Fakultät für Informatik
Technische Universität München

Lehrstuhl für Wirtschaftsinformatik und Entscheidungstheorie

Core Pricing and Spectrum Auction Design

Andor Roger Goetzendorff

Vollständiger Abdruck der von der Fakultät für Informatik der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwisschenschaften (Dr.rer.nat.)

genehmigten Dissertation.

Vorsitzender:
Prof. Dr.-Ing. Georg Carle

Prüfende der Dissertation:
1. Prof. Dr. Martin Bichler
2. Prof. Dr. Florian Matthes

Die Dissertation wurde am 08.11.2017 bei der Technischen Universität München eingereicht und durch die Fakultät für Informatik am 15.05.2018 angenommen.

Abstract

Auctions are useful in settings when fixed price sourcing is not optimal because of an information asymmetry between buyers and sellers in a given market. Combinatorial auctions can increase social welfare by allowing bidders to express substitute or synergistic values for combinations of items. Considering synergistic valuations turns the allocation problem into a computationally hard optimization problem that generally cannot be approximated to a constant factor in polynomial time. We analyze specific scenarios, where the optimality cannot be guaranteed, either because of the computational hardness of the underlying winner determination and pricing problem, or because inefficiencies that were due to the auction formats used. We concentrate our research on complex markets in general and focus on spectrum auctions specifically because of its pivotal position in combinatorial auction design.

We first concentrate on the computation hardness of large markets. Here we introduce an auction design framework for large markets with hundreds of items and complex bidder preferences. Such markets typically lead to computationally hard allocation problems. Our new framework consists of compact bid languages for sealed-bid auctions and methods to compute second-price rules such as the Vickrey-Clarke-Groves or bidder-optimal, core-selecting payment rules when the optimality of the allocation problem cannot be guaranteed. For realistic instances of the respective winner determination problems found in the analyzed markets, very good solutions with a small integrality gap can be found quickly. Closing the integrality gap to find marginally better solutions or prove optimality can take a prohibitively large amount of time, however. Our subsequent adaptation of a constraint-generation technique for the computation of bidder-optimal core payments to this environment is a practically viable paradigm by which core-selecting auction designs can be applied to large markets with potentially hundreds of items. We complement our computational experiments in the context of TV ad markets with additional results for volume discount auctions in procurement in order to illustrate the applicability of the approach in different types of large markets.

The spectrum market is another prominent example where complex bidder valuations can be found. Ascending auction designs such as the Simultaneous Multiple Round Auction (SMRA) and the single-stage or two-stage Combinatorial Clock Auction (CCA), the two most dominant spectrum auction formats, can be seen as simple heuristic algorithms to solve this problem. Such heuristics do not necessarily compute the optimal solution, even if bidders are truthful. We study the average efficiency loss that can be attributed to the simplicity of the auction algorithm with different levels of synergies. Our simulations are based on realistic instances of bidder valuations we inferred from bid data from the 2014

Canadian 700MHz auction. We compare the SMRA and CCA under different synergy conditions when bidders maximize payoff in each round. With "linear" synergies, a bidder's marginal value for a license grows linearly with the total number of licenses won, while with the "extreme national" synergies, this marginal value is independent of the number of licenses won, unless the bidder wins all licenses in a national package. We find that with the extreme national synergy model, the CCA is indeed more efficient than SMRA. However, for the more realistic case of linear synergies, SMRA outperforms various versions of CCA that have been implemented in the field including the one used in the Canadian 700MHz auction. Overall, the efficiency loss of all ascending auction algorithms is small even with high synergies, which is remarkable given the simplicity of the algorithms.

Optimal bidder behavior cannot be guaranteed, however: Although the CCA draws on a number of elegant ideas inspired by economic theory which provide incentives to bid truthfully, this is not always the case. Bidders might not respond to these incentives due to strategic reasons or practical limitations. We introduce metrics based on Afriat's Efficiency Index to analyze straightforward bidding and report on empirical data from the lab and from the field in the British 4G auction in 2013 and the Canadian 700 MHz auction in 2014, where the bids were made public. The data provides evidence that bidders deviate significantly from straightforward bidding in the clock phase, which can restrict the bids they can submit in the supplementary phase. We show that such restrictions can have a significant negative impact on efficiency and revenue.

Contents

List of Figures

List of Tables

1 Introduction

Auctions have been used for at least 2 millenia to exchange goods or services between parties. The usage of auctions as a market mechanism has been steadily increasing, with numerous examples: In business to business markets, from perishable goods, such as fresh flowers, to commodities and public utilities, auctions are now in use in more and more markets (Krishna, 2009). Also private sellers started to adopt auctions since the start of the 21st century using platforms such as ebay.com as an alternative to fixed-price sourcing.

But auctions are used for more than just selling single items: Electronic markets often allow market participants to express rich information about their preferences for the offered goods or services beyond single item valuations. More complex valuations such as complementarities between goods or stepwise discounts for large volumes of items are readily found in multiple markets. The possibility to allow the direct expression of these valuations can lead to an increase in allocative efficiency and social welfare, a target often desired by regulators and market designers.

Combinatorial auctions are a pivotal example of such smart markets (Cramton et al., 2006). By allowing bidders to bid on packages (i.e., combinations of items) instead of single items only, combinatorial auctions protect bidders against the well-known "exposure problem". This effect is present if items were to be auctioned separately and a bidder is unable to win all items of a highly valued combination (causing him to pay more for the subset than the subset's worth to him), or when a bidder wins too many items which he considers to be substitutes.

Markets successfully using combinatorial auctions include logistics (Caplice, 2007), energy exchanges (Meeus et al., 2009) and industrial procurement (Bichler et al., 2006). For several years regulators started to sell spectrum auctions around the world using combinatorial auction designs (Cramton, 2013). In all of these cases the market design has a profound effect on the bidder behavior and efficiency. Being inherently multi-disciplinary, the design of such markets created an interest in academia, with significant contributions coming from researchers from various fields including economics, operations research and computer science. Cramton et al. (2006) provides a first overview of the different application areas.

Auction theory provides a basic framework to think about strategies and efficiency of auctions, but with the increase in computing power and improvements in optimization algorithms, new problems arose which are complementary to those discussed in microeconomics. As an example, the required amount of communication needed for an efficient solution and the computational complexity of markets with multiple items has been a topic of interest in operations research and computer science (Lehmann et al., 2006; Nisan and Blumrosen, 2007). In the information systems literature, contributions regarding the decision support and information feedback (Adomavicius and Gupta, 2005), the analysis of bidder behavior (Scheffel et al., 2011), as well as the design for specific domains (Bapna et al., 2007; Guo et al., 2007) can be found. As we will see in the following chapters, all of these issues will also be present in our main contributions.

Spectrum Auction design is one of the most challenging and visible application domains for combinatorial auctions. It is often seen as a pivotal example for the design of multi-object markets and successful auction designs are likely role-models for other markets in areas such as procurement and logistics.

Radio spectrum is a key resource in the digital economy. With numerous applications important for society and its dramatic increase in demand in the last years due to the rapid increase of wireless devices, the US Federal Communication Commission (FCC) decided in 1994 to adopt a market-based approach to assign spectrum instead of previously established methods such as "beauty contests", in which the candidates' offers are matched against a weighted list of criteria. The Simultaneous Multiple Round Auction (SMRA) was the first auction used by the FCC and has since then been adopted successfully by regulators in multiple national markets around the world, generating hundreds of billions of dollars worldwide.

Despite this success, the SMRA has also led to a number of strategic challenges for bidders. Telecom operators often have preferences for certain packages of licenses, but are unable to express them directly in the SMRA. This leads to the aforementioned exposure problems: if bidders who compete aggressively for a certain package risk ending up with only a subset, they risk paying more more than what this subset is worth to them. This adds a possibly unwanted strategic complexity to the auction and is a source of inefficiency in the SMRA. To mitigate this, regulators started incorporating combinatorial auctions also for selling spectrum. In theory, the Vickrey-Clarke-Groves (VCG) auction is the only strategy-proof (i.e., truthful) and efficient auction, but for multiple practical reasons it has rarely been used so far (see Section 2.3). This brought up several other combinatorial auction formats, some of them specifically designed for the spectrum market. In 2008, the FCC decided to adopt an extension of the SMRA format (see Goeree and Holt (2010)) that allows the auctioneer to specify packages on which bidders can bid on. Meanwhile, Ofcom, UK's telecom regulator, adopted the Combinatorial Clock Auction (CCA) in it's L-Band auction (Cramton, 2008). The format rapidly gained popularity and variations of it have been used by regulators around the world including Austria, Australia, Canada, Denmark, Ireland, Mexico, the Netherlands, Switzerland and the UK (Mochon and Saez, 2017).

Efficiency, revenue, and strategic simplicity for bidders are typical design goals that a regulator has in mind. Choosing the best auction design is not always trivial: In auctions with a purely additive bidding language, bids for each package are just the sum of the bids on the individual items. This allows for a potentially higher efficiency in larger markets than some combinatorial auction designs due to its simplicity, but, as mentioned earlier, bidders risk of winning only parts of a bundle of interest and therefore having to pay more than this subset of items is worth to them. This can introduce an unwanted level of strategic complexity and potential inefficiency into the auction. If on the other hand bidders have to explicitly enumerate all packages they are interested in, bidders can typically only specify a small proportion of these bids in larger auctions. Missing package bids are usually treated as if a bidder had no value for the package; an unlikely scenario. As we will see in Paper B and C, lab experiments, field data and simulations have shown that this “missing bids problem” can lead to substantial efficiency losses.

Combinatorial auction design therefore faces a natural trade-off between the possibly higher efficiency due to the option of directly expressing preferences on packages, and the efficiency losses due to the introduced complexity. This observation has caused a debate on the design of spectrum auctions, but the discussion goes beyond this application and begs the question how large markets with many items can be designed so that bidders are incentivized and able to express their preferences truthfully while allowing auctioneers to achieve allocations with high efficiency.

1.1 Outline and Contributions

The debate is still not solved in many cases: In spectrum auctions, regulators still tend to choose from a variety of possible auction formats. None of these formats has emerged as a clear winner, as all used formats have different advantages and drawbacks. In the paper *Synergistic Valuations and Efficiency in Spectrum Auctions* we try to quantify the differences in efficiency and revenue of these formats by simulating truthful and rational bidders. Instead of creating synthetic scenarios, we based both our ruleset and bidder valuations on the Canadian 700 Mhz Auction of 2014, including all the intricacies that are often ignored in simulations or laboratory experiments due to the added complexity. As of today, no other simulation framework is known to us that allows us to generate results with the desired fidelity. The paper was created in collaboration with Martin Bichler and Jakob K. Goeree and was published in “Telecommunication Policy”.

But what if bidders do not behave rationally? This is analyzed in *(Un)expected Bidder Behavior in Spectrum Auctions*: Using the bidding data from laboratory experiments and two recent spectrum auctions in the field, we demonstrate that bidders do indeed not always display rational behavior in the Combinatorial Clock Auction. This often has unexpected consequences, as even small deviations from truthful behavior can cause severe limitations to the bidder's bidding possibilities due to the used activity rules. The impact this has on the efficiency and revenue was analyzed in conjunction to the "missing bids" phenomenon. The paper is based on the work done in collaboration with Martin Bichler and Christian Kroemer and was published in "Group Decision and Negotiation".

In our third contribution we concentrate on very large markets with potentially hundreds of items and complex bidder preferences. In such a setting, solving the allocation problem to optimality cannot be guaranteed, although near-optimal solutions can be found relatively quickly. As several payment schemes, such as the VCG auction (see Section 2.3) depend on strict optimality, using these payment schemes is not always possible. In the paper *Compact Bid Languages and Core Pricing in Large Multi-item Auctions*, we develop an auction framework for such markets, allowing us to apply these payment schemes even under near-optimality. To demonstrate the efficacy, we introduce a compact bidding language for TV advertising markets and complement our computational experiments with additional results for volume discount auctions to illustrate the applicability of the approach in different types of large markets. The paper was done in collaboration with Martin Bichler, Pasha Shabalin and Robert W. Day and was published in "Management Science".

The dissertation is structured as follows:

- Chapter 1 introduces the reader to auctions, combinatorial auctions, and the spectrum auction market and its challenges.
- Chapter 2 gives a concise overview over the main concepts of auction theory.
- Paper A depicts *Compact Bid Languages and Core Pricing in Large Multi-item Auctions*.

- Paper B compares the *Efficiency of Spectrum Auctions: SMRA versus CCA*.
- Paper C analyzes the *(Un)expected Bidder Behavior in Spectrum Auctions*.
- Chapter 3 concludes by discussing the selected methods and by giving an outlook into further research possibilities.

If the reader is already familiar with the basics of auction theory and spectrum auctions, the paper A, B and C can be read on their own. Concluding remarks and appendices can be found at the end of these papers, respectively.

2 Theoretical Framework

Before going deeper into the main contributions of this dissertation, a short overview of basic notation and concepts typically used in auction theory is presented in this chapter. The topics here are presented on their own, as they are of interest to the reader for all of the following work. For a more detailed review please refer to Nisan and Blumrosen (2007) and Cramton et al. (2006), the main references for this chapter.

2.1 Auctions

The term “auction” stands for any mechanism that collects bids of potential buyers and where the outcome is determined solely by the information provided by these bidders. While a wide variety of mechanisms fall under this category, in its most basic form, an auction can be broadly characterized by its running procedure, its pricing rule and its direction. An auction can last only one or multiple rounds (a sealed-bid or iterative auction, respectively). In standard auctions, the prices winning bidders have to pay are typically at most the value of their winning bid (a first-price auction) and at least the marginal harm they cause to the other participants with their bid (a second-price auction). This coincides with the bid of the second-highest bidder in single-item single-unit auctions. Finally, an auction can be either a forward auction, a reverse auction, or an exchange. A forward auction is the auction you would expect when going to an auction house. Here the auctioneer is the seller. A reverse auction is typically used by an auctioneer to procure services or goods from suppliers instead. In an exchange both buyers and sellers participate in the price finding process of the auction to reach an equilibrium.

Using an auction makes sense if it is difficult for the auctioneer to determine an optimal price point for the goods or services to sell in order achieve an optimal allocation, but the participating bidders already have or can form an opinion about the value of the auctioned goods. If each bidder knows only his value of the items to be sold, and each of the bidders' types is independent from one another, we refer to this as *independent private values*. Formally, we can define the valuation function of a bidder j as

$$v_j : X \to \mathbb{R}_{\geq 0} \tag{2.1}$$

where $X = (S_1, \ldots, S_n)$ stands for an outcome chosen by the auctioneer. In this case, $v_j(S)$ denotes the value bidder j has for the allocation or bundle S[1]. Throughout this dissertation we will adopt the standard assumption of quasilinear utility functions, meaning that the utility of bidder j is

$$\pi_j(S) = v_j(S) - p_j(S) \tag{2.2}$$

where the term p_j represents the monetary amount the bidder has to pay for bundle S.

We further adopt the assumption that all bidder valuations satisfy *free disposal*

$$v_j(S) \geq v_j(T) \;\forall\; T \subseteq S \tag{2.3}$$

i.e., receiving an additional item never reduces the valuation of bidder j.

A combinatorial auction allows a bidder to communicate valuations for whole bundles, compared to single items. This can cause a higher efficiency if the bidder's valuation function is nonadditive: Objects are *complements* if the marginal value of receiving an additional item i is larger if the set of objects already received is bigger, i.e. we have a

[1] Note that, with independent private values, the bidder's valuation only depends on the bundle he receives.

superadditive valuation function:

$$v_j(T_s \cup \{i\}) - v_j(T_s) \leq v_j(T_b \cup \{i\}) - v_j(T_b) \; \forall \; T_s \subseteq S, T_b \subseteq S, T_s \subseteq T_b \tag{2.4}$$

Likewise, we say objects are *substitutes* if the marginal value of receiving i is less in this setting. In this case, we have subadditive valuations.

By introducing a binary decision variable $x_j(S)$ for each set S a bidder j demands, a *winner determination problem* (WD) can be formulated as an integer linear program (ILP). This ILP will form the basis of many of the allocation problems presented in this dissertation. Note that in the current formulation, a bidder j can win at most one bundle $S \subseteq I$. This restriction can be easily removed by omitting Equation (2.7).

$$\begin{aligned}
\max \quad & \sum_{S \subseteq I} \sum_{j \in J} x_j(S) v_j(S) & & & (2.5)\\
\text{s.t.} \quad & \sum_{j \in J} \sum_{S: i \in S} x_j(S) \leq & q_i \quad & \forall i \in I & (2.6)\\
& \sum_{S \subseteq I} x_j(S) \leq & 1 \quad & \forall j \in J & (2.7)\\
& x_i(S) \in & \{0,1\} \quad & \forall S \subseteq I, \forall j \in J &
\end{aligned}$$

A bid submitted to an auction is a representations of a bidder's communicated valuation $v_j(S)$ for a specific bundle S. While solving the WD to optimality is an NP-hard problem (Sandholm, 2002), finding a solution is possible even for instances with hundreds of bids (Ibid.).

In order to quantify the 'goodness' of an allocation generated by an auction, two metrics will be used throughout the dissertation. The allocative efficiency (or simply efficiency) of an auction is measured by comparing the outcome of the algorithm against the social welfare generated by the optimal allocation X^*, i.e.

$$E(X) = \frac{\sum_{j \in J} v_j(S_j)}{\sum_{j \in J} v_j(S_j^*)} \tag{2.8}$$

We also measure the revenue distribution $R(X)$, which compares the auctioneers revenue (i.e., the sum of the bidders' prices p_j) against the optimally achievable surplus.

$$R(X) = \frac{\sum_{j \in J} p_j(S_j)}{\sum_{j \in J} v_j(S_j^*)} \tag{2.9}$$

To make the WD usable in practice, two issues have to be addressed before putting it to use: With an exponential number of possible bids (there are 2^m possible bundles for m items), it would be impractical to ask all bidders to submit a bid for each combination. This will be addressed in the next session. Another point of interest is the strategy-proofness of the proposed mechanism: As of now, it is unclear how much a winning bidder should pay for his winning bundle. This will be addressed in Section 2.3.

2.2 Bid Languages

Bid languages are used to allow participants to limit the amount of information they have to communicate, while still being expressive enough. Designing a fitting language is a nontrivial task for all but smaller markets, and often makes use of domain knowledge to reduce the amount of information needed (see Paper A). The two most basic operators used to form a collection of bids are OR and XOR bids. In a bidder's submitted XOR bid collection, only at most one of the bids in such a collection can be won by the bidder. In an OR bid collection, on the other hand, it is possible for the bidder to win more than one bid (Nisan and Blumrosen, 2007) of such a collection. Both operators have very different levels of compactness and expressiveness: To allow a bidder to express any arbitrary valuation function through atomic bids, an XOR language is needed. In this case, the bidder has to submit up to $2^m - 1$ bundle bids (for m items). An OR language is much more compact, as the bidder only has to submit up to m bids. In this case however, the bidder only has the possibility to express superadditive valuations.

As an example, take the bids $((\{a, b\}, 3)\ XOR\ (\{c\}, 5))$ and $((\{a, b\}, 3)\ OR\ (\{c\}, 5))$. Each of them has a value of 0 for $\{b\}$ and values the bundle $\{a, b\}$ at 3. The difference shows up in the bundle $\{a, b, c\}$. Here, the XOR bid has a value of 5 for it, whereas the OR bid values the bundle at 8.[2]

Often, a hierarchy of super- and subadditive valuations is specific to the market. If this is known to the auctioneer, a specific structure of OR and XOR bids or the introduction of dummy bids can help to achieve a good balance between compactness and expressiveness (Bichler et al., 2011; Nisan and Blumrosen, 2007). In more complex settings and larger markets, domain specific compact bid languages such as the ones found in Boutilier and Hoos (2001) and Paper A might be needed to lower the communication complexity.

Another way to lower the needed communication between bidders and auctioneer will be introduced in Section 2.5, where we will be discussing iterative combinatorial auctions. Before depicting this concept, we will introduce two different payment schemes that are essential if an auctioneer is interested in inducing truthful behavior or a stable outcome in sealed-bid auctions.

2.3 The VCG mechanism

Efficiency, revenue, and strategic simplicity for bidders are typical design goals that an auctioneer has in mind. The Vickrey-Clarke-Groves (VCG) auction is the *only* strategy-proof and efficient auction (Krishna, 2009). Even though it is seldom used in the field for several reasons (Rothkopf, 2007), it often serves as a baseline to compare other auction formats against, or as a starting point for further price adjustments.

In a VCG auction, bidders report their valuations for all relevant packages through the used bidding language. The goods are then assigned optimally, maximizing the total value. The requested payment for a winning bid is not necessarily equal to the clearing price, however. Instead, it is equal to the marginal harm this bid inflicts to the other bidders (i.e., the combined opportunity cost). For a set of bidders J and a winner

[2] Note that, if a bidder submitted an OR bid for two packages S_1 and S_2 with $S_1 \cap S_2 \neq \emptyset$, only at most one of those bids would win.

determination algorithm $WD(J)$, the price a winning bidder j has to pay for his winning bid amount $b_j(S_j)$ can therefore be formulated as:

$$p_j^{\mathrm{VCG}} = b_j(S_j) - (WD(J) - WD(J \setminus \{j\})) \tag{2.10}$$

The VCG auction is the only strategy-proof and efficient auction, but it is rarely used for practical reasons. If there is at least one bidder who does not have substitute preferences, the revenue a VCG auction generates is no longer monotonous regarding the set of bidders and the amounts bid, and the mechanism is no longer group-strategy-proof. But even if bidders bid truthfully, the revenue of a VCG auction can be as low as 0 even for high bids and ample competition.

For example, Ausubel and Milgrom (2006) provide a classic setup where VCG payments total zero for two bidders, despite a competitor's bid to pay the seller a large amount for their combined winnings, and show that these payments of zero can be achieved through group manipulation or the use of false-name (i.e., shill) bids. This occurs when the first two bidders bid M for their disjoint respective bundles of interest, with a losing third competitor offering exactly M for the union of these bundles.

This makes the auction unattractive, even if the auctioneer's primary interest is social welfare maximization and not revenue maximization. In a public setting, a scenario as the one just described would be most likely unacceptable for the general population (as the revenue would be zero) and the losing bidder (as he would be willing to pay up to M for the bundle).

2.4 The Core

The constructed example mentioned in Section 2.3 is quite extreme: The seller earns zero even if the goods in the auctions were in demand. The result does seem neither fair nor stable: Specifically, compared to the computed outcome, the seller and bidder 3 would be better off settling for any price $\leq M$. The concept introduced in this section allows a market designer to determine if prices are 'too low' or 'too high' by looking at the stability of the outcome. For this we draw upon a concept used in cooperative game theory.

Auctions can be seen as cooperative games with transferable utility. Whenever a subcoalition of the grand coalition (i.e., the auctioneer and all bidders) can form an outcome that increases the payoff of all participants of this subcoalition, we say the outcome is not in the *core*, and that such a coalition is *blocking* the proposed outcome. Hence, in the example above, the coalition {seller, bidder 3} is a blocking coalition.

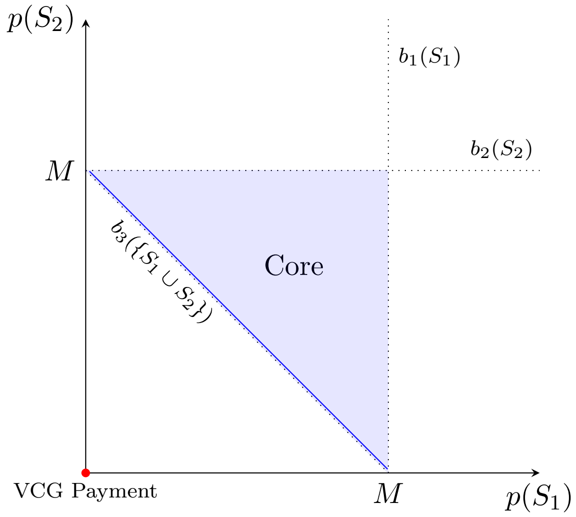

Figure 2.1: VCG Payments and the Core

We base our more formal definition of the core on Milgrom (2004). We define the set of players J_0 as the set of all bidders J in addition to the auctioneer (with index 0). The seller's utility π_0 is equal to the revenue the auction generates, i.e. $\pi_0 = \sum_{j \in J} p_j(S_j)$. We further define a *coalitional value* function $w(C)$ for a coalition C as the maximum sum of the participants' bundles S_j^x that they receive in the alternative S out of the set of alternatives $\mathcal{X}$.

$$w(C) = \begin{cases} \max_{X \in \mathcal{X}} \left\{ \sum_{j \in C \setminus \{0\}} v_j(S_j^x) \right\} & \text{if} \quad 0 \in C \\ 0 & \text{if} \quad 0 \notin C \end{cases} \tag{2.11}$$

The set of core payoffs can then be defined as:

$$Core(C, w) = \left\{ \pi : \sum_{j \in C} \pi_j = w(C), (\forall D \subseteq C) \; w(D) \leq \sum_{j \in D} \pi_j \right\} \tag{2.12}$$

I.e., if a set of payoffs is not in the core, a (sub-)coalition of C has a higher coalitional value than the summed up payoffs. For the aforementioned example, a simple computation shows that the payoff vector $\pi \notin Core(J_0, w)$. Figure 2.1 shows the same setting graphically.

Core payments are often seen as more practical compared to the VCG auction because of their stability, but as we can see in Figure 2.1, the core is often not unique. A number of authors provide valuable insights into the possibilities to manipulate core outcomes and the perceived fairness of outcomes (Day and Raghavan, 2007; Day and Cramton, 2012; Lubin et al., 2015). Bidder-Pareto-Optimal core (BPOC) outcomes, i.e., a core outcome for which there is no other core outcome that is weakly preferred by every player[3], minimize the incentives to bid strategically among all core outcomes, and are used in several markets such as spectrum auctions. Paper A will extend one of these payment schemes in order to be able to compute core-like payments also for problems whose size allows only the computation of near-optimal solutions.

2.5 Iterative Combinatorial Auction Formats

The main advantage of combinatorial auctions is the ability for bidders to express preferences on packages, but this makes it also extremely difficult for bidders, as they have to

[3]In Figure 2.1, the thickened edge of the core polyhedron marks all BPOC outcomes in this example.

provide a valuation function over the space of all possible bundles of items. (Compact) bidding languages (see Section 2.2) can be a remedy for this, but they do not address the issue of the cost of precisely computing the parameters of the model in question to fit their valuation function. Another possibility to tackle the inherent communication complexity of combinatorial auctions is by introducing rounds to an auction, and querying the bidders each round. The introduction of iterations makes the bidders' strategy space richer (Sandholm and Boutilier, 2006), which can introduce unwanted strategic complexity. At the same time, intelligent querying can reduce the needed communication between bidders and auctioneers (Ibid.). Another reason to opt for an iterative auction format are privacy concerns: Here bidders only have to reveal partial and indirect information about their (private) valuation function. Furthermore, while not yet proven for combinatorial auctions, the dynamic nature of iterative auctions is known to enhance revenue and efficiency in single item auctions with interdependent valuations (Milgrom and Weber, 1982). Before introducing the main families of iterative combinatorial auctions found to date, we outline the preliminary concepts needed to reason about this type of auction. We base our description on Krishna (2009) and Cramton et al. (2006).

A first idea about how prices are formulated can be readily seen by extending the English (open outcry) auction to multiple items. This concept is a stark contrast compared to the result of the sealed-bid payment rules such as VCG or BPOC auctions, introduced in Sections 2.3 and 2.4, where the resulting prices are typically not comparable between bidders. In order to reason more about the payments bidders are required to make, we can categorize an auction's (ask) prices as follows:

$$\begin{aligned} &\text{linear prices} && p(i) \geq 0 && \forall i \in I && (2.13)\\ &\text{non-linear prices} && p(S) \geq 0 && \forall S \subseteq I && (2.14)\\ &\text{personalized, non-linear prices} && p_j(S) \geq 0 && \forall S \subseteq I, j \in J && (2.15) \end{aligned}$$

	A	B	AB
b_1	a	b	$a+b+c$
b_2	$a+d$	$b+d$	$a+b+d$

Table 2.1: No equilibrium with linear prices

Each introduced price level allows the market designer to discriminate on a specific additional dimension. With linear prices (2.13), bundle prices are simply the sum of the item prices, i.e. $p(S) = \sum_{i \in S} p_i$. For non-linear prices (2.14), the statement $p(S) \neq p(S_1) + p(S_2)$ is legal even if $S = S_1 \cup S_2$ and $S_1 \cap S_2 = \emptyset$. Finally, prices can also be discriminatory towards bidders, i.e. it can be that, under eq. (2.15), $p_j(S) \neq p_{j'}(S)$ for $j \neq j'$, in addition to all the characteristics of (2.14).

Linear prices are certainly the easiest concept to grasp for participants in an auction, but for general valuations, defining an allocation function that is guaranteed to result in an equilibrium, a necessary condition for an efficient iterative auction, is not possible with linear prices. As an example, Milgrom (2000) showed that with at least three bidders and at least one non-substitutes valuation, i.e., this bidder regards at least one package of items as complements, it is not guaranteed that an equilibrium with linear prices (a Walrasian equilibrium) exists. An intuition can be given by a two-item, two-bidder example, based on (Ibid.).

We characterize this example by defining two bidders, b_1 and b_2, and two items, A and B. The valuations for the individual items and the package AB can be found in Table 2.1. Here, the valuations $a, \ldots, d \geq 0$ and $\frac{c}{2} < d < c$, i.e., A and B are substitutes for b_2 but complements b_1. In this case, a social welfare maximizing function such as WD (see eq. (2.5)) can only allocate AB to b_1. In order to achieve an equilibrium with linear prices, $p(A)$ and $p(B)$ must be so that b_2 is not demanding these items anymore, i.e., $p(A) \geq a + d, p(B) \geq b + d$. At these prices, b_1 is unwilling to acquire them, however: No equilibrium with linear prices can be formed.

The concept of Competitive Equilibrium extends the notion of a Walrasian equilibrium to a combinatorial auction. A competitive equilibrium (p, X) is such that an allocation

$X = (S_1, \ldots, S_n)$ maximizes the payoff of every bidder and the seller given prices p, i.e. prices p and allocation X are in competitive equilibrium if:

$$\pi_j(S_j, p) = \max_{S \subseteq I}(v_j(S) - p_j(S), 0) \qquad \forall\, j \in J \tag{2.16}$$

$$\pi_0(X, p) = \max_{S \in \Gamma} \sum_{j \in J} p_j(S_j) \tag{2.17}$$

where $\Gamma \in \mathcal{X}$ denotes the set of all feasible allocations.

We note that the allocation X is supported in competitive equilibrium if and only if X is an efficient allocation (Parkes and Ungar, 2000). Furthermore, all core outcomes (see Section 2.4) can be priced with a set of prices $p_j(S_j)$[4], and all competitive equilibrium prices have corresponding core payoffs (Bikhchandani and Ostroy, 2002).

As we will see, many iterative combinatorial auctions are designed to converge to minimal competitive equilibrium prices, i.e., a set of prices which minimizes the auctioneer's total revenue $\pi_0(X, p)$ on the efficient allocation X^* across all competitive equilibrium prices. Also, minimal competitive equilibrium prices have a corresponding bidder Pareto-optimal core payoff vector, see Section 2.4.

The design space for iterative auctions is quite large, with possibilities to decide on different dimensions such as timing issues, information feedback, bidding rules, termination conditions and the use of proxy agents (Parkes and Ungar, 2000). The biggest challenge in iterative combinatorial auction design is to support focused bidding by reducing the communication needed between bidders and auctioneer, without allowing new strategic behavior to compromise efficiency.

In the following, we will introduce a short number of iterative auction formats used to sell spectrum worldwide, with a focus on price-based and discrete-round-based iterative combinatorial auctions with a rolling termination rule. In addition to the bidding language (see Section 2.2), and ask price formation, the following additional criteria have to be specified:

[4] A trivial example would be $p_j(S_j) = v_j(S_j)$.

- The information feedback bidders receive in each round. This includes for example the provisional allocation and price feedback, and can have an important strategic impact on bidders.

- The (inter-round) bidding rules. These typically are used to put an incentive on straightforward bidding, or to speed up the convergence of an auction toward an equilibrium. They can range from simple monotonicity rules, i.e., bidders can only express downward sloping demand curves over a series of rounds, as prices increase, to revealed preference constraints, which are thoroughly analyzed in Paper C.

As we will see in Paper B, details such as the ones listed can cause a significant change on the perceived fairness and efficiency of an auction. They are however highly market specific and therefore omitted in the following survey.

Although not a combinatorial auction format, simultaneous ascending auctions such have been used successfully for numerous applications, from the commissioning of divisible goods like electricity to spectrum auctions. One prominent example here is the *Simultaneous Multiround Auction* (SMRA), a simultaneous auction format used by regulators to auction off spectrum. The widespread adoption of this auction format in spectrum auctions was one of the reasons we included the SMRA in our analysis in Paper B. The SMRA is a generalization of the English open-outcry auction when taking into consideration many goods. Here, all items are offered at the same time, each item having a prices associated with it. Bidders can then submit bids for each of these items, surpassing the existing price, with the highest bidder provisionally winning these items. This continues until no bidder is willing to submit any new bids. The auction terminates by asking the bidders to pay the price they bid on each of their provisionally winning items. The auction suffers from the aforementioned exposure problem, but the procedure and termination rule is easily understood by bidders, which might offset the inherent inefficiency of this auction format (see Paper B).

Hierarchical Package Bidding (HPB) is an extension to the SMRA format by Goeree and Holt (2010). Instead of bidding on items only, bidders are allowed to bid on a series of packages that were defined by the auctioneer. The hierarchical, tree-like structure of the defined packages allows for an optimization in polynomial time, while still allowing bidders to potentially overcome the exposure problem associated with a simultaneous ascending auction. Defining a suitable hierarchy of packages is challenging however, especially when the participating bidders would like to express heterogeneous complementarities or if the complementarities are unknown to the auctioneer.

While all of the previous auction formats raise prices in their iterative phase by taking into account the overdemand and increasing the item's ask price by a specified amount (typically between 5% and 10%), Primal-Dual Auction Formats such as RAD and *ALPS(m)* (see Bichler et al. (2009)) take a different approach: After each round, a mixed integer linear program is solved with the aim to generate a set of linear prices by taking into account the current allocation computed by a WD (see eq. (2.5)). While these pseudo-dual linear price iterative combinatorial auctions have performed well in laboratory experiments and simulations, they are not guaranteed to converge to a competitive equilibrium and don't guarantee monotonicity on the generated linear prices throughout the rounds (Bichler et al., 2009; Parkes and Ungar, 2000).

Similar to the SMRA, the *Clock Auction* format is, in its most basic form, an extension of the Japanese Auction to multiple items. Here, each item is associated to "price clock", which determines the price of the respective item. Bidders can then in each round decide which items they would like to acquire at the current prices and express this demand set to the auctioneer. Contrary to the SMRA format, bidders are not able to specify a custom price for each of the demanded items, i.e., jump bidding is not allowed. After collecting the demands of each item, the auctioneer increases the prices of all items with overdemand, followed by another round, where all bidders have, again, to submit their demand sets at the current prices. The auction is finished when no item is overdemanded. Being a simultaneous ascending auction, the clock auction also bears the potential risk of the exposure problem for bidders.

The Clock auction format is an important precursor for two other combinatorial auction formats prominent in spectrum auctions, the single stage and two stage *Combinatorial Clock Auction* (SCCA and CCA, resp.). The SCCA allows bidders to express their demand on packages of items instead of items alone, eliminating the exposure problem if a fully expressive language is used[5]. The auction formats combines the intuitivity of the clock auction format with monotonically increasing linear prices and a simple termination rule: In case of overdemand, prices are increased as in the clock auction. In a situation with oversupply, the auctioneer solves a winner determination problem (see eq. (2.5)) to determine, whether there is an excess demand by also including bids from previous rounds. If this is not the case, or if supply equals demand, the auction terminates, and bidders pay the respective sum of ask prices of their winning packages.

The (two stage) CCA extends the clock auction format was developed to overcome the shortcomings of a linear price combinatorial auction. Here a supplementary, sealed-bid phase is added to the initial clock auction. This supplementary phase allows bidders to specify package bids with a fully expressive bid language, while the first clock phase is used to determine the bid price range. Prices are then determined by using a bidder-Pareto optimal core selecting rule (see also Section 2.4 and Paper A).

[5] In its original design, Porter et al. (2003) use an OR-bidding language, but already mentions mutually exclusive bids. Extensions with a more expressive bidding language were already used in the field, for example in the Danish Spectrum Auction of 2016.

Compact Bid Languages and Core Pricing in Large Multi-item Auctions

Peer-reviewed Journal Paper

Title: Compact Bid Languages and Core Pricing in Large Multi-item Auctions

Authors: A. Goetzendorff, M. Bichler, R. Day, P. Shabalin

In: Management Science 61.7 (2015): 1684-1703

Abstract: We introduce an auction design framework for large markets with hundreds of items and complex bidder preferences. Such markets typically lead to computationally hard allocation problems. Our new framework consists of compact bid languages for sealed-bid auctions and methods to compute second-price rules such as the Vickrey–Clarke–Groves or bidder-optimal, core-selecting payment rules when the optimality of the allocation problem cannot be guaranteed. To demonstrate the efficacy of the approach for a specific, complex market, we introduce a compact bidding language for TV advertising markets and investigate the resulting winner-determination problem and the computation of core payments. For realistic instances of the respective winner-determination problems, very good solutions with a small integrality gap can be found quickly, although closing the integrality gap to find marginally better solutions or prove optimality can take a prohibitively large amount of time. Our subsequent adaptation of a constraint-generation technique for the computation of bidder-optimal core payments to this environment is a practically viable paradigm by which core-selecting auction designs can be applied to large markets with potentially hundreds of items. Such auction designs allow bidders to express their preferences with a low number of parameters, while at the same time providing incentives for truthful bidding. We complement our computational experiments in the context of TV advertising markets with additional results for volume discount auctions in procurement to illustrate the applicability of the approach in different types of large markets.

Contribution of thesis author: Methodology, results (except for proof of Theorem 2), implementation, mathematical model, presentation, project and paper management

MANAGEMENT SCIENCE
Vol. 61, No. 7, July 2015, pp. 1684–1703
ISSN 0025-1909 (print) | ISSN 1526-5501 (online)
http://dx.doi.org/10.1287/mnsc.2014.2076

Compact Bid Languages and Core Pricing in Large Multi-item Auctions

Andor Goetzendorff, Martin Bichler, Pasha Shabalin
Department of Informatics, Technische Universität München, 85748 Garching, Germany
{goetzend@in.tum.de, bichler@in.tum.de, shabalin@mytum.de}
Robert W. Day
Operations and Information Management, University of Connecticut, Storrs, Connecticut 06269,
bob.day@business.uconn.edu

We introduce an auction design framework for large markets with hundreds of items and complex bidder preferences. Such markets typically lead to computationally hard allocation problems. Our new framework consists of compact bid languages for sealed-bid auctions and methods to compute second-price rules such as the Vickrey–Clarke–Groves or bidder-optimal, core-selecting payment rules when the optimality of the allocation problem cannot be guaranteed. To demonstrate the efficacy of the approach for a specific, complex market, we introduce a compact bidding language for TV advertising markets and investigate the resulting winner-determination problem and the computation of core payments. For realistic instances of the respective winner-determination problems, very good solutions with a small integrality gap can be found quickly, although closing the integrality gap to find marginally better solutions or prove optimality can take a prohibitively large amount of time. Our subsequent adaptation of a constraint-generation technique for the computation of bidder-optimal core payments to this environment is a practically viable paradigm by which core-selecting auction designs can be applied to large markets with potentially hundreds of items. Such auction designs allow bidders to express their preferences with a low number of parameters, while at the same time providing incentives for truthful bidding. We complement our computational experiments in the context of TV advertising markets with additional results for volume discount auctions in procurement to illustrate the applicability of the approach in different types of large markets.

Data, as supplemental material, are available at http://dx.doi.org/10.1287/mnsc.2014.2076.

Keywords: TV ads; core-selecting auction; market design
History: Received February 9, 2013; accepted June 30, 2014, by Lorin Hitt, information systems. Published online in *Articles in Advance* March 16, 2015.

1. Introduction

Electronic markets allow market participants to express rich information about their preferences for different goods or services beyond a price quote for individual items only. It is easy for an auctioneer to elicit complementarities, synergies, or volume discounts for large volumes of items. More comprehensive information about cost structures or utility functions of market participants can increase allocative efficiency and lead to higher economic welfare.

In recent years, a growing body of literature in the management sciences is devoted to the design of such smart markets (Gallien and Wein 2005), with combinatorial auctions (CAs) emerging as a pivotal example (Cramton et al. 2006). A CA allows bidders to bid on combinations of items, offering protection against the well-known "exposure problem" present in simultaneous auctions for heterogeneous items, in which a bidder is exposed to winning too few complementary goods to realize synergies at a high price or winning too many items that she considers substitutes at a high cost. Logistics markets have used combinatorial auctions for several years (Caplice 2007), industrial procurement is a large field of application (Bichler et al. 2006), energy exchanges are using bundle bidding in day-ahead markets (Meeus et al. 2009), and more recently spectrum auctions across the world have started using combinatorial auction designs (Cramton 2013). In all these cases, the market design has a profound effect on bidder behavior and efficiency, and many market designs originated from academic research.

Mechanism design and auction theory provide a basic framework to think about strategies and efficiency of auctions, but the design of multi-item markets has led to many new problems complementary to those discussed in microeconomics. For example, the computational complexity of allocation problems in multi-item markets has been a topic of interest in operations research and computer science Lehmann et al. (2006). The information systems literature has made contributions on decision support, pricing, and information feedback (Xia et al. 2004, Adomavicius and Gupta 2005, Bichler et al. 2009), the analysis of bidder behavior (Scheffel et al. 2011, Adomavicius et al. 2012), and the design of markets for specific domains (Guo et al. 2007, Bapna et al. 2008).

Combinatorial auctions have been used for increasingly large markets. For example, in some spectrum auctions there are around 100 licenses for sale—i.e., 2^{100} packages, which is on the order of 1.267×10^{30}.[1] As a comparison, 3×10^{23} is the number of stars in the observable universe. It is clear that larger bidders can only specify a small proportion of their bids of interest. Note that the winner-determination problem for auctions with a fully expressive XOR bid language treats missing package bids as if a bidder had no value for the package. Recent lab experiments have shown that this "missing bids problem" can lead to substantial efficiency losses, even with a much lower number of possible packages compared with a simultaneous multiround auction where bids can only be submitted on individual items (Bichler et al. 2013). In a simultaneous multiround auction, the bids are additive (OR bid language), and for each package there is an estimate of the valuations for this package, which is just the sum of the bids on the individual items. This allows for higher efficiency in larger markets than some combinatorial auction designs, even though bidders cannot express their complementarities without the risk of winning only parts of a bundle of interest and having to pay more than this subset of items is worth to the bidder. CA designs, therefore, face a natural trade-off between the efficiency gains of allowing bids on packages and the efficiency losses resulting from missing bids. This observation has caused a debate on the design of spectrum auctions, but the debate goes beyond this application and asks the question how large markets with many items can be designed such that bidders are incentivized and able to express their preferences truthfully and that auctioneers achieve allocations with high efficiency.

Our paper provides a contribution to the design of large markets with dozens or hundreds of items. In general, the term "large markets" can be used to refer to those in which the number of packages in a fully enumerative XOR bid language is more than a few hundred bids, making it clear that bidders cannot be expected to submit bids on all possible packages. Simplification was suggested as a theoretical concept to reduce the message space without losing efficient equilibria (Milgrom 2010), but apart from this, there has been little research in the design of large markets. We outline a framework that includes the design of compact bid languages and computational techniques to determine VCG and bidder-Pareto-optimal core payments, which provide incentives for bidders in a sealed-bid auction to submit their preferences truthfully.

There are many potential applications of this approach, ranging from the allocation of a large number of resources with potentially varying quality, such as the capacity on flexible manufacturing machines, the procurement of a large number of raw materials, or the sale of both TV and Internet banner advertisement with varying quality in terms of customer reach. To prove the computational applicability to realistic markets, we focus on two particular markets of interest: (1) multi-item procurement auctions with economies of scale and discounts for large quantities of each item and (2) markets for TV advertising (hereinafter, TV ad) slots. We emphasize that the techniques we propose are quite general and applicable to a number of markets, but because they involve heuristics, it is necessary to show that the methods function effectively on a few realistic implementations in terms of size and complexity.

The procurement context involves several types of raw materials that need to be procured, where the procurement manager needs hundreds or thousands of tons of each. (For example, Bichler et al. 2006 report on real-world procurement auctions with 10 to 90,000 line items. See also Sandholm 2007 for another discussion of advanced procurement auctions in recent decades). Such markets could be organized as a combinatorial procurement auction where bidders can win one out of many package bids they submit. But with only 10 material types (items) and six units of each, however, a bidder faces more than 284 million packages to consider under an XOR bidding language, which requires a unique bid for any package that might be won to verify full efficiency. Similarly, markets for TV ad slots involve the sale of air time in hundreds of different time slots, weekly or biweekly, where bidding advertisers have differing preferences over slots, based on varying audience demographics and firm-specific needs for sufficient ad reach or coverage. Although these are only examples of large markets, the size of these markets, in terms of the large number of distinct interrelated goods and heterogeneity of bidder preferences, prohibits the application of existing CA designs, making it a good test bed for the methods proposed here.

[1] The spectrum auction in Canada in 2014 included 97 licenses (Government of Canada 2009). Planned spectrum auctions in Canada and the United States have even more items for sale.

1.1. Compact Bid Languages and Allocation Rules

First, we will discuss compact bid languages as a remedy for efficiency losses resulting from missing bids. Several generic logic-based bid languages have been discussed in the literature on CAs (Boutilier and Hoos 2001). But for large markets like the ones discussed here, even these could require too many bids to be submitted. Often, prior knowledge about bidder preferences and market nuances allow for compact bid languages with a very low number of parameters that bidders need to specify to describe their preferences. For the procurement markets discussed here, the various discount policies that are regularly used in pricing can be elements of a bid language as described in Goossens et al. (2007) or Bichler et al. (2011), who substantially extend the expressiveness of a bid language for markets with economies of scale and scope. These bid languages follow established market practices, and bidders do not need to change their established discount policies. In a similar way, we will introduce a bid language for TV ad markets, which is natural to media agencies, allowing them to express their preferences with a few parameters only by describing substitutes in a succinct way. Such *domain-specific* bid languages require adequate optimization models to compute cost-minimal allocations in procurement or revenue-maximal allocations in forward TV ad auctions.

Advanced mixed-integer programming solvers allow for the computation of allocations of large TV ad markets with hundreds of ad slots and procurement markets with dozens of items and several quantity schedules to near optimality. Although such near-optimal solutions can typically be found in minutes, finding (or proving) the exact solution might take hours or even be intractable. This is a widespread pattern in combinatorial optimization. An integrality gap of a few percent would be considered acceptable in the types of large-scale private-sector markets that we discuss in this paper. Even the recent design of incentive auctions for the Federal Communications Commission in the United States includes allocation problems that are too large and difficult to be solved to full optimality.[2]

1.2. Payment Rules

Second, we will discuss payment rules to encourage truthful bidding in large markets. The celebrated Vickrey–Clarke–Groves (VCG) payment rule charges each bidder the harm they cause to other bidders and ensures that the dominant strategy for a bidder is to bid her true valuation of the items. The VCG outcome can be "outside the core" leading to low revenue and possibilities for shill bidding among other problems (Ausubel and Milgrom 2006). Intuitively, VCG provides discounts to ensure that an *individual* cannot benefit from *unilateral* deviation from truth-telling, but the resulting discounts can be so large that payments are absurdly low and remain manipulable by *groups* of bidders. For example, Ausubel and Milgrom (2006) provide a classic setup where VCG payments total zero for two bidders, despite a competitor's bid to pay the seller a large amount for their combined winnings, and they show that these payments of zero can be achieved through group manipulation or the use of false-name (i.e., shill) bids. This occurs when the first two bidders bid for their disjoint respective bundles of interest, with a losing competitor offering exactly the same amount for the union of these bundles.

Core-selecting auctions were introduced in recent years (Day and Raghavan 2007) to combat these weaknesses of VCG. As a payment paradigm for multi-item markets in general, they were designed to balance the incentives of bidders to reveal bids truthfully (achieved by making the bidders pay the least amount possible) against the perceived fairness of payments (such that payments are adequately large to preclude any set of losing bids from become winning). This auction design computes prices that are "in the core" with respect to submitted bids, stating roughly that no coalition of bidders could claim that their bids offered a mutually preferable outcome that would also raise seller revenue. Thus in the example above, the winners will always combine to pay at least the same amount in a core-selecting auction.

The game-theoretical properties of bidder-Pareto-optimal core (BPOC) auctions have been discussed extensively in the recent years (Day and Milgrom 2007, Goeree and Lien 2015), and core-selecting auction rules have been adopted for spectrum license auctions around the world, including Australia, Austria, Canada, Denmark, Ireland, Portugal, the Netherlands, and the Unites Kingdom. The approach in these spectrum auctions has been to use a combinatorial clock auction (CCA) with bidding in an iterative auction, in response to rising price clocks for each item, and finishing with a sealed-bid core-selecting auction using all bids from these iterative rounds, as well as additional combinatorial bids submitted in a sealed-bid round subject to activity rules. Thus, even if one were to argue to use the CCA format as is used in spectrum auctions for our applications (although we do not) the auctioneer would still need to run the winner-determination and core-pricing algorithm, and the algorithmic contributions of this paper would still be relevant as the auction gets large.

Although BPOC payment rules are not strategy-proof, the incentives for manipulation can be considered minimal in most large-scale markets, where

[2] http://www.fcc.gov/topic/incentive-auctions (accessed February 25, 2015).

typically neither the number of bidders nor their exact preferences are known. Note that existing game-theoretical models assume that bidders are all interested in only a single package and that all bidders know which packages their competitors bid on, in order to keep the analysis tractable (Goeree and Lien 2015). Although these analyses are insightful and illustrate situations where bidders would not bid truthful in equilibrium in a BPOC auction, such information is rarely available in large real-world markets. The number of packages that bidders could bid on can serve as a proxy for how much information would be needed by a bidder to profitably manipulate a market. Still, a simple pay-as-bid rule sets strong incentives for bid shading, while the benefits of bid shading are greatly reduced under VCG or BPOC payment rules in large markets with little or no prior distributional information. Recent lab experiments comparing a BPOC payment rule with a pay-as-bid payment rule provide evidence for this hypothesis (Bichler et al. 2014).

1.3. Relationship to Approximation Mechanisms

Recent research in computer science has explored whether strategy-proofness can be maintained by giving up on optimal social welfare and using approximation algorithms with provable approximation ratios on the quality of the allocation as an allocation rule (Lavi 2007). Unfortunately, in spite of the theoretical value of results in this field, the approximation ratios of algorithms for most combinatorial optimization problems are often not acceptable for real-world market design, and often no such approximation algorithms are available for specific problems. The approximation ratio of approximation algorithms to solve the winner-determination problem in CAs with general valuations is $O(\sqrt{N})$ (Halldorsson et al. 2000), where N is the number of items. No strategy-proof approximation mechanism can have a better ratio than this algorithmic bound. This means in an auction with 25 items only, the solution can be five times worse than the optimal solution in the worst case. Randomized approximation mechanisms with the same approximation ratio have already been found (Lavi and Swamy 2011, Dobzinski et al. 2012). However, the best deterministic truthful approximation guarantee known for general combinatorial auctions is $O(N/\sqrt{\log N})$ (Holzman et al. 2004). Note that much of the literature on approximation mechanisms relies on randomized mechanisms, which also lead to somewhat weaker notions of truthfulness than strategy-proofness with deterministic mechanisms.

We consider this literature as complementary to our research. Although there are no provable guarantees to solve certain problem sizes of combinatorial optimization problems, experiments typically lead to high confidence about the problem sizes that can be solved in due time in practice. There is a great deal of literature with various benchmark problems analyzing the empirical hardness of certain optimization problems in operations research,[3] which practitioners rely on for scheduling, vehicle routing, or other types of resource allocation problems. We do this as well. We also give up on strategy-proofness in the strong sense. Strategy-proofness is a powerful but also restrictive concept, which is why we instead focus on the weaker notion of core-selecting payments.

It is worth noting that the VCG mechanism is no longer strategy-proof if the allocation does not necessarily maximize social welfare. The simple proof showing that the VCG mechanism leads to a dominant-strategy equilibrium for each individual bidder (see, e.g., Shoham and Leyton-Brown 2008, p. 276) relies on the argument that the auctioneer chooses the allocation that maximizes the *coalitional value* based on the reported bids of all bidders. So, if the allocation cannot be computed optimally, then the VCG mechanism also loses this strong game-theoretical properties. We refer to coalitional value as the result of the allocation problem assuming bidders report their true valuations.

Still, the basic concept of a second-price rule can encourage truthful bidding because shading one's bids might not increase profit, but it might increase the risk of losing in the auction or getting a less desired outcome. Note that the information a bidder would need to manipulate grows exponentially with the number of items in a combinatorial auction. With many bidders and many items but little distributional information about all possible combinations, profitable manipulation becomes almost impossible. The amount of information required by a bidder to profitably manipulate in a specific auction could well serve as an alternative way to characterize markets, different from the game-theoretical solution concepts that are typically used in auction theory. On one hand, dominant strategies restrict the auction designer to the VCG mechanism (Green and Laffont 1979), only applicable with optimal allocation rules. On the other hand, Nash equilibria are computationally hard to compute in general (Daskalakis et al. 2009), and Bayes–Nash equilibria of combinatorial auctions require a very large number of distributional assumptions, rendering this solution concept intractable in large markets as discussed in our paper. Given the lack of sufficient information about other bidders' valuations or the specific packages they are interested in, and the hardness of computing Bayes–Nash equilibrium strategies in large markets, we argue that a second-price rule

[3] See "OR Library" at http://people.brunel.ac.uk/~mastjjb/jeb/info.html for different types of discrete optimization problems.

such as in BPOC payments offers a compelling compromise, encouraging truthful bidding with substantial discounts, rather than guaranteeing it.

The application of second-price payment rules such as BPOC or VCG rules with near-optimal rather than exact solutions to the allocation problem in our framework is not without challenges, however. For example, the coalitional value of a coalition without one of the winners (required to compute the VCG payments) might return a higher value than the coalitional value with all bidders. Our adaptation of a constraint-generation technique for the computation of BPOC payments by Day and Raghavan (2007) to large-scale markets is a new, practically viable paradigm by which core-selecting auction designs with good incentive properties can be applied to large markets. It is the combination of the compact bid language and the payment rule that allows bidders to express their complementarities, but at the same time, it provides incentives for truthful bidding. Our experiments help understand how the near optimality of the allocations impacts the payments of bidders. Overall, this can be a recipe for many large-scale markets beyond the ones discussed in this paper.

1.4. Contributions and Outline

In summary, our contributions are as follows: First, motivated by work with industrial partners, we propose a compact bid language for the TV ad market. The TV ads application will be our leading example because it provides a rich test bed to demonstrate our ideas with a realistic valuation model and a type of winner-determination problem that would benefit most readily from the approach. The bid language allows the expression of preferences for a large number of packages with only a few parameters. The winner-determination problem in such markets is $\mathcal{NP}$-hard and cannot always be solved optimally. However, as is typical in many combinatorial optimization problems, near-optimal solutions can be found within a few minutes for limited problem sizes, which is promising, but previous algorithms for computing core payments break down when solutions are not exact. We therefore propose two algorithms to deal with markets where the winner determination might not be solved to optimality. The approaches are evaluated in an extensive set of experiments and their properties are characterized. In addition to the TV ad market, we also analyze the two algorithms in the context of volume-discount auctions to show that the basic framework and the results carry over to other large markets. Here we draw on a compact bid language for procurement markets with economies of scale introduced by Bichler et al. (2011).

Overall, we show that the dynamic reuse algorithm we develop is not too slow relative to the quicker trimming algorithm, so the time cost of computing more accurate and more fair outcomes should not be out of reach for this or similar applications. Although the trimming algorithm is quicker and generates more revenue for a fixed set of bids, its inferior efficiency and incentive properties make its use harder to justify, particularly for government applications where efficiency concerns naturally dominate. The market design can serve as a template for other large markets with many items and complex bidder preferences, an area of research with many applications but little attention in the literature so far.

In §2, we discuss related work. Section 3 introduces the market design including the allocation and the pricing rules. In §4.2, we propose two algorithms to deal with nonoptimal solutions to the winner-determination problem. Section 5 summarizes the results of our experiments, and §6 concludes with a summary and future outlook. A complete glossary of symbols is also provided in the appendix for quick reference.

2. BPOC Auctions

In this section, we want to provide further background on auctions with bidder-Pareto-optimal payment rules because they are central to this paper. The concept of the core has a long history in economics, and indeed a mechanism that selects a core outcome based on submitted preferences was the foundation for the 2012 Nobel Prize in Economics, although that stable-matching market involved allocations that do not allow for monetary transfers. The extension of these ideas to the auction context (with payments) began indirectly in the work of Parkes and Ungar (2000), Ausubel and Milgrom (2002), with explicit computation of core outcomes and formal treatment of the core-selecting approach coming later in Day and Milgrom (2007), Day and Raghavan (2007), and Day and Cramton (2012). Core-selecting auctions have been suggested as an alternative to the VCG mechanism, which suffers from a number of problems such as low seller revenue (Ausubel and Milgrom 2006). VCG solutions outside the core, where a subset of bidders has indicated to be able to pay more than what the winners paid, is often seen as undesirable.

Day and Raghavan (2007) showed that under semi-sincere bidding strategies and perfect information, every BPOC price vector forms a Nash equilibrium. Thus, assigning BPOC payments based on bids being true values only makes a bidder pay what she should have bid in equilibrium if she had expertly anticipated the true values (and bids) of others. Day and Milgrom (2007) show that bidder-Pareto optimality implies optimal incentives for truthful revelation over all core-selecting auctions, among other supporting results, including decreased vulnerability to false-name bidding and collusive behavior relative to other auction formats discussed in the literature, in particular the VCG mechanism. Day and Raghavan also note that *total-payment minimizing* BPOC points are further resistant to certain forms of collusion with side payments, and this minimum revenue condition has been implemented in all core-selecting spectrum auctions to date. Selecting a BPOC point is further supported by the fact that, if the truth-revealing VCG vector is in the core, then any BPOC algorithm will produce VCG as its output.

A core-selecting auction only provides a dominant strategy if the VCG outcome is in the core. Goeree and Lien (2015) actually show that no Bayesian incentive-compatible core-inducing auction exists when the VCG outcome is not in the core. In specific settings, where the VCG outcome is outside the core, the equilibrium bidding strategy is to shade bids below one's true valuation, speculating that the reduced bid can lower one's payment without affecting the bundle of goods awarded. Simple threshold problems where multiple local bidders only interested in one item compete against a global bidder, who is interested in all items, provide an illustrative example. Local bidders could try to free ride on each other. However, Bayesian analyses of such markets assume that bidders are interested in a single item and that they know what other bidders bid on and have commonly known prior distributions available about other bidders' valuations for their bundles of interest. In large TV ad markets, such information is not available to bidders and bidders are multiminded. Often bidders do not even know how many competitors exist in the market, making speculation quite risky. The same assumptions hold in procurement markets such as the volume-discount auctions analyzed in this paper. In such situations, manipulation comes at a high risk of winning nothing.

A potential alternative to the approach proposed here is the use of proxy agents bidding in multiple rounds of an ascending auction until an equilibrium is reached, as in the ascending proxy auction (Ausubel and Milgrom 2002) or iBundle (Parkes and Ungar 2000). Unfortunately, this approach requires a very large number of auction rounds (unless the problem size is quite small) and the auctioneer needs to solve a winner-determination problem in each round (Schneider et al. 2010). By trying to avoid unnecessary winner-determination optimizations, constraint generation after a sealed-bid auction is a more effective and practical method for price generation, especially when considering cases of hard winner-determination problems.

3. Compact Bid Languages and Allocation Problems

The TV ad slot market will serve as the main example in our paper, which shares many of the features of other large markets for the sale of spectrum licenses, in logistics, or in industrial procurement. Also, we will briefly review volume-discount auctions designed for industrial procurement to illustrate that the framework outlined in this paper can easily be applied to other large-scale markets. For this latter setting, we are able to find exact solutions in our computational experiments for smaller instances, allowing for direct benchmarking against optimality. These benchmarks are provided to demonstrate the approach, but clearly we propose the near-optimal approach to be relevant in practice to larger instances where exact optimality is out of reach.

3.1. TV Ad Markets

In what follows, we provide a brief overview of the essential requirements. Parts of the advertisement capacity of a typical TV station are sold via specially negotiated, large, long-term contracts of about a year and are not considered in our study. We focus instead on the sale of the remaining ad slot inventory to specific marketing campaigns that run in the short term, which in Europe are typically sold via posted prices, in advance of airing. Prices for different slots can range from 6,000€ to 50,000€ for a duration of 30 seconds and are set by the TV station based on historical demand. Buyers are large media agencies that purchase a set of slots with the intent to procure the best slots for each of their customers' campaigns. The number of agencies in a particular market depends on the country and the particular station, but a typical short-term market for a TV station in Germany, for example, consists of approximately 50 media agencies, booking slots for several hundred customers in a particular channel.

Because the amount of airtime filled by long-term customers varies, the length of a slot available in the short-term market can vary between 2 and 5 minutes, while the length of an ad also varies considerably, lasting up to 1 minute. For a particular channel in the markets we investigated, there are on the order of 150 short-term slots available during the program per week.

Different slots have a different reach for different customer segments or the population overall. The reach of a particular slot varies over time, but there are estimates based on historical panel data available to clients of the media agencies. Clients use reach per segment (based on gender, age, or other demographics) or per population to determine their willingness to pay for different slots. Clearly, the value of some slots, such as those during the finals of the national soccer league, may be difficult to estimate and their valuation varies considerably depending on the target market of an advertiser. Apart from these high-value slots, there is also typically a segment of low-value slots, which are also difficult to price because the demand is hard to predict. This difficulty in demand or valuation prediction together with limited supply suggests that an auction market would outperform the existing posted-price mechanism.

The allocation of TV ad slots can be modeled as a multiple knapsack problem, in which each time slot i in the set $I = \{1, 2, \ldots, N\}$ is treated as a knapsack with a maximum capacity/duration of c_i, which cannot be exceeded. As mentioned above, each slot can potentially hold a number of ads, although some may have been previously allocated to larger customers, so we assume that c_i reflects only short-term capacity in the current market, making for a potentially heterogeneous list of c_i values, even for a TV station with slots of the same size when considering all ads aired. We also assume that each slot i has a reservation value or minimum price per time unit r_i, which reflects the station's ability to off-load excess capacity at a low price to existing customers if needed. Station call signs and other brief announcements can also be used to fill any excess unused time.

Each bidding advertiser k in the set $K = \{1, 2, \ldots, \mathcal{K}\}$ has an ad of duration d_k to be shown repeatedly (at most once per time slot) if she wins a bundle of time slots in the auction. To ensure adequate reach, each bidder specifies an abstract "priority vector" or "weight vector" W_k, containing an arbitrary weight value w_{ik} for each time slot. These "weights" conveying "strength of priority" could specifically represent the expected viewership, expected viewership of a particular demographic, viewership weighted by expected sales, etc., reflecting the advertiser's performance metric of choice. She can then bound the total priority value in the auction outcome to be greater than or equal to a minimum amount to qualify bids of various levels.

Thus, after specifying the priority vector and ad duration, a bidder places one or more tuples $(w_j^{\min}, b_j)$ containing the desired sum of priority values $w_j^{\min}$ necessary to justify a monetary bid b_j. At most, one of the bids placed by a bidder can win, making the bidding language an XOR of "weight threshold levels." For example, if the bidder sets the priority weights w_{ik} at the expected viewership of each slot i, the XOR structure lets her set an exact price for any particular price point of interest. She can set a price for a total of $w_j^{\min} = 1$ million viewers, a price for $w_j^{\min} = 2$ million viewers, etc., regardless of which slots are chosen to reach this total viewership. This price-point structure reflects the ability of the language to represent the fundamental complementarity in this type of market; a small number of ad slots (or small reach, etc.) may have little or no value, but several of them together are worth more than the sum of the parts.

The set J_k contains all bid indexes j by a bidder k, and the superset J is defined as $J := \bigcup_{k\in K} J_k$. We assume these bids are submitted in a sealed-bid format, consistent with the timing of Google's auction, in which bids were submitted once to a proxy. In such markets, it is not practical for media agencies to participate in an ascending auction every week or two. After the bids are submitted, the market is cleared at a particular point in time, and the allocation is determined for some period for a time (e.g., two weeks) in the future.

Formulation WD maximizes the value of accepted bids given that ad durations do not exceed capacity in any slot (1a), the bid values are not less than the seller's reservation values (1b), the priority threshold level $w_j^{\min}$ of a bid j is met if and only if that bid is accepted (1c,d), and at most one bid j is accepted for each bidder k (1e). Decision variables x_{ij} and y_j indicate time if slot i is assigned to bid j and bid j itself is accepted, respectively, while M is a sufficiently large positive constant parameter. WD(K) indicates that all bidders $k \in K$ are included. Later we refer to WD(C) for a coalition $C \subseteq K$, which indicates the same formulation but with all bids made $k \notin C$ set to zero. Overall, we use the term *coalitional value* to describe the objective function value of formulation WD.

$$\mathrm{WD}(K) = \max \sum_{j\in J} b_j y_j \tag{WD}$$

$$\text{subject to} \quad \sum_{j\in J} d_k x_{ij} \le c_i \quad \forall\, i \in I, \tag{1a}$$

$$d_k \sum_{i\in I} r_i x_{ij} \le b_j \quad \forall\, j \in J, \tag{1b}$$

$$\sum_{i\in I} w_{ik} x_{ij} \le M y_j \quad \forall\, j \in J, \tag{1c}$$

$$w_j^{\min} - \sum_{i\in I} w_{ik} x_{ij} \le M(1 - y_j) \quad \forall\, j \in J, \tag{1d}$$

$$\sum_{j\in J_k} y_j \le 1 \quad \forall\, k \in K, \tag{1e}$$

$$x_{ij} \in 0, 1 \quad \forall\, i \in I, j \in J, \tag{1f}$$

$$y_j \in 0, 1 \quad \forall\, j \in J. \tag{1g}$$

The priority vector W_k provides quite a bit of flexibility to the bidders in expressing their preferences over ad slots, and we propose that this novel bidding language could be relevant in a number of other areas. Indeed, the language captured in this formulation is quite general and includes the "k-of-singletons" expressions described in Hoos and Boutilier (2000), which were shown to be difficult to express succinctly with more fundamental logical operators and result in hard optimizations. For example, a bidder in the ad slot auction might want her ad to be on the air at least five times within one week between 8 P.M. and 10 P.M. That is, all ad slots between 8 P.M. and 10 P.M. are substitutes, but the bidder needs at least five of them, a complementarity valuation for a sufficient volume from a group of substitutes. The priority-vector format would then have weights equal to one for the selected set of substitute times and $w_j^{\min} = 5$ playing the role of the k term in Hoos and Boutilier (2000).

Theorem 1. *The decision version of the* WD *problem is strongly* $\mathcal{NP}$*-complete.*

The proof is by reduction from the multiple knapsack problem, and it can be found in the appendix. The decision version of the multiple knapsack problem is strongly $\mathcal{NP}$-complete (Chekuri and Khanna 2006). Whereas weakly $\mathcal{NP}$-complete problems may admit efficient solutions in practice as long as their inputs are of relatively small magnitude, strongly $\mathcal{NP}$-complete problems do not admit efficient solutions in such cases. Unless $\mathcal{P} = \mathcal{NP}$, there is no fully polynomial-time approximation scheme (FPTAS) for strongly $\mathcal{NP}$-complete problems (Garey and Johnson 1979). Even if we cannot hope for FPTAS, we can get near-optimal solutions with standard mixed-integer programming solvers for practically relevant problem sizes as we will show.

3.2. Volume-Discount Auctions

Volume discounts are in widespread use in markets with economies of scale. Davenport and Kalagnanam (2000) were among the first authors to discuss volume-discount bids in an auction. Their bid language requires suppliers to specify continuous supply curves for each item. They apply discounts only to additional units above a threshold of a specific price interval. In contrast to these incremental volume-discount bids, Goossens et al. (2007) proposed tiered bids, which they refer to as total-quantity bids. The latter are valid for the entire volume of goods purchased after a certain quantity threshold. For example, a supplier charges \$4 per unit for up to 1,500 units, but only \$2.50 per unit for the entire quantity if the purchasing manager were to buy 1,500–2,000 units. In practice, suppliers employ various types of such discount policies in different settings. In addition to volume-discount bids and total-quantity bids one can find lump-sum rebates on total spend, and such discounts can be based on the quantity or spend of one or a few items that are being auctioned off. Bichler et al. (2010) introduced a comprehensive bid language that allows for different types of discount policies, including volume-discount bids, total-quantity bids, and lump-sum rebates. They propose a mixed-integer program to solve problems of up to 30 suppliers, 30 items, and five quantity schedules to near optimality in less than 10 minutes. Because of space constraints, we refer the interested reader to Bichler et al. (2010) for a detailed description of the bid language and the experimental setup and results. An analysis of various heuristics and metaheuristics to solve the problem can be found in Hass et al. (2013). Even though such near-optimal solutions were always possible with these problem sizes, proving the optimality of a solution was typically intractable and even after hours there would be a small integrality gap. This phenomenon is widespread in combinatorial optimization overall. In our experiments, we will use the compact bid language introduced by Bichler et al. (2010) and their experimental setup to compute VCG and BPOC payments for near-optimal allocations.

4. Payment Rules

If we only aim for near-optimal solutions, not for exact solutions to the winner-determination problem, some computational issues can arise. For example, the objective function value of the best allocation with one winner excluded might be higher than that of the best allocation with all bidders included when computing VCG payments. We will first revisit BPOC payments, before we discuss different algorithms how to compute these payments with near-optimal solutions to the winner-determination problem. We will use the terms payments and prices interchangeably.

4.1. Bidder-Pareto-Optimal Core Payments

We will determine BPOC payments in the following treatment and compare them with the VCG payments in our experiments. The technique of using constraint generation to find the coalitions defining the core was designed to work in any context where the winner-determination problem could be solved exactly. Here, we quickly reiterate that approach before extending it to situations of near-optimal winner determination in the next section.

The approach discussed in the literature is to find core prices by iteratively creating new price vectors p^t for the winning coalition W and then checking at each iteration t whether there is an alternative outcome that generates strictly more revenue for the seller and for which every bidder in this new outcome weakly prefers to the current outcome. If such a coalition exists, the alternative winning coalition C is called a *blocking coalition*, and a constraint is added to a partial representation of the core in payment space until no further blocking coalitions can be found. To discover the most violated blocking coalition C^t relative to the

Figure 1 **(Color online) A Typical Instance Showing the Reduction of the Integrality Gap Over Time**

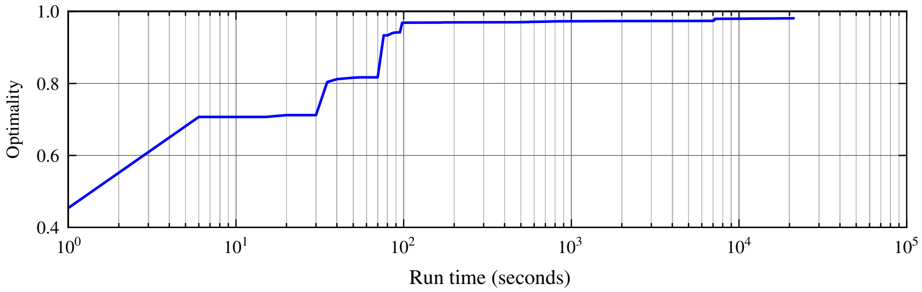

current payments at iteration t, the WD is extended as in the separation problem SEP^t.

$$z(p^t) = \max \sum_{j\in J} b_j y_j - \sum_{k\in W} (b_k^* - p_k^t)\gamma_k \qquad (\mathrm{SEP}^t)$$

$$\text{subject to } (1a)-(1g),$$

$$\sum_{j\in J_k} y_j \le \gamma_k \quad \forall k \in W,$$

$$\gamma_k \in [0,1] \quad \forall k \in W.$$

Here, W is the set of winners from the solution of WD(K), and b_k^* represents bidder k's winning bid. If the sum of the current payments p^t is less than the solution to (SEP^t), then a violated core constraint has been found, and we must add a constraint to our partial representation of the core. Following Day and Raghavan (2007), this partial representation is given in the following linear program to find equitable bidder-Pareto-optimal (EBPO) payments. The program is then solved to find the next tentative set of payments p^{t+1} until no further constraints can be found.

$$\theta(\epsilon) = \min \sum_{k\in W} p_k + \epsilon m \qquad (\mathrm{EBPO}^t)$$

$$\text{subject to } \sum_{k\in W\setminus C^\tau} p_k \ge z(p^\tau) - \sum_{k\in W\cap C^\tau} p_k^t \quad \forall \tau \le t, \qquad (\mathrm{EBPO}^t.1)$$

$$p_k - m \le p_k^{vcg} \quad \forall k \in W,$$

$$p_k \le b_k^* \quad \forall k \in W,$$

$$p_k \ge p_k^{vcg} \quad \forall j \in W.$$

As in SEP^t, b_k^* is the winning bid for k, the parameters $p_k^{vcg} = b_k^* - (\mathrm{WD}(K) - \mathrm{WD}(K_{-k}))$ represent VCG payments, and m represents a maximum deviation from VCG, which is minimized as a secondary objective after minimizing total payments.[4] We then use the value of each p_k in the solution for the next iteration (i.e., set $p_k^{t+1} = p_k$).

[4] In practice, these two minimizations can be handled as separate optimizations, but they are presented here as a single optimization using a sufficiently small ϵ for the sake of concise exposition.

4.2. Core Payments with Near-Optimal Allocations

For many combinatorial optimization problems, good solutions often can be found quickly, even though finding a provably optimal solution may take a very long time. Figure 1 shows a typical example of WD with 336 slots and 50 bidders, where a feasible solution with 95% optimality could be reached within a few minutes. This is a common phenomenon in many combinatorial optimization problems.

Without the ability to guarantee an optimal solution quickly enough for a practical application, one would naturally consider a provably high-quality solution that can be found quickly. Most industrial mixed-integer programming solvers (e.g., CPLEX, Gurobi) provide absolute and relative worst-case optimality gap parameters, allowing the optimization routine to terminate if the optimality gap (the difference between the best feasible solution and the theoretical bound) falls below some target or is a small enough percentage of the best feasible solution, respectively. For now, we leave the exact specification of how a "good enough" approximate solution is qualified, but motivated by Figure 1, the reader may assume a 5% optimality gap or an at-least-95%-optimal solution for concreteness. We will thus write WD^a for any qualified approximation of a WD value and consider an implementation using these approximations in place of true WD values. Similarly, we will write $z^a(p^t)$ for separation problem values found using near-optimal solutions.

Problems can arise, however, during the VCG and core price calculation when accepting these approximate or near-optimal solutions. For example, under truly optimal solutions, with the standard assumption of free disposal, WD(K_{-k}) is always at most the value of WD(K). But with a series of near-optimal computations this is not guaranteed, opening up the possibility that one might compute an approximate VCG payment with $b_k^* - (\mathrm{WD}^a(K) - \mathrm{WD}^a(K_{-k})) > b_k^*$. Similarly, under near-optimal computation the coalitional value of SEP^t can be higher than the value of the WD. If this happens, the newly generated constraint added to EBPO^t can cause an infeasibility if

$\sum_{k\in W\setminus C^t} b_k^* < z^a(p^t) - \sum_{k\in W\cap C^t} p_k^t$. Two different solutions are presented to address this problem that can potentially arise during computations.

4.2.1. The TRIM Algorithm. With known b_k^* values determining individual rationality (IR) constraints (i.e., payments must not exceed bids), a natural first approach is to adjust each WD-based result so that it fits into the IR region.

For the VCG prices, this technique makes use of the fact that

$$b_k^* \geq p_k^{vcg} \geq 0 \quad \forall k \in W \tag{2}$$

whereas for the (constant) right-hand side (RHS) of the constraints in the EBPO^t, we must always have:

$$\sum_{k\in W\setminus C^\tau} b_k^* \geq z(p^\tau) - \sum_{k\in W\cap C^\tau} p_k^\tau \quad \forall \tau \leq t \tag{3}$$

Thus, our first algorithm[5] is to simply trim the infeasibilities based on known bids, represented in Algorithm 1 in the two steps using min functions.

Algorithm 1 (Core Price Calculation—TRIM)
Solve: $\mathrm{WD}^a(K)$;
for $k \in W$ **do**
  Solve: $\mathrm{WD}^a(K_{-k})$;
  $p_k^{vcg} \leftarrow \min(b_k^*, b_k^* - (\mathrm{WD}^a(K) - \mathrm{WD}^a(K_{-k})))$;
$p^1 \leftarrow p^{vcg}$;
$\theta^0 \leftarrow \sum_{k\in W} p_k^{vcg}$;
while *true* **do**
  Solve: SEP^t
  if $z^a(p^t) \leq \theta^{t-1}$ **then**
    Break: 'core' price vector found;
  Generate RHS of new constraint:
    $\alpha^t \leftarrow \min(\sum_{k\in W\setminus C^t} b_k^*, z^a(p^t) - \sum_{k\in W\cap C^t} p_k^t)$;
  Add constraint $\sum_{k\in W\setminus C^t} p_k \geq \alpha^t$ to EBPO;
  $p^t, \theta^t \leftarrow$ Solve: EBPO;
  if $(C^t, \theta^t) = (C^{t-1}, \theta^{t-1})$ **then**
    Break: no better price vector possible;
  Iterate: $t \leftarrow t+1$.

4.2.2. The REUSE Algorithm. An alternative to trimming infeasibilities is based on the observation that whenever an infeasibility is found, the validity of expressions (2) and (3) imply that an update can be made to an approximate WD value, from a previously best-known feasible solution to a new tentatively optimal feasible solution. To implement this change, the storage of any value based on a winner-determination solution can no longer be treated as constant and must be regenerated at each iteration based on current WD^a values. This includes VCG price estimations and the RHS values for generated core constraints, whose definitions must be reformulated based on current WD^a values.

Thus, our second approach is to store a list of all discovered $\mathrm{WD}^a(C)$ values, reusing all coalitions found so far and reformulating the entire separation problem and EBPO problem at each iteration, noting that the set of relevant core constraints, and indeed the set of winners itself, may be changing as new information becomes available. Whenever we run WD^a (the first time, to compute each VCG price, and inside each run of SEP), we get a new collection of feasible bids, representing a coalition of bidders, and we check these values against our current list of coalitions and WD^a values. If the coalition has not been found before, our list is extended to include it as among the "potentially important" coalitions to consider. If any superset coalition has been listed previously but with a lower coalitional value, we can update it to the current $\mathrm{WD}^a(C)$ value, as a new better approximation has been found.

Because we will now store the blocking coalitions C^t and its value instead of $z(p^t)$ after each SEP^t has been solved, we are forced to work with a reformulation of core constraints based on WD^a values rather than separation levels. For a general winner-determination problem (i.e., with respect to any alternative bidding language), the core constraints can be expressed with the alternative, equivalent expression that can be derived by substitution:

$$\sum_{k\in W\setminus C^\tau} p_k \geq \mathrm{WD}(C^\tau) - \sum_{k\in W\cap C^\tau} b_k^* \quad \forall \tau \leq t \tag{EBPO^t'.1}$$

Using this formulation of the core, we can generate a constraint in EBPO for any C^τ found so far using the current best-approximation $\mathrm{WD}^a(C^\tau)$ in place of $\mathrm{WD}(C^\tau)$. For bidding languages with only one relevant bid b_k per bidder (as is the case in the scenarios presented in §5), this constraint can be further simplified, resulting in the following formulation (4) in place of EBPO^t'.1.

$$\sum_{k\in W\setminus C^\tau} p_k \geq \sum_{k\in C^\tau\setminus W} b_k \quad \forall \tau \leq t \tag{4}$$

This new set of constraints provides an intuitively pleasing interpretation of core constraints in the TV ad context: Any subset of winners pays enough to match the counteroffer including a set of losing bidders that would otherwise benefit the seller, a direct analogue to second prices.

Although it is not guaranteed that each stored coalition C^t provides a potential maximally violated coalition at the end of the constraint generation process, the addition of all constraints found at any point

[5] In all algorithmic implementations that follow, we assume that all feasible integer solutions are stored by the optimizer and used to generate bounds on subsequent optimizations using alternate objective functions.

drastically improves the overall performance of the algorithm when compared with having to completely rebuild a set of blocking coalitions after a change in W. That is, it is better to reuse potentially redundant constraints than to start over, looking for relevant constraints from scratch each time the set of winners is updated. Also, because the core pricing is computed as an LP (without integer constraints), it is not computationally expensive to have redundant constraints.

The formulation of the separation problem as an altered WD problem has the additional benefit that all feasible solutions remain feasible for a WD or SEP^t instance. MIP solvers store feasible (integer) bases internally, and if the separation problem is implemented as the same problem instance with some changes to objective coefficients, all feasible bases (stored as a branch-and-bound tree) remain feasible and thus provide immediate bounds on the SEP^t problem, making efficient use of all information found by the solver. This makes it progressively more difficult to find relevant feasible solutions. Therefore, in our experimental evaluation, we allowed for longer time limits on the optimization routine for later SEP^t instances.

Function (Core Price Calculation—REUSE—Update Coalitions(*optwinners*, *bidsum*))
for $C \in$ *Coalitions* **do**
  if $C \supseteq$ *optwinners* and $\mathrm{val}(C) <$ *bidsum* **then**
    $\mathrm{val}(C) \leftarrow$ *bidsum*;
    if $C = K$ **then**
      $W \leftarrow$ *optwinners*;
      EBPO $\leftarrow$ null;
      reset $\leftarrow$ *true*;
for $k \in W$ **do**
  $p_k^{vcg} \leftarrow b_k^* - \mathrm{val}(K) + \mathrm{val}(K_{-k})$.

Algorithm 2 as presented below keeps a list *Coalitions*, each element being a list of winners under some feasible integer solution to WD. For each coalition $C \in$ *Coalitions*, we also store the best known value $\mathrm{val}(C)$, which can be revised as the algorithm progresses. Furthermore, for ease of exposition, these algorithms refer to the winning bidders from the most recent optimization run as *optwinners*, and the sum of the (actual, i.e., unaltered) bids of these winners is given as *bidsum*.

Algorithm 2 (Core Price Calculation—REUSE)
Solve: $\mathrm{WD}^a(K)$;
Coalitions $\leftarrow \{K\}$;
$\mathrm{val}(K) \leftarrow$ *bidsum*;
$W \leftarrow$ *optwinners*;
$t \leftarrow 0$;
EBPO $\leftarrow$ null;
reset $\leftarrow$ *false*;
while *true* **do**
  if EBPO = *null* **then**
    ComputeVCG:
    for $k \in W$ **do**
      Solve: $\mathrm{WD}^a(K_{-k})$;
      Coalitions $\leftarrow$ *Coalitions* $\cup \{K_{-k}\}$;
      $\mathrm{val}(K_{-k}) \leftarrow \mathrm{WD}^a(K_{-k})$;
      UpdateCoalitions(*optwinners, bidsum*);
      if *reset* = *true* **then**
        Break k loop;
    $p^t \leftarrow p^{vcg}$;
    $\theta^t \leftarrow \sum_{k\in W} p_k^t$;
  if *reset* = *true* **then**
    reset $\leftarrow$ *false*;
    Continue;
  Solve: $\mathrm{SEP}^t(p^t)$;
  if $z(p^t) \le \theta^t$ **then**
    Break: 'core' price vector found;
  Coalitions $\leftarrow$ *Coalitions* $\cup$ {*optwinners*};
  val(*optwinners*) $\leftarrow$ *bidsum*;
  UpdateCoalitions(*optwinners, bidsum*);
  if *reset* = *true* **then**
    reset $\leftarrow$ *false*;
    Continue;
  if *EBPO* = *null* **then**
    build EBPO with constraints $\mathrm{EBPO}^{t\prime}.1$ for all $C \in$ *Coalitions* with $\mathrm{val}(C)$ as a best approximation of $\mathrm{WD}(C)$;
  else
    add constraint $\mathrm{EBPO}^{t\prime}.1$ to EBPO with $C =$ *optwinners* and with $\mathrm{val}(C)$ as a best approximation of $\mathrm{WD}(C)$;
  $p^{t+1}, \theta^{t+1} \leftarrow$ Solve: EBPO;
  Iterate: $t \leftarrow t+1$.

The difference in the quality of the TRIM and REUSE approaches, in terms of closeness to core-selecting prices, can be described as follows. Let ζ_C^t represent the amount that the final prices p^t and allocation x^t violate the core-defining constraint (with respect to submitted bids) indexed by coalition C. Let $trim^t$ denote the amount "trimmed" in the final iteration of the TRIM algorithm—i.e., $trim^t = \max(0, z^\alpha(p^t) - \sum_{k\in W\cap C^t} p_k^t - \sum_{k\in W\setminus C^t} b_k^*)$. Finally, let gap^t represent the final absolute optimality gap when solving SEP^t. Theorem 2 provides simple bounds on possible deviation from optimality-based core-selecting prices. This result indicates that the optimality gap (in absolute rather than relative terms) of the final separation measures the potential for violation of core selection under a near-optimal approach, with any trimming performed by the TRIM algorithm translating one-to-one into further potential for core violation.

Theorem 2. *For a fixed set of bids,* $\zeta_C^t \le gap^t \ \forall\, C \subseteq K$ *under REUSE, while* $\zeta_C^t \le gap^t + trim^t \ \forall\, C \subseteq K$ *under TRIM.*

Proof. Core constraints are most often written as $\sum_{k\in C}\pi_k \geq \mathrm{WD}(C)$, where C must include the seller and π represents each player's payoff. Suppose such a constraint is violated and replace payoffs with surplus for bidders (i.e., $b_k^* - p_k^t$) and total payments for the seller. We get that for some positive value ζ_C^t:

$$\sum_{k\in C} p_k^t + \zeta_C^t = \mathrm{WD}(C) - \sum_{k\in C}(b_k^* - p_k^t).$$

Under the REUSE algorithm, this expression becomes

$$z^a(p^t) + \zeta_C^t = z(p^t, C)$$

because the final separated cut must be tight. Here $z(p^t, C)$ represents the true value of the separation objective function evaluated at the feasible solution implied by WD(C). Because that same feasible solution was a candidate when solving SEP^t approximately, however, by the definition of the optimality gap we must have

$$z^a(p^t) + gap^t \geq z(p^t, C)$$

with the desired result for REUSE following by substitution. The result follows analogously for TRIM, with the difference that the second line becomes

$$z^a(p^t) - trim^t + \zeta_C^t = z(p^t, C). \quad \square$$

5. Experimental Evaluation

In this section, we examine the solution quality of the presented algorithms under constrained computation time. Using the simulations provided below, we analyzed a number of primary attributes to measure overall performance, such as allocative efficiency, revenue, and speed. To directly compare the quality of the generated prices, a series of secondary metrics were computed. These values allow a comparison of how much a bidder could possibly gain by shading her bid by comparing the ratios of the bids with the BPOC prices and to the VCG prices, respectively.

- Primary metrics

1. The efficiency in terms of the coalitional value achieved by the set of winners computed after a restricted time compared with that of the optimal allocation (E).
2. The overall auctioneer's revenue based on the computed payments compared with that in the optimal allocation (R).
3. The duration of the computation (D).

- Secondary metrics

4. The ratio of the BPOC payments p_k to the bids b_k (core/bid).
5. The ratio of the VCG payments p_k^{vcg} to the bids b_k (VCG/bid).
6. The ratio of the VCG payments p_k^{vcg} to the BPOC payments p_k (VCG/core).

Figure 2 provides values of different instances for the maximum revenue and the sum of payments achieved with TRIM and REUSE. These absolute values are difficult to compare because the instances are based on different value draws. We are interested in the comparison between the performance of the algorithms across different instances, so we require a baseline for a sensible comparison of different payment schemes. A potential baseline is the optimal revenue of the winner-determination problem, against which we could compare the value of the winning coalition after a restricted solving time, the VCG payments, and the solutions by TRIM or REUSE. In the volume discount auction, we could select instance sizes for which we could compute the optimal solution with more time allotted. Achieving optimality took a prohibitively large amount of time for the TV ad market experiments, however. All instances could be solved to near optimality, but not to optimality in several hours. For these experiments, we used the objective function value or optimal revenue of the best linear programming relaxation (LPR) of the winner determination problem as an upper bound for the optimal integer solution.

An example of a typical instance of WD is shown in Table 1. REUSE runs longer and generates less

Figure 2 (Color online) Coalitional Values in Different Experiments

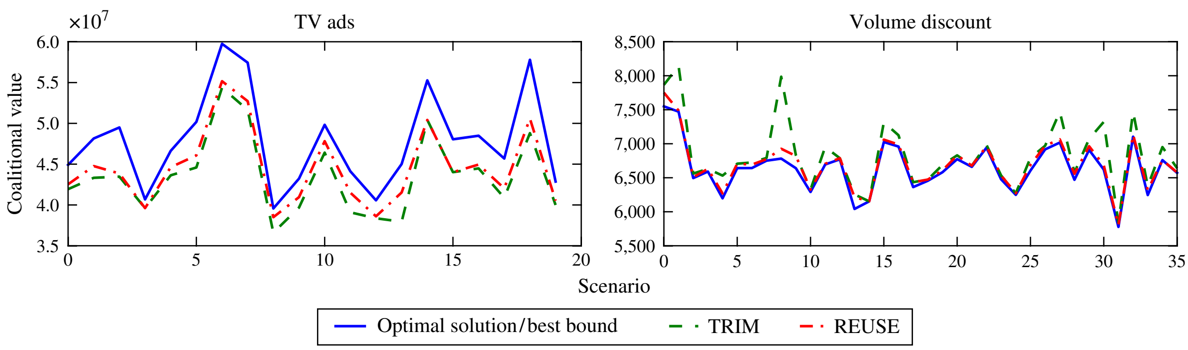

Table 1 Comparison of a Representative Experiment

Example instance	REUSE	TRIM	Optimal revenue (LPR)
Coalitional value	46,073,899	44,590,749	50,387,546
Revenue	36,569,158	42,766,735	—
Run time (h)	2.8	2.0	—
Median(p_k/b_k)	0.79	1.00	—
Median(p_k^{vcg}/b_k)	0.58	1.00	—
Median(p_k^{vcg}/p_k)	0.78	1.00	—

revenue than TRIM, but it is more efficient. The TRIM technique, on the other hand, often results in pay-as-bid pricing where bids and payments coincide. We will see that this pattern also emerges in a larger set of experiments.

All experiments were run on Dual Socket Octo Core AMD Opteron 2.4 GHz computers running the Linux operating system with 8 GB DDR2 RAM. All optimization problems were solved with the Gurobi 5.5 mixed-integer programming solver using the default parameters. The time limit for each single optimization was set to 300 seconds for TRIM and REUSE, unless stated otherwise. An optimization describes the process of solving a single mixed-integer program such as the initial WD, the WD_{-j} needed to compute the VCG payments and the discovery of blocking coalitions in SEP^t. Overall, this can lead to a solution time of two to three hours for each experiment because many of these optimization problems need to be solved for a single experiment. A time limit of 300 seconds for one optimization problem allowed us to conduct a larger number of experiments and get statistically significant results. For significance tests, we will provide the *p*-values of a Wilcoxon signed-rank test throughout.

5.1. Research Design for the TV Ad Market

For the generation of sample instances for the TV ad market, we could draw on data from a booking system of an industry partner. This provides us with a distribution of prices paid in this market. We will briefly summarize the main characteristics of the generated data. The distributions of all relevant random variables in the experiments can be found in Table 2.

- A typical campaign duration is from one to four weeks, averaging two weeks.
- An advertisement slot is 120 seconds long, but it can be prefilled before the auction starts because of the existing booking system (effectively reducing capacity available in a slot).
- The duration of an ad is at most 40 seconds long.
- Up to 50 different bidders (media agencies) are interested in placing ads during the average campaign time span.

Table 2 Parameters for the Experiments

	Name	Parameters $\{\mu;\sigma\}$ or $\{\lambda\}$	Distribution
I	Number of slot	336	—
J	Number of bids	50	—
K	Number of bidders	50	—
c_i	Slot duration	{60; 30}	Normal
r_i	Slot reserve price steps (in €/s) [1, 2, 5, 10, 50, 75]	{1.2}	Poisson
d_k	Ad duration	{20; 10}	Normal
β_j	Bid base price (in €/s)	{50; 25}	Normal
$w_{j,\mathrm{rel}}^{\min}$	Min $\sum$ of campaign priorities (in %)	{30; 20}	Normal
—	Correlation of priority to slot reserve price	—	Linear
—	Distribution of priorities around the priority/price value	—	Normal

- Each bidder has its own budget and target customer group that defines the slots she is interested in.
- The reserve price per second during a particular time is set by the TV station, which puts different slots into sets with different reserve prices.

Although the reported evaluation concentrates on a biweekly market (336 slots), a series of experiments with 168 slots (one week) and 504 slots (three weeks) was also performed to verify the robustness of the results presented here. The integrality gap was small in these cases as well (Table 3). We will therefore only report the detailed results for experiments with the biweekly market and 336 slots. The following parameters and distributions were used for the random variables of our experiments (see Table 2).

The normal distributions are truncated to an interval $[0; 2\mu]$. The Poisson distribution models the frequency of the six discrete reservation prices $[1, 2, 5, 10, 50, 75]$, which follows the empirical distribution that we observed in the field. The bid base price β_j can be interpreted as how much a bidder would spend at a maximum to obtain the right to reach one priority point with her ad for one second. The actual bid price for a campaign is then computed as $b_j = d_k \beta_j w_j^{\min}$. This means, for the bid price we multiply the duration of the ad, the base price of the bidder for one viewer, and the minimum reach or viewership the bidder wants to achieve. Based on the parameters of Table 2, we generated 20 scenarios used in our experiments.

Table 3 Average Integrality Gaps for 168, 336, and 504 Ad Slots After 300 and 3,600 Seconds

	Time limit (%)	
Slots	300 s	3,600 s
168	0.04	0.01
336	0.10	0.01
504	0.20	0.02

Table 4 Efficiency E, Revenue R, and Duration D in Minutes for the TV Ad Market Experiments

	TRIM			REUSE			REUSE⁺			Optimal (LPR)		
Iteration	E	R	D	E	R	D	E	R	D	E	R	D
1	0.93	0.78	115	0.95	0.65	290	0.96	0.60	580	1.00	—	—
2	0.90	0.83	105	0.93	0.70	225	0.93	0.72	890	1.00	—	—
3	0.88	0.68	95	0.89	0.67	205	0.91	0.65	480	1.00	—	—
4	0.97	0.65	95	0.97	0.66	310	0.98	0.65	410	1.00	—	—
5	0.94	0.78	105	0.96	0.63	245	0.96	0.65	460	1.00	—	—
6	0.89	0.85	85	0.92	0.73	150	0.94	0.68	450	1.00	—	—
7	0.91	0.76	75	0.92	0.66	100	0.93	0.71	340	1.00	—	—
8	0.90	0.76	100	0.92	0.70	245	0.92	0.67	490	1.00	—	—
9	0.93	0.92	105	0.97	0.71	205	0.98	0.71	390	1.00	—	—
10	0.92	0.81	105	0.95	0.66	245	0.96	0.66	560	1.00	—	—
11	0.93	0.79	90	0.96	0.65	185	0.97	0.66	400	1.00	—	—
12	0.89	0.85	110	0.94	0.68	295	0.95	0.65	480	1.00	—	—
13	0.95	0.80	85	0.95	0.75	120	0.96	0.77	330	1.00	—	—
14	0.84	0.84	90	0.92	0.68	220	0.92	0.69	590	1.00	—	—
15	0.91	0.66	95	0.91	0.64	195	0.93	0.63	680	1.00	—	—
16	0.92	0.76	90	0.92	0.77	300	0.93	0.74	610	1.00	—	—
17	0.92	0.78	100	0.93	0.73	170	0.94	0.67	470	1.00	—	—
18	0.89	0.79	95	0.92	0.64	300	0.93	0.60	430	1.00	—	—
19	0.84	0.83	60	0.88	0.74	90	0.90	0.71	370	1.00	—	—
20	0.93	0.79	95	0.95	0.72	225	0.95	0.73	410	1.00	—	—
μ	0.91	0.79	95	0.93	0.69	216	0.94	0.68	491	1.00	—	—
σ	0.03	0.06	12	0.03	0.04	65	0.02	0.04	129	0.00	—	—

5.2. Results of the TV Ad Market Experiments

We will first report the efficiency E, revenue R, and duration D in minutes in Table 4. The REUSE algorithm is able to improve the winning coalition if a coalition of bidders with higher revenue is found, and therefore the REUSE efficiency is higher than TRIM's (p-value < 0.001). However, the actual revenue R generated by the payments from the REUSE algorithm is consistently lower than that of TRIM (p-value < 0.001), despite the fact that a higher coalitional value was achieved. REUSE$^+$ describes the results of running the REUSE algorithm with a time limit of 3,600 seconds for the first winner-determination problem and 600 seconds for every subsequent optimization problem. This helps understand the impact of allowing for a longer computation time. This impact is low as the numbers in Table 4 show, illustrating that even twice the computation time has little impact on efficiency and revenue.

To understand why the overall revenue is significantly higher for TRIM in spite of the lower efficiency, we compared the ratios between the bids submitted and the resulting VCG and BPOC payments. Table 5 provides an overview of the average secondary metrics across all experiments. The ratios were all significantly higher for TRIM than for REUSE (p-value < 0.001).

In addition to these aggregate secondary metrics, we provide a more detailed summary in Figure 3, where a single frame groups an algorithm and a metric. For each algorithm and metric, we aggregated the individual (i.e., bidder-wise) ratios for all bidders in small box plots. In each of the frames of Figure 3, the light gray area of the box plot marks the interquartile range for a specific metric and the line the median for one of the 20 experiments. This provides an overview of the ratio distribution for all 50 bidders. Finally, the solid line across the box plots in each frame marks the overall mean of all ratios in all scenarios. What follows is a brief interpretation of these values.

The core/bid ratio for TRIM shows that the core prices are close to the bid prices submitted and different from the results seen in the second row in Figure 3 for the REUSE algorithm. Multiple factors influence this high ratio for TRIM: As seen in §4, the VCG payment vector sets the lower bound for the BPOC payment computation. Because of the possibility of switching to coalitions with a higher coalitional value, the VCG payment vector computed by the REUSE algorithm is always at most as high as the TRIM algorithm relative to the coalitional value of the respective winner coalition. If the coalitional value computed with TRIM $\mathrm{WD}^a(W)$ is smaller than all $\mathrm{WD}^a(W_{-k})$ for all winners k, the VCG payments are equal to the winning bid prices. In contrast, REUSE would switch the winning coalition in such a case,

Table 5 Average Ratios for Secondary Metrics

	TRIM			REUSE		
Secondary metrics	μ	σ		μ	σ	
Core/bid	0.88	0.18	▲	0.69	0.25	▼
VCG/bid	0.75	0.31	▲	0.39	0.29	▼
VCG/core	0.84	0.30	▲	0.53	0.37	▼

Note. ▼ and ▲ values are significantly lower or higher, respectively, compared with the other algorithm.

Figure 3 (Color online) Secondary Metrics of the TV Ad Market Experiments

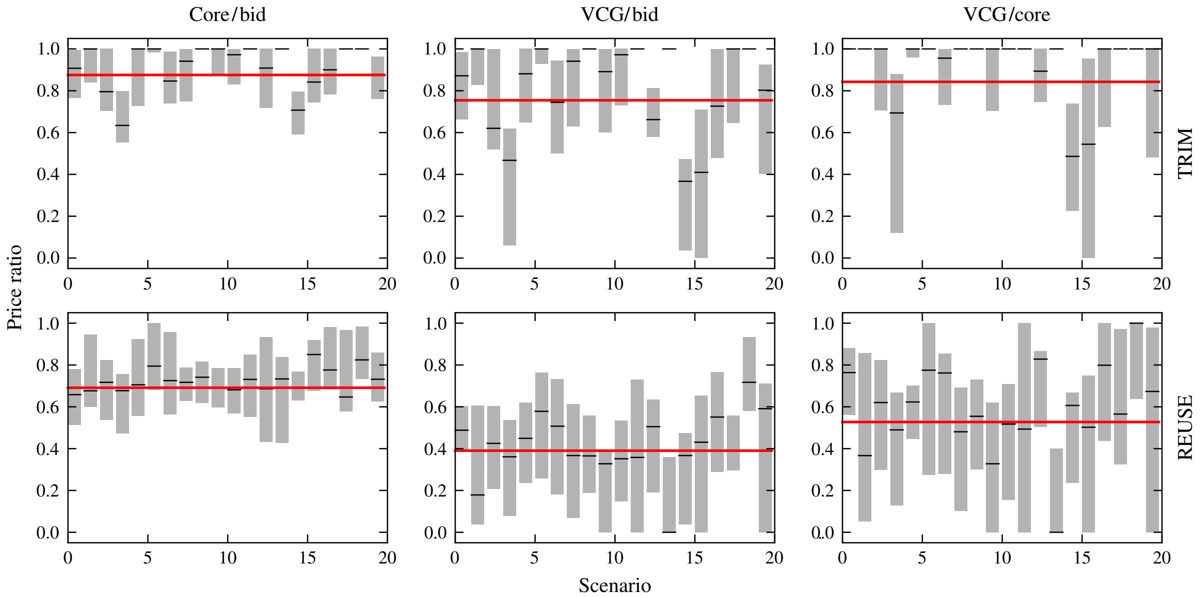

effectively increasing the difference between $\mathrm{WD}^a(W)$ and $\mathrm{WD}^a(W_{-k})$. This in turn increases the VCG discount of individual bidders and hence decreases the VCG payments for each bidder.

Similarly, if a blocking coalition is found during the BPOC payment computation in TRIM, where the coalitional value is higher than the coalitional value of the winning coalition at this point, then this coalition remains blocking even if all BPOC payments of the winners are at their bid price. In contrast, REUSE would switch the winning coalition, effectively raising the upper bound on the BPOC payments, in addition to lowering all VCG payments as described in the last paragraph. In the plots describing the VCG/bid ratio and those describing the core/bid ratio, the values for TRIM are higher than those for REUSE. In many TRIM instances, the median ratios are 1.0—i.e., for most winners the VCG and BPOC payments correspond to their bid price.

The overall duration D needed to compute using the REUSE algorithm takes significantly longer than when using TRIM (p-value < 0.001), even if each optimization run is restricted to the same limit of 300 seconds. The REUSE algorithm updates the winning coalition 3.5 times on average, whereas the TRIM algorithm will always maintain the initial coalition. Updating the winning coalition also initiates new VCG computations, which explains why the REUSE algorithm takes 63% longer than the TRIM algorithm.

5.3. Research Design for the Volume Discount Auctions

To generate bid data for the volume discount auction format, we draw on the cost function and the experiments base on Bichler et al. (2011). Being a procurement auction, all primary and secondary metric ratios (i.e., the E, R, core/bid, VCG/bid, VCG/core) have to be reversed to allow an easier comparison between the two auction formats. The multiproduct cost function $c_s(x_1, \ldots, x_{|I|})$ used by Bichler et al. draws on econometric literature and enables a systematic evaluation of markets with different economies of scale and scope. Based on these cost functions, incremental volume discount bids are generated approximating the cost curve. We could fortunately use the very same bid generation as described in Bichler et al. (2011). These bids serve as an input to the winner determination problem.

We will briefly introduce the cost function and the main parameters. There are $s \in S$ suppliers competing for a fixed quantity W_i of one or more items $i \in I$, and x_i describes the quantity produced of each item.

$$c_s(x_1, \ldots, x_{|I|}) = \sum_{i\in I} B_{i,s}\lceil x_i/z_i\rceil + \sum_{i\in I} \beta_{i,s}(x_i/\gamma_{i,s})^{\rho}.$$

The function allows us to model very different shapes with convex and concave sections. The item-specific stepwise fixed cost of supplier s for item i is denoted by $B_{i,s}$. The parameter z_i models the capacity bound, after which an additional machine or plant is needed, adding an additional $B_{i,s}$ fixed costs. Note that with the inclusion of stepwise fixed costs, the cost functions are no longer continuous. The term $\beta_{i,s}$ describes the slope of a variable cost function for product i, and the exponent ρ is the nonlinear element in the cost function, representing economies (or diseconomies) of scale. The distribution for ρ was truncated at zero such that only positive values were drawn. Parameter $\gamma_{i,s}$ moderates the economies of

Table 6 Parameters of the Cost Curve $c_s(x_1, \ldots, x_I)$

Parameter	Description	μ	σ_{item}	σ_{supplier}
$B_{i,s}$	Per item (stepwise) fixed costs	100.0	0.0	0.0
z_i	Capacity of production line (%)	100	0	0
ρ	Power of the variable cost function	0.0	0.1	1.0
$\beta_{i,s}$	Slope of the variable cost function	1,000.0	0.0	0.0
$\gamma_{i,s}$	Slope delay of the variable cost function	100.0	0.0	0.0

scale. A brief summary of all relevant variables and their distributions can be found in Table 6.

An example of an average cost function based on Table 6 can be found in Figure 4. The winner determination problem minimizes total costs of the procurement manager and is formally described in Bichler et al. (2011). This mixed-integer program allows for various allocation constraints. The only constraint used in our experiments was that a supplier could only win a certain percentage of the volume of each item (20%, 40%, 60%, and 80%), but nothing in between. This requirement can be found in the field where procurement managers try to avoid odd quantity splits.

5.4. Results of the Volume Discount Auction Experiments

In each of our experiments, 14 bidders submit volume discount bids for eight items. We chose this problem size because it is a realistic problem size, but at the same time, the exact solution can be computed within a few hours at most. No time constraints were imposed. This allows us to report the ratio of the optimal (OPT) to the near-optimal integer solution as efficiency E, in contrast to the TV ads problem where we could only use the LPR, because the optimal integer solution of the winner-determination problem proved intractable for all but unrealistically small problem instances. While the optimal allocation was indeed found for all instances, 4 out of the 40 experiments were aborted because of out of memory exceptions during the BPOC payment computation, as seen in lines 37–40 of Table 7. However, even in such cases, the integrality gap of the near-optimal solution was low, on the order of 4% at most.

Table 7 Efficiency E, Revenue R, and Duration D in Minutes for the Volume Discount Auction Experiments

	TRIM			REUSE			OPT		
Iteration	E	R	D	E	R	D	E	R	D
1	0.96	0.90	2.21	0.97	0.85	3.06	1.00	0.88	100
2	0.92	0.90	2.56	1.00	0.77	2.22	1.00	0.86	195
3	0.99	0.83	3.06	1.00	0.77	2.56	1.00	0.85	47
4	0.99	0.82	2.21	0.99	0.82	2.39	1.00	0.84	56
5	0.95	0.86	2.72	0.99	0.74	2.72	1.00	0.79	62
6	0.99	0.80	2.41	0.99	0.80	2.04	1.00	0.85	42
7	0.99	0.77	2.21	0.99	0.73	2.22	1.00	0.78	62
8	0.99	0.82	2.90	1.00	0.80	2.22	1.00	0.84	86
9	0.85	0.85	2.38	0.98	0.82	2.56	1.00	0.91	61
10	0.97	0.81	2.56	0.97	0.80	3.23	1.00	0.80	125
11	0.99	0.78	2.04	0.99	0.78	1.87	1.00	0.79	44
12	0.96	0.91	2.21	1.00	0.80	2.22	1.00	0.82	42
13	1.00	0.76	2.89	1.00	0.78	3.06	1.00	0.87	70
14	0.97	0.83	2.22	0.97	0.81	2.22	1.00	0.80	43
15	1.00	0.74	2.89	1.00	0.74	2.73	1.00	0.77	37
16	0.96	0.88	2.89	1.00	0.78	2.22	1.00	0.80	108
17	0.98	0.88	2.55	1.00	0.85	2.39	1.00	0.89	144
18	0.99	0.74	2.55	0.99	0.75	2.55	1.00	0.76	77
19	1.00	0.77	2.21	1.00	0.79	2.55	1.00	0.82	53
20	0.98	0.78	2.38	1.00	0.74	2.90	1.00	0.80	68
21	0.99	0.82	2.21	0.99	0.82	2.21	1.00	0.83	77
22	1.00	0.79	3.57	1.00	0.77	2.72	1.00	0.81	183
23	0.99	0.81	3.57	1.00	0.73	2.55	1.00	0.83	53
24	0.99	0.79	2.56	0.99	0.79	2.39	1.00	0.82	52
25	1.00	0.77	2.90	1.00	0.73	2.55	1.00	0.81	46
26	0.97	0.90	2.56	0.99	0.86	2.38	1.00	0.86	37
27	0.99	0.80	2.56	0.99	0.82	2.38	1.00	0.82	41
28	0.94	0.85	3.06	0.99	0.80	3.23	1.00	0.81	89
29	0.98	0.78	2.21	0.99	0.77	2.21	1.00	0.76	73
30	0.98	0.91	3.23	0.99	0.85	2.38	1.00	0.85	75
31	0.90	0.90	2.56	0.99	0.85	2.56	1.00	0.83	57
32	0.99	0.75	2.38	1.00	0.75	2.73	1.00	0.78	35
33	0.95	0.86	2.72	1.00	0.78	2.21	1.00	0.82	71
34	0.98	0.84	2.38	0.99	0.81	2.89	1.00	0.82	41
35	0.97	0.76	2.89	1.00	0.71	3.07	1.00	0.77	68
36	0.99	0.78	2.55	1.00	0.73	2.38	1.00	0.79	50
37	0.99	0.76	2.62	0.99	0.77	2.63	1.00	—	—
38	0.99	0.77	2.63	0.99	0.81	2.80	1.00	—	—
39	0.96	0.85	2.93	0.99	0.79	3.12	1.00	—	—
40	0.96	0.88	2.97	0.98	0.79	2.95	1.00	—	—
μ	0.97	0.82	2.63	0.99	0.79	2.56	1.00	0.82	71
σ	0.03	0.05	0.36	0.01	0.04	0.34	0.00	0.04	38

As in the previous section, the experiments with TV ads, the efficiency of REUSE is significantly higher than for the TRIM algorithm (p-value < 0.001). Additionally, the efficiency of TRIM and REUSE with a time limit of 300 seconds per optimization

Figure 4 (Color online) Average Costs for One Unit Produced and Increasing Quantity Sold

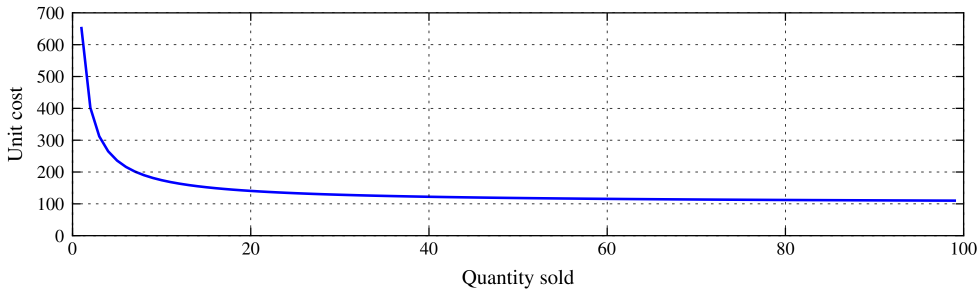

Table 8 Average Ratios for Secondary Metrics

	TRIM			REUSE			OPT	
Secondary metrics	μ	σ		μ	σ		μ	σ
Bid/core	0.85	0.15	▲	0.81	0.15	▼	0.82	0.11
Bid/vcg	0.80	0.18	▲	0.72	0.16	▼	0.82	0.11
Core/vcg	0.93	0.13	▲	0.90	0.14	▼	1.00	0.00

Note. ▼ and ▲ values are significantly lower or higher, respectively, compared with the competing BPOC algorithm.

are significantly lower than the optimal efficiency (p-value < 0.001). Also the results on revenue are in line with the TV ad market experiments: The revenue achieved with REUSE is significantly lower than with TRIM.

The secondary metrics are illustrated in Table 8 and Figure 5 similar to what we have reported for the TV ad market experiments in §5.1. However, now we are also able to compare the ratios of payment vectors using near-optimal solutions with those if the problems are solved optimally. Note that in the optimal solution the core payment vector coincides with the VCG payments. We conjecture that this is because we did not have economies of scope in our cost functions. Table 8 shows that again the ratios for TRIM are higher than those for REUSE. The difference is again significant (p-value < 0.001 for bid/core and bid/VCG, p-value of 0.02 for core/VCG), but not as high as in the TV ad market experiments. This is because the integrality gap was lower for the instance sizes chosen.

A comparison with the price vectors based on the optimal solution, OPT, shows that the core payments achieved with TRIM are indeed very high. The difference of bid/core between TRIM and OPT was significant (p-value < 0.001). The difference between the bid/core ratios of OPT and reuse REUSE was not significant (p-value of 0.4). The bid/VCG ratios of TRIM and OPT were not significantly different (p-value of 0.965). However, the core/VCG ratio was significantly higher for OPT than for TRIM and REUSE (p-value $<$ 0.03). This is because there was no difference in OPT, i.e., both payment vectors coincided, while there was a difference with the near-optional winner determination in TRIM and REUSE.

Figure 5 again provides a more detailed view of the secondary metrics. As we have seen in the TV ad mar-

Figure 5 (Color online) Secondary Metrics of the Volume Discount Auction Experiments

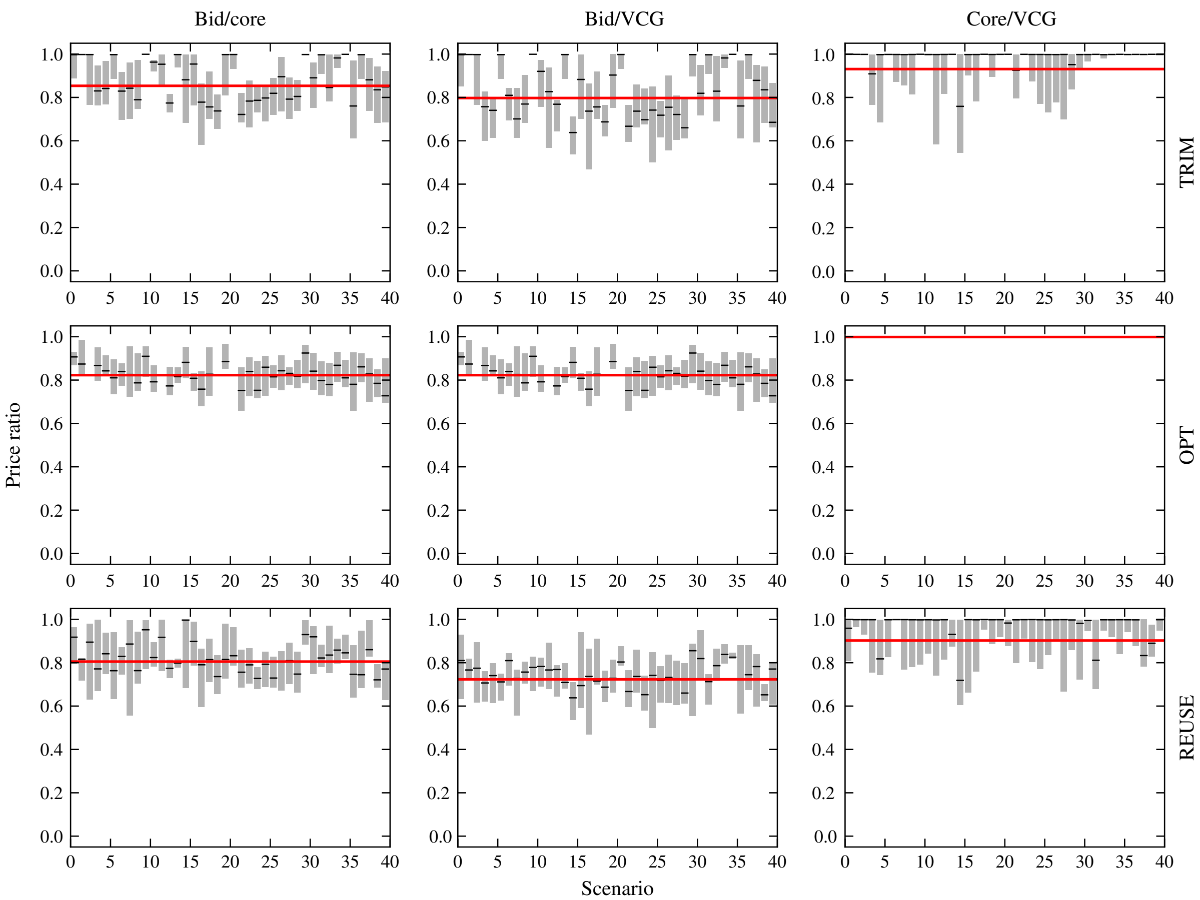

Figure 6 (Color online) Allocative Distance and Relative Payoff Difference Between TRIM or REUSE and OPT for All Bidders in a Single Auction Instance

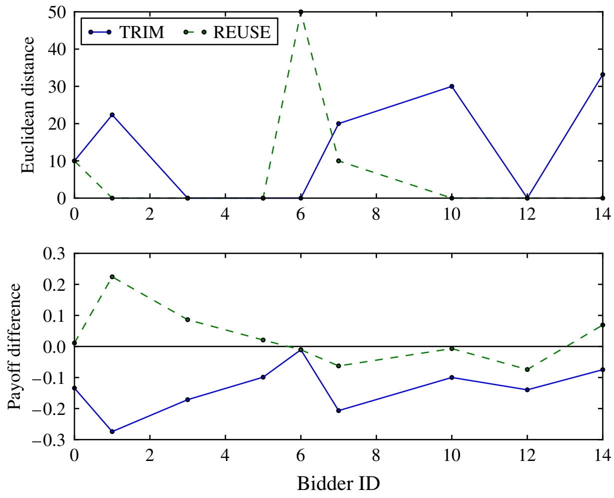

ket experiments, bidders' final BPOC and VCG payments are nearer to the bid price than necessary on an aggregate level, often as low as their bid prices for the volume discount auctions if the TRIM algorithm is used.

The availability of optimal solutions in these experiments allows us to compare how the allocation and the payments for individual bidders would differ in OPT, TRIM, and REUSE. These differences are described in Figure 6 for a specific auction. The upper part of the figure describes how the final allocation of TRIM or REUSE differs from the one in OPT by treating each bidder's allocation as a vector and computing the Euclidean distance between the different allocations. For example, suppose there were just two different items, A and B, and two different allocations X_{i1} and X_{i2} describing the quantity of each item that the bidder i has to provide. Then, for $X_{i1} = (10, 10)$ and $X_{i2} = (0, 15)$, the Euclidean distance is

$$\sqrt{((10-0)^2 + (10-15)^2)} \approx 11.2.$$

The lower part of Figure 6 shows how the payoff in TRIM or REUSE differs from the one in OPT. To gain insight about the magnitude of the payoff difference, the values are normalized against the average payment the auctioneer has to provide to the winners. In this example, many bidders have a lower payoff in TRIM because the payments in this reverse auction are lower in TRIM, which is indicated by a high bid/core ratio. For REUSE, the individual payoff can even surpass the payoff obtained when computations are solved to optimality. Also, a change in the allocative efficiency does not necessarily cause drastic changes in every bidder's payoff, as we see a payoff difference close to zero for bidder 6, for example. The

Table 9 Summary of the Primary Metrics Comparing TRIM and REUSE

		TRIM			REUSE			Baseline
		μ			μ			μ
TV ad market								LPR
Efficiency	E	0.914	▼	∘	0.928	▲	∘	1.000
Revenue	R	0.788	▲	—	0.680	▼	—	—
Run time (minutes)	D	95	▼	—	222	▲	—	—
Volume discount auction								OPT
Efficiency	E	0.988	▼	∘	0.994	▲	∘	1.000
Revenue	R	0.807	▲		0.786	▼	∘	0.818
Run time (minutes)	D	3		∘	3		∘	54

Note. ▼ and ▲ indicate significant difference compared with the competing BPOC algorithm; ∘ indicates significant difference from the baseline.

computational hardness of these problems is such that these differences from the optimal benchmark can in some cases not be avoided, but the overall efficiency of the near-optimal solutions is still very high in all experiments.

Table 9 provides a summary of the primary metrics to compare TRIM and REUSE, showing that the main results are the same in the TV ad market experiments and the volume discount auction experiments. Overall, if speed and revenue are primary concerns, then TRIM may be the right approach. In other situations, where incentives for truthful bidding and high efficiency are a concern, the REUSE algorithm is preferable.

6. Conclusions

The design of large-scale markets where bidders have complex preferences has been given little attention in the literature as of yet. In several countries, regulators sell dozens or hundreds of licenses to telecom companies. The incentive auctions in the US are another example where complex bidder preferences and allocation constraints lead to computationally hard allocation problems. Similar examples can be found in many other domains including the sale of TV ads to media agencies or multi-item and multiunit industrial procurement auctions. Much research in market design has focused on ascending combinatorial auctions with a fully expressive XOR bid language, and such designs have recently been used for selling spectrum (Cramton 2013, Bichler et al. 2013) and in logistics and procurement (Bichler et al. 2006). Such designs do not scale to large markets because of the exponential growth in the number of package bids that can be submitted.

We describe an auction design framework using compact bid languages and payment rules that incentivize truthful bidding. In markets where bidders have independent private values, which is the standard assumption in auction theory, this can yield highly efficient allocations. Compact bid languages can often draw on domain specifics and allow bidders to describe their preferences with a low number of parameters that they have to specify, as the TV ad market and the volume discount auctions in this paper illustrate. Commercial off-the-shelf mixed-integer programming solvers can now solve large and realistic instances of such problems to near optimality on standard hardware, which allows us to use such bid languages in real-world markets. Such compact bid languages, however, defy the ask pricing rules typically used in ascending combinatorial auctions (Scheffel et al. 2011), but they can easily be used in sealed-bid auctions.

In sealed-bid auctions, second price rules such as VCG or BPOC payment rules can be used to provide incentives for truthful bidding. In many markets, auctioneers would prefer core pricing to VCG mechanisms to avoid noncore outcomes where the bids of losing bidders are higher than the payments of the winners. With the introduction of core-selecting auctions for spectrum licenses in recent years, stakeholders have developed software to determine winners and core prices based on the use of integer programming to solve a series of winner-determination problems. Extending the use of this software to larger and more complex markets (such as the TV ads and procurement contexts we address here) cannot be accomplished by merely specifying time limits or optimality-gap thresholds to the solver engine because it could for the more simple case of a single optimization problem. Doing so would often result in an infeasible pricing problem. This general problem exists for all larger markets with near-optimal winner determination.

We compared two potential algorithms for dealing with these infeasibilities, finding one faster and higher revenue method (for a fixed set of bids) and one slower but more efficient method. Our results show that the former TRIM algorithm may be suited to a fast-clearing market in which speculation to lower bids is offset by uncertainty about the competition. For other applications, such as government spectrum auctions the goal of public efficiency might outweigh the computational costs and suggest an advantage for the latter REUSE algorithm.

Further study may improve the application of core-selecting auction algorithms to large and complex markets like the TV ad and volume-discount markets, but we have provided the first steps to the extension of the core-selecting auction paradigm beyond provably optimal winner-determination settings. The paper shows that the overall auction design framework using compact bid languages and second-price payment rules provides a computationally feasible approach to achieve high efficiency in large-scale markets with dozens or hundreds of items.

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/mnsc.2014.2076.

Appendix

Proof of Theorem 1. The reduction is from the decision version of the strongly $\mathcal{NP}$-hard multiple knapsack problem: given a set of n items and a set of m knapsacks ($m \leq n$), with a profit b_j and a weight d_j for each item j, and a capacity c_i of each knapsack i, can you select m disjoint subsets of items such that the total profit of the selected items exceeds a given target profit T, with each subset assigned to a knapsack, and the total weight of any subset not exceeding the capacity of the assigned knapsack?

To see that this problem is a special instance of the WD problem, let the minimum price per unit $r_i = 0$; also let each bidder only bid with a single bid (item) j with a bid price of b_j, and each bidder's priority vector $W_k = \{1, \ldots, 1\}$ with a $w_j^{\min} = 1$. This means, he wants his ad with a length (weight) d_j to be assigned to one out of all slots (knapsacks) i with a duration (capacity) c_i. The multiple knapsack decision problem can be answered affirmatively if and only if this specific WD instance has an optimal objective value greater than or equal to T. The problem is in $\mathcal{NP}$ because it is straightforward to check for a given solution, whether it is correct. □

List of Symbols

b_j	bid amount for bid j
b_k^*	bid amount of bidder k's winning bid
C	coalition of bidders. $C \subseteq K$
C^t	most violated coalition relative to the current payment vector at iteration t
I	set of goods offered/requested.
i	index of a single good i
J	set of bids containing all bid indexes
j	index of a single bid j
J_k	set of bids containing all bid indexes j by a bidder k
K	set of bidders
k	index of a single bidder k
p^t	payment vector at iteration t
p_k^t	payment amount for bidder k at iteration t
p_k	shorthand notation for p_k^t with the currently highest t
p_k^{vcg}	VCG payment amount for bidder k
$w_j^{\min}$	minimum sum of weight values to justify a monetary bid b_j
W_k	weight vector for bidder k
w_{ik}	weight value for time slot i and bidder k
x_{ij}	decision variable indicating that bid j is assigned to time slot i
y_j	decision variable indicating that the bid j is accepted
$z(p^t)$	computed value of SEP^t
LPR	solution of the linear programming relaxation of the winner determination problem
OPT	optimal solution of the winner determination problem
SEP^t	core separation problem

WD winner determination problem, optimally solved
WD^a winner determination problem, approximately solved

References

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

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

Ausubel L, Milgrom P (2002) Ascending auctions with package bidding. *Frontiers Theoret. Econom.* 1(1):1–42.

Ausubel L, Milgrom P (2006) The lovely but lonely Vickrey auction. Cramton P, Shoham Y, Steinberg R, eds. *Combinatorial Auctions* (MIT Press, Cambridge, MA), 57–95.

Bapna R, Das S, Garfinkel R, Stallaert J (2008) A market design for grid computing. *INFORMS J. Comput.* 20(1):100–111.

Bichler M, Gupta A, Ketter W (2010) Designing smart markets. *Inform. Systems Res.* 21(4):688–699.

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

Bichler M, Shabalin P, Wolf J (2013) Do core-selecting combinatorial clock auctions always lead to high efficiency: An experimental analysis of spectrum auction designs. *Experiment. Econom.* 16(4):511–545.

Bichler M, Davenport A, Hohner G, Kalagnanam J (2006) Industrial procurement auctions. Cramton P, Shoham Y, Steinberg R, eds. *Combinatorial Auctions* (MIT Press, Cambridge, MA), 1116–1147.

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

Bichler M, Schneider S, Guler K, Sayal M (2011) Compact bidding languages and supplier selection for markets with economies of scale and scope. *Eur. J. Oper. Res.* 214:67–77.

Boutilier C, Hoos HH (2001) Bidding languages for combinatorial auctions. *Proc. 17th Internat. Joint Conf. Artificial Intelligence* (Morgan Kaufmann, San Francisco), 1211–1217.

Caplice C (2007) Electronic markets for truckload transportation. *Production Oper. Management* 16(4):423–436.

Chekuri C, Khanna S (2006) A polynomial time approximation scheme for the multiple knapsack problem. *SIAM J. Comput.* 35:713–728.

Cramton P (2013) Spectrum auction design. *Rev. Indust. Organ.* 42(2):161–190.

Cramton P, Shoham Y, Steinberg R (2006) Introduction to combinatorial auctions. Cramton P, Shoham Y, Steinberg R, eds. *Combinatorial Auctions* (MIT Press, Cambridge, MA), 29–55.

Daskalakis C, Goldberg PW, Papadimitriou CH (2009) The complexity of computing a Nash equilibrium. *SIAM J. Comput.* 39(1):195–259.

Davenport A, Kalagnanam J (2000) Price negotiations for procurement of direct inputs. *IMA "Hot Topics" Workshop: Math. Internet: E-Auction and Markets, Minneapolis, MN*, Vol. 127, 27–44.

Day RW, Cramton P (2012) The quadratic core-selecting payment rule for combinatorial auctions. *Oper. Res.* 60(3):588–603.

Day RW, Milgrom P (2007) Core-selecting package auctions. *Internat. J. Game Theory* 36:393–407.

Day RW, Raghavan S (2007) Fair payments for efficient allocations in public sector combinatorial auctions. *Management Sci.* 53(9):1389–1406.

Dobzinski S, Nisan N, Schapira M (2012) Truthful randomized mechanisms for combinatorial auctions. *J. Comput. System Sci.* 78(1):15–25.

Gallien J, Wein LM (2005) A smart market for industrial procurement with capacity constraints. *Management Sci.* 51(1):76–91.

Garey MR, Johnson DS (1979) *Computers and Intractability: A Guide to the Theory of NP-Completeness* (Springer, New York).

Goeree JK, Lien Y (2015) On the impossibility of core-selecting auctions. *Theoret. Econom.* Forthcoming.

Goossens DR, Maas AJT, Spieksma F, van de Klundert JJ (2007) Exact algorithms for procurement problems under a total quantity discount structure. *Eur. J. Oper. Res.* 178(2):603–626.

Government of Canada (2009) 700 MHz spectrum auction-process and results. Date modified February 19, 2014, http://news.gc.ca/web/article-en.do?nid=816869.

Green J, Laffont J (1979) *Incentives in Public Decision Making* (North-Holland, Amsterdam).

Guo Z, Koehler GJ, Whinston AB (2007) A market-based optimization algorithm for distributed systems. *Management Sci.* 53(8):1345–1358.

Halldorsson MM, Kratochvil J, Telle JA (2000) Independent sets with domination constraints. *Discrete Appl. Math.* 99(1–3):39–54.

Hass C, Bichler M, Guler K (2013) Optimization-based decision support for scenario analysis in electronic sourcing markets with volume discounts. *Electronic Commerce Res. Appl.* 12(3):152–165.

Holzman R, Kfir-Dahav N, Monderer D, Tennenholtz M (2004) Bundling equilibrium in combinatorial auctions. *Games Econom. Behav.* 47(1):104–123.

Hoos HH, Boutilier C (2000) Solving combinatorial auctions using stochastic local search. *Proc. 17th Natl. Conf. Artificial Intelligence and 12th Conf. Innovative Appl. Artificial Intelligence* (AAAI Press, Palo Alto, CA), 22–29.

Lavi R (2007) Computationally efficient approximation mechanisms. *Algorithmic Game Theory* 301–329.

Lavi R, Swamy C (2011) Truthful and near-optimal mechanism design via linear programming. *J. ACM (JACM)* 58(6):25.

Lehmann D, Mueller R, Sandholm T (2006) The winner determination problem. Cramton P, Shoham Y, Steinberg R, eds. *Combinatorial Auctions* (MIT Press, Cambridge, MA), 555–596.

Meeus L, Verhaegen K, Belmans R (2009) Block order restrictions in combinatorial electric energy auctions. *Eur. J. Oper. Res.* 196(3): 1202–1206.

Milgrom P (2010) Simplified mechanisms with applications to sponsored search and package auctions. *Games Econom. Behav.* 70(1):62–70.

Parkes DC, Ungar LH (2000) Iterative combinatorial auctions: Theory and practice. *Proc. 17th Natl. Conf. Artificial Intelligence and 12th Conf. Innovative Appl. Artificial Intelligence* (AAAI Press, Palo Alto, CA), 74–81.

Sandholm T (2007) Expressive commerce and its application to sourcing: How we conducted $35 billion of generalized combinatorial auctions. *AI Magazine* 28(3):45–58.

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

Schneider S, Shabalin P, Bichler M (2010) On the robustness of nonlinear personalized price combinatorial auctions. *Eur. J. Oper. Res.* 206(1):248–259.

Shoham Y, Leyton-Brown K (2008) *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations* (Cambridge University Press, New York).

Xia M, J KG, B WA (2004) Pricing combinatorial auctions. *Eur. J. Oper. Res.* 154(1):251–270.

Synergistic Valuations and Efficiency in Spectrum Auctions

Peer-reviewed Journal Paper

Title: Synergistic Valuations and Efficiency in Spectrum Auctions

Authors: A. Goetzendorff, M. Bichler, J. Goeree

In: Telecommunications Policy (2017)

Abstract: In spectrum auctions, bidders typically have synergistic values for combinations of licenses. This has been the key argument for the use of combinatorial auctions in the recent years. Considering synergistic valuations turns the allocation problem into a computationally hard optimization problem that generally cannot be approximated to a constant factor in polynomial time. Ascending auction designs such as the Simultaneous Multiple Round Auction (SMRA) and the single-stage or two-stage Combinatorial Clock Auction (CCA) can be seen as simple heuristic algorithms to solve this problem. Such heuristics do not necessarily compute the optimal solution, even if bidders are truthful. We study the average efficiency loss that can be attributed to the simplicity of the auction algorithm with different levels of synergies. Our simulations are based on realistic instances of bidder valuations we inferred from bid data from the 2014 Canadian 700MHz auction. The goal of the paper is not to reproduce the results of the Canadian auction but rather to perform "out-of-sample" counterfactuals comparing SMRA and CCA under different synergy conditions when bidders maximize payoff in each round. With "linear" synergies, a bidder's marginal value for a license grows linearly with the total number of licenses won, while with the "extreme national" synergies, this marginal value is independent of the number of licenses won unless the bidder wins *all* licenses in a national package. We find that with the extreme national synergy model, the CCA is indeed more efficient than SMRA. However, for the more realistic case of linear synergies, SMRA outperforms various versions of CCA that have been implemented in the field including the one used in the Canadian 700MHz auction. Overall, the efficiency loss of all ascending auction algorithms is small even with high synergies, which is remarkable given the simplicity of the algorithms.

Contribution of thesis author: Methodology, results, implementation, mathematical model, project and paper management

Synergistic Valuations and Efficiency in Spectrum Auctions

Abstract

In spectrum auctions, bidders typically have synergistic values for combinations of licenses. This has been the key argument for the use of combinatorial auctions in the recent years. Considering synergistic valuations turns the allocation problem into a computationally hard optimization problem that generally cannot be approximated to a constant factor in polynomial time. Ascending auction designs such as the Simultaneous Multiple Round Auction (SMRA) and the single-stage or two-stage Combinatorial Clock Auction (CCA) can be seen as simple heuristic algorithms to solve this problem. Such heuristics do not necessarily compute the optimal solution, even if bidders are truthful. We study the average efficiency loss that can be attributed to the simplicity of the auction algorithm with different levels of synergies. Our simulations are based on realistic instances of bidder valuations we inferred from bid data from the 2014 Canadian 700MHz auction. The goal of the paper is not to reproduce the results of the Canadian auction but rather to perform "out-of-sample" counterfactuals comparing SMRA and CCA under different synergy conditions when bidders maximize payoff in each round. With "linear" synergies, a bidder's marginal value for a license grows linearly with the total number of licenses won, while with the "extreme national" synergies, this marginal value is independent of the number of licenses won unless the bidder wins *all* licenses in a national package. We find that with the extreme national synergy model, the CCA is indeed more efficient than SMRA. However, for the more realistic case of linear synergies, SMRA outperforms various versions of CCA that have been implemented in the field including the one used in the Canadian 700MHz auction. Overall, the efficiency loss of all ascending auction algorithms is small even with high synergies, which is remarkable given the simplicity of the algorithms.

Key words: Market design, Spectrum auctions, Combinatorial auctions, Simulation experiments

1. Introduction

Radio spectrum is a key resource in the digital economy. Recognizing its enormous value for society, the US Federal Communication Commission (FCC) decided in 1994 to replace their bureaucratic process ("beauty contest") with a market-based approach to assign spectrum: the Simultaneous Multiple Round Auction (SMRA).

Since then the SMRA has successfully been used by many regulators and has generated hundreds of billions of dollars worldwide. Despite this success, the SMRA has also led to a number of strategic problems for bidders. Telecom operators typically have preferences for certain packages of licenses. In the SMRA this leads to the so-called *exposure problem*: bidders who compete aggressively for a certain package risk ending up with only a subset, possibly paying more than what this subset is worth to them. The inability to express preferences for packages directly adds strategic complexity for bidders and is a source of inefficiency in the SMRA. The exposure problem that is inherent to item-by-item competition has stirred interest in combinatorial auctions, which allow bidders to submit preferences for combinations or packages directly. The design of combinatorial spectrum auctions has drawn significant attention from researchers from various fields including economics, game theory, operations research, and computer science, see e.g. Cramton et al. (2006).

Combinatorial auctions have also drawn interest from regulators. For their 2008 700MHz auction, the FCC decided to augment the SMRA with the possibility to bid on a national package. This simple combinatorial auction was based on the Hierarchical Package Bidding (HPB) format that had been designed and tested by Goeree and Holt (2010). In the same year, the British regulator Ofcom pioneered the Combinatorial Clock Auction (Cramton, 2008) and regulators world-wide have since followed their example by adopting different versions of the CCA. The single-stage CCA (Porter et al., 2003) is a simple ascending format where bidders can submit multiple package bids in each round and prices are increased on items for which there is excess demand. Variants of single-stage CCA have been used in Romania in 2012 and in Denmark in 2016. The single-stage CCA creates incentives for demand reduction (Ausubel et al., 2014), a problem which the two-stage CCA tries to address. The two-stage CCA allows for only a single package bid in each round and adds a sealed-bid "shoot out" phase and a core-selecting payment rule (Cramton, 2013). Both phases are governed by a revealed-preference activity rule. The two-stage CCA has been used in many countries including Austria, Australia, Canada, Ireland, the Netherlands, Slovakia, Switzerland, and the UK (Mochon and Saez, 2017; Cave and Nicholls, 2017; Bichler and Goeree, 2017).

One takeaway message from the recent literature is that the design of spectrum auctions is still a topic of intense debate. Another is that it requires different approaches – theory, laboratory experiments, and simulations – to understand the properties of alternative formats. Mechanism design theory has identified the unique efficient auction in which bidding truthfully is a (weakly) dominant strategy so that bidders do not require information about their rivals' valuations. Despite its desirable features, the Vickrey-Clarke-Groves (VCG) mechanism is rarely used in the field for various practical reasons (Ausubel and Milgrom, 2006). Bayesian-Nash implementation allows for a broader class of auction formats but imposes a strong common-prior assumption. Moreover, recent game-theoretical models of spectrum auction formats typically make simplifying assumptions about bidders' valuations (Goeree and Lien, 2014), focus on small environments with a few items and players (Levin and Skrzypacz, 2016), or assume complete information to highlight strategic problems (Janssen and Karamychev, 2016). Laboratory experiments provide valuable insights as well, e.g. Goeree and Holt (2010), but the size of the markets that can be organized in an economic experiment is typically limited. And, like in theoretical analyses, experimental designs often ignore or simplify complicated institutional details (e.g. activity rules or spectrum caps).

In contrast, simulations allow one to analyze realistic market sizes and to take institutional details into account. As such they can provide complementary insights (Consiglio and Russino, 2007). Interestingly, there are no published simulation studies about spectrum auction markets that we are aware of, even though simulations are regularly used by consultants and telecoms to explore different bidding strategies.

1.1. Auctions as Algorithms

Spectrum auctions can be seen as large games with many bidders, licenses, and additional rules such as spectrum caps and activity rules. For example, the 2014 Canadian 700MHz auction allowed bidders to bid on 18 packages in 14 regions leading to 18^{14} possible packages. In large games like this, bidders need a lot of information about competitors to bid strategically, and one might argue that strategic manipulation is less of a concern. But even if we ignore strategic bidding, it is far from obvious that auctions yield efficient outcomes. It is well-known that the allocation problem in a combinatorial auction where bidders have preferences for combinations of licenses is an NP-hard optimization problem (Cramton et al., 2006). The SMRA and different versions of the CCA can be interpreted as algorithms to solve this problem, and it is important to understand the approximation ratios of these algorithms (Domowitz and Wang, 1994). It is unclear whether to expect efficient outcomes even when bidders bid straightforwardly in each round of the auction.

Theoretical results on the allocation problem in combinatorial auctions are not encouraging. There is no polynomial-time algorithm that guarantees an approximate solution to the winner determination problem within a factor of $l^{1-\varepsilon}$ from the optimal allocation, where l is the number of submitted bids and ε a small number (Pekec and Rothkopf, 2003). The problem is APX-hard and the worst-case approximation ratio of any polynomial-time algorithm for the allocation problem in combinatorial auctions is in $O(\sqrt{m})$, where m is the number of objects to be sold. In large spectrum auctions with many licenses such as the Canadian auction in 2014, this lower bound on efficiency is obviously very low[1] and provides no practical guidance. There are also results on the worst-case efficiency of the single-stage CCA with bidders who truthfully reveal their preferences (i.e., bid straightforward) in each round. This can also be seen as an algorithm to solve the allocation problem. Unfortunately, the

[1] $\sqrt{18 * 14} = 15.87$, such that the worst-case approximation of any polynomial-time algorithm might be $OPT/15.87$, where OPT is the optimal solution to the allocation problem.

worst-case approximation ratio can be $2/(m+1)$ with m being the number of licenses (Bichler et al., 2013b).

Worst-case bounds might be too pessimistic, and it is interesting to understand the average-case approximation ratio to the fully efficient solution of different auction types based on realistic problem instances. Numerical experiments are widely used in operations research and computer science to analyze the average-case solution quality of an algorithm, and they help to understand aspects of the algorithm that do not lend themselves to theoretical analysis or lab experiments. In particular, we want to study the average-case efficiency of different auction algorithms under the assumption of straightforward bidding. This provides an estimate of the efficiency that can be achieved by the auction algorithms that are commonly employed in the field *for realistic market sizes and considering all details of the auction design.*

Another important element of realism in our simulations stems from the fact that we estimate bidder value models from bids observed in the Canadian 700MHz auction. Unlike other regulators, the Canadian regulator, Industry Canada, revealed detailed bid data.[2] We use the supplementary bids from the Canadian 700MHz auction to estimate individual license values for each bidder.[3] Our goal is not to provide the most precise estimates of bidders' private valuations and then reproduce the outcomes of the Canadian auction. This would lead to over-fitting problems in the counterfactuals we have in mind. Our goal instead is to robustly infer reasonable instances of bidder valuations and then compare the performance of different auction formats under various "out-of-sample" synergy conditions.

In particular, we compare efficiency of the SMRA to that of the single-stage and two-stage CCA using extensive numerical experiments. The experiments employ a simulation framework that implements the exact rules of the 2014 Canadian auction, i.e. band plan, spectrum caps, regional interests of bidders and rules of the two-stage CCA. In addition, we implemented the rules of SMRA and the single-stage CCA as they were used in other countries. Since the Canadian auction shares similarities with other large spectrum auctions, e.g., in Australia, India, and the US, the results of our study are of interest beyond the Canadian market.

1.2. Outline

The paper is organized as follows. The next section provides details about the SMRA and CCA formats. Section 3 covers the experimental design, including details about the Canadian 700MHz auction. The results of the simulation experiments can be found in Section 4. Section 5 concludes and the Appendices provide further details about the value model we estimate (Appendix A), the optimization algorithms used

[2] Besides Industry Canada, the UK's Ofcom is the only regulator that disclosed bids data for their CCA. However, the auctions in the UK have only national licenses and there are too few bidders and objects to conduct a simulation study as in this paper.

[3] The bid data used in this paper are available for replication studies.

in the simulations (Appendix B), our modeling assumptions (Appendix C), and the results (Appendix D).

2. The Auctions

In this section, we briefly summarize the SMRA and different versions of combinatorial clock auctions and provide references to the details of the implementations, which follow the very auction rules used in Canada or other countries.

2.1. The Simultaneous Multiple Round Auction (SMRA)

The SMRA is an extension of the English auction to more than one license. All licenses are sold at the same time, each with a price associated with it, and the bidders can bid on any one of the licenses. The auction proceeds in rounds, which is a specific period of time in which all bidders can submit bids. After the round is closed, the auctioneer discloses who is winning and the prices of each license, which coincide with the highest bid submitted on each license. There are differences in the level of information revealed about other bidders' bids. Sometimes all bids are revealed after each round, sometimes only prices of the currently winning bids are published.

The bidding continues until no bidder is willing to raise the bid on any of the licenses any more. In other words, if in one round no new bids are placed, the bidders receive the spectrum for which they hold the highest active bid, then the auction ends with each bidder winning the licenses on which he has the high bid, and paying its bid for any license won.

SMRA uses simple *activity rules* which enforce bidder activity throughout the auction. Monotonicity rules are regularly used, where bidders cannot bid on more licenses in later rounds. This forces bidders to be active from the start. Typically, bidders get eligibility points assigned at the start of the auction, which define the number of licenses they are allowed to bid on maximally. If the number of licenses they win in a round and the new bids they submit require less eligibility points than in the last round, then they risk losing points, which limits the number of items they can bid on in future rounds.

Apart from the activity rules, there are typically additional rules that matter. Auctioneers set *reserve prices* for each license, which describe prices below which an license will not be sold. They need to define *bid increments* and how bid increments might change throughout the auction. A bid increment is the minimum amount by which a bidder needs to increase his bid beyond the ask price in the next round. Sometimes, auctioneers allow for *bid withdrawals* and sometimes bidders get bid waivers, which allow bidders not to bid in a round without losing eligibility points. Finally, auctioneers often set *bidding floors and caps*, which are limits on how much a winner in the auction needs to win at a minimum and how much he can win at most. These rules should avoid unwanted outcomes such as a monopoly after the auction or a winner who wins so little spectrum that it is not sufficient for a viable business.

The auction format is popular because it is easy to implement and the rules are simple. If the valuations of all bidders were additive, the properties of a single-object ascending auction carry over. Unfortunately, this is rarely the case and bidders have often synergies for specific licenses in a package or their preferences are substitutes. Only if bidders have substitutes preferences and bid straightforwardly, then the SMRA terminates at a Walrasian equilibrium, i.e., an equilibrium with linear prices (Milgrom, 2000). *Straightforward bidding* means that a bidder bids on the bundles of licenses, which together maximize the payoff at the current ask prices in each round. Milgrom (2000) also showed that with at least three bidders and at least one non-substitutes valuation (for example super-additive valuations for a package if licenses) no Walrasian equilibrium exists.

We assume a simple straightforward bidding strategy in each round where bidders submit bids on their payoff maximizing package. The exposure problem is a key strategic challenge in the SMRA. In our simulations we assume that bidders take different levels of exposure. Either they only bid up to the additive values of a package and ignore the synergies or they exceed the additive value of the licenses in a package and take some exposure risk. This is a treatment variable in the experiments. More details can be found in Section 3.5 and Appendix C.

2.2. Alternative Versions of the Combinatorial Clock Auction

There are various versions of combinatorial clock auctions that have been implemented in the field: the *two-stage combinatorial clock auction (CCA)* that was used in e.g. the Canadian 700MHz auction, the *two-stage combinatorial clock auction with base and OR bids in the second stage (CCA+)* as it was used in the Canadian 2.5GHz spectrum auction, and the *single-stage combinatorial clock auction (SCCA)*, which was recently used in Romania and Denmark.

2.2.1. The Two-Stage Combinatorial Clock Auction (CCA)

The two-stage combinatorial clock auction was introduced by Ausubel et al. (2006). In contrast to SMRA, the auction avoids the exposure problem by allowing for bundle bids. Maldoom (2007) describes a version as it has been used in spectrum auctions across Europe. In a two-stage combinatorial clock auction, bids for bundles of licenses are made throughout a number of sequential, open rounds (the primary bid rounds or clock phase) and then a final sealed-bid round (the supplementary bids round). In the primary bid rounds the auctioneer announces prices and bidders state their demand at the current price levels. Prices of licences with excess demand are increased by a bid increment until there is no excess demand anymore. Jump bidding is not possible. In the primary bid rounds, bidders can only submit a bid on one bundle per round. This rule is different to the initial proposal by Ausubel et al. (2006). If bidders bid straightforward on their payoff maximizing bundle in each round and all goods get sold after the clock phase, allocation and prices would be in competitive equilibrium. It might well be that there is excess supply after the clock phase, however. The sealed-bid supplementary bids phase and a Vickrey-closest core-selecting payment rule try to induce truthful bidding and avoid incentives for demand reduction. This is because core payments in Day and Cramton (2012) are computed such that a losing bid of a winner does not increase his payment for his winning bid. The winner determination after the supplementary bids round considers all bids, which have been submitted in the primary bid rounds and the supplementary bids round and selects the revenue maximizing allocation. The bids by a single bidder are mutually exclusive (i.e., the CCA uses an XOR bidding language).

Activity rules should provide incentives for bidders to reveal their preferences truthfully and bid straightforwardly already in the primary bid rounds. Bidders should not be able to shade their bids and then provide large jump bids in the supplementary bids round. An eligibility-points rule is used to determine activity and eligibility to bid in the primary bid rounds. Each license in a band requires a certain number of eligibility points, and a bidder cannot increase his activity across rounds. In the supplementary bids round, revealed preferences during the primary bid rounds are used to derive relative caps on the supplementary bids that impose consistency of preferences between the primary and supplementary bids submitted. The consequence of these rules is that all bids are constrained relative to the bid for the final primary package by a difference determined by the primary bids. This should set incentives for straightforward bidding in the primary bid rounds.

2.2.2. The Two-Stage Combinatorial Clock Auction with OR Bids (CCA+)

The two-stage combinatorial clock auction reaches full efficiency if bidders bid on their payoff maximizing package during the primary phase and truthfully on *all* packages in the supplementary phase. The number of packages increases exponentially in the number of items and in the Canadian auction this would not be possible for national bidders. To counter this "missing bids" problem (see Bichler et al. (2013a)), Industry Canada introduced the possibility to submit one or more mutually exclusive collections of OR bids[4] in the supplementary round, in addition to the mutually exclusive XOR bids. These OR bid collections (see Footnote 4) are used in conjunction with the bidder's final primary package and can be used to express additive valuations on top of the final primary package bid in a concise manner.

2.2.3. The Single-Stage Combinatorial Clock Auction (SCCA)

We also analyze the single-stage combinatorial clock auction as it has been described by Porter et al. (2003). This format was used in Romania[5] in 2012, and in

[4] http://www.ic.gc.ca/eic/site/smt-gst.nsf/eng/sf10730.html#aD-s10

[5] http://www.ancom.org.ro/en/uploads/forms_files/terms_of_reference1331893175.pdf

the recent the Danish 1800Mhz auction[6] in 2016. Here, bidders can place a bid on one or multiple package bids at the current ask prices in each round. The pricing rule is simple: As long as at least one item is over-demanded, the prices for these over-demanded items increase by a bid increment. If at some point the supply equals demand, the auction terminates and assigns the items to all bidders. In the case of excess supply, the auctioneer considers all bids, including those of the previous rounds, and solves a winner determination problem. If the all winning bidders in the current round are in the winner determination problem's solution, the auction terminates. Otherwise, a new round begins with increased prices on all items that were not allocated in the previous round. Note that the Danish 1800MHz auction had some additional rules about the prices one can specify for package bids in each round that were not considered in our simulations. We argue that these differences do not influence the results of the research question in this paper significantly.

3. Experimental design

Our experiments were conducted by using an auction framework which allows the run of all of the major spectrum auction formats, i.e. the SMRA, the single-stage and two-stage CCA (SCCA, CCA), and the two-stage CCA with OR bids (CCA+). The implementation follows the very rules specified in the documents that we referenced in the previous sections. In addition, we used the exact band plan, the licenses and caps of the Canadian 700MHz auction. Next, we introduce the value model and the strategies of the automated bidders, before we summarize the experimental design.

3.1. Market Environment

We will briefly summarize the environment of the Canadian 700 MHz auction in 2014. We used the very same band plan, the same start prices, and the same spectrum caps as in this auction.[7]

Licenses

The band plan consists of five paired spectrum licenses (A, B, C, C1, and C2), and two unpaired licenses (D, E) in 14 service areas. B and C as well as C1 and C2 were treated as generic licenses, i.e., substitutes. Although the licenses are all in the 700MHz band, they are technically not similar enough to sell all of them as generic licenses of one type.

[6] https://ens.dk/sites/ens.dk/files/Tele/information_memorandum_june_2016.pdf

[7] The detailed auction rules of the 2014 Canadian 700Mhz auction can be found at http://www.ic.gc.ca/eic/site/smt-gst.nsf/eng/h_sf01714.html.

Market Participants

The auction was dominated by three national carriers Bell, Rogers, and Telus. Rogers was the strongest bidder and contributed 62.45% to the overall revenue, while Telus paid 21.69% and Bell 10.73%. Rogers did not bid on C1/C2 and aimed for licenses in A, and B/C throughout the auction, while Bell and Telus also bid on C1/C2 in certain service areas. The smaller bidders mainly bid on remaining C1/C2 licenses. Bell and Telus had to coordinate and find an allocation such that they both got sufficient coverage in the lower 700 MHz band (A, B and C licenses), which explains much of the bid data. The bid data generated for our experiments was based on the field data.

Caps

In order to facilitate the market entry for new entrants, Industry Canada set up spectrum caps. All eight bidders were restricted to at most 2 paired frequency licenses in each service area. Large national wireless service providers such as Rogers, Bell, and Telus were further limited in that they could only bid on one paired license in each service area among licenses B, C, C1 and C2. This cap on large wireless service providers did not, however, include license A. Still, the national bidders could bid on $2 * 3 * 3 = 18$ packages per region including the empty package, which leads to $18^{14} \approx 3.75 * 10^{17}$ packages across all regions. In contrast, in the Canadian 700MHz auction, Rogers submitted 12 supplementary bids, Bell 543 and Telus 547 bids.

Activity Rules

All implemented iterative auction formats use an eligibility point or, in all CCA formats, a revealed preference/eligibility point activity rule during the rounds. Bidders begin each round with a number of eligibility points and they can only bid on a set of licenses that in sum requires a less or equal number of eligibility points. The rounds' activity requirement is 100%, i.e., a bidder loses all eligibility points that he does not use in a round. Industry Canada allowed bidders to bid on packages even beyond the current eligibility point limit, as long as this choice is consistent with the bidder's revealed preferences up to this point.

The two-stage CCA formats apply this rule also to all supplementary bids: All bids submitted in this stage have to be consistent with the bidder's revealed preferences, taking into account his bid in the final clock round and all eligibility-reducing rounds starting from the last round in which the bidder's eligibility was at least as high as the total points associated with the current bid.[8]

[8] Details of the activity rule can also be found at `http://www.ic.gc.ca/eic/site/smt-gst.nsf/eng/h_sf01714.html`.

3.2. Estimates of Base Valuations

In order get reasonable problem instances, we constructed bidder-specific value models for all participants from the auction data. The Canadian regulator published all bids from the 700Mhz spectrum auction, which allowed us to estimate realistic valuations. We do *not* aim to get precise value estimates with the goal to reproduce the outcomes of the Canadian 700MHz auction. Rather we aim to obtain realistic valuations from the observed package bids in the supplementary stage of the Canadian auction which have a reasonable order of magnitude with respect to the starting and end prices in the auction. For this, we fitted a L1 linear regression to the supplementary bids. As an abstract example, suppose there are three items, A, B, and C, and a bidder submits the following package bids: $b(AB) = 10$, $b(AC) = 10$, $b(BC) = 10$, and $b(ABC) = X$. Then an L1 regression that explains these bids in terms of underlying license values yields $v_A = v_B = v_C = 5$ *irrespective* of X.[9] The estimated valuations of course depend on the estimation technique, but our main simulation results are robust and hold for very different samples.

Interestingly, we find that the linear model fits the data well (see A for details), which suggests that the regression results provide reasonable estimates of the bidders' license values. As in the abstract example, there were some non-negligible residuals for larger packages but this is to be expected given the synergistic nature of bidders' values. We purposefully do not try to estimate these package synergies. First, this would be difficult given that only a small proportion from the exponential set of possible package bids were submitted.[10] More importantly, we are not interested in reproducing the outcomes of one particular auction, i.e. the Canadian 700MHz, but want to robustly compare the performance of the SMRA and CCA under various "out-of-sample" synergy conditions.

3.3. Synergy Model

We will consider two synergy models: with "linear" synergies the marginal value of each license rises linearly with the total number of licenses won. In contrast, with "extreme national synergies" the marginal value of a license is independent of the number of licenses won unless *all* licenses in a national package are won. To illustrate, consider an abstract example with K items that all have one eligibility point. Denote a bidder's individual license values by ϑ_k for $k = 1, \ldots, K$. When there are synergies, the individual license values go up when they are part of a larger package. Suppose license k is part of a set of L licenses that the bidder wins then

$$\vartheta_k(L) \;=\; \vartheta_k \left(1 + (\alpha - 1)\frac{L}{K}\right) \qquad (1)$$

[9] An L2 regression would instead yield $v_A = v_B = v_C = (X + 20)/7$.

[10] Telus submitted the most supplementary bids: 547. But this is only 0.00000000000015% of all possible package bids.

where the synergy coefficient, $\alpha \geq 0$, determines the scale of package valuations. We are mainly interested in the case where licenses are complements ($\alpha > 1$), but the value parametrization in (1) also applies to the substitutes case ($\alpha < 1$). When $\alpha = 1$, there are no synergies (neither positive nor negative) and values are simply additive.

With the "linear synergy" model that license values rise linearly with the number of licenses won. In the "extreme national synergy" model, non-additive valuations are limited to

$$\vartheta_k(L) = \begin{cases} \vartheta_k & \text{if} \quad 1 \leq L < K \\ \alpha\vartheta_k & \text{if} \quad L = K \end{cases}$$

With extreme national synergies the marginal value of each license is independent of the number of licenses won unless *all* licenses in the national package are won, see also Goeree and Lien (2014). The fact that synergies are only applied for a selected few out of all possible packages is an extreme case. In the field we have seen similar models where super-additive valuations were only determined for a small number of packages as it is often difficult to determine the right polynomial. However, there is no public information about the structure of such valuation models and the level of synergies that we are aware of. We will see that with such a synergy model, package auctions do achieve higher efficiency, but this is not the case with the linear synergy value model.

The experiments use the following parameters: $\alpha \in \{1, 2, 2.5\}$ and both synergy models. A difference with the abstract example above is that in the Canadian 700MHz auction not all licenses had the same eligibility points, but it in the case of the linear model it is straightforward to account for those.[11] Another difference is that in the Canadian auction there were "national bidders", i.e. Telus, Bell, and Rogers, and "regional bidders". In the experiments we assume that only the national carriers enjoy synergies. A final detail is that national bidders could cover the country once or twice with a specific set of paired licenses: In the case of covering the country twice, the synergy coefficient for the national package was α, whereas if a bidder covers the nation within a specific paired frequency block once[12], his synergy coefficient is $\alpha - 0.5$.

3.4. Efficient Allocation

Given the large number of licenses and bidders the computation of the optimal allocation is challenging. Obviously, we cannot enumerate all valuations for all possible packages of all bidders. Instead, we solve a mixed integer program, which

[11] Let e_k denote the eligibility points associated with license $k = 1, \ldots, K$. Suppose license k is part of a winning set S. Define $e = \sum_{k=1}^{K} e_k$ and $e_S = \sum_{k \in S} e_k$ then $\vartheta_k(L) = \vartheta_k\Big(1 + (\alpha - 1)\frac{e_S}{e}\Big)$.

[12] As an example, winning one licence in each region in the A frequency block would cause a higher surplus, whereas covering the nation once with a mix of A and BC licences would not.

leverages the structure of the valuation models described above without having to enumerate all possible packages. The model is described in Appendix B.[13] A version of this model is also used to compute the payoff-maximizing package in each round for bidders.

3.5. Bidding Strategies

There are no closed-form equilibrium bidding strategies for the SMRA or the CCA in complex environments such as the Canadian 700MHz auction. But strategic manipulation would be difficult given the large number of licenses and packages bidders are interested in, which is why we assume bidders naively optimize in each round of the auction. Such *straightforward bidding* means that a bidder bids on the package with the highest net value taking into account current prices. Formally, a straightforward bid $\beta_j^{\mathrm{SF}}(v_j, p_j)$ is defined as

$$\beta_j^{\mathrm{SF}}(v_j, p_j) \in \arg\max_{S \subseteq I}(v_j(S) - p_j(S))$$

where S is the set of items bidder j wants to bid on from all possible items I, $v_j(S)$ is the bidder's valuation for the package S and $p_j(S)$ the price he has to pay for it. In the case of an exposure limit λ, we limit the value $v_j(S)$ so that $v_j^{\lambda}(S) = min\left(v_j(S), \lambda \sum_{i \in S} v_j(i)\right)$ for all S where $p_j(S) > \sum_{i \in S} v_j(i)$.[14]

3.6. Experimental Design

The main treatment variables in our experiments are the auction format, the synergies in the two value models (α_l or α_n), and the exposure risk that bidders are willing to take in SMRA. We also consider limits on the number of bids that bidders are willing to submit.

We compare SMRA against various versions of the combinatorial clock auction format, which include the single stage CCA where bidders submit one or two bids per round, SCCA(1) and SCCA(2), the two stage CCA with zero[15] or 200 additional bids, CCA(0), CCA(200). We also analyze the CCA with an OR bid language in the supplementary stage, in addition to 200 supplementary bids, CCA+(200).

We draw 100 sets of random valuations based on the derived value models from a uniform distribution of ±5% around the item valuations and subsequently use these

[13]The optimization model is an effective way to determine payoff-maximizing packages in each round with both synergy models. Unfortunately, with non-linear, e.g. quadratic, synergies the integer programming problem becomes intractable.

[14]As an example, assume a bidder with $v(i_1) = 5, v(i_2) = 5, v(i_3) = 20$ and $v(\{i_1, i_2\}) = 100$. Assume further a fixed number $5 < \rho < 10$, $p(i_1) = p(i_2) = \rho$ and $p(i_3) = 2\rho$, and a winning cap of two items per bidder. With $\lambda < 2$, β_j^{SF} is $\{i_3\}$, whereas with a $\lambda > 2$, the bidder strictly prefers the package $\{i_1, i_2\}$.

[15]This means that bidders submit only supplementary bids on combinations they already bid for in the clock phase. The result of the clock phase of the two stage CCA is also added for comparison.

Variables	Values
Auction formats	SMRA, SCCA(1), SCCA(2), CCA(0), CCA(200), CCA+(200)
Synergy model	No, National-2.0, National-2.5, Linear-2.0, Linear-2.5
Exposure	No, +50%, +100%, Full

Table 1: Treatment variables

valuations for all treatment combinations, which are described in Table 1. In total, we have $9*4+6 = 42$ treatment combinations and 4,200 simulation runs with approx. 1,500 hours of simulation run time.

We use efficiency $E(X)$ as the primary aggregate measure.[16] Let the optimal allocation be denoted X^* then:

$$E(X) = \frac{\sum_{j\in J} v_j(X)}{\sum_{j\in J} v_j(X^*)}$$

We also measure the revenue distribution $R(X)$, which compares the auctioneers revenue against the optimally achievable surplus.

$$R(X) = \frac{\sum_{j\in J} p_j(X)}{\sum_{j\in J} v_j(X^*)}$$

4. Results

In the following, we summarize the main results of the numerical experiments. The efficiency of the auction formats is quite high and with a synergy value of less or equal to 2, efficiency is typically higher than 80% of the optimal surplus. This is surprising given that the auctions are simple heuristics and the worst-case approximation ratio for the allocation problem are low, as we discussed in the introduction.

Result 1. *The two synergy models lead to significantly different results:*

- *With linear synergies, the efficiency level of the SMRA exceeds those of the various CCA versions. This is true even for a high value of the synergy coefficient α_l. Among the CCA versions, the two stage CCA is more efficient than the single stage SCCA.*
- *With the extreme national synergies, the single- and two-stage CCA yield comparable efficiency levels, which exceed those of the SMRA. CCA+ has a significantly higher effiency than all other formats. Efficiency in the SMRA and in the CCA formats decreases for higher levels α_n, but not for the CCA+.*

[16] A formulation of the mixed integer program can be found in Appendix B.

For the national synergy model, efficiency levels are ranked as follows: CCA+ $\succ^{***}$ SCCA $\succ^{**}$ CCA $\succ^{***}$ CCA (clock) $\succ^{***}$ SMRA. For the linear synergy model, the ranking is: SMRA $\succ^{***}$ CCA+ $\sim$ CCA $\succ^{***}$ SCCA $\succ^{***}$ CCA (clock).

We provide box plots with the efficiency levels in the appendix in Figures 1 to 10, which show the results for different levels of synergy. Figure 1 shows box plots in a market with purely additive valuations where the efficiency of all auction formats is very high. SMRA achieves full efficiency while the combinatorial clock auction formats have a slightly lower efficiency with a median above 97%. This efficiency loss might be due to the missing bids problem, as bidders can only reveal a small subset of all their bundle preferences in a combinatorial auction of this size. SMRA$_{No}$ refers to an implementation of SMRA with bidding agents who do not take any exposure risk, while they bid up to their full package valuation in SMRA$_{Full}$.

Figure 2 shows the box plots for the extreme national synergy model with a synergy coefficient of $\alpha_n = 2.0$ while Figure 3 shows the box plots for $\alpha_n = 2.5$. With national synergies, the median efficiency of the SMRA with bidders who do not take exposure risk (SMRA$_{No}$) is 0.77 or 0.69 respectively. If bidders would take full exposure, some of the efficiency would be recovered. The efficiency of the CCA formats is high at around 0.96 and 0.85 for the CCA+ and CCA, respectively.

Figures 4 and 5 show efficiency for the linear synergy model with synergy coefficients of $\alpha_l = 2.0$ and $\alpha_l = 2.5$. Interestingly, efficiency of SMRA remains very high for both synergy levels. An explanation for this is the auction format and the caps in the Canadian auction. The caps limited the size of the packages any national bidder could win. None of the bidders could win all a huge package with all licenses. All national bidders were also able to win larger packages in SMRA, which all result in some level of complementarity in the linear synergy model. The average efficiency of the combinatorial clock auctions is substantially lower.

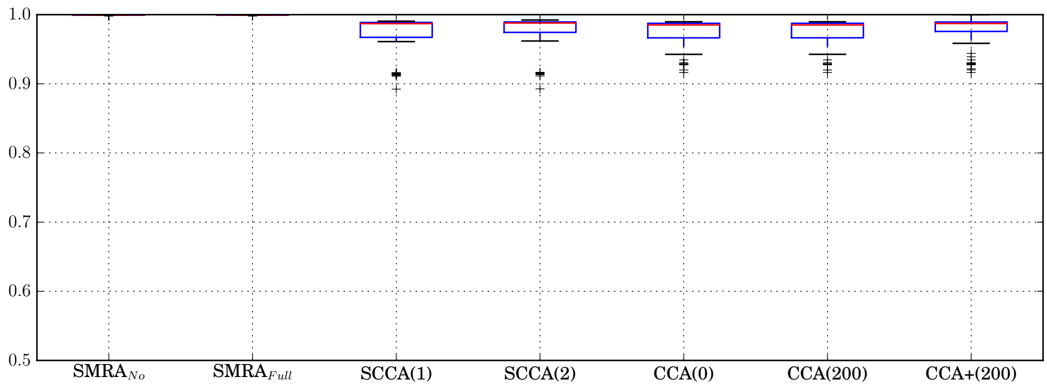

Figure 1: Efficiency in the additive value model

The pattern in the box plots is reflected in the regression analysis. Table 2 summarizes the OLS estimates for the extreme national synergy model and Table 3

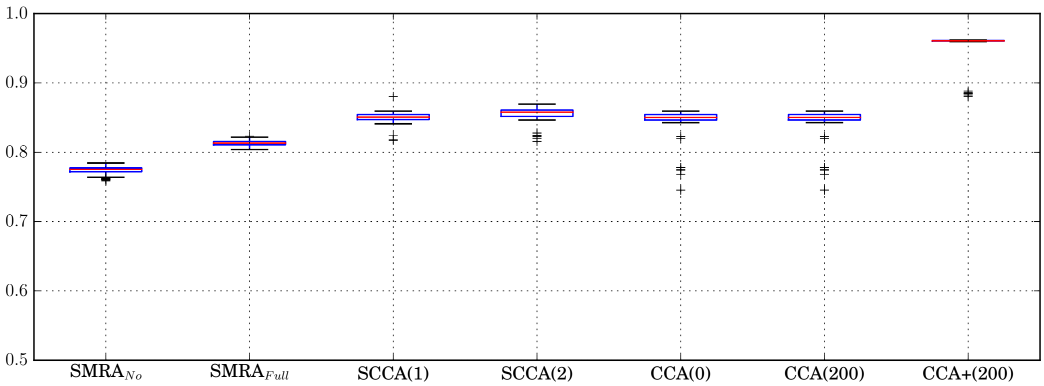

Figure 2: Efficiency in the National-2.0 value model

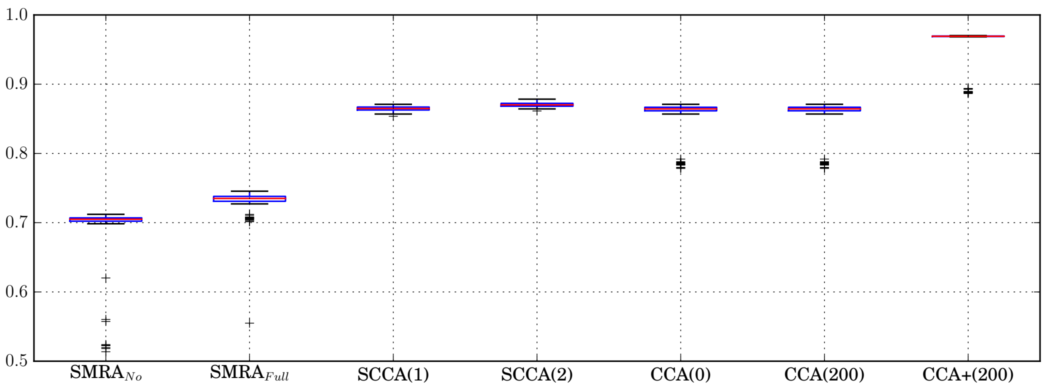

Figure 3: Efficiency in the National-2.5 value model

does the same for the linear synergy model (the dependent variable in each case is efficiency). SMRA describes the baseline auction format in this regression and no exposure the baseline for the various exposure levels that the bidders take. In Table 2 all combinatorial clock auction formats have a positive and significant influence on efficiency compared to SMRA. Differences among the combinatorial clock auction formats are small. In contrast, in Table 2 all the regression coefficients are negative indicating a negative impact of alternative auction formats on efficiency compared to SMRA. The CCA auction with the OR bid language (CCA+) was best among the combinatorial clock auction format as bidders revealed a significantly larger number of package valuations.

One of the specifics of the Canadian value model is the number of three national competitors. We ran also simulations with four national bidders[17], but the main

[17] A fourth national bidder was created by first averaging the estimated valuations of the three national bidders. We then drew 100 sets of random valuations based on this model from a uniform

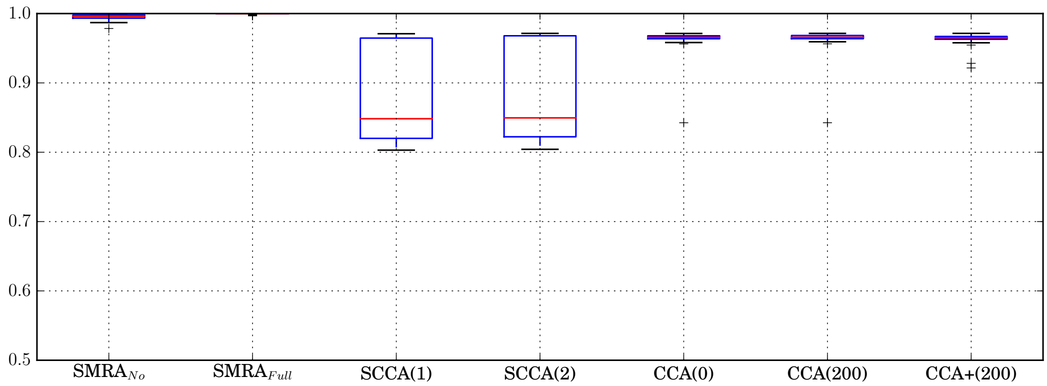

Figure 4: Efficiency in the Linear-2.0 value model

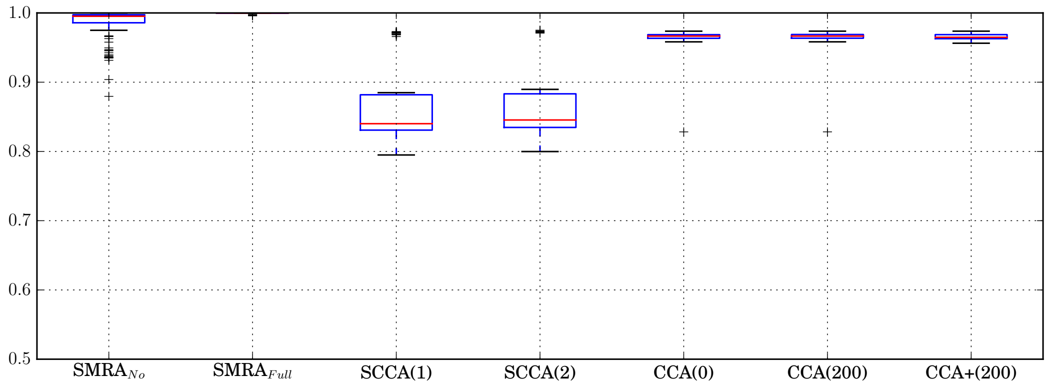

Figure 5: Efficiency in the Linear-2.5 value model

results regarding the comparison of auction formats remain.

distribution of ±5%, identically to the procedure used for our main simulations.

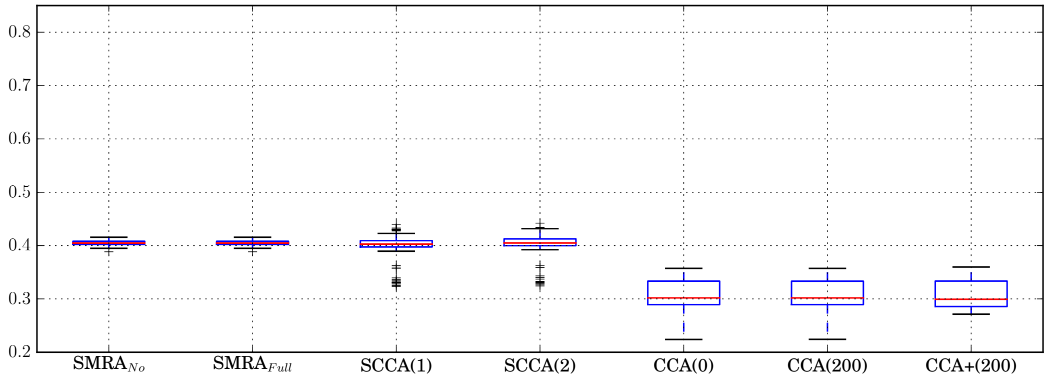

Figure 6: Revenue in the additive value model

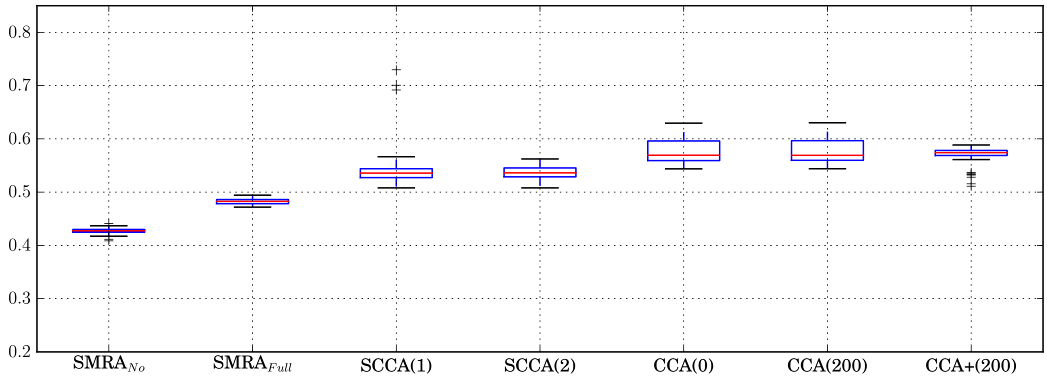

Figure 7: Revenue in the National-2.0 value model

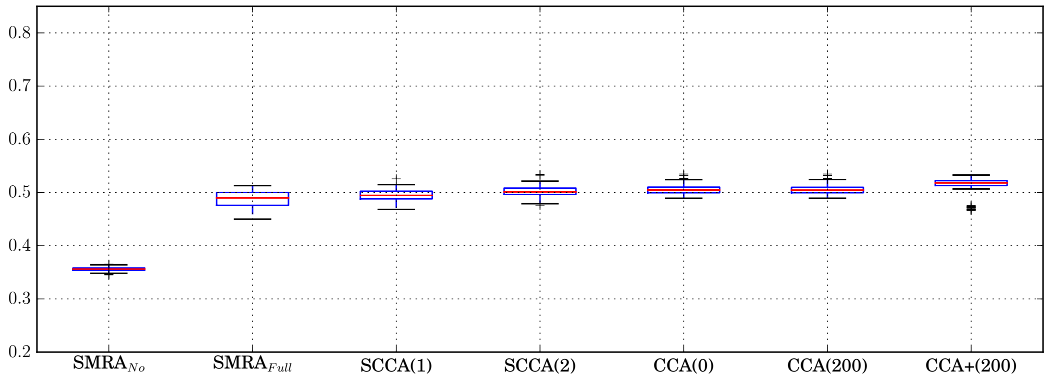

Figure 8: Revenue in the National-2.5 value model

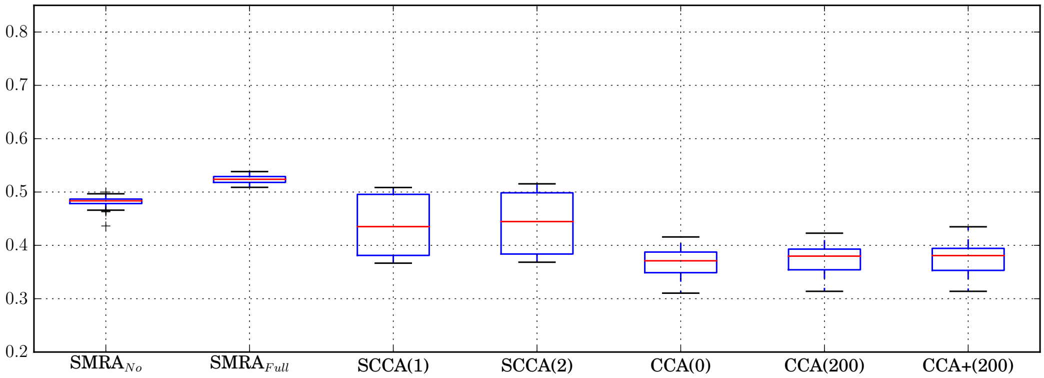

Figure 9: Revenue in the Linear-2.0 value model

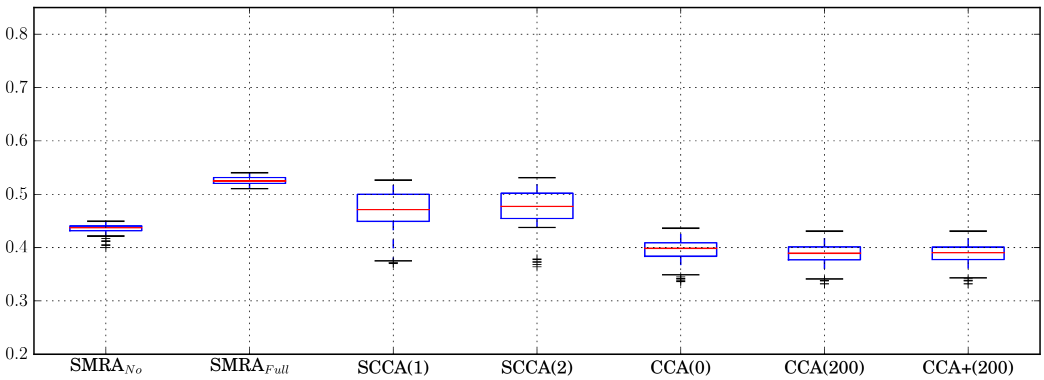

Figure 10: Revenue in the Linear-2.5 value model

Coefficients	$E(X)$	p-value	$R(X)$	p-value
Intercept	0.8568	< 0.001	0.5965	< 0.001
Auction Format				
CCA (clock)	0.1050	< 0.001	0.0989	< 0.001
CCA(0)	0.1176	< 0.001	0.1497	< 0.001
CCA(200)	0.1177	< 0.001	0.1499	< 0.001
CCA+(200)	0.2253	< 0.001	0.1505	< 0.001
SCCA(1)	0.1256	< 0.001	0.1266	< 0.001
SCCA(2)	0.1310	< 0.001	0.1282	< 0.001
Synergy	-0.0556	< 0.001	-0.0911	< 0.001
Exposure				
+50%	0.0344	< 0.001	0.0506	< 0.001
+100%	0.0395	< 0.001	0.0933	< 0.001
Full	0.0390	< 0.001	0.0931	< 0.001
Adjusted R^2	0.79		0.84	

Table 2: Regression results for the extreme national synergy model (base: SMRA_{No}, no exposure) with efficiency and revenue as dependent variables.

Coefficients	$E(X)$	p-value	$R(X)$	p-value
Intercept	1.0125	< 0.001	0.4941	< 0.001
Auction Format				
CCA (clock)	-0.1425	< 0.001	-0.0796	< 0.001
CCA(0)	-0.0246	< 0.001	-0.0784	< 0.001
CCA(200)	-0.0244	< 0.001	-0.0781	< 0.001
CCA+(200)	-0.0243	< 0.001	-0.0784	< 0.001
SCCA(1)	-0.1185	< 0.001	-0.0037	-
SCCA(2)	-0.1158	< 0.001	-0.0006	-
Synergy	-0.0104	< 0.001	-0.0159	< 0.001
Exposure				
+50%	0.0062	< 0.1	0.0030	-
+100%	0.0084	< 0.05	0.0272	< 0.001
Full	0.0106	< 0.01	0.0663	< 0.001
Adjusted R^2	0.74		0.66	

Table 3: Regression results for the linear synergy model (base: SMRA_{No}, no exposure) for efficiency and revenue as dependent variables.

Result 2. *Allowing OR bids in the the supplementary phase of the two-stage CCA significantly raises efficiency and revenue compared to the CCA with XOR bid language (p-value:* < 0.001*) in the extreme national synergy model, but not in the linear synergy model. Bidding up to 200 additional supplementary bids does not significantly improve efficiency compared to a two-stage CCA without additional package bids.*

Straightforward bidding is fully efficient in the two stage CCA if bidders are able to provide their preference for all packages. In larger combinatorial auctions such as the Canadian auction, this is not possible with a fully combinatorial XOR bid language. Bidding only on a small fraction of the 10^{17} possible packages in the supplementary phase does not improve the allocation significantly, even if bidders bid on up to 200 of their most valuable packages. The addition of OR bids in the supplementary phase does improve the efficiency of the auction outcome, but this can only be observed in the extreme national synergy model.

Result 3. *With national synergies, revenue is lowest in the SMRA when bidders do not take any exposure. With linear synergies, revenue of the two-stage CCA (and CCA+) is lowest. The single-stage CCA achieves higher revenue than the two-stage CCA on average in the linear synergy model.*

Figure 6 shows the revenue results of the additive model, where the two-stage CCA, both CCA(0) and CCA(200), is worst. Figures 7 to 8 provide box plots for the extreme national synergy model, while Figures 9 to 10 shows those for the linear synergy model. Since efficiency is low with extreme national synergies, also revenue is low. The opposite is true for linear synergies. The revenue of the single-stage CCA, SCCA(1) and SCCA(2), is always higher on average than that of the two-stage CCA, CCA(0) and CCA(200).

Table 2 summarizes the robust regression estimates for the extreme national synergy model and Table 3 shows those for the linear synergy model (with revenue as the dependent variable). We see a similar pattern as for efficiency. With national synergies, revenue is significantly higher in the CCA formats. But with linear synergies it is significantly higher in SMRA. A summary of the results for efficiency and revenue across all auction formats and value models can be found in Table 5 and Table 6.

5. Conclusions

Much has been written about the design of efficient spectrum auctions in the past two decades (see e.g. Bichler and Goeree (2017) for an up-to-date overview). Traditionally, game-theoretic analyses and laboratory experiments have been used to analyze different auction formats. These methods have their limitations. In particular, spectrum auctions in the field typically have many licenses, complex activity rules, and spectrum caps. Such design elements are important, but typically ignored in theoretical and lab studies. Simulation studies complement theory and experiments. They allow economists to study alternative auction formats under exactly the same rules as used in the field. One contribution of the paper is the implementation of the SMRA and various versions of the CCA formats in a unified simulation framework, as well as an instance generator that estimates bidder valuations based on drop out bids from the Canadian 700MHz spectrum auction.

The economic environment in this simulation mirrors the Canadian market with all its institutional details, and it allows us to study the efficiency of wide-spread spectrum auction formats with different levels of synergies in the valuations of bidders. We assume that bidders maximize payoff in each round. This can serve as a reasonable approximation of bidder behavior in larger markets such as the Canadian auction. In any case, it is important to understand the average approximation ratio of simple auction algorithms in realistic environments when bidders bid straightforward.

The main results of the experiments go against wide-spread wisdom. Even high synergies do not always lead to higher efficiency in the combinatorial clock auctions compared to SMRA, and the relative efficiency ranking depends on the type of synergies. We analyzed two types of synergies motivated from observations in the field. In the "extreme national" synergy model synergies only occur when a bidder wins *all* licenses in a national package (and not if the bidder wins, say, 99% of the licenses). The extreme national synergies create the largest possible risk for a bidder who wants to aggregate licenses in the SMRA and, not surprisingly, the SMRA results in low efficiencies in this model. More moderate synergies occur when the marginal value of a license rises linearly with the number of licenses won. Surprisingly, under this assumption of "linear" synergies, the SMRA outperforms various versions of the CCA, in terms of efficiency as well as revenue. Overall, it is interesting to observe that the average efficiency loss in both models is remarkably low considering the simplicity of the algorithms and the worst-case approximation ratio of the allocation problem in combinatorial auctions.

References

Andersen, R., 2008. Modern methods for robust regression. No. 152. Sage.

Ausubel, L., Cramton, P., Milgrom, P., 2006. The clock-proxy auction: A practical combinatorial auction design. In: Cramton, P., Shoham, Y., Steinberg, R. (Eds.), Combinatorial Auctions. MIT Press, Cambridge, MA.

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

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, rdu023.

Bichler, M., Goeree, J., 2017. Handbook of Spectrum Auction Design. Cambridge University Press, Cambridge, UK.

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, 1–35.

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

Cave, M., Nicholls, R., 2017. The use of spectrum auctions to attain multiple objectives: Policy implications. Telecommunications Policy 41 (5–6), 367 – 378, optimising Spectrum Use.

Consiglio, A., Russino, A., 2007. How does learning affect market liquidity? a simulation analysis of a double-auction financial market with portfolio traders. Journal of Economic Dynamics and Control 31 (6), 1910 – 1937, tenth Workshop on Economic Heterogeneous Interacting AgentsWEHIA 2005.

Cramton, P., 2008. A review of the l-band auction. Tech. rep.
URL `http://stakeholders.ofcom.org.uk/binaries/spectrum/spectrum-awards/completed-awards/714758/Report_by_Professor_Peter_C1.pdf`

Cramton, P., 2013. Spectrum auction design. Review of Industrial Organization 42.
URL `http://ideas.repec.org/p/pcc/pccumd/09sad.html`

Cramton, P., Shoham, Y., Steinberg, R., 2006. Introduction to combinatorial auctions. In: Cramton, P., Shoham, Y., Steinberg, R. (Eds.), Combinatorial Auctions. MIT Press, Cambridge, MA.

Day, R., Cramton, P., 2012. The quadratic core-selecting payment rule for combinatorial auctions. Operations Research 60, 588–603.

Domowitz, I., Wang, J., 1994. Auctions as algorithms. Journal of Economic Dynamics and Control 18 (1), 29 – 60.

Goeree, J. K., Holt, C. A., 2010. Hierarchical package bidding: A paper & pencil combinatorial auction. Games and Economic Behavior 70 (1), 146 – 169.

Goeree, J. K., Lien, Y., 2014. An equilibrium analysis of the simultaneous ascending auction. Journal of Economic Theory 153, 506–533.

Janssen, M., Karamychev, V., 2016. Spiteful bidding and gaming in combinatorial clock auctions. Tech. rep.

Levin, J., Skrzypacz, A., 2016. Properties of the combinatorial clock auction. American Economic Review 106, 2528–2551.

Maldoom, D., 2007. Winner determination and second pricing algorithms for combinatorial clock auctions. Discussion paper 07/01, dotEcon.

Milgrom, P., 2000. Putting auction theory to work: The simultaneous ascending auction. Journal of Political Economy 108 (21), 245–272.

Mochon, A., Saez, Y., 2017. A review of radio spectrum combinatorial clock auctions. Telecommunications Policy 41 (5), 303 – 324, optimising Spectrum Use.

Pekec, A., Rothkopf, M. H., 2003. Combinatorial auction design. Management Science 49 (11), 1485–1503.

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

A. The Bidders' Value Models

We estimate value models for the individual bidders from the bid data in an attempt to get realistic valuations. While we do not aim to estimate the true valuations of bidders from the data, we want to get valuations which resemble those in the field. In an initial step, we preprocessed the data and based our estimates on the supplementary bids. We also removed a few outliers, bids that were substantially higher or lower than the other bids of a bidder, and which might have had strategic reasons.

Instead of a standard L2 linear regression, we used an L1 linear regression, which is constrained to have no intercept. L1 regressions are more robust against outliers (Andersen, 2008). In this regression, the package bid is the dependent variable, the vector of licenses included in a package bid describes exogenous variables. Such regressions were run for each bidder. The R^2 of the resulting models for the three national bidders Rogers, Bell, and Telus are 0.60, 0.89, and 0.96, respectively. We provide the value models in our data companion. These estimated valuations provide the means of a distribution describing the value of each type of license. For each valuation set used in a simulation, we draw different valuations for the individual licenses of different bidders. Synergies for packages were then modeled on top as described in our experimental design.

B. The Efficient Allocation

In what follows, we provide a quick introduction into the mixed integer program (MIP) that is used to compute the efficient allocation in all our experiments. A similar formulation (albeit only for a single bidder) can be used to compute a bidder's best response and is used in the simulation framework to compute the bidders' payoff-maximizing package in their straightforward bidding strategy.

The allocation of licences can be modeled as a multi-knapsack problem, in which each licence $i \in I$ with capacity or quantity $q(i)$ will be assigned to a bidder $j \in J$. Part of the objective function is the sum of the bidder's additive valuations $v_j^A(i)$ for a licence i times whether he actually won the licence $x_{i,j}$ (3) or not. Assignining at most $q(i)$ licences for each licence type is checked at (7). Each licence also has $e(i)$ eligibility points associated with it and which will be tested against bidder's maximum allowed eps $e_j^{\max}$ (10). Equations (8) and (9) apply the capping rules defined by the Canadian regulator. Note that the national bidders $J^{\mathrm{N}} \subset J$ are slightly more constrained across the available regions R. Each bidder can potentially have one or multiple non-additive demand vectors $r_{j,k}^N \in (\{-1\} \cup \mathbb{N}_0)^{|I|}$, where -1,0 and $\mathbb{N}$ signals indifference, should not receive, and the minimum quantity, respectively. If a demand vector is fulfilled, the equations (13) through (16) make sure the associated non-additive valuation $v_j^N(k)$ is added to the objective (4). This part of the algorithm is used to model the National Synergy Value Model. Similarly, a bidder can have an incremental bonus that reaches its maximum relative complementarity b_j^T, when all of the bidders relevant bonus items I_j^B reach the required quantity $r_j^B(i)$ (equation (6)). In order to stay in linear space, the relative factors were linearized, as seen in (6) and (17) to (19). Please refer to Table 4 for a summary of all involved parameters and variables.

Parameters	
I	Items
J	Bidders
J^{N}	national bidders
b_j^T	added complementarity (1.0 = 100 % = additivity)
e_j^T	total eps of the bidder's bonus package
$e(i)$	eps of item
$e_j^{\max}$	maximum eps of the bidder
$r_j^B(i)$	quantity j requires for bonus of i
$\mathrm{band}(i)$	band of item
R	available regions
$v_j^A(i)$	additive surplus
$v_j^N(k)$	nonadditive surplus
$r_{j,k,i}^N$	required quantity of i in bidder j's package k
$q(i)$	available quantity of i
Variables	
$x_{i,j} \in \mathbb{N}$	quantity bidder j receives of item i
$y_{j,k} \in \mathbb{B}$	the bidder won package k
$y_{j,k,i} \in \mathbb{B}$	variable for non-additive packages, item-wise
$b_{j,i_1,i_2} \in \mathbb{B}$	bonus active for i_1, i_2

Table 4: Sealed Model Parameters and Variables

$$
\begin{aligned}
\max \quad & \sum_{j \in J} \left(z_j^{add} + z_j^{na} \right) && && (2)\\
\text{s.t.} \quad & \sum_{i \in I} v_j^A(i) \cdot x_{i,j} = z_j^{add} && \forall j \in J && (3)\\
& \sum_{k \in K_j} v_j^N(k) \cdot y_{j,k} + && && (4)\\
& \frac{b_j^T - 1}{\mathrm{ep}_j^T} \sum_{i_1 \in I_j^B} \sum_{i_2 \in I_j^B} r_j(i_1) && && (5)\\
& \cdot e(i_1) \cdot r_j(i_2) \cdot v_j^A(i_2) \cdot b_{j,i_1,i_2} \geq z_j^{na} && \forall j \in J && (6)\\
& \sum_{j \in J} x_{i,j} \leq q(i) && \forall i \in I && (7)\\
& \sum_{i \in I_r : \mathrm{band}(i) != \text{'DE'}} x_{i,j} \leq 2 && \forall r \in R, \forall j \in J && (8)\\
& \sum_{i \in I_r : \mathrm{band}(i) \in \{\text{'BC'},\text{'C1C2'}\}} x_{i,j} \leq 1 && \forall r \in R, \forall j \in J^{\mathrm{N}} && (9)\\
& \sum_{i \in I} e(i) \cdot x_{i,j} \leq e_j^{\max} && \forall j && (10)\\
& \sum_{i} y_{j,k,i} \leq |I| \cdot (1 - y_{j,k}) && \forall j, \forall k && (11)\\
& 1 - \sum_{i} y_{j,k,i} \leq |I| \cdot y_{j,k} && \forall j, \forall k && (12)\\
& r_{j,k,i}^N - x_{i,j} \leq q(i) \cdot y_{j,k,i} && \forall j \in J, \forall k \in K_j, \forall i \in I : r_{j,k,i}^N > 0 && (13)\\
& 1 - r_{j,k,i}^N + x_{i,j} \leq q(i) \cdot (1 - y_{j,k,i}) && \forall j \in J, \forall k \in K_j, \forall i \in I : r_{j,k,i}^N > 0 && (14)\\
& x_{i,j} \leq q(i) \cdot y_{j,k,i} && \forall j \in J, \forall k \in K_j, \forall i \in I : r_{j,k,i}^N = 0 && (15)\\
& 1 - x_{i,j} \leq q(i) \cdot (1 - y_{j,k,i}) && \forall j \in J, \forall k \in K_j, \forall i \in I : r_{j,k,i}^N = 0 && (16)\\
& r_j^B(i) \cdot b_{j,i,i} \leq x_{i,j} && \forall j \in J, \forall i \in I && (17)\\
& b_{j,i_1,i_2} \leq b_{j,i_1,i_1} && \forall j \in J, \forall i_1 \in I_j^B, \forall i_2 \in I_j^B : i_1 \neq i_2 && (18)\\
& b_{j,i_1,i_2} \leq b_{j,i_2,i_2} && \forall j \in J, \forall i_1 \in I_j^B, \forall i_2 \in I_j^B : i_1 \neq i_2 && (19)
\end{aligned}
$$

$$
\begin{aligned}
x_{i,j} \in &\quad \mathbb{N}_0 && \forall j \in J, \forall i \in I \\
y_{j,k} \in &\quad \mathbb{N}_0 && \forall j \in J, \forall k \in K \\
y_{j,k,i} \in &\quad \mathbb{N}_0 && \forall j \in J, \forall k \in K, \forall i \in I
\end{aligned}
$$

C. Modeling the Exposure Risk

A limit to how much a bidder should expose himself can be directly included in the bidders individual bidding selection mixed integer problem. We do this by first deciding on an exposure factor exposure_j he should not surpass, i.e.:

$$\text{exposure}_j \geq \frac{z_j^{add} + z_j^{na}}{z_j^{add}}$$

We take the individual parts of the bidders target function z_j of the bidder, i.e., its additive part, z_j^{add} and the non-additive part z_j^{na}, which we want to limit. After reformulating the formula above we get the following inequality

$$z_j^{na} \leq (\text{exposure}_j - 1) \left(\sum_{i \in I} v_j^A(i) \cdot x_{i,j} \right)$$

which can then be integrated into the selection MIP to further limit the value of z_j^{na}.

D. Results Tables

What follows are the mean and standard deviation for all aggregate metrics.

	Additive	National-2.0	National-2.5	Linear-2.0	Linear-2.5
$SMRA_{No}$	0.9994 (0.0003)	0.7735 (0.0060)	0.6902 (0.0484)	0.9952 (0.0037)	0.9830 (0.0251)
$SMRA_{+50\%}$	0.9994 (0.0003)	0.8130 (0.0040)	0.7194 (0.0319)	0.9954 (0.0033)	0.9952 (0.0035)
$SMRA_{+100\%}$	0.9994 (0.0003)	0.8130 (0.0040)	0.7296 (0.0210)	0.9998 (0.0004)	0.9952 (0.0035)
$SMRA_{Full}$	0.9994 (0.0003)	0.8130 (0.0040)	0.7286 (0.0214)	0.9998 (0.0005)	0.9997 (0.0007)
SCCA(1)	0.9737 (0.0258)	0.8502 (0.0076)	0.8647 (0.0030)	0.8733 (0.0629)	0.8679 (0.0530)
SCCA(2)	0.9754 (0.0260)	0.8554 (0.0098)	0.8702 (0.0031)	0.8763 (0.0641)	0.8703 (0.0509)
CCA (clock)	0.9674 (0.0281)	0.8296 (0.0284)	0.8441 (0.0343)	0.8596 (0.0622)	0.8338 (0.0576)
CCA(0)	0.9742 (0.0192)	0.8445 (0.0225)	0.8544 (0.0272)	0.9642 (0.0132)	0.9648 (0.0148)
CCA(200)	0.9742 (0.0192)	0.8446 (0.0226)	0.8544 (0.0272)	0.9644 (0.0132)	0.9649 (0.0148)
CCA+(200)	0.9790 (0.0215)	0.9554 (0.0195)	0.9590 (0.0268)	0.9641 (0.0067)	0.9654 (0.0042)

Table 5: Efficiency - mean, (std. deviation)

	Additive	National-2.0	National-2.5	Linear-2.0	Linear-2.5
VCG	0.3820 (0.0038)	0.6182 (0.0078)	0.5634 (0.0071)	0.4879 (0.0043)	0.4946 (0.0043)
SMRA_{No}	0.4048 (0.0049)	0.4272 (0.0052)	0.3559 (0.0040)	0.4821 (0.0090)	0.4347 (0.0094)
$\text{SMRA}_{+50\%}$	0.4048 (0.0049)	0.4822 (0.0052)	0.4019 (0.0119)	0.4840 (0.0057)	0.4390 (0.0056)
$\text{SMRA}_{+100\%}$	0.4048 (0.0049)	0.4822 (0.0051)	0.4874 (0.0183)	0.5235 (0.0066)	0.4482 (0.0073)
SMRA_{Full}	0.4048 (0.0049)	0.4823 (0.0053)	0.4869 (0.0165)	0.5235 (0.0071)	0.5258 (0.0075)
SCCA(1)	0.3974 (0.0263)	0.5408 (0.0315)	0.4954 (0.0101)	0.4379 (0.0492)	0.4713 (0.0348)
SCCA(2)	0.3998 (0.0266)	0.5374 (0.0109)	0.5019 (0.0104)	0.4438 (0.0497)	0.4717 (0.0380)
CCA (clock)	0.3912 (0.0257)	0.5060 (0.0274)	0.4748 (0.0342)	0.3915 (0.0631)	0.3662 (0.0590)
CCA(0)	0.3098 (0.0256)	0.5769 (0.0211)	0.5055 (0.0084)	0.3665 (0.0269)	0.3932 (0.0229)
CCA(200)	0.3098 (0.0256)	0.5772 (0.0212)	0.5056 (0.0084)	0.3741 (0.0266)	0.3863 (0.0218)
CCA+(200)	0.3097 (0.0271)	0.5713 (0.0136)	0.5127 (0.0174)	0.3739 (0.0278)	0.3860 (0.0217)

Table 6: Revenue - mean, (std. deviation)

(Un)expected Bidder Behavior in Spectrum Auctions

Peer-reviewed Journal Paper

Title: (Un)expected Bidder Behavior in Spectrum Auctions: About Inconsistent Bidding and Its Impact on Efficiency in the Combinatorial Clock Auction.

Authors: C. Kroemer, M. Bichler, A. Goetzendorff

In: Group Decision and Negotiation 25.1 (2016): 31-63

Abstract: The combinatorial clock auction is a two-stage auction format, which has been used to sell spectrum licenses worldwide in the recent years. It draws on a number of elegant ideas inspired by economic theory. A revealed preference activity rule should provide incentives to bid straightforward, i.e., consistent with the bidders' valuations on a payoff-maximizing package, in each round of the clock phase. A second-price rule should set incentives to bid truthfully in both phases. If bidders respond to these incentives and bid straightforward in the clock phase and truthful in the second sealed-bid stage, then the auction is fully efficient. Unfortunately, bidders might neither bid straightforward in the clock phase nor truthful on all packages in the second sealed-bid stage due to strategic reasons or practical limitations. We introduce metrics based on Afriat's Efficiency Index to analyze straightforward bidding and report on empirical data from the lab and from the field in the British 4G auction in 2013 and the Canadian 700 MHz auction in 2014, where the bids were made public. The data provides evidence that bidders deviate significantly from straightforward bidding in the clock phase, which can restrict the bids they can submit in the supplementary phase. We show that such restrictions can have a significant negative impact on efficiency and revenue.

Contribution of thesis author: Results (Canada 700 Mhz Auction of 2014), implementation (Canada 700 Mhz Auction of 2014), presentation, guidance and joint paper management

Group Decis Negot
DOI 10.1007/s10726-015-9431-0

(Un)expected Bidder Behavior in Spectrum Auctions: About Inconsistent Bidding and Its Impact on Efficiency in the Combinatorial Clock Auction

Christian Kroemer · Martin Bichler · Andor Goetzendorff

Abstract The combinatorial clock auction is a two-stage auction format, which has been used to sell spectrum licenses worldwide in the recent years. It draws on a number of elegant ideas inspired by economic theory. A revealed preference activity rule should provide incentives to bid straightforward, i.e., consistent with the bidders' valuations on a payoff-maximizing package, in each round of the clock phase. A second-price rule should set incentives to bid truthfully in both phases. If bidders respond to these incentives and bid straightforward in the clock phase and truthful in the second sealed-bid stage, then the auction is fully efficient. Unfortunately, bidders might neither bid straightforward in the clock phase nor truthful on all packages in the second sealed-bid stage due to strategic reasons or practical limitations. We introduce metrics based on Afriat's Efficiency Index to analyze straightforward bidding and report on empirical data from the lab and from the field in the British 4G auction in 2013 and the Canadian 700 MHz auction in 2014, where the bids were made public. The data provides evidence that bidders deviate significantly from straightforward bidding in the clock phase, which can restrict the bids they can submit in the supplementary phase. We show that such restrictions can have a significant negative impact on efficiency and revenue.

Keywords Market design · Spectrum auctions · Activity rules

C. Kroemer · M. Bichler (✉) · A. Goetzendorff
Department of Informatics, Technische Universität München, Boltzmannstr. 3, 85748 Garching, Munich, Germany
e-mail: bichler@in.tum.de

C. Kroemer
e-mail: kroemer@in.tum.de

A. Goetzendorff
e-mail: goetzend@in.tum.de

Published online: 24 February 2015

1 Introduction

The design of auction protocols and systems has received considerable academic attention in the recent years and found application in industrial procurement, logistics, and in public tenders (Airiau and Sen 2003; Bellantuono et al. 2013). Spectrum auction design is one of the most challenging and visible applications. It is often seen as a pivotal example for the design of multi-object markets and successful auction designs are likely role-models for other markets in areas such as procurement and logistics.

Efficiency, revenue, and strategic simplicity for bidders are typical design goals that a regulator has in mind. In theory, the Vickrey–Clarke–Groves (VCG) auction is the only strategy-proof and efficient auction but for practical reasons, it has rarely been used so far (Rothkopf 2007). Several other auction formats have been designed and used for selling spectrum. The most prominent example is the Simultaneous Multi-Round Auction (SMRA) which has been used since the mid-90s to sell spectrum licenses world-wide. The more recent Combinatorial Clock Auction (CCA) is a two-phase auction format with an initial ascending clock auction and a sealed-bid supplementary bid phase afterward. It has lately been used to sell spectrum in countries such as Australia, Austria, Canada, Denmark, Ireland, the Netherlands, Slovenia, and the UK.

The CCA draws on a number of elegant ideas inspired by economic theory. A revealed preference activity rule should provide incentives for bidders to *bid straightforward* or *consistent*, i.e., to bid truthfully on one of the payoff-maximizing packages in each round of the clock phase. If bidders fail to maximize utility and bid on a package with a less than optimal payoff, we will also refer to this as *inconsistent bidding behavior*, i.e., bids which are not consistent with the assumption of utility maximization. A second-price rule should set incentives to bid all valuations truthfully in the second sealed-bid phase. It can be shown that if bidders respond to these incentives in both phases of the CCA, then the outcome is efficient and in the core (Ausubel et al. 2006). However, bidders might not have incentives to bid truthful in both phases, and this can lead to inefficiencies.

1.1 Reasons for Inefficiency in the CCA

The CCA is used in high-stakes auctions and much recent research tries to better understand when it is efficient in theory and in the lab. For the former, Goeree and Lien (2013) highlight possibilities for profitable manipulation and deviations from truthful bidding in core-selecting auctions in a market with several local and one global bidder. They show that the Bayesian Nash equilibrium outcome in this market can be further from the core than that of the VCG auction in a sealed-bid auction, and that in their model truthful bidding is never an equilibrium in a core-selecting auction. Sano (2012) analyzes the same market situation and shows that in ascending auctions a core-selecting payment rule can lead to an inefficient perfect Bayesian equilibrium where local bidders drop out at the start. Janssen and Karamychev (2013) and later Levin and Skrzypacz (2014) provide a complete information analysis of the CCA rules considering the activity rules of the CCA and show that there are multiple equilibria with no guarantee for efficiency. The equilibria depend on assumptions about bidders' incentives to drive up prices of competitors, which is risk-free in the CCA as was shown in Bichler et al. (2013a) (see Sect. 2.4).

Lab experiments yielded low revenue and low efficiency for the CCA in a market with a larger number of licenses (Bichler et al. 2013a). Interestingly, also the CCA conducted in the UK in 2013 achieved a revenue below the expectations, leading to an investigation by the UK National Audit Office (Arthur 2013), whereas some other CCAs such as the one in Austria in 2013 achieved high revenue. It turns out that one reason for low efficiency and revenue in the experiments was that bidders submitted only a small subset of the thousands or millions of packages they could bid on. This can have strategic but also very practical reasons. In larger combinatorial auctions such as the Canadian 700 MHz auction in 2013 with 98 licenses, national bidders could potentially bid up to 18^{14} packages. It will only be possible to submit bids on a small subset of all possible packages for any bidder. All other packages are treated by the winner determination in the CCA as if bidders had no valuation for these combinations, which is unlikely.

In contrast, the SMRA uses an "OR" bidding language, where bidders can have multiple winning bids. During the winner determination, bids on different items provide an estimate for the value that a bidder has for every possible combination of bids on individual items. Also in the British auction in 2013 only a low number of package bids was submitted. Problems due to the exponential growth in the number of packages can sometimes be addressed by a compact bid language, as was discussed in Bichler et al. (2014). The recently released rules for the upcoming CCA in Canada in 2015 try to address this problem by allowing for restricted OR bids in the supplementary stage. Of course, the bid language does not solve the strategic reasons for bidders to bid on many or only a few packages in the supplementary stage. We will discuss some of these reasons in Sect. 2.4.

1.2 Contribution of this Paper

In this paper, we show that apart from missing bids in the supplementary phase, also inconsistent bidding in the clock phase can be a source of inefficiency. We show that bidders in the lab and in the field (Canada and UK) do not bid straightforward in the clock phase. There are actually several reasons for inconsistent bidding behavior. For example, bidders might have budget constraints (Shapiro et al. 2013) or values might be interdependent, which can lead to inconsistent bidding as bidders revise their valuations when they learn about other bidders' valuations during the auction. Even if bidders have independent and private values without budget constraints, there can be incentives to reduce or inflate demand in the clock phase in order to drive up payments of competitors (Bichler et al. 2013a; Janssen and Karamychev 2013; Levin and Skrzypacz 2014).

However, the revealed preference activity rule prohibits bidders from bidding truthfully up to their valuation in the supplementary phase, if they do not bid straightforward in the clock phase, and this can lead auctioneers to select an inefficient allocation. We provide evidence from the lab and from the field showing that the resulting inefficiencies can be significant, while being much less obvious at the same time. Measuring this inefficiency due to restrictions on the supplementary bid prices is straightforward in the lab, where the values of bidders are available. But also the analysis of the field data from the British LTE auction in 2013 and from the Canadian 700 MHz auction in 2014 suggests that inconsistent bidding was an issue. We introduce metrics based on Afriat's Efficiency Index, which allow measuring the level of inconsistency. Numerical simulations based on data from the lab and from the UK indicate that the impact of inconsistent bidding on efficiency can be substantial. In the lab we found an overall efficiency loss of around 5 %, which can be attributed to inconsistent bidding in the clock phase. In the data from the field, where we don't know the bidders' true valuations, we also found a surprising large number of supplementary stage bids at the bid price limit imposed by the clock phase. This can be seen as an indication that these bids were also below the true valuation, although one can assume that bidders in these these countries tried to bid up to their true valuation.

A strong activity rule, which forces bidders to be consistent across auction rounds, appears to be an intuitive solution to fix the problems discussed in this paper. However, in the conclusions we will outline issues which arise when a regulator tries to force bidders to bid straightforward.

1.3 Outline

The remainder of this paper is structured as follows: After briefly introducing the rules of the CCA in Sect. 2, we will discuss Afriat's Efficiency Index to analyze whether bidders in the CCA are bidding straightforward in Sect. 3. We will use this metric to analyze bidders in the lab in Sect. 4 and bidders from the British and the Canadian auction in Sect. 5. Finally, we will use computer simulations to analyze the impact of these deviations on the auction's final outcome in Sect. 6. Section 7 discusses stronger activity rules to force consistent bidding in the clock phase and potential problems arising from such rules.

2 The Combinatorial Clock Auction

Used for the first time in 1994, the SMRA has been the de facto standard auction format for spectrum sales for almost 20 years (Milgrom 2000). A number of well-known strategic problems have led to substantial research on alternative auction formats. In particular, the exposure problem turned out to be central. Bidders are often interested in specific combinations or packages of licenses. Their value for these packages can be much higher than the sum of the individual license values in this package. As the SMRA allows only bidding on single items, a bidder risks winning only part of his package, having to pay more than what the subpackage is worth to him. Combinatorial auctions address this problem by allowing bidders to submit bids on packages rather than on single items. In 2008 the British regulator Ofcom decided on the two-stage Combinatorial Clock Auction (CCA) (Ausubel et al. 2006), a format which has been used in many countries world-wide in the last 5 years.

First, we briefly describe the overall auction process that was the same in the recent auctions. Then we discuss the activity rules, and draw on the latest version used in Canada in 2014.[1]

2.1 The CCA Auction Process

In the *clock phase*, the auctioneer announces ask prices for all licenses at the beginning of each round. In every round bidders communicate their demand for each item at the current prices. At the end of a round, the auctioneer determines a set of over-demanded licenses for which the bidders' demand exceeds the supply. The price for all over-demanded lots is increased by a bid increment for the next round. This clock phase continues until there are no over-demanded lots left.

The *supplementary stage* is designed to eliminate incentives for demand reduction and other inefficiencies in the combinatorial clock auction due to the limited number of bids that bidders can submit in the first phase. In this sealed-bid stage bidders are able to increase bids from the clock phase or submit bids on bundles they have not bid on so far. Bidders can submit as many bids as they want, but the bid price is restricted subject to the CCA activity rule (see next subsection). Finally, all bids from both phases of the auction are considered in the winner determination and the computation of payments for the winners. The winner determination is an NP-hard combinatorial optimization problem (Lehmann et al. 2006). For the computation of payments, a Vickrey-nearest bidder-optimal core-pricing rule is used (Day and Cramton 2012).

With certain assumptions on the bidders' valuations it is possible to determine the efficient allocation and the VCG payments, even if bidders do not bid up to their true valuation in the supplementary stage. For example, if bidders have independent and decreasing marginal valuations for homogeneous items and all bidders bid straightforward then it is possible to determine Vickrey payments even bidders would not increase their bids after the clock phase. Under these assumptions bidders have strong incentives to bid truthful as the clock auction is ex post incentive compatible. However, combinatorial auctions are typically used when bidders have complementary valuations and this is when the clock auction loses its favorable properties. Without substitutes valuations an efficient outcome can not be guaranteed in a clock auction, not even with fully straightforward bidding by all participants. Actually, simple examples show that the clock phase can have very low efficiency, if all bidders bid straightforward (see Sect. 7). Actually, even if valuations were gross substitutes no ascending auction can always impute Vickrey prices (Gul and Stacchetti 1999), i.e., payments for which bidders have no incentives to shade their true valuations.[2]

[1] The auction rules of the Canadian 700 MHz auction in 2014 can be found at http://www.ic.gc.ca/eic/site/smt-gst.nsf/eng/sf10583.html. The auction rules of the British auction in 2013 can be found at http://stakeholders.ofcom.org.uk/spectrum/spectrum-awards/awards-archive/completed-awards/800mhz-2.6ghz/.

[2] Ausubel (2006) showed that there is an ascending auction with multiple price trajectories and item-level prices, which is efficient and yields the VCG allocation and payments. The auction runs one ascending auction with all bidders, and one with each bidder excluded in turn. However, this auction format is quite different from the clock auctions used in the field so far.

Apart from the observation that bidders in spectrum auctions often have complementary valuations, a number of other reasons can cause differences between the true VCG payments and the payments computed in the CCA. For example, in larger auctions with many licenses bidders might be unable to submit supplementary bids on all possible packages. However, such missing bids can have an impact on the payments of others. There are also differences to the VCG payments, if bidders bid higher or lower than their valuation for strategic reasons, and there can be multiple non-truthful equilibria in this auction (Levin and Skrzypacz 2014).

2.2 Activity Rules in the CCA

The CCA combines two auctions in the clock and in the supplementary phase. This requires additional rules setting incentives to bid truthfully in both phases. Without activity rules, bidders might not bid actively in the clock phase, but wait for the other bidders to reveal their preferences, and only bid in the supplementary phase. Originally, the clock phase of the CCA employed a simple monotonicity rule which does not allow to increase the size of the package in later rounds as prices increase. It has been shown that with substitutes preferences straightforward bidding is impossible with such an activity rule (Bichler et al. 2011, 2013a). Later versions use a hybrid activity rule using a monotonicity rule and a revealed preference rule (Ausubel et al. 2006). Revealed preference rules allow bidders to bid straightforward in the clock phase. If they do, then bidders are able to bid on all possible packages up to their true valuation in the supplementary stage (Bichler et al. 2013a). In the following we describe the latest version of the activity rules as they have been used in the Canadian 700 MHz auction in 2014. These rules have also been used in our simulations in Sect. 6.

First, an eligibility points rule is used in the clock phase to enforce activity in the primary bid rounds. The number of bidder's eligibility points is non-increasing between rounds, such that bidders cannot bid on more licenses when the prices rise. A bidder may place a bid on any package that is within its current eligibility. Second, in any round, the bidder is also permitted to bid on a package that exceeds its current eligibility provided that the package satisfies revealed preference with respect to each prior eligibility-reducing round. Bidding on a larger package does not increase the bidder's eligibility in subsequent rounds.

The revealed preference rule works as follows: A package in clock round t satisfies revealed preference with respect to an earlier clock round s for a given bidder if the bidder's package x_t has become relatively less expensive than the package bid on in clock round s, x_s, as clock prices have progressed from the clock prices in clock round s to the clock prices in clock round t. x_s and x_t are vectors where each component describes the number of licenses demanded in the respective category, i.e., region or spectrum band. The revealed preference constraint is:

$$\sum_{i=1}^{m}(x_{t,i} \times (p_{t,i} - p_{s,i})) \leq \sum_{i=1}^{m}(x_{s,i} \times (p_{t,i} - p_{s,i}))$$

where:

- i indexes the licenses;
- m is the number of licenses;
- $x_{t,i}$ is the quantity of the ith license bid in clock round t;
- $x_{s,i}$ is the quantity of the ith license bid in clock round s;
- $p_{t,i}$ is the clock price of the ith license bid in clock round t; and
- $p_{s,i}$ is the clock price of the ith license bid in clock round s.

A bidder's package, x_t, of clock round t is consistent with revealed preference in the clock rounds if it satisfies the revealed preference constraint with respect to all eligibility-reducing rounds prior to clock round t for the given bidder.

2.3 Activity Rules in the Supplementary Phase

Under the activity rule for the supplementary round, there is no limit on the supplementary bid amount for the final clock package. All supplementary bids on packages other than the final clock package must satisfy revealed preference with respect to the final clock round regardless of whether the supplementary bid package is smaller or larger, in terms of eligibility points, than the bidder's eligibility in the final clock round. This is referred to as the *final cap rule*.

In addition, supplementary bids for packages that exceed the bidder's eligibility in the final clock round must satisfy revealed preference with respect to the last clock round in which the bidder was eligible to bid on the package and every subsequent clock round in which the bidder reduced eligibility. This is also called the *relative cap rule*.

Let x denote the package on which the bidder wishes to place a supplementary bid. Let x_s denote the package on which the bidder bid in clock round s and let b_s denote the bidder's highest monetary amount bid in the auction on package x_s, whether the highest amount was placed in a clock round or the supplementary round.

A supplementary bid b on package x satisfies revealed preference with respect to a clock round s, if b is less than or equal to the highest monetary amount bid on the package bid in clock round s, that is, b_s plus the price difference in the respective packages, x and x_s, using the clock prices of clock round s. Algebraically, the revealed preference limit is the condition that:

$$b \leq b_s + \sum_{i=1}^{m} (p_{s,i} \times (x_i - x_{s,i}))$$

where:

- x_i is the quantity of the ith license in package x;
- b is the maximum monetary amount of the supplementary bid on package x; and
- b_s is the highest monetary amount bid on package x either in a clock round or in the supplementary round.

In addition, for supplementary bid package x, let $t(x)$ denote the last clock round in which the bidder's eligibility was at least the number of eligibility points associated with package x.

A given bidder's collection of supplementary bids is consistent with the revealed preference limit if the supplementary bid for package x, with a monetary amount b for the given bidder satisfies the following condition: for any package x, the monetary amount b must satisfy the revealed preference constraint, as specified above with respect to the final clock round and with respect to every eligibility-reducing round equal to $t(x)$ or later.

Note that, in the application of the formula above, the package x_s may itself be subject to a revealed preference constraint with respect to another package. Thus, the rule may have the effect of creating a chain of constraints on the monetary amount of a supplementary bid for a package x relative to the monetary amounts of other clock bids or supplementary bids.

2.4 Incentives for Strategic Manipulation and the CCA's Prisoner's Dilemma

These activity rules have strategic implications, which have been analyzed in a number of papers. Possibilities for spiteful bidding have been shown in Bichler et al. (2011) and later in Bichler et al. (2013a), who show that standing bidders after the clock phase can determine bid prices in the supplementary round (aka. safe supplementary bids) such that their standing bid from the clock phase becomes winning with certainty. Consequently, the allocation cannot change anymore after the clock phase providing little incentives for bidding truthful in the second phase assuming independent and private values.

However, in reality bidders might often care about the prices others have to pay and consequently their payoff, i.e., bidders might be spiteful. Since the allocation cannot change anymore, the CCA provides possibilities for supplementary bids which drive up the competitors' payments, but at no risk of losing the standing bid from the clock phase (Bichler et al. 2013a). Also, they cannot pay more for this bid than what they have bid. In recent spectrum auction implementations, the regulator decided not to reveal excess supply in the last round, in order to make spiteful bidding risky. It depends on the market specifics, if this risk is high enough to eliminate spiteful bidding.

Another issue in both the VCG auction and the CCA is that they violate the law of one price. This means, two bidders might win identical allocations at different prices. We introduce a brief example following Bichler et al. (2013a) to illustrate this point: Suppose there are two bidders and two homogeneous units of one item. Bidder 1 and bidder 2 both have preferences for only one unit and a standing bid of \$5 on one unit after the clock phase. If both bidders only bid on one unit, they both pay zero. Now, according to the CCA activity rules, the allocation cannot change any more. Suppose, bidder 2 also bids \$9 for two units in the CCA, although he does not have such a valuation for two units. As a consequence, bidder 2 would still pay zero, while bidder 1 would pay \$4. However, outcomes where bidders get the same allocation at very different prices are typically perceived as problematic (see Sect. 5.3), no matter if they are due to spiteful bids or truthful bidding.

Violations of the law of one price and possibilities for riskless spiteful bidding introduce a situation much like in a prisoner's dilemma: If a bidder does not want to pay more for his allocation relative to competitors, he can bid high on losing package bids to drive up payments of competitors after the clock phase. If all bidders follow this strategy, then the payments will be at their bid prices. Often there is excess supply after the clock phase, and the standing clock bids need to be increased by the price of the unsold licenses in the final clock round to win with certainty (Bichler et al. 2013a). This, of course, can also drive up their own payments to the level of this safe supplementary bid.

Janssen and Karamychev (2013) shows in a complete information analysis that bidders with an incentive to raise rivals' costs can submit large final round bids and aggressive bids in the clock phase. Levin and Skrzypacz (2014) recently provided an elegant complete information model characterizing the ex post equilibria and resulting inefficiencies that can arise in the CCA. First, they show that the CCA can have many ex post equilibria if bidders have independent private values. If several bidders try to raise each others payments spitefully, then they show that there are again multiple equilibria featuring demand reduction in the clock phase with no guarantee of efficiency. Knapek and Wambach (2012) discuss strategic complexities partly related to an earlier version of the CCA activity rule.

3 Revealed Preference Theory and Straightforward Bidding in Auctions

As outlined earlier, straightforward bidding is a central assumption for the two-stage CCA to be efficient (Ausubel et al. 2006). Note that the revealed preference activity rules in the CCA are such that bidders can be limited in the amount they bid in the supplementary round if they do not bid straightforward in the clock phase (Bichler et al. 2011, 2013a). This can also lead to inefficiency, as we will show. Ausubel and Baranov (2014) draw on the theory of revealed preference as a rationale for the activity rules used in the latest version of the CCA in the Canadian 700 MHz auction and for future versions. They show that the current version is based on the Weak Axiom of Revealed Preference (WARP), while future versions should be based on the General Axiom of Revealed Preference (GARP) and eliminate eligibility-point-based activity rules. In what follows, we will revisit important concepts of revealed preference theory and then discuss how they relate to straightforward bidding in an auction. We will also introduce a version of Afriat's Efficiency Index, which allows us to measure straightforward bidding in empirical bid data.

The concept of revealed preferences was originally introduced by Samuelson in order to describe rational behavior of an observed individual without knowing the underlying utility function. He described the simple observation that "if an individual selects batch one over batch two, he does not at the same time select two over one" (Samuelson 1938). The term "select over" relates to a concept which is nowadays known as "revealed preferred to" and can be defined as follows:

Definition 1 Given some vectors of prices and chosen bundles (p_t, x_t) for $t = 1, \ldots, T$, x_t is directly revealed preferred to a bundle x $(x_t R_D x)$ if $p_t x_t \geq p_t x$. Furthermore, x_t is strictly directly revealed preferred to x $(x_t P_D x)$ if $p_t x_t > p_t x$. The relations R and P are the transitive closures of R_D and P_D, respectively.

Intuitively, a selected bundle x_1 is directly revealed preferred to bundle x_2 if given x_1 and x_2, both at price p, x_1 is chosen. This definition implies some sort of budget (or income) for each observation. Consider a world with only two bundles x_1 and x_2, x_1 being the more expensive one. If an individual chooses to consume x_1 nevertheless, we know that she prefers it over x_2 such that $x_1 R_D x_2$. This implies that as a rational utility maximizer, she will never strictly prefer x_2 when x_1 is affordable at the same time. More formally, this is known as the Weak Axiom of Revealed Preference (WARP).[3] If she chooses x_2, though, we do not know if that decision is due to an actual preference or a budget constraint below the price of x_1. Hence, there is also no way to predict which choice will be made in another observation where she might have a higher income or face different prices as we have learned nothing about the relation R_D.

In a setting with more than two bundles, WARP is not enough to determine if a consumer is a rational utility maximizer. A set of choices $\{x_1 R_D x_2, x_2 R_D x_3, x_3 R_D x_1\}$ is not violating WARP but is possibly irrational. In order to detect this inconsistency, we need to consider the transitive closure R which also includes $x_1 R x_3$, possibly contradicting $x_3 R_D x_1$. Therefore, in a world with more than two bundles the consumption data of a rational utility maximizer needs to satisfy the Strong Axiom of Revealed Preference (SARP)[4] or, if indifference between distinct bundles is valid, the Generalized Axiom of Revealed Preference (GARP).[5] Varian (2006) provides an extensive discussion of WARP, SARP, and GARP.

Applying these axioms to the clock phase of the CCA is straightforward: In each clock round (observation), there is a single known price vector for which each bidder submits a single demand vector. Hence, we can easily build the revealed preference relation R_D and its transitive closure R for every bidder. For the supplementary round S, we know the bid prices $p^S x$ even without an explicit price vector p^S, as bidders bid on bundles instead of single items. As only at most one of the bidder's bids will win, for any pair of supplementary bids $\{x_1^S, x_2^S\}$, the bidder reveals her preference for the higher bid. This allows us to infer $x_1^S R_D x_2^S$ if the bid on x_1^S is higher or equal to the bid x_2^S, or vice versa. A bid in the clock phase x and a supplementary bid x^S will be treated as the same observation if both bids have identical demand vectors.

Example 1 Table 1 provides a simple example of CCA bidding data for an auction with 3 clock rounds and a supplementary phase. In each round of the clock phase, the considered bidder reveals her preference of the chosen bundle over all other affordable bundles. In the supplementary phase, bundles with higher bids are preferred over those with lower bids. The given data is consistent with a set of valuations such as (85, 75, 55) for the three bundles. However, it is not consistent with the assumed actual valuations (100, 100, 100) that would require to always choose the cheapest of the three packages. When using the actual valuations to infer revealed preferences as well,[6] the resulting relation violates GARP in this case, but it cannot be detected without knowing the true valuations.

3 If $x_t R_D x_s$ then it must not be the case that $x_s P_D x_t$ for WARP to be satisfied.

4 If $x_t R x_s$ then it must not be the case that $x_s R x_t$ for SARP to be satisfied.

5 If $x_t R x_s$ then it must not be the case that $x_s P x_t$ for GARP to be satisfied.

6 $x_i R_D x_j$ for any pair (i, j) as all valuations are equal.

Table 1 An example applying the revealed preference theory to a CCA

Round t	Prices p_t		Bundle prices $p_t x$			Revealed preference
	A	B	$x_1 = (2, 2)$	$x_2 = (2, 1)$	$x_3 = (1, 1)$	
1	10	10	40*	30	20	$x_1 P_D x_2, x_1 P_D x_3$
2	20	20	80	60*	40	$x_2 P_D x_3$
3	30	20	100	80	50*	
S			85		55	$x_1 P_D x_3$
Valuations			100	100	100	

Bundles x_t, which are selected by the bidder, are marked with *

Table 2 Example for non-straightforward bidding behavior with GARP

Round t	Prices p_t		Bundle prices $p_t x$		Revealed preference
	A	B	$x_1 = (1, 0)$	$x_2 = (0, 1)$	
1	10	50	10*	(50)	
2	30	80	30	80*	$p_2 x_2 > p_2 x_1 \Rightarrow x_2 P_D x_1$

Afriat's Theorem says that a finite set of data is consistent with utility maximization (i.e., straightforward bidding) if and only if it satisfies GARP (Afriat 1967). However, GARP allows for changes in income or budget across different observations (see Table 2) as traditional revealed preference theory is based on the assumption of an idealized individual who "confronted with a given set of prices and with a given income [...] will always choose the same set of goods" (Samuelson 1938).

The auction literature typically assumes that bidders have quasi-linear utility functions such that they maximize their payoff given the prices. Quasi-linear utility functions imply that there are no binding budget constraints or "infinite income." Ausubel and Baranov (2014) argue that a GARP-based activity rule would require GARP and quasi-linearity. Also, the efficiency results for the CCA in Ausubel and Milgrom (2002) and Ausubel et al. (2006) only hold if bidders are quasi-linear and they bid straightforward. Unfortunately, Table 2 shows that the traditional definition of GARP allows for changes in income and therefore allows substantial deviations from straightforward bidding if we assume quasi-linear utility functions.

Example 2 The example in Table 2 is no violation of GARP. It can be explained by an increase in income from $t = 1$ to $t = 2$.

Therefore, we aim for a stronger definition of revealed preference with non-binding budgets, as they are assumed in theory. With this assumption, the different bids in an auction also reveal how much one bundle is preferred to another one:

Definition 2 Given some vectors of prices and chosen bundles (p_t, x_t) for $t = 1, \ldots, T$ and a constant income, we say x_t is revealed preferred to a bundle x by amount c (written $x_t R_c x$) if $p_t x_t \geq p_t x + c$.

Intuitively, $x_t R_c x$ can be interpreted as "x_t is chosen over x if it costs no more than the price of x plus c". We will refer to this definition of revealed preference as *GARP with quasi-linear utility (GARPQU)*. Note that c will be negative in all cases where x is more expensive than x_t, which would be ignored in the traditional definition of revealed preferences (see Definition 1). The result of applying this definition to a set of bid data will be a family of relations R_c instead of a single revealed preference relation R. R_c has several properties:

- $x_1 R_c x_2$ implies $x_1 R x_2$ if $c \geq 0$ (definition)
- $x_1 R_c x_2$ implies $x_1 P x_2$ if $c > 0$ (definition)
- $x R_c x$ for all $c \leq 0$ (reflexivity)
- $x_1 R_{c_1} x_2$ and $x_2 R_{c_2} x_3$ imply $x_1 R_{c_1+c_2} x_3$ (transitivity)
- $x_1 R_{c_1} x_2$ implies $x_1 R_{c_2} x_2$ if $c_1 > c_2$ (derived from transitivity and reflexivity of $R_{c_1-c_2}$)

These properties are sufficient to derive a contradiction $x R_c x$ with $c > 0$ ("$u(x) > u(x)$") for any non-straightforward bidding behavior that can be detected without knowing the actual utility function u. For example, it is easy to see that the choices in Table 2 do not describe straightforward bidding because they are not consistent under the above properties of R_c: $(x_1 R_{-40} x_2 \wedge x_2 R_{50} x_1 \Rightarrow x_1 R_{10} x_1)$.

The clock stage and the supplementary stage lead to different questions to the bidders. In the clock stage a straightforward bidder is asked to indicate which bundle has the highest payoff given some vector of prices. In contrast, a bidder should submit his true valuations for all packages. Therefore, a bidder who submits bids on the packages of the last round at the clock prices in the last round does not necessarily satisfy GARPQU.

Example 3 Let's assume there are two lots A and B. At a price of (\$100, \$50) for both lots, a bidder demands a quantity vector of (0, 3). In the next round prices increase to (\$100, \$100), and the bidder demands (1, 2). Let $v(\cdot)$ be the value of a package. In the first round, the bidder revealed that $v(0, 3) + \$50 \geq v(1, 2)$. In the second round, he reveals that $v(1, 2) \geq v(0, 3)$. The auction stops and the bidder submits exactly the same prices for the packages as supplementary bids: $\langle (0, 3), \$150 \rangle$ and $\langle (1, 2), \$300 \rangle$. The differences in supplementary bids are interpreted as differences in the valuations, such that $v(1, 2) - 150 \geq v(0, 3)$. Together with the revealed preferences from the clock phase, this leads to a violation of GARPQU $v(0, 3) + \$50 \geq v(1, 2) \geq v(0, 3) + \$150$. If the bidder revealed his valuations truthfully in the supplementary stage, he would not submit the very same bid as in the clock phase. With a bid of $\langle (0, 3), \$270 \rangle$ GARPQU will not be violated.

Note that the result of such an analysis of a series of bids is always binary: either a set of data satisfies GARPQU or it does not. In revealed preference theory, measures such as Afriat's Efficiency Index (AI) were developed to describe how well a set of consumer choices conforms to utility maximization. The AI is a goodness of fit metric that spans the range [0; 1] with 1 indicating perfect compliance with a tested axiom (Afriat 1973). It requires a variable e in all revealed preference inequations (see Definitions 1, 2):

- $x_t R_D x$: $p_t x_t \geq p_t x$ becomes $e \cdot (p_t x_t) \geq p_t x$
- $x_t P_D x$: $p_t x_t > p_t x$ becomes $e \cdot (p_t x_t) > p_t x$
- $x_t R_c x$: $p_t x_t \geq p_t x + c$ becomes $e \cdot (p_t x_t) \geq p_t x + c$

Applying the axioms with $e < 1$ leads to a relaxed version that is easier to satisfy. For instance, assume $e = 0.9$: If bundle x_t was chosen for a price of \$100 the pair (x_t, x) will only be included in R_D if $p_t x \leq \$90$. The AI is equal to the maximum value of e which satisfies the tested axiom. We will use a graph-based algorithm based on Smeulders et al. (2012) for computing the AI. There are related metrics such as the Varian Index (VI) which follows the same principle as the AI but uses a vector instead of a single constant value e (Varian 1990). Unfortunately, the computation of VI is *NP*-hard (Smeulders et al. 2012).

4 Evidence from the Lab

In a lab experiment we cannot only observe the bids, but also know the induced valuations of bidders. In what follows, we will analyze straightforward bidding in the lab and draw on the data from experiments conducted by Bichler et al. (2013a). We will focus on 16 auctions with 4 bidders in a multi-band value model with 24 blocks in 4 different bands. This means, bidders could submit up to 2400 package bids. This experimental setup is comparable to multi-band auctions with national licenses as they were conducted in Austria, Ireland, the UK, and Switzerland, although the number of bands differed from country to country.

4.1 Missing Bids

The auctions in the lab suffered from the missing bids problem with only 8.3 supplementary bids per bidder on average. Bichler et al. (2013a) argue that this has contributed to the low efficiency of only 89.3 % observed in the CCA, which was substantially lower than that of the auctions with SMRA, which achieved an average efficiency of 98.5 %. In comparison with the standing clock bids, the allocation changed after the supplementary phase in 14 auctions by 34.9 % of all licenses on average. Significant changes in the allocation could also be observed in the British auction after the supplementary stage.

4.2 Inconsistent Bidding

Figure 1 shows the AI based on GARPQU for all 64 bidders participating in a CCA in the lab experiments. The left-hand box plot describes bids from the clock phase only, the middle box plot the bids submitted in both phases, and the right box plot the clock bids and all true valuations for all packages of a bidder. A median AI of 1.000 for clock bids shows that there is no evidence for significant deviations from straightforward bidding in the bids during the clock rounds. When including data from the supplementary round, however, the median AI drops to 0.938, indicating inconsistencies between the two phases. The AI with truthful supplementary bids for

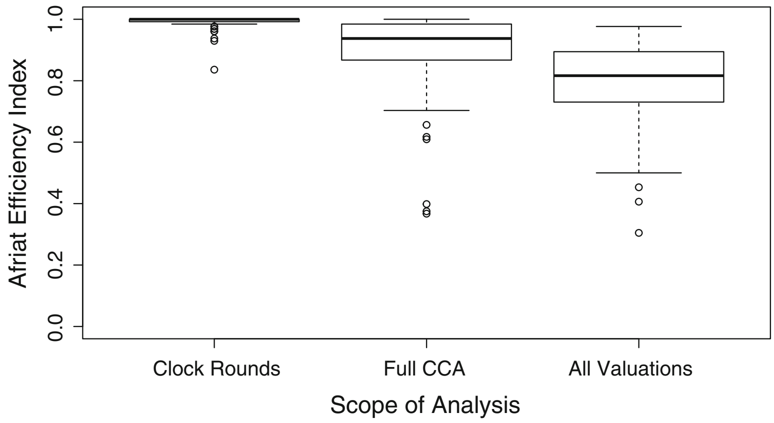

Fig. 1 *Boxplot* for AI of 16 · 4 bidders from Lab auctions

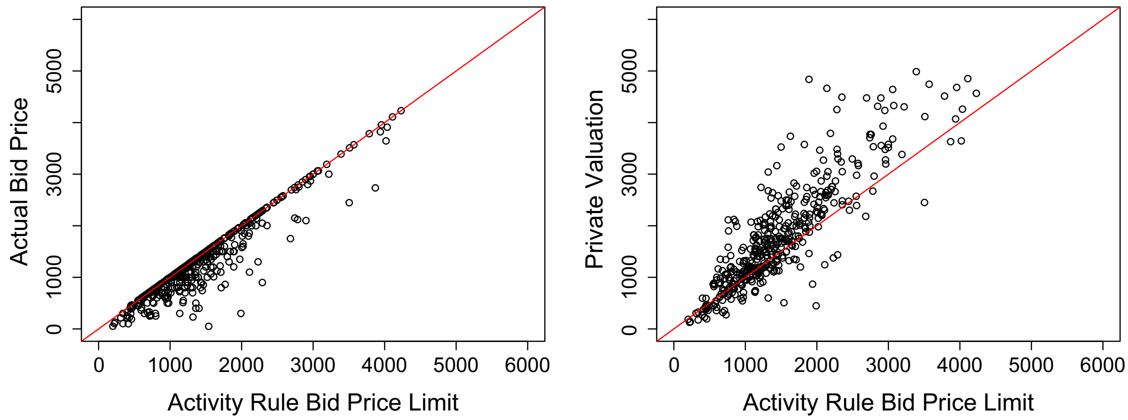

Fig. 2 *Scatterplots* for supplementary bids of 4 bidders in 16 auctions in the lab (*left*) in comparison to their induced private valuations (*right*)

all possible bundles, which is described in the third boxplot (All valuations) drops to 0.816 and suggests that bidders did indeed not bid straightforward with respect to their true valuations in the clock phase.

Deviations from straightforward bidding such as those indicated by boxplot 3 can limit the possible bid amount in the supplementary phase substantially. For the lab data we can see how high bidders have bid in the supplementary phase relative to their bid price limit.

The left scatter plot in Fig. 2 shows that bidders often bid close to the bid price limit (Pearson correlation coefficient of 0.9448). The right scatter plot illustrates the private valuations with respect to the bid price limit imposed by the activity rule and their behavior in the clock phase. For 57.2 % of all submitted supplementary bids, the bid price limit was lower than their valuation for the corresponding bundle and hence it did not allow bidders to bid their valuation truthfully in the supplementary phase.

Figure 2 deserves further explanation. As described in Sect. 2, if bidders had independent and decreasing marginal valuations, then they would not need to bid up to their true valuation in the second phase and even if bidders did not bid at all after the clock phase, the auctioneer could compute the correct Vickrey payments. The valuations of bidders in the lab were complements and there was often excess supply after the clock phase. Given the uncertainty that bidders faced in the lab, their most likely strategy was to bid truthful on their supplementary packages if possible. Bidders in the lab knew in which order they had to submit supplementary bids such that they could maximize the bids for supplementary packages. However, there were significant differences between the final payments of bidders in the clock stage and the payments one would get if bidders submitted all their valuations truthfully in a sealed-bid auction. In Sect. 6 we analyze the impact that inconsistent bidding has on the efficiency of these auctions.

4.3 Clock Prices

It would be helpful for bidders, if there was some connection between the final clock prices and the core payments, because this could give bidders a useful hint on how high they need to bid in the second phase. However, the final prices from the clock phase can differ substantially from the payments. We compared the core payments of all winning bids with the corresponding linear bundle prices in the final clock round and found that the average payment was only 59.1 % of the last clock price. The standard deviation of this ratio in the lab was 22.6 %. For the British LTE auction in 2013 this average payment was at 56.5 % of the final clock prices. Also in simulations with straightforward bidders who bid truthful in the supplementary round the clock prices do not necessarily provide an indication for payments or winning supplementary bids.

5 Evidence from the Field

The British regulator Ofcom was the first to publish the bid data on a CCA in 2008 and 2013 (Ofcom 2013a). We will primarily focus on the 2013 multi-band spectrum auction as it is closest to auctions in other countries and similar to the environment analyzed in the lab (Bichler et al. 2013a). Then we will discuss the Canadian 700 MHz auction in 2014, where bid data was revealed as well, before we summarize public information about CCA applications in some other countries. Although, all these auctions used a CCA there are important differences in the caps used, in the licenses and the band plan, and in details of the auction rules, which requires caution in the comparison of the results. Of course, we cannot know the true valuations of bidders in these auctions, however, we highlight some patterns which are similar to what we found in the lab data. In particular, bidders only bid on a small subset of all possible packages and there was a very high number of supplementary package bids at the bid price limit and not below, which can be seen as an indication of bidders over-constraining themselves in the supplementary phase due to inconsistent clock bids.

5.1 The British LTE Auction in 2013

In the British auction in 2013, 28 licenses in the 800 MHz and 2.6 GHz bands were sold, and the bid data was released to the public. There were 4 A1 blocks of paired spectrum in 800 MHz and another A2 block with a coverage obligation. In addition, there were 14 blocks of paired spectrum in the 2.6 GHz band, and another 9 blocks of unpaired spectrum in the 2.6 GHz band. The unpaired spectrum was considered less valuable than paired spectrum bands. There were seven bidders, Vodafone, Telefonica, Everything Everywhere, Hutchinson, Niche, HKT, and MLL. A spectrum cap was put on the 800 MHz band for Vodafone and Telefonica, who are considered large bidders. The detailed rules can be found at (Ofcom 2013b).

The bid data reveals the main interests of these seven bidders. Vodafone and Telefonica bid on 800 MHz and both 2.6 GHz bands. They consistently bid on two 2×5 MHz blocks in 800 MHz spectrum throughout the clock phase and both won two blocks. Everything Everywhere and Hutchinson also bid on the valuable 800 MHz spectrum, but ceased to bid on 800 MHz in the clock phase. Niche, MLL, and HKT can be considered smaller players. MLL and HKT only bid on the unpaired spectrum in 2.6 GHz and they did not win anything. Niche bid on both 2.6 GHz bands and also won blocks in both bands. More details on the auction can be found in “Appendix”, where we describe the valuations of bidders for our numerical experiments.

5.1.1 Missing Bids

Let us now provide some statistics to shed light on the missing bids problem in the British auction, which might be one of the reasons for the low revenue encountered (Arthur 2013; Smith 2013), before we discuss straightforward bidding. With all the caps considered, larger bidders such as Vodafone and Telefonica could bid on 750 packages in this auction. However, after 52 clock rounds in which the seven bidders selected 7.7 distinct bundles on average, they submitted only 39.6 supplementary bids per bidder on average (277 bids in total). Bidders always submitted higher bids on the packages submitted in the clock phase, but bid on average on 31.9 new bundles only in the supplementary phase. Telefonica submitted no more than 11 supplementary bids, while Vodafone submitted 94 of 750 supplementary bids mostly covering combinations of licenses with 20 MHz in low frequency bands. Everything Everywhere submitted 84 supplementary bids, and Hutchinson only 17 bids.

Note that the winner determination treats a missing bid as if the valuation of a bidder for this package was zero in a CCA. It is questionable if bidders had no value for all the other packages or a value below the reservation prices. In this case, the missing bids problem appears to have been an issue.

The total revenue from the bidder-optimal core prices of £2.23 bn is equivalent to the Vickrey payments in this auction, which is also due to the low number of supplementary bids which led to a lower number of core constraints when computing the bidder-optimal core payments (Day and Cramton 2012). Consequently, the discounts were very high. The sum of the bids in the revenue maximizing allocation amounts to £5.25 bn.

It is interesting to note that with only the bids from the clock phase and without the supplementary phase the auction had a revenue of £1.92 bn, which is only 13.9 % less than the final result including the bids of the supplementary phase. The supplementary phase did change the allocation considerably, however, which might have come as a surprise to some bidders. 19.3 % of all winning licenses from the clock phase (weighted by their eligibility points) were re-allocated after the supplementary round.

Table 3 Afriat's Index (AI) based on bids of the clock phase and based on all bids including the supplementary stage in the British LTE auction

Bidder	AI in clock phase	AI of all bids	# of suppl. bids
EE	0.811	0.811	48
H3G	0.988	0.845	17
HKT	1	0.626	9
MLL	1	0.619	9
Niche	0.995	0.814	89
Telefonica	1	0.467	11
Vodafone	0.946	0.946	94

5.1.2 Inconsistent Bidding

Next, we analyze straightforward bidding in the British auction using Afriat's index as we have discussed it in Sect. 3. Table 3 shows the AI per bidder for the clock phase only and for all bids including the supplementary phase. Although the median AI is high (0.995) for bids in the clock phase only, it decreases to 0.811 when we also consider supplementary bids. Note that this value is lower than what we have found in the lab auctions even though it is an upper bound for the "true AI". If the true valuations of each bidder are taken into account the AI can be considerably lower as we have seen in Sect. 4 and in the example in Table 1.

The reason for low auctioneer revenue after the supplementary phase might, however, also have been due to limits on the bid prices imposed by the activity rule. Figure 3 compares the supplementary bid prices and the corresponding bid price limit imposed by the activity rule and shows that the bids of many bidders are very close to this limit. The bid data is highly correlated with the bid price limit imposed by the activity rule

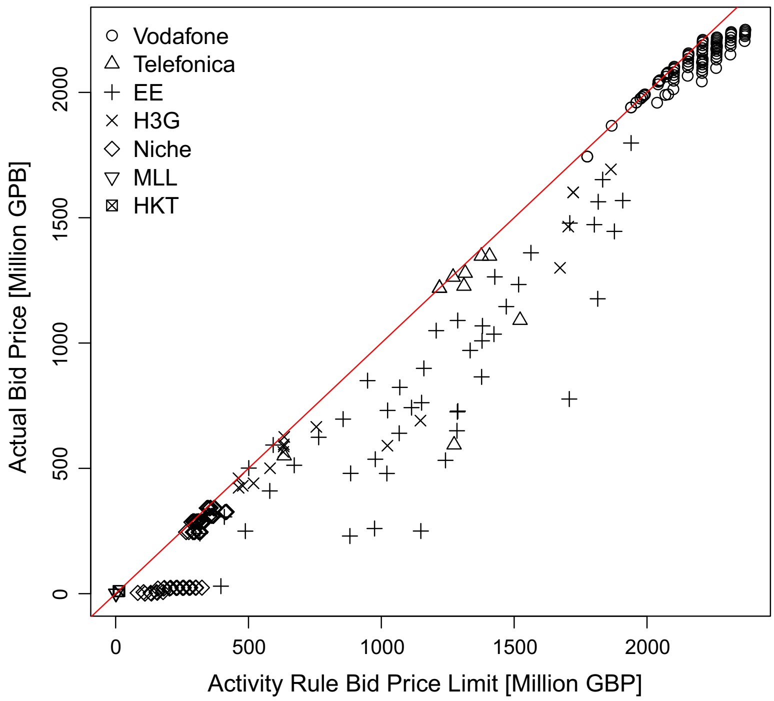

Fig. 3 Scatterplot for supplementary bids (UK)

(Pearson correlation coefficient of 0.9824) and yields a median ratio of the bid price to the bid price limit of 92.3 % (mean: 80.5 %). Interestingly, this ratio was particularly high for the big bidders Vodafone and Telefonica with a median of 98.1 and 96.5 % respectively. For supplementary bids of the remaining five bidders, the median ratio was only 83.0 %, which might be due to the fact that these were financially weaker bidders. In this auction spiteful bidding (bidders submitting high losing package bids to drive up payments of competitors) did not seem to be an issue such that the more likely explanation is that bidders could not bid up to their valuations.

5.2 The Canadian 700 MHz Auction in 2014

The Canadian 700 MHz auction in 2014 comprised 5 paired spectrum licenses (A, B, C, C1, and C2), and two unpaired licenses (D, E) in 14 service areas. B and C as well as C1 and C2 were treated as generic licenses. Although the licenses are all in the 700 MHz band, they are technically not similar enough to sell all of them as generic licenses of one type.

The total revenue of \$5.27 bn from the bidder-optimal core prices was 32.4 % less than \$7.14 bn, the sum of provisionally winning bids after the final clock round. The sum of the bids in the revenue maximizing allocation was \$9.13 bn. Again, the clock prices provided little guidance for what might constitute a winning bid in the supplementary phase.

The auction was dominated by three national carriers Bell, Rogers, and Telus. Rogers was the strongest bidder and contributed 62.45 % to the overall revenue, while Telus paid 21.69 % and Bell 10.73 %. Rogers did not bid on C1/C2 and aimed for licenses in A, and B/C throughout the auction, while Bell and Telus also bid on C1/C2 in certain service areas. The smaller bidders mainly bid on remaining C1/C2 blocks. Bell and Telus had to coordinate and find an allocation such that they both got sufficient coverage in the lower 700 MHz band (A, B and C blocks), which explains much of the bid data. There was a disparity in how much bidders had to pay for different packages, which can be explained by different valuations that bidders placed on packages and the payment rule. Still, due to the high competition and revenue the auction is considered successful.

5.2.1 Missing Bids

Overall, the high competition among the three national telecoms Bell, Rogers, and Telus and the clever spectrum caps for them explains much of the result. All eight bidders were restricted to at most 2 paired frequency blocks in each service area. Large national wireless service providers such as Rogers, Bell, and Telus were further limited in that they could only bid on one paired license in each service area among licenses B, C, C1 and C2. This cap on large wireless service providers did not, however, include block A. Still, the national bidders could bid on $2 \times 3 \times 3 = 18$ packages per region including the empty package, which leads to $18^{14} \approx 3.75 \times 10^{17}$ packages in all regions. Rogers submitted 12 supplementary bids, Bell 543 and Telus 547 bids, which suggests that there was a missing bids problem as it is questionable if all other

Table 4 Afriat's Index (AI) based on bids from the clock phase and based on all bids including the supplementary stage in the Canadian 700 MHz auction

Bidder	AI in clock phase	AI of all bids	# of suppl. bids
Bell	0.871	0.159	544
Bragg	0.722	0.151	14
Feenix	0.873	0.730	16
MTS	0.893	0.450	32
Novus	0.949	0.379	4
Rogers	0.977	0.454	13
SaskTel	0.499	0.384	12
TbayTel	0.857	0.716	2
Telus	0.930	0.235	548
Videotron	0.728	0.493	24

packages had no valuation for the bidders. Note that only one license remained unsold after the auction. Rogers bid consistently on the A licenses and one license in B/C, such that the coordination problem was largely solved by Bell and Telus, who split the regional service areas on B/C and C1/C2.

5.2.2 *Inconsistent Bidding*

It is interesting to understand straightforward bidding in the Canadian 700 MHz auction as well. Although the regulator disclosed all the bid data, the clock prices were not made public. We used a linear program which helped us reconstruct clock prices from the bid data. There are some assumptions in this linear program and we cannot compute the price trajectories and the resulting AI with certainty, such that the numbers in Table 4 are only estimates. However, the order of magnitude in the AI was similar for different price trajectories that we could derive. The numbers suggest that bidders deviated substantially from straightforward bidding. One explanation is that bidders such as Bell and Telus actively tried to coordinate and agree on non-overlapping packages of licenses. It is also interesting to note that some small local bidders bid on competitive service areas in A and B/C outside the service area in which they operate. One conjecture is that this was done in an attempt to park eligibility rights and keep clock prices low in their own service area. As the regulator did not disclose the excess supply after the clock phase and due to the uncertainty in this large scenario, it is not unreasonable to believe that bidders tried to bid up to their true valuation in the supplementary stage. Actually, in Canada the supplementary bid on the final clock package was substantially higher than the final clock round bid for many bidders. Figure 4 shows that, again, a very large proportion of the other supplementary bids are exactly at their bid price limit indicating that they might have been truncated due to restrictions imposed by the activity rule. The low AIs for the different bidders provide further evidence.

5.3 Observations from Other Countries

Apart from Canada and the UK, bids were not made public in other countries. As mentioned earlier, the UK also released data for two earlier CCAs in 2008, the L-band

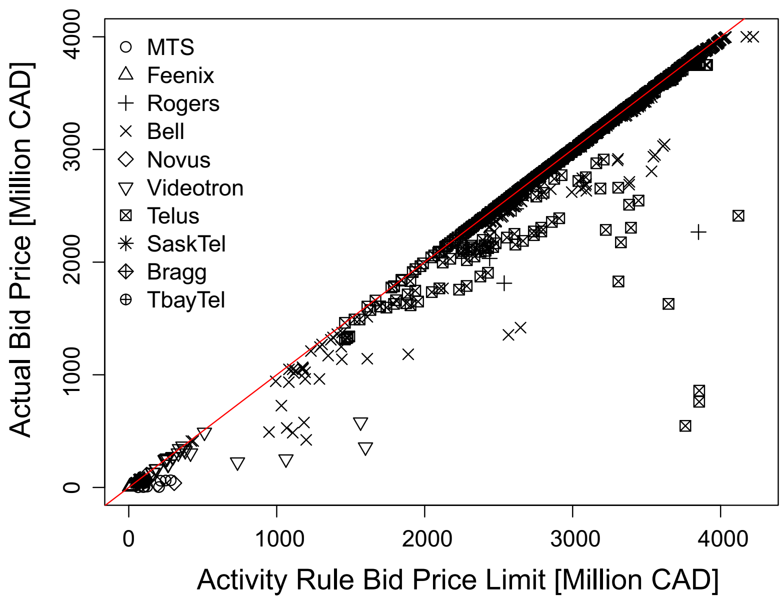

Fig. 4 Scatterplot for supplementary bids (Canada)

auction with 17 licenses, and the 10–40 GHz auction with 27 licenses. In the L-band auction bidders submitted between 0 and 15 bids in the supplementary phase also indicating missing bids from at least some of the bidders. In this auction with much less valuable spectrum than in 2013, one bidder won all 17 lots with a bid of £20 m. The bidder only had to pay £8.334 m, which was the revenue of the best coalition of bidders without the winner (Cramton 2008). In the 10–40 GHz auction all but one bidder made their highest supplementary bid either on the final clock package, or on a subset thereof (Jewitt and Li 2008).

The Swiss auction in 2012 was remarkable, because one bidder payed almost 482 million Swiss Francs, while another one payed around 360 million Swiss Francs for almost the same allocation. This can happen in a Vickrey auction as well as in a CCA when one bidder contributes more to the overall revenue with his bids than another bidder (see Sect. 2.4).

The Austrian Auction in 2013 on the 800, 900, and 1800 MHz bands is another interesting case. Bidders could potentially submit up to 12,810 package bids. The regulator reported that the three bidders actually submitted 4000 supplementary bids in total. The regulator also disclosed that most of these bids were submitted on very large packages (RTR 2013). This large number of supplementary bids can be seen as one reason for the high prices paid in Austria. The attempt to drive up prices of other bidders and avoid having to pay more for an allocation than ones competitors, as it happened in Switzerland, can serve as one explanation for this bidding behavior. However, it leads to the Prisoner's dilemma discussed in Sect. 2.4.

6 Estimating the Impact of Missing Bids and Inconsistent Bidding

We performed computer simulations of the CCA for the lab value model as well as for the British 4G auction. For the latter, we estimated valuations for the seven bidders from the bid data with base valuations, intra-band and inter-band synergies.

The estimated valuations are described in "Appendix". We did not perform this analysis for the large Canadian auction with 98 licenses, because this would require many more assumptions due to the regional structure and the many licenses involved. Data from 16 auctions in the lab and 10 sets of synthetic valuations for the British scenario were used. All significance tests reported in this section are using a Wilcoxon signed-rank sum test.

Efficiency and revenue of an auction are typically used as primary metrics. Throughout the rest of this paper, we will use the terms *allocative efficiency*:

$$E = \frac{\text{actual surplus}}{\text{optimal surplus}} \times 100\,\%$$

and auctioneer's *revenue share*:

$$R = \frac{\text{auctioneer's revenue}}{\text{optimal surplus}} \times 100\,\%$$

The revenue share shows how the resulting total surplus is distributed between the auctioneer and the bidders. Optimal surplus describes the resulting revenue of the winner-determination problem if all valuations of all bidders were available, while actual surplus considers the true valuations for those packages of bidders selected by the auction. In contrast, auctioneer's revenue describes the cumulative payments of the bids selected by the auction, not their underlying valuations.

In the following subsections we analyze the impact of missing bids and inconsistent bidding in the clock phase. The main results are summarized in Table 5. A baseline for this analysis are the simulations with truthful bidders, i.e., bidders who bid straightforward in each clock round and submit truthful supplementary bids on all bundles. As expected, all simulations where bidders submitted all package bids truthfully were 100 % efficient in contrast to the efficiency of 89.3 % we measured in the lab.

6.1 Impact of Missing Bids in the Supplementary Phase

We first evaluate the impact of the missing bids problem. In this set of simulations, the simulated bidders bid straightforward in the clock phase (see Fig. 5 and the first two column-pairs in Table 5), such that they could bid up to their valuation in the supplementary phase. As human bidders only submit a small subset of possible supplementary bids, there are just a few core constraints leading to lower prices and hence lower auctioneer revenue. In order to better understand the impact of this effect, we restricted our bidders in the number N of additional packages they can bid on in the second phase.

More precisely, bidders always started the supplementary phase with truthful bids on all clock bundles in reverse order of submission which allows them to maximize the amount they can bid on other packages without violating the activity rule. Then they submitted additional truthful bids on bundles chosen after a heuristic which we

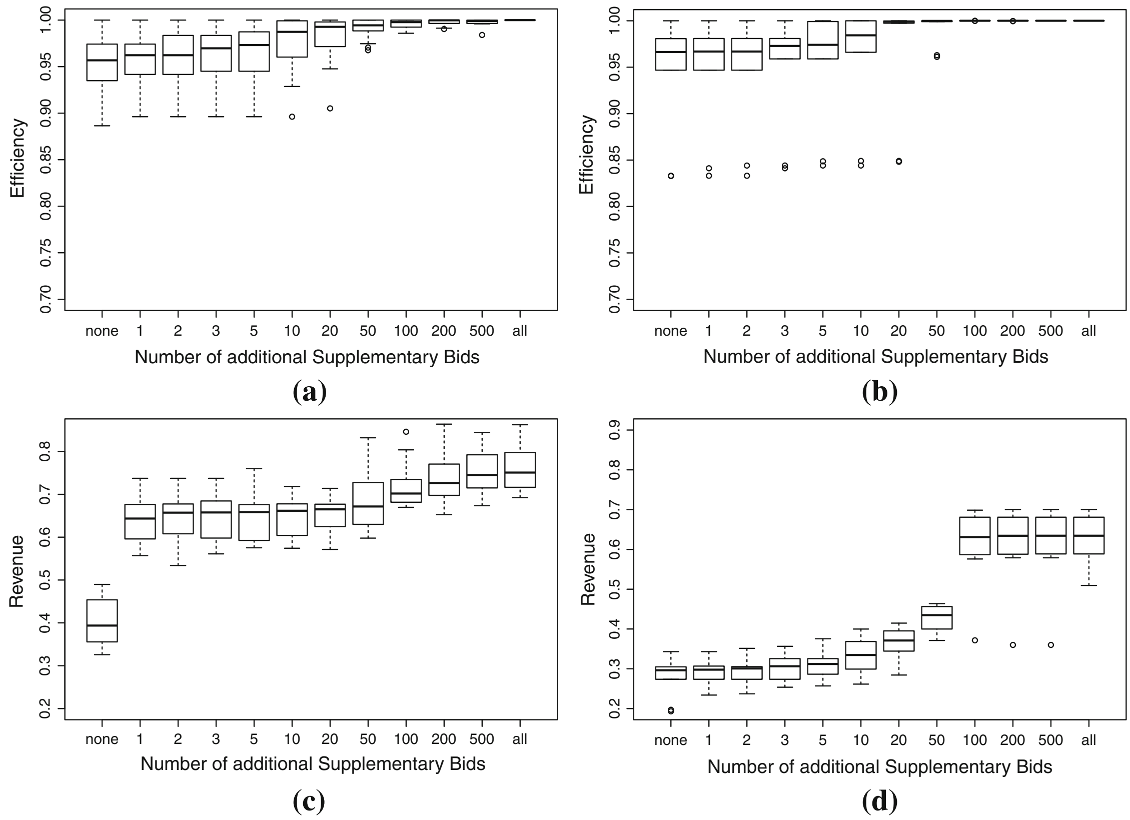

Fig. 5 Mean efficiency and revenue for the British and the lab value models with the CCA. **a** Efficiency (lab). **b** Efficiency (UK). **c** Revenue (lab). **d** Revenue (UK)

observed in the British auction. First, bidders do not demand more units of a certain band than they did in the clock phase, and second, the bidders Telefonica and Vodafone do not submit any bids without two A blocks. Out of this pre-selection, bidders selected up to N of their $2N$ strongest bids. We define the strength of a bid as the valuation divided by the bundle size in terms of the corresponding eligibility points. We have also tested different bundle selection heuristics, but the differences in efficiency were minor. The artificial bidders were bidding truthful as far as they could, such that only limitations in the number of bids submitted matter.

For the simulations with the lab value model, there is a significant difference in revenue between no supplementary bids at all (first line in Table 5) and the submission of one new bid. This is due to the fact that in the treatment without additional supplementary bids, we just evaluated the bids submitted in the clock phase. In the treatment with one additional bid, clock bids were updated to their true valuation. The number of bids affects auctioneer revenues in both value models. Even for 50 additional bids, the revenue share is still significantly lower than with supplementary bids on all bundles (p value $=$ .0000), which is due to missing bids.

With 50 additional bids, the efficiency was beyond 99 % in both value models. While we found an average efficiency of 89.3 % for the CCA in the lab with 8.3 supplementary bids on average, simulated auctions with no bids in the second phase at all yielded an efficiency of 95.3 %, which is significantly higher than in the lab (p value $=$ .0052). A substantial part of this difference can be attributed to inconsistent bidding in the clock phase which we will discuss next.

Table 5 Mean efficiency and revenue for the British and the lab value models with the CCA and bidders bidding truthful on a subset of the packages in the supplementary round up to their true valuation or a restricting bid price limit, resp

Supplementary bids in simulations	British VM (Straightforward)		Lab VM (Straightforward)		Lab VM (Actual clock bids)	
	Efficiency (%)	Revenue (%)	Efficiency (%)	Revenue (%)	Efficiency (%)	Revenue (%)
None	94.7	28.1	95.3	40.5	86.3	37.2
1	94.8	29.3	95.7	63.8	89.2	54.0
2	94.8	29.5	96.0	64.5	90.4	54.5
3	95.2	30.2	96.2	64.3	91.7	54.0
5	95.5	31.2	96.4	64.7	92.8	53.3
10	96.1	33.3	97.5	64.8	92.5	54.4
20	96.9	36.5	98.1	65.5	93.7	56.1
50	99.2	42.8	99.1	68.5	95.1	60.0
100	100	61.3	99.6	71.8	94.3	63.4
200	100	61.4	99.8	73.6	94.5	65.6
500	100	61.4	99.8	75.0	94.3	66.6
all	100	62.9	100	75.7	95.0	71.4
Human subjects					89.3	41.0
	39.6 suppl. bids/bidder				8.3 suppl. bids/bidder	
	Estimated values				Induced values	

6.2 The Impact of Inconsistent Bidding in the Clock Phase

As we have discussed in Sect. 4, bidders in the lab and in the field did not bid straightforward in the clock phase and were therefore limited by the activity rule in the supplementary phase. Now, we want to understand how much efficiency loss can be attributed to these limitations in the simulations. In the British value model we only have the bid data of a single instance, which is why we only report on the lab value model.

For all 16 instances we replicated the bids of human bidders in the clock phase. In the supplementary phase the agents tried to bid their true valuations on additional bundles like in the previous subsection. If this was impossible due to the revealed preference activity rule, they chose the highest possible bid price instead.

The third column-pair of Table 5 and Fig. 6 summarizes the results. Even for supplementary bids on all possible bundles, the efficiency was only 95.0 % on average. This is not significantly different (p value $>$ .95) to the mean efficiency of 95.3 % that we measured with straightforward bidders but without any supplementary bids. These findings provide evidence that non-straightforward bidding in the first phase reduces efficiency of the final outcome. The auctioneer revenue share was significantly lower as well. For some simulations the average differences in revenue share were more than 10 %, which was only due to inconsistent bidding in the clock phase.

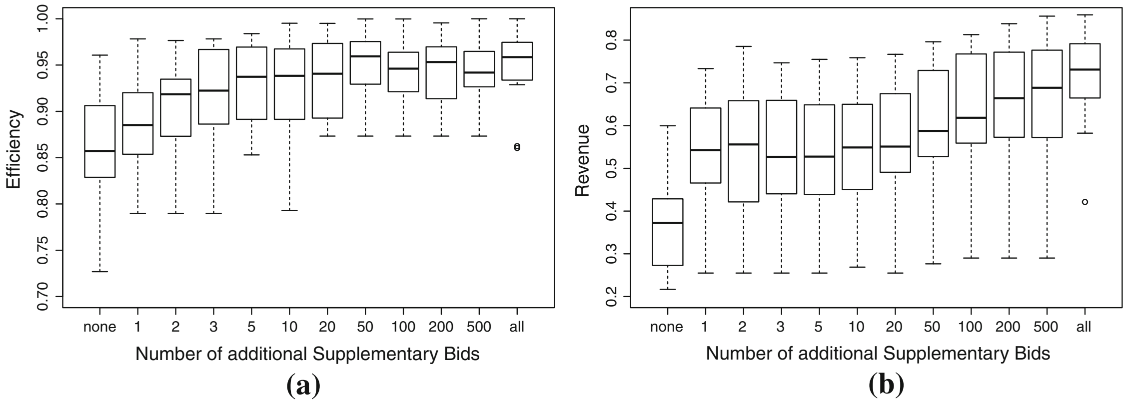

Fig. 6 Efficiency and revenue for lab instances with different numbers of (as possible) truthful supplementary bids after actual clock round bids. **a** Efficiency. **b** Revenue

7 Can Strong Activity Rules Serve as a Remedy?

Our analysis of bids in a CCA in the lab and in the recent British and Canadian spectrum auctions indicates that bidders do not bid straightforward in the clock phase of the CCA. This inconsistent bidding with respect to their true valuations can lead to inefficiencies, because the deviations from straightforward bidding in the clock phase restricts bidders from bidding up to their true valuations in the supplementary phase. The difference in efficiency and revenue in simulations with bidders bidding on their bid price limit induced by the activity rule and bidders bidding truthful is substantial, even if we assume the same number of supplementary bids being submitted by the bidders. If bidders do not bid up to their true valuations in the supplementary stage, this can have an impact on payments and the allocation of bidders as simulations show. Both, the missing bids problem and restrictions due to inconsistent bidding can lead to payments in the CCA, which are quite different from the VCG or core payments if bidders submitted their valuations truthfully.

Efficiency, simplicity, transparency, and robustness against manipulation are often considered design goals for spectrum auctions. No auction format is perfect and there are always trade-offs that an auctioneer needs to make. For example, a Vickrey auction exhibits dominant strategies, but the payments of bidders are not anonymous and it can happen that two bidders with similar allocations pay vastly different prices, which can cause envy. In a similar way, non-core outcomes can be considered unfair, however, core-selecting auctions cannot have dominant strategies for general valuations. For regulators it is important to understand the properties of different auction formats and make an informed choice. Giving up anonymous prices and the transparency of a simple ascending auction format should only be done if the resulting auction achieves higher efficiency and has stronger incentives for bidders to bid truthful.

The CCA has developed over the recent years and a number of suggestions have been picked up to improve the design. For example, new versions of the CCA will allow for a restricted set of OR bids to address the missing bids issue in large auctions. There have also been suggestions to address problems such as dead ends arising from the current activity rule (Ausubel and Baranov 2014) via stronger activity rules in the clock phase, which enforce straightforward bidding. While the current activity rules can be derived from the Weak Axiom of Revealed Preference (WARP), future activity rules should be based on the General Axiom of Revealed Preference (GARP), which checks for consistency throughout the entire bidding history of a bidder. Such strong activity rules would also avoid problems due to inconsistent bids in the clock phase. However, there are a number issues that need to be considered.

- First, straightforward bidding with a larger number of licenses is challenging for human bidders and probably requires automated bidding agents or decision support for larger auctions with dozens or hundreds of licenses, let alone that there are reasons for bidders not to bid straightforward, such as budget constraints mentioned in the introduction or interdependencies in the valuations of bidders. One might be able to address budget constraints during the auction such that automated agents could be a remedy. However, they would effectively turn the clock phase into a sealed-bid auction, which is then followed by another supplementary sealed-bid stage in the current CCA design. The advantages of such a two-stage design compared to ascending auctions deserve some discussion.
- Second, Bichler et al. (2013b) show that the efficiency of a clock auction with certain types of bidder valuations and straightforward bidding can be close to zero.[7] Not only that the standing bids after the final clock round do not provide an indication for the efficient allocation, also the clock prices do not provide helpful information about the final payments, as can be seen in data from the field and the lab. At least, it is not obvious how bidders should use these price signals from the clock phase.

Both points raise the question, which added value the clock phase provides. One argument in favor of an ascending or dynamic multi-object auction is that bidders do not need to provide all their valuations on exponentially many packages in one step. Levin and Skrzypacz (2014) write that ”economists think of dynamic auctions as having an advantage in this regard because bidders can discover gradually how their demands fit together.” Although the single-stage combinatorial clock auction was shown to be highly efficient in lab experiments apparently helping bidders to find efficiency-relevant bundles, bidders in the lab did not bid straightforward (Scheffel et al. 2012). Overall, using GARP with a traditional clock auction exhibits some challenges.

Many regulators have adopted an ascending auction over sealed-bid alternatives for efficiency reasons. Evan Kwerel, senior economist at the FCC, explained the decision of the US Federal Communications Commission (FCC) to adopt an ascending auction format for selling spectrum licenses by saying: “In the end, the FCC chose an ascending bid mechanism, largely because we believed that providing bidders with more information would likely increase efficiency and, as shown by Paul Milgrom and Robert J. Weber, mitigate the winner’s curse” (Milgrom 2004). The argument draws

[7] Let’s introduce a simple example to better illustrate how straightforward bidding can lead to inefficiency in the clock auction: Consider a market with two items $\{A, B\}$ and three bidders. Bidder 1 has a value of \$10 for A, bidder 2 has a value of \$4 for B and \$10 for $\{A, B\}$, and bidder 3 only has a value of \$10 for the package $\{A, B\}$. If all bidders bid straightforward starting with prices of zero and unit increments, then bidder 2 will never reveal his valuation for A, leading to 71% efficiency. Bidders 2 and 3 would actually drop out at a price of \$5 for both items in the clock stage, which is when bidder 1 still bids on item A. It is easy to extend the example and achieve very low revenue.

on the *linkage principle*, which implies that ascending auctions generally lead to higher expected prices than sealed-bid auctions with interdependent bidder valuations (Milgrom and Weber 1982). In contrast to bidders with independent values, bidders with interdependent values might not always bid consistent as their valuations can change and GARP can be too strong to allow for these changes.

Transparency is also an important argument for ascending auctions as a bidding team needs to set expectations throughout the auction and inform stakeholders. Bidders in SMRA see the final allocation and prices develop throughout the auction, which typically takes several weeks. However, this type of transparency is much reduced in the CCA. How much bidders finally have to pay depends on the bids submitted in the supplementary stage and is a result of a quadratic optimization problem which is almost impossible to predict given the many possible packages bidders can bid on and the missing bids problem. If they are unable to submit a safe supplementary bid, then the allocation can change substantially after the clock phase, as it has happened in the British LTE auction. This makes the outcome of the CCA hard to predict during the auction.

One advantage that an ascending auction still has over a sealed-bid auction is the fact that winners do not need to reveal their valuation for the winning package to the regulator. Regulators need to decide whether this feature outweighs the added complexity stemming from a two-stage CCA. Ascending combinatorial auctions can certainly be of help for bidders in coordinating with other bidders and finding a feasible allocation among the many possible ones. However, if an activity rule enforces straightforward bidding, the possibilities for such coordination will be much reduced.

Designing efficient multi-item auctions is difficult when a regulator needs to consider conflicting design goals such as incentive-compatibility, simplicity, efficiency, and the law-of-one-price. The bid language, the payment rule, and the decision to use a sealed-bid or an ascending format are design choices, which all have significant impact on efficiency and revenue of an auction. A simple bid language can have a substantial positive impact on the efficiency of an auction as was shown in lab experiments (Bichler et al. 2014), and it is not unreasonable to assume similar effects in the field. The pros and cons of different activity rules considering realistic assumptions about bidder preferences in a spectrum auction are still a fruitful area for future research.

Appendix: Details on the Value Model and the Simulations based on the British 4G Auction

The value model used in our simulations in Sect. 6 is based on the British 4G auction in 2013 in which the 800 MHz as well as the 2.6 GHz band were sold (Ofcom 2013b). We will provide a brief description of the British auction and how we derived the value model for each bidder in our simulations, mirroring the main characteristics of this market. The valuations can be made available upon request.

Licenses Up for Sale

Table 6 illustrates the lots used in the auction. We simplified this band plan to allow for an easier analysis. The 800 MHz spectrum was split into two generic lots A(i) and

Table 6 Overview of auctioned lots in the UK 4G auction (Ofcom (2012))

Lot	Amount	Description	EPs	Start price
A(i)	4	2 × 5 MHz paired spectrum in the 800 MHz band	2250	\$225 mn
A(ii)	1	2 × 10 MHz paired spectrum in the 800 MHz band with coverage obligation	4500	£250 mn
C	10/12/14	2 × 5 MHz paired spectrum in the 2.6 GHz band	150	£15 mn
D(i)	≤10	2 × 10 MHz paired spectrum in the 2.6 GHz band (shared low power)	30	£3 mn
D(ii)	≤10	2 × 20 MHz paired spectrum in the 2.6 GHz band (shared low power)	60	£6 mn
E	9	5 MHz unpaired spectrum in the 2.6 GHz band	1	£0.1 mn

Table 7 Overview of auctioned lots for the simplified UK 4G auction scenario

Lot	Amount	Description	EPs	Start price
A	6	2 × 5 MHz paired spectrum in the 800 MHz band	2250	£225 mn
B	14	2 × 5 MHz paired spectrum in the 2.6 GHz band	150	£15 mn
C	9	5 MHz unpaired spectrum in the 2.6 GHz band	1	£0.1 mn

A(ii) where A(ii) has twice as much bandwidth and eligibility points. Furthermore, the winner of A(ii) is obliged to use his spectrum to build a nationwide network. For simplicity, we neglected these legal details in our experiments and combined A(i) and A(ii) into one generic lot A with 6 licenses and the specifications of A(i).

The paired 2.6 GHz spectrum was split into three generic lots C, D(i), and D(ii) with amounts dependent on the bids submitted in the auction. D(i) and D(ii) are shared low power lots of different bandwidth whose winners will jointly use the same frequencies. The British auction rules allowed three different outcomes: First, up to 10 units of D(ii) are sold along with 10 units of C; second, up to 10 units of D(i) are sold along with 12 units of C, or third, the entire paired 2.6 GHz spectrum is sold in 14 units of C. Whichever allocation maximizes revenues wins. Based on the fact that 14 units of C were sold in the British auction and almost no bids containing shared low power lots were submitted, we discarded D(i) and D(ii) in our numerical experiments, as they did not seem to be important for this market.

For the unpaired 2.6 GHz spectrum (band E), only one lot with 9 units was used. However, the number of licenses that can actually be used is lower and dependent on the number of winners since one reserved block per winning package is required as a protection ratio between any two different users. In our simulations, we ignored this limitation and assumed 9 fully useable blocks as C is the least important band in the auction. The resulting list of bands used in our simulations can be found in Table 7.

In addition, a number of spectrum caps were imposed, some of which were also based on existing spectrum holdings in the British auction. For simplicity, we only used one simple spectrum cap that limits the amount of 800 MHz spectrum assigned

Table 8 Final allocation of the British 4G Auction in the simplified band plan

	A (800 MHz paired)	B (2.6 GHz paired)	C (2.6 GHz unpaired)	
Vodafone	2	4	5	Primary bidder
Telefonica	2	–	–	
Everything everywhere	1	7	–	Secondary bidder
Hutchison 3G	1	–	–	
Niche spectrum ventures	–	3	4	2.6 GHz bidder
MLL telecom	–	–	–	Small bidder
HKT company	–	–	–	

to a single bidder to no more than 4 blocks or 2×20 MHz which is common in similar auctions.[8]

Bidders in the British 4G Auction

Seven bidders participated in the auction and five of them won at least one license. As expected, the most valuable target—a pair of A blocks for building a nationwide network with maximum reach—was won by the two big incumbents Vodafone and Telefonica. Table 8 shows the results of the British auction using the simplified lots introduced in the previous section.

Since all auction data has been made publicly available, the segmentation of participants into four generic bidder types is based not only on the results, but also on actual bidding behavior throughout the auction. As illustrated in Fig. 7, there are four bidders who competed in all three bands: the primary bidders Vodafone and Telefonica as well as the secondary bidders Everything Everywhere and Hutchison 3G. The reason for separating them into two groups is the obvious strength of the primary bidders with regard to lot A. For these lots, the primary bidders' bids were much higher than the final clock prices in the supplementary round, compared to both secondary bidders.[9] The 2.6 GHz Bidder Niche Spectrum Ventures was focused on the 2.6 GHz lots only while the small bidders MLL Telecom and HKT Company only competed for the licenses in the C band.

A Value Model for the Simulations

Based on the public bid data, we derived a value model, i.e., valuations for each bidder, which allowed us to run simulations and estimate the impact of different

[8] E.g. German 800 MHz/1.8 GHz/2.0 GHz/2.6 GHz auction in 2010 (Bundesnetzagentur 2010), Danish 800 MHz auction in 2008 (Danish Business Authority 2012).

[9] Since EE is the largest mobile service provider in the UK (Ofcom 2011), it might be surprising to describe them as secondary bidders. However, the classification was solely made based on the bids in this particular auction.

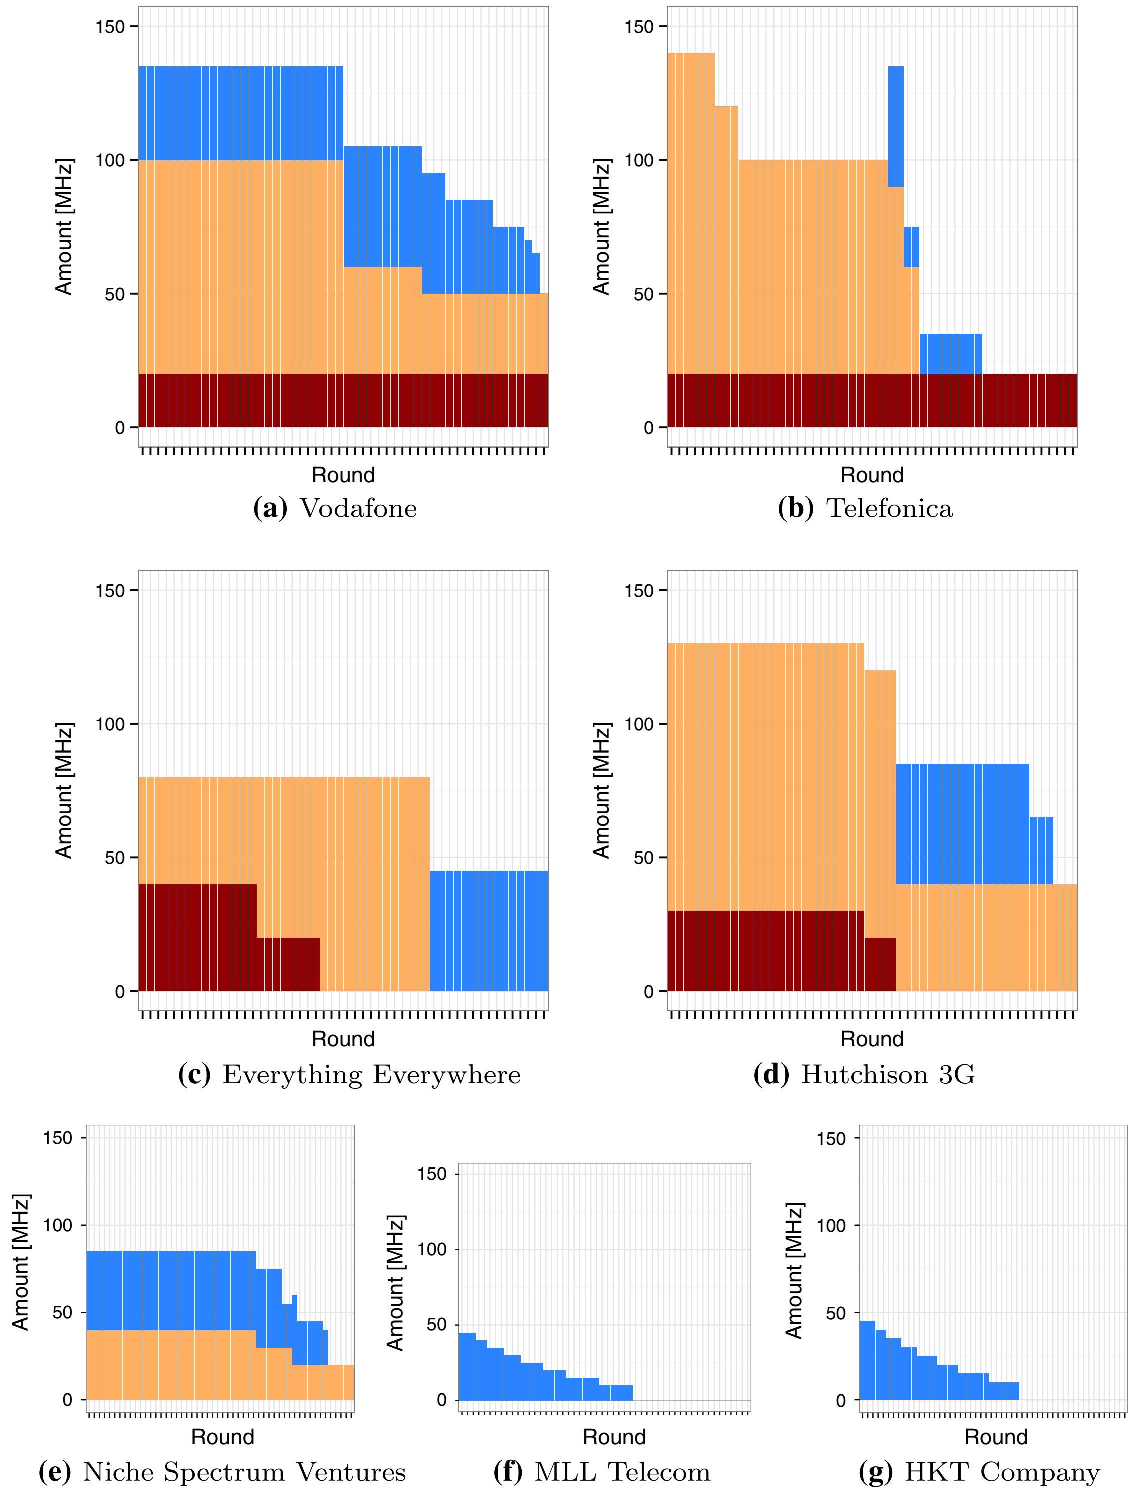

Fig. 7 Visualization of bidding behavior throughout the clock rounds of the British 4G auction with a simplified band plan (A = *red*, B = *apricot*, C = *blue*). **a** Vodafone. **b** Telefonica. **c** Everything everywhere. **d** Hutchison 3G. **e** Niche spectrum ventures. **f** MLL Telecom. **g** HKT company. (Color figure online)

bidding strategies in Sect. 6. First, we defined individual base valuations for each bidder indicating how much he is willing to pay for a single license in a band. Second, intra-band synergies were defined for any package with more than one license within

Table 9 Mean base valuations for the national licenses scenario

	A (800 MHz paired)	B (2.6 GHz paired)	C (2.6 GHz unpaired)
Primary bidder	£300 mn	£70 mn	£15 mn
Secondary bidder	£200 mn	£70 mn	£15 mn
2.6 GHz bidder	–	£70 mn	£15 mn
Small bidder	–	–	£8 mn

the same band up to a certain limit (e.g., 2 blocks of A, 4 blocks of B or C). More licenses of the same band exhibit decreasing marginal value beyond these limits. For the expensive A blocks we even assumed that no bidder is interested in winning more than two licenses. Third, inter-band synergies were defined increasing the value of a bundle comprising licenses from bands A and B. The mean base valuations were defined based on the final clock prices and the supplementary bids and can be found in Table 9. The primary bidders had a much higher valuation in the A band compared to other bidders, while we assumed similar valuations for the B and C bands. Even though the true valuations of bidders are unknown, this allowed for a reasonable sensitivity analysis in Sect. 6.

Based on the mean base valuations v in Table 9 the valuations for each simulation were determined based on two parameters, the relative strength of a bidder s_i and a random influence r_i. Both values are drawn from a uniform distribution in the interval $[v \times 0.75, v \times 1.25]$ and multiplied with the mean base valuations v for each band. For example, consider a primary bidder with relative strength 0.8 and random influence for blocks A and B of 1.1 and 1.0, resp. His valuations are $v_i(A) = 0.8 \cdot 1.1 \cdot £300\,\text{mn} = £264\,\text{mn}$ and $v_i(B) = 0.8 \cdot 1.0 \cdot £70\,\text{mn} = £56\,\text{mn}$. Then we determined the valuation $v_i(nX)$ for different bundles with n licenses within a band X.

$$v_i(nA) = \left(\min\{2, n\} + \min\left\{\frac{1}{2}, \frac{n-1}{n}\right\} \cdot syn_i(A) + \max\{0, ln(n-1)\}\right) \cdot v_i(A) \tag{1}$$

$$v_i(nB) = \left(\min\{4, n\} + \min\left\{\frac{3}{4}, \frac{n-1}{n}\right\} \cdot syn_i(B) + \max\{0, ln(n-3)\}\right) \cdot v_i(B) \tag{2}$$

$$v_i(nC) = \left(\min\{4, n\} + \min\left\{\frac{3}{4}, \frac{n-1}{n}\right\} \cdot syn_i(C) + \max\{0, ln(n-3)\}\right) \cdot v_i(C) \tag{3}$$

The first and second summand correspond to the linear increase in value when adding more blocks and the synergies on top of that. Both only increase in value until a threshold is reached. The third summand is only relevant when additional blocks are added, but with decreasing marginal value. The final valuation for a bundle within a band is computed as the sum of the licenses within a band multiplied by $(1 + syn_i)$. All intra-band synergies $syn_i(A)$ are drawn from a uniform distribution [1.75; 2.25].

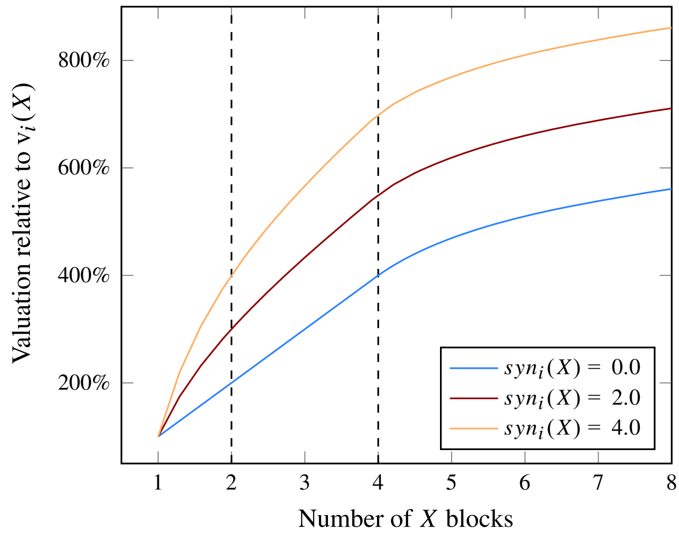

Fig. 8 Plot of intra-band valuations

Only the synergies for the primary bidders in the A band were higher and drawn from a uniform distribution in the interval [3.75; 4.25] assuming that two blocks in A was their primary target. Figure 8 illustrates the resulting valuation function which is only valid for A for up to 2 blocks.

Finally, a uniformly distributed parameter is drawn for each bidder to determine his inter-band synergies for bands A and B. Synergies across these bands can be assumed to be much lower than intra-band synergies, and we use a uniform distribution in the interval [0.0; 0.2]. The valuation for a bundle containing licenses from bands A and B is now computed as the sum of the valuations for inter-band bundles multiplied with $(1 + syn_i)$. For example, a bidder's valuation for a bundle $ABBC$ comprised of one block of A, two blocks of B, and one license in band C is $v_i(ABBC) = (v_i(A) + v_i(2B) + v_i(C)) \cdot (1 + syn_i)$. Based on the random variables above, we generated 10 different instances of the value model.

The correlation between the supplementary bids in the British 4G auction and the valuations generated with the above model is 0.957, indicating that the generated valuations are a reasonable basis for our simulation study.

References

Afriat SN (1967) The construction of utility functions from expenditure data. Int Econ Rev 8(1):67–77

Afriat SN (1973) On a system of inequalities in demand analysis: an extension of the classical method. Int Econ Rev 14(2):460–472

Airiau S, Sen S (2003) Strategic bidding for multiple units in simultaneous and sequential auctions. Group Decis Negot 12(5):397–413. doi:10.1023/B:GRUP.0000003741.29640.ac

Arthur C (2013) 4G auction to be investigated by audit office after poor return. http://www.theguardian.com/technology/2013/apr/14/4g-auction-national-audit-office

Ausubel L, Baranov O (2014) Market design and the evolution of the combinatorial clock auction. Am Econ Rev Pap Proc 104(5):446–451

Ausubel LM (2006) An efficient dynamic auction for heterogeneous commodities. Am Econ Rev 96(3):602–629

Ausubel LM, Milgrom PR (2002) Ascending auctions with package bidding. Front. Theor. Econ. 1(1):1–42. http://www.ausubel.com/auction-papers/ascending-proxy-auctions.pdf

Ausubel LM, Cramton P, Milgrom P (2006) The clock–proxy auction: a practical combinatorial auction design. In: Cramton P, Shoham Y, Steinberg R (eds) Combinatorial auctions, chap 5. MIT Press, Cambridge, pp 115–138. The ISBN-10 is 0262514133

Bellantuono N, Ettorre D, Kersten G, Pontrandolfo P (2013) Multi-attribute auction and negotiation for e-procurement of logistics. Group Decis Negot 1–21. doi:10.1007/s10726-013-9353-7

Bichler M, Shabalin P, Wolf J (2011) Efficiency, auctioneer revenue, and bidding behavior in the combinatorial clock auction. In: Second conference on auctions. Market mechanisms and their applications (AMMA), New York, NY, USA

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. Exp Econ 16:511–545

Bichler M, Shabalin P, Ziegler G (2013b) Efficiency with linear prices? A theoretical and experimental analysis of the combinatorial clock auction. In: INFORMS information systems research, pp 394–417

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

Bundesnetzagentur (2010) Frequency Award 2010. http://www.bundesnetzagentur.de/cln_1931/EN/Areas/Telecommunications/Companies/FrequencyManagement/ElectronicCommunicationsServices/FrequencyAward2010_Basepage.html

Cramton P (2008) A review of the L-band auction. Technical report August. http://works.bepress.com/cramton/11/

Danish Business Authority (2012) Information memorandum—800 MHz auction. http://erhvervsstyrelsen.dk/file/251159/information-memorandum-800mhz-auction.pdf

Day RW, Cramton P (2012) Quadratic core-selecting payment rules for combinatorial auctions. Oper Res 60(3):588–603

Goeree JK, Lien Y (2013) On the impossibility of core-selecting auctions. Theor Econ (forthcoming). http://econtheory.org/

Gul F, Stacchetti E (1999) Walrasian equilibrium with gross substitutes. J Econ Theory 87(1):95–124

Janssen M, Karamychev V (2013) Gaming in combinatorial clock auctions. Technical report, Tinbergen Institute

Jewitt I, Li Z (2008) Report on the 2008 UK 10–40 GHz spectrum auction. Technical report. http://stakeholders.ofcom.org.uk/binaries/spectrum/spectrum-awards/completed-awards/jewitt.pdf

Knapek S, Wambach A (2012) Strategic complexities in the combinatorial clock auction. CESifo Working Paper Series 3983, CESifo Group Munich. http://ideas.repec.org/p/ces/ceswps/_3983.html

Lehmann D, Müller R, Sandholm T (2006) The winner determination problem. In: Cramton P, Shoham Y, Steinberg R (eds) Combinatorial auctions. MIT Press, chap 12. http://www.cs.cmu.edu/~sandholm/winner-determination-final.pdf

Levin J, Skrzypacz A (2014) Are dynamic vickrey auctions practical? Properties of the combinatorial clock auction. Stanford University working paper September

Milgrom P (2000) Putting auction theory to work: the simultaneous ascending auction. doi:10.1086/262118

Milgrom P (2004) Putting auction theory to work. Cambridge University Press, Cambridge

Milgrom PR, Weber RJ (1982) A theory of auctions and competitive bidding. Econometrica 50(5):1089–1122

Ofcom (2011) Everything Everywhere becomes the UKs largest network in terms of revenue. http://stakeholders.ofcom.org.uk/market-data-research/market-data/communications-market-reports/cmr11/telecoms-networks/5.48

Ofcom (2012) The wireless telegraphy (licence award) regulations 2012. http://www.legislation.gov.uk/uksi/2012/2817/contents/made

Ofcom (2013a) 800 MHz & 2.6 GHz auction data. http://stakeholders.ofcom.org.uk/spectrum/spectrum-awards/awards-archive/completed-awards/800mhz-2.6ghz/auction-data/

Ofcom (2013b) 800 MHz & 2.6 GHz combined award. http://stakeholders.ofcom.org.uk/spectrum/spectrum-awards/awards-archive/completed-awards/800mhz-2.6ghz/

Rothkopf MH (2007) Thirteen reasons why the vickrey-clarke-groves process is not practical. Oper Res 55(2):191–197

RTR (2013) Multiband Auction 800/900/1800 MHz. Rundfunk und Telekom Regulierungs-GmbH. https://www.rtr.at/en/tk/multibandauktion

Samuelson PA (1938) A note on the pure theory of consumer's behaviour. Economica 5(17):61–71. http://www.jstor.org/stable/10.2307/2548836

Sano R (2012) Non-bidding equilibrium in an ascending core-selecting auction. Games Econ Behav 74: 637–650

Scheffel T, Ziegler A, Bichler M (2012) On the impact of package selection in combinatorial auctions: an experimental study in the context of spectrum auction design. Exp Econ 15:667–692

Shapiro R, Holtz-Eakin D, Bazelon C (2013) The economic implications of restricting spectrum purchases in the incentive auctions. Georgetown University Washington working paper

Smeulders B, Spieksma FCR, Cherchye L, De Rock B (2012) Goodness of fit measures for revealed preference tests: complexity results and algorithms. aghedupl V(212):1–16. http://home.agh.edu.pl/~faliszew/COMSOC-2012/proceedings/paper_11.pdf

Smith C (2013) Did Ofcom's 4G auction rules cost the UK an extra 3 billion? http://www.techradar.com/news/phone-and-communications/did-ofcom-s-4g-auction-rules-cost-the-uk-an-extra-3-billion-1138302

Varian HR (1990) Goodness-of-fit in optimizing models. J Econom 46:125–140. http://www.sciencedirect.com/science/article/pii/030440769090051T

Varian HR (2006) Revealed preference. In: Szenberg M, Ramrattan L, Gottesman AA (eds) Samuelsonian economics and the twenty-first century, January 2005, chap 6. Oxford University Press, Oxford, pp 99–115

3 Discussion and Conclusion

Although much has been written about the design of efficient spectrum auctions in the past two decades, the design of large-scale markets with complex bidder preferences has been given little attention in the literature as of yet. Traditionally, game-theoretic analyses and laboratory experiments have been used to analyze different auction formats. While these methods are important to gain insights about possible equilibria and the behavior of the participants, these methods have their limitations. In particular, auctions in the field often have many licenses, complex activity and bidding rules. Such design elements are important, but are typically ignored in most theoretical studies due to the added complexity. Simulation studies complement experiments and theoretical work, as they allow economists to study complex market designs under exactly the same rules as used in the field. As we have seen in Papers A to C, replicating realistic bidder behavior can potentially be nontrivial, depending on the market size and auction format.

Even though the used payment rules are not strategy-proof, we argue that incentives for strategic behavior is minimized in large-scale markets such as the ones presented in this dissertation due to the amount of information needed to estimate correctly the valuations or the underlying distribution. Differing from smaller markets or stylized settings (see Goeree and Lien (2014)), such information is rarely available in large real-world markets. The number of packages that bidders could bid on can serve as a guideline for how much information would be needed by a bidder to profitably manipulate a market.

One contribution of this dissertation is the implementation of the SMRA, HPB and various versions of the CCA formats in a unified simulation framework, as well as an instance generator that estimates bidder valuations based on drop out bids from the Canadian 700MHz spectrum auction.

The economic environment in this simulation mirrors the Canadian market with all its institutional details, and it allows us to study the efficiency of wide-spread spectrum auction formats with different levels of synergies in the valuations of bidders. We assume that bidders maximize payoff in each round. This can serve as a reasonable approximation of bidder behavior in larger markets such as the Canadian auction, as explained above. In any case, it is important to understand the average approximation ratio of simple auction algorithms in realistic environments when bidders bid straightforward.

The results are surprising: Even high synergies do not always lead to higher efficiency in the combinatorial clock auctions compared to SMRA, and the relative efficiency ranking depends on the type of synergies. We analyzed two types of synergies motivated from observations in the field. In the "extreme national" synergy model synergies only occur when a bidder wins all licenses in a national package within a specific band. The extreme national synergies create the largest possible risk for a bidder who wants to aggregate licenses in the SMRA and, not surprisingly, the SMRA results in low efficiencies in this model. More moderate synergies occur when the marginal value of a license rises linearly with the number of licenses won. Under this assumption of "linear" synergies, the SMRA outperforms various versions of the CCA, in terms of efficiency as well as revenue. Overall, it is interesting to observe that the average efficiency loss in both models is remarkably low, considering the simplicity of the algorithms and the worst-case approximation ratio of the allocation problem in combinatorial auctions.

Laboratory and field data can give insights into the sometimes irrational behavior of market participants, with sometimes surprising consequences. By developing a metric based on Afriat's Critical Cost Efficiency Index, we demonstrate that bidders did indeed deviate from straightforward bidding in the clock phase in both the British 4G auction in 2013 and the Canadian 700 Mhz auction in 2014. This, in turn, can limit the amounts bidders can bid on in the supplementary phase of these auctions, with significant negative consequences on both efficiency and revenue.

Apart from spectrum auctions, other examples where complex bidder preferences and allocation constraints lead to computationally hard allocation problems can be found in several industries, including the sale of TV ads to media agencies or multi-item and multi-unit industrial procurement auctions. Because of their full expressiveness, the XOR bid languages and similar designs have been used in markets such as spectrum and procurement, but this does not scale to very large markets due to the exponential growth in the number of package bids that can be submitted.

To address this, we describe an auction design framework using compact bid languages and payment rules which incentivize truthful bidding. Compact bid languages can often draw on domain specifics and allow bidders to describe their preferences with a low number of parameters that they have to specify as the TV ads market and the volume discount auctions in Paper A illustrate. Commercial off-the-shelf mixed integer programming solvers can now solve large and realistic instances of such problems to near optimality on standard hardware, which allows us to use such bid languages in real-world markets.

In sealed-bid auctions second price rules such as VCG or BPOC payment rules can be used to provide incentives for truthful bidding. In many markets, auctioneers would prefer core pricing to VCG mechanisms, in order to avoid non-core outcomes where the bids of losing bidders are higher than the payments of the winners. With the introduction of core-selecting auctions (such as the CCA in spectrum auctions) and algorithms (see Day and Raghavan (2007); Day and Cramton (2012)), stake-holders have developed software to determine winners and core prices based on the use of integer programming to solve a series of winner-determination problems.

Extending the use of this software to larger and more complex markets (such as the TV ads and procurement contexts we presented) cannot be accomplished by merely specifying time limits or optimality-gap thresholds to the solver engine, as it could for the more simple case of a single optimization problem. Doing so would often result in an infeasible pricing problem. This general problem exists for all larger markets with near-optimal winner determination. We developed and compared two potential algorithms for dealing with these infeasibilities, finding one faster and higher revenue method (for a fixed set of bids) and one slower but more efficient method.

The toolset and analyses provided by this dissertation allow a further exploration of complex markets along several dimensions. The study on the extension of core-selection auctions beyond provably-optimal winner determination settings as the ones shown in Paper A might enable the usage of sophisticated pricing rules and smart market clearing mechanisms for markets even more complex as the ones we studied.

At the same time, the ever evolving spectrum auction formats such as the Combinatorial Clock Auction allow for a deeper investigation: During the last years, the CCA's bidding rules were adopted and refined, leading to new activity rules which in theory should induce truthful behavior. We have demonstrated that this did not always succeed. The introduction of planned stricter activity rules might limit bidders even further, while more sophisticated bidding languages might counter the bidders' tendency to only select a small subset of valuable packages. Bidders might also not bid truthfully due to strategic reasons. As shown in the analyses by Goeree and Lien (2014); Knapek and Wambach (2012) and Janssen and Karamychev (2016), core-selecting auctions in general and certain implementations of the Combinatorial Clock Auction in particular can enhance a bidder's payoff by deviating from truthful bidding. This is especially so if bidders have budget constraints (Janssen and Karamychev, 2013) or a lexicographic preference to raising their rivals' costs. An extension of the provided analyses might shed light on the changes in efficiency and revenue if this strategic behavior is included in computational experiments similar to our own.

We show that the frameworks and analyses provided in this dissertation offer a first approach in faithfully recreating auction formats used in complex settings and spectrum auctions, including all intricacies. Extending the research by testing new auction formats and bidder valuation functions will allow for further insights into complex markets such as the selling of wireless spectrum.

Bibliography

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

L. Ausubel and P. Milgrom. The lovely but lonely vickrey auction. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. MIT Press, Cambridge, MA, 2006.

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

M. Bichler, A. Davenport, G. Hohner, and J. Kalagnanam. Industrial procurement auctions. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. MIT Press, 2006.

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

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

Sushil Bikhchandani and Joseph M Ostroy. The package assignment model. *Journal of Economic theory*, 107(2):377–406, 2002.

C. Boutilier and H. H. Hoos. Bidding languages for combinatorial auctions. In *17th International Joint Conference on Artificial Intelligence (IJCAI) 2001*, pages 1211–1217, Washington, USA, 2001.

C. Caplice. Electronic markets for truckload transportation. *Production and Operations Management*, 16(4):423–436, 2007. ISSN 1937-5956. doi: 10.1111/j.1937-5956.2007.tb00270.x. URL `http://dx.doi.org/10.1111/j.1937-5956.2007.tb00270.x`.

P. Cramton. A review of the L-band auction. Technical report, 2008. URL `http://stakeholders.ofcom.org.uk/binaries/spectrum/spectrum-awards/completed-awards/714758/Report_by_Professor_Peter_C1.pdf`.

P. Cramton. Spectrum auction design. *Review of Industrial Organization*, 42, 2013. URL `http://ideas.repec.org/p/pcc/pccumd/09sad.html`.

P. Cramton, Y. Shoham, and R. Steinberg. Introduction to combinatorial auctions. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. MIT Press, Cambridge, MA, 2006.

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

R. W. Day and P. Cramton. Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research*, 60(3):588–603, 2012.

J. Goeree and C. Holt. Hierarchical package bidding: A paper & pencil combinatorial auction. *Games and Economic Behavior*, 70(1):146 – 169, 2010. ISSN 0899-8256. doi: 10.1016/j.geb.2008.02.013.

J. Goeree and Y. Lien. On the impossibility of core-selecting auctions. *Theoretical Economics*, 2014.

A. Goetzendorff. An auction format for the award of TV advertisement rights, 2013.

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

M. Janssen and V. Karamychev. Gaming in combinatorial clock auctions. Technical report, Tinbergen Institute, February 2013.

M. Janssen and V. Karamychev. Spiteful bidding and gaming in combinatorial clock auctions. Technical report, 2016.

Stephan Knapek and Achim Wambach. Strategic Complexities in the Combinatorial Clock Auction. CESifo Working Paper Series 3983, CESifo Group Munich, 2012. URL `http://ideas.repec.org/p/ces/ceswps/_3983.html`.

Vijay Krishna. *Auction Theory*. Academic Press, 2009.

D. Lehmann, R. Müller, and T. Sandholm. The Winner Determination Problem. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 12. MIT Press, 2006.

B. Lubin, B. Bünz, and S. Seuken. New core-selecting payment rules with better fairness and incentive properties. In *Proceedings of the 3rd Conference on Auctions, Market Mechanisms and Applications (AMMA)*, 2015.

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

P. Milgrom. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108(21):245–272, 2000.

P. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.

P. R. Milgrom and R. J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, 1982.

Asuncion Mochon and Yago Saez. A review of radio spectrum combinatorial clock auctions. *Telecommunications Policy*, 41(5):303 – 324, 2017. ISSN 0308-5961. Optimising Spectrum Use.

Noam Nisan and Liad Blumrosen. Combinatorial auctions. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani, editors, *Algorithmic game theory*, volume 1. Cambridge University Press Cambridge, 2007.

D. Parkes and L. H. Ungar. Iterative combinatorial auctions: Theory and practice. In *17th National Conference on Artificial Intelligence (AAAI-00)*, 2000.

A. Pikovsky. Pricing and Bidding Strategies in Iterative Combinatorial Auctions, 2008.

D. Porter, S. Rassenti, A. Roopnarine, and V. Smith. Combinatorial auction design. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 100:11153–11157, 2003.

Michael H Rothkopf. Thirteen reasons why the vickrey-clarke-groves process is not practical. *Operations Research*, 55(2):191–197, 2007.

T. Sandholm and C. Boutilier. Preference elicitation in combinatorial auctions. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. MIT Press, Cambridge, MA, 2006.

Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial intelligence*, 135(1-2):1–54, 2002.

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

S. Schneider. Compact Bidding Languages and Supplier Selection for Markets with Economies of Scale and Scope, 2011.

A General Auction Simulator

For static scenarios, the implementation of a small number of mixed integer linear problems (MILP) and the computation of these MILPs with a linear solving library is often a feasible approach. All presented experiments in this dissertation required user input (in our lab experiments) or were highly dynamic simulations, however. The presented research therefore required adequate computer aided experimental frameworks. While Paper A and C of this dissertation leverages in part prior software[1], the simulation and analysis of large regional spectrum auctions as presented in Paper B required the creation of an entirely new simulation framework.

We designed the *General Auction Simulator* (GAS) to allow us the conduction of computational experiments of this scale. The software forms a central research tool for the simulations taken in this dissertation. The main requirements we elicited for the software are as follows:

- Spectrum auction formats – In order to be able to compare spectrum auction formats in terms of efficiency and revenue, a selected number of auction formats has to be supported. This includes the Simultaneous Multi Round Auction, the HPB auction, the single stage Combinatorial Clock Auction (CCA), and the two stage CCA (see also Paper B).

[1] See Goetzendorff (2013); Schneider (2011) for the *TvAuction* and *Supplier Selection* software (Paper A) and Pikovsky (2008) for *MarketDesigner* framework extended for Paper C.

- Bidding behavior – The simulation framework should include the possibility to easily load and configure the valuation model and bidding behavior for a bidder at runtime.
- Idempotency – To guarantee replicability, the framework should always behave deterministically: Given the same input, the same result should be returned.
- Efficiency – Solving the Winner Determination Problem and deciding on which package to bid are both *NP*-hard problems (see Paper B). As such, the software should be as efficient as possible, regarding the usage of CPU, RAM and time. This becomes especially important in our computational experiments, as we have to simulate a high number of different scenarios.
- Extensibility – The design of spectrum auction formats is steadily evolving. As such, it was not possible for us to know all auction formats GAS should support in advance. A modular and extensible architecture is therefore needed, to allow a fast addition of new auction formats and bidding languages.
- KISS ("keep it small and simple") – The software should be designed to be as simple as possible. Specifically, adding a new feature should only be done if there is a real need for it.

The GAS framework is designed as a multi-tiered application based on several loosely coupled Python packages, each tier leveraging modules of the lower tiers. External dependencies are kept to a minimum. In order to understand the basic architecture of the framework, we will give a short overview of the framework's structure, its internal dependencies, and demonstrate a typical simulation flow.

For the structural dependencies between the different classes, please refer to Figure A.1. The basic item of both every auction format and bidder valuation is the *Item*. An *Item* represents the good the auctioneer wants to sell and contains properties that are inherent to it, such as how many of it are available, and what its initial price is. In order to reason about a collection of items, it sometimes makes sense to introduce more sophisticated

Tier	Package	Description
1	value-models auction-formats	base item structure, bidder valuation functions base auction format structure, auction formats
2	behavior auction	bidding behaviors and strategies auction format and bidder front controllers
3	runners simulation	synchronous, async. and distributed simulation drivers simulation scenario entry and persistence strategy

Table A.1: *GAS* packages

data structures as e.g. a list. The *ItemStructure* allows exactly this: All implementations of this interface provide methods that allow to select and filter the underlying data by a variety of criteria, such as the filtering by frequency band.

An *AuctionFormat* makes use of the *ItemStructure* and implements the set of rules defining the specific auction format. Data objects representing the bidders' state, the *BidderInfo*, can be registered to the auction format. Whenever the *AuctionFormat*'s `clear()` method is called, the underlying winner determination problem is solved and its result are written into the AuctionFormat's *RoundInfo* object. At the same time, all *BidderInfo* objects are updated.

In order for a bidder to reason about on which packages to bid, a bidder's *Strategy* can consult its *ValueModel*. The *Selector* allows the *Strategy* to find the most valuable (i.e., payoff maximizing) bundles given a set of prices.

Communication between the clients (i.e., the different *Strategy* objects) and the server (i.e., the *AuctionFormat*) is done by the two classes *Agent* and *Auction*, respectively. By registering an *Agent* to the *Auction*, the *Agent* gets notified on all relevant state changes (see also fig. A.2) and is able to respond with *Action*s generated by its *Strategy*.

A minimal code example for an auction driver can be found in Listing A.1. Note that the interaction between auction and agents is not dependent on the driver, but is handled automatically as soon as the auction's `open()` method is called. How the observing *Agents* are acting when this happens is shown in the sequence diagram in Figure A.3:

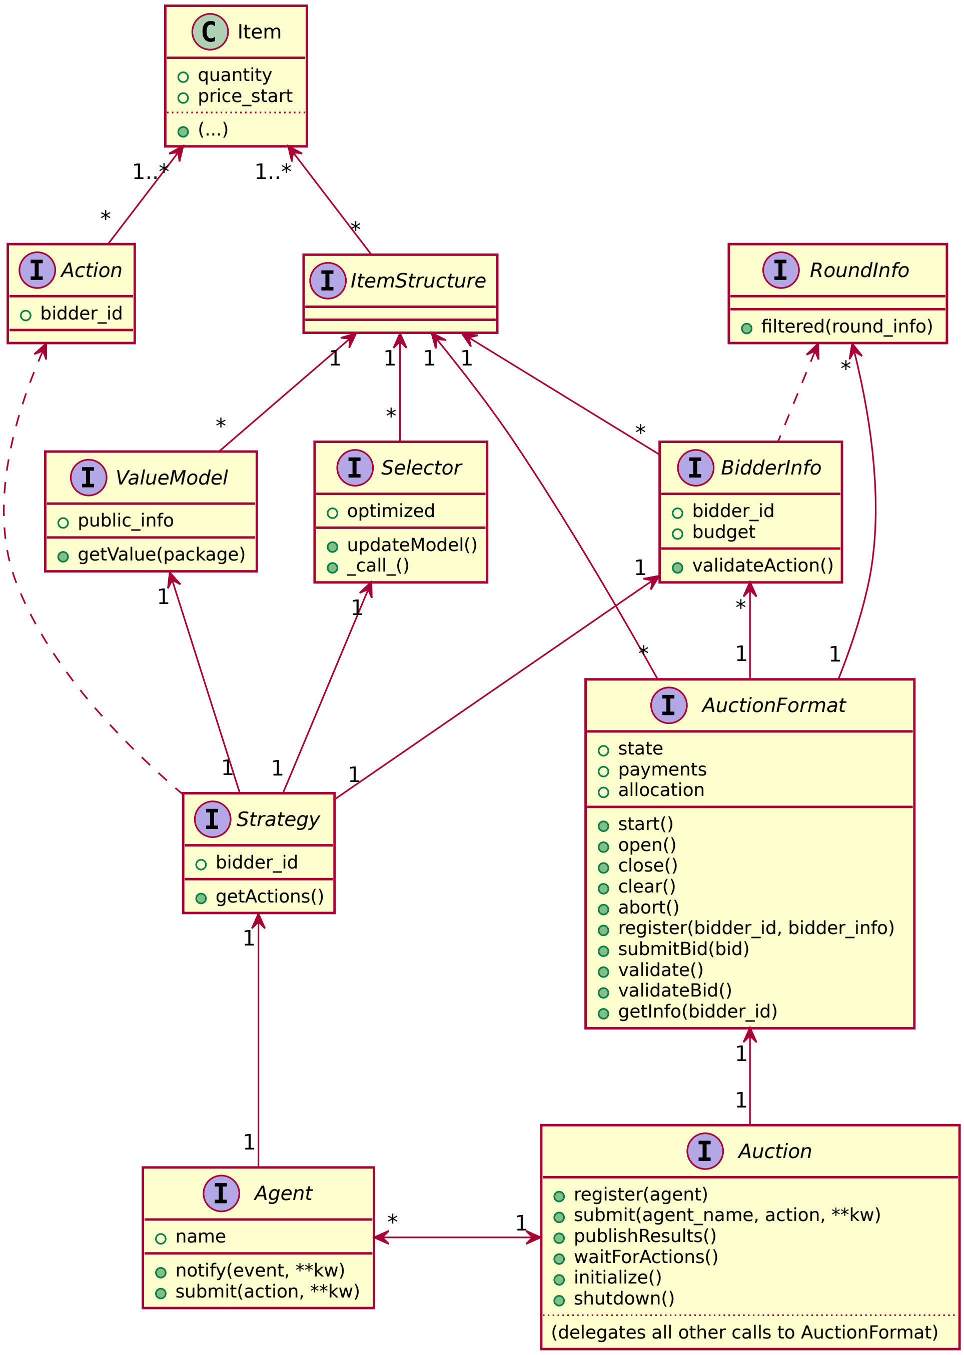

Figure A.1: *GAS* Structure

After initializing all relevant objects, registering the agents (**1**, **2**), and starting the auction, the auction's `open()` method is called (**3**). This causes the auction to `notify()`

```python
def runAuction(auction, agents):
    for agent in agents.itervalues():
        auction.register(agent)
    auction.start()
    while auction.state != STATE_FINISHED:
        auction.nextRound()
        auction.open()
        auction.close()
        auction.clear()
    auction.publishResults()
    auction.shutdown()
    result = (auction.allocation, auction.payments)
    return result
```

Listing A.1: A Minimal Auction Run Example

all listening agents with the `STATE_CHANGE` event and the additional named values {`state=STATE_OPEN, action_requested=True`}, among others.

This causes the agent in this example to query its strategy what kind of actions the agent should perform (**5**). At this point, the strategy updates the selector with the most current round information (**6**) gets the current most valued demand set (**7** to **10**). The strategy then uses this combination to construct a bid and validates it (**11**), before submitting it to the auction (**14**). The submitted bid then gets validated on the auction server (**16**) before issuing the response.

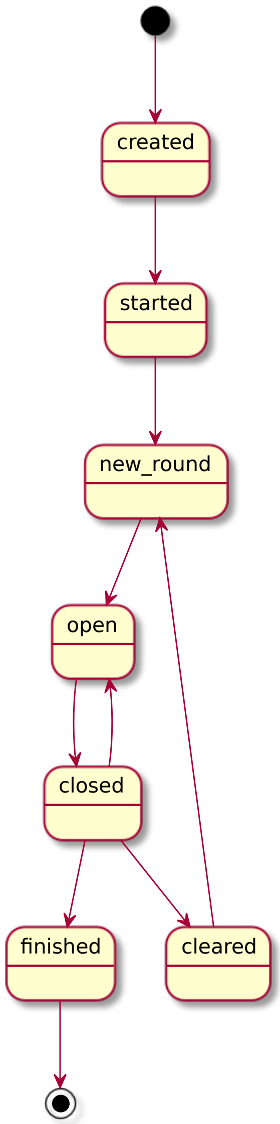

Figure A.2: Auction States

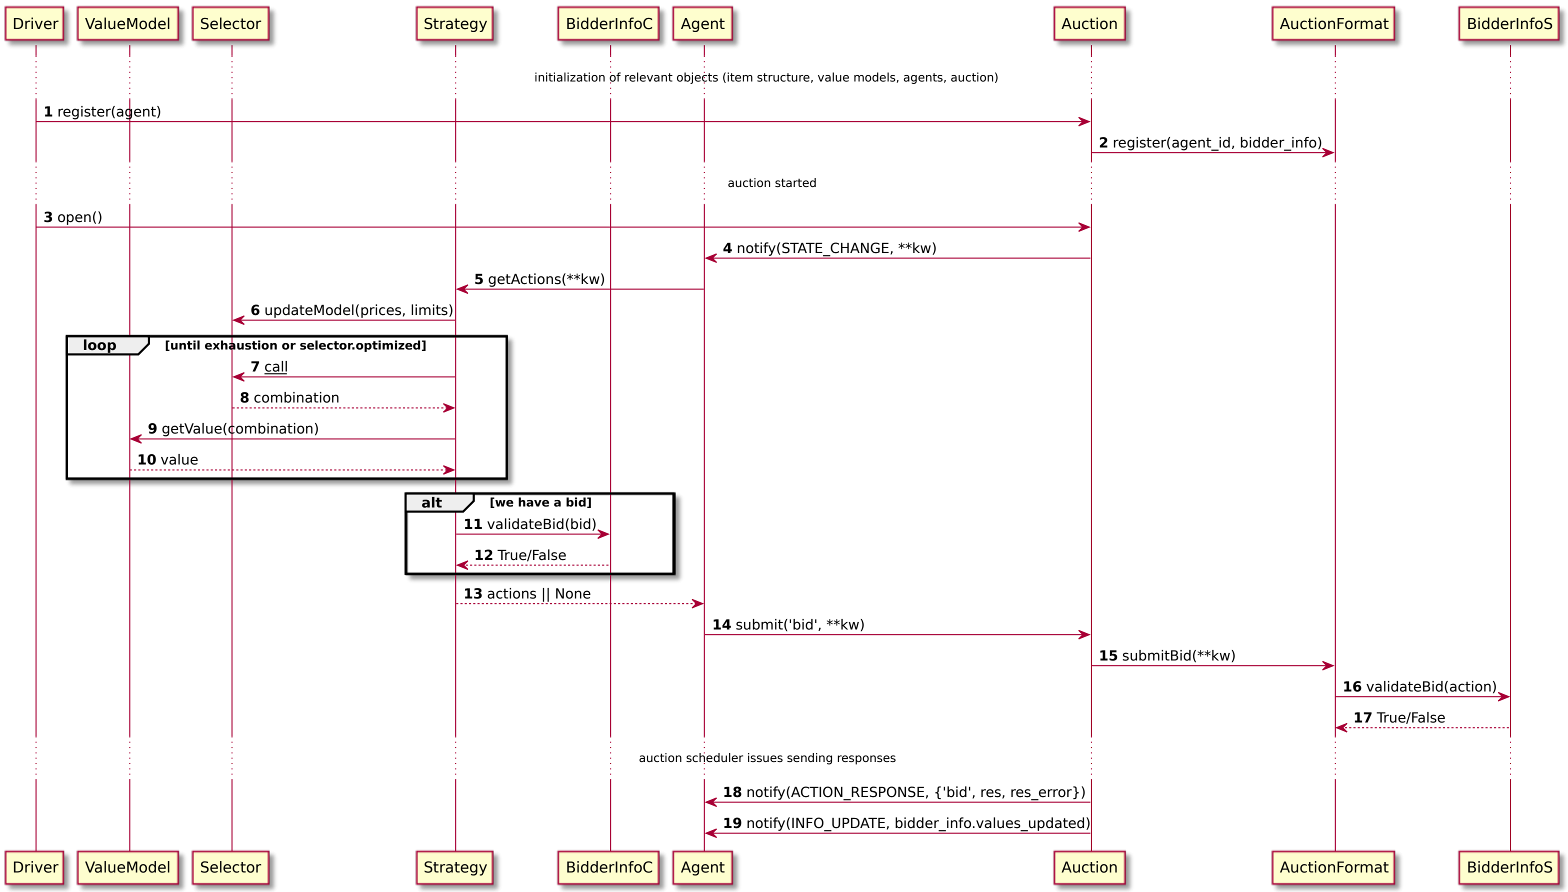

Figure A.3: Bid Submission Sequence