



Sealed-bid versus ascending spectrum auctions

Nicholas C. Bedard¹ · Jacob K. Goeree² · Philippos Louis³ · Jingjing Zhang⁴

Received: 25 February 2022 / Revised: 8 May 2023 / Accepted: 9 May 2023 /

Published online: 26 May 2023

© The Author(s) 2023

Abstract

We compare the first-price sealed-bid (FPSB) auction and the simultaneous multiple-round auction (SMRA) in an environment based on the recent sale of 900 MHz spectrum in Australia. Three bidders compete for five indivisible items. Bidders can win at most three items and need to obtain at least two to achieve profitable scale, i.e. items are complements. Value complementarities, which are a common feature of spectrum auctions, exacerbate the “fitting problem” and undermine the usual logic for superior price discovery in the SMRA. We find that the FPSB outperforms the SMRA across a range of bidding environments: in terms of efficiency, revenue, and protecting bidders from losses due to the exposure problem. Moreover, the FPSB exhibits superior price discovery and almost always results in competitive (“core”) prices unlike the SMRA, which frequently produces prices that are too low because of demand-reduction or too high because of the exposure problem. We demonstrate the robustness of our findings by considering two-stage variants of the FPSB and SMRA as well as environments in which bidders know their own values but not the distributions from which values are drawn.

Keywords Spectrum auctions · Laboratory experiments · Price discovery · Exposure problem · Market design

JEL Classification C78 · C92 · D47

We would like to thank the Australian Research Council (DP190103888 and DP220102893) for financial support. We thank staff of the Australian Communication and Media Authority (ACMA) and participants at the 2018 BEAT market-design conference in Beijing for helpful comments. The replication material for the study is available at DOI 10.17605/OSF.IO/NTVPB.

✉ Philippos Louis
phlouis@gmail.com

¹ University of Western Ontario, 1151 Richmond St, London, ON N6A 3K7, Canada

² AGORA Center for Market Design, UNSW, Sydney 2052, Australia

³ Department of Economics, University of Cyprus, 1678 Nicosia, Cyprus

⁴ School of Economics, Finance, and Marketing, RMIT, 445 Swanson Street, Melbourne, Australia

1 Introduction

Since pioneered by the US Federal Communication Commission in 1994, the simultaneous multiple-round auction (SMRA) has come to dominate the allocation of radio spectrum, earning hundreds of billions of dollars for treasuries around the world (Milgrom, 2004). Its procedure is simple: all items are put up for sale simultaneously with a separate price associated with each item. Bidders can bid on any subset of items they wish and the auction ends only when no new bids are made on *any* of the items. At the end of the auction, bidders win the items they bid highest on and pay the price they bid. Analysis of the SMRA for allocating spectrum licenses since 1994 suggests that it has been exceedingly successful, with allocations generally thought to be efficient and revenues high (Cramton, 1997).

In computational terms, the SMRA can be seen as a “greedy algorithm” used to solve the following “fitting problem:” items need to be allocated to the bidders but it is not clear who gets what and at what prices. The algorithm starts at zero (or low reserve) prices at which aggregate demand does not fit into supply. As prices rise, the algorithm picks off demand that is no longer profitable for bidders and increases prices until excess demand is eliminated, i.e. until demand and supply fit.¹

The use of the SMRA is typically motivated by its ability to produce correct prices, i.e. competitive prices that clear the market and allocate goods efficiently across agents (Ausubel & Cramton, 2004; Milgrom, 2004).² This motivation, however, rests on the assumption that goods are substitutes so that bidders reduce their demands over the rounds of the auction as prices rise until supply equals demand.³ This logic breaks down when goods are complements, as they are in many spectrum auctions. At high prices, bidders may only be interested in *large* packages. As a result, the fitting problem the SMRA is intended to solve becomes more severe as the auction proceeds—at low prices, although demand exceeds supply, bidders may be willing to buy packages of various sizes, but at high prices, bidders may only be willing to buy large packages that cannot be fit to match supply. Moreover, the SMRA assigns provisional winners for each item in each round, which exposes

¹ The algorithm is greedy in this sense, which makes it ideally suited to auctions with many items, such as many of the spectrum auctions in the US and Canada (often with hundreds of items). Greediness is less of an advantage when few items are being sold, as in the recent spectrum auctions in Australia and across Europe.

² When competitive prices do not exist, recent research suggests using the *core* as a proper benchmark for “reasonably” competitive outcomes. Core outcomes are reasonably competitive in the sense that there does not exist a coalition of players (including the seller) that would be better off exiting the market and trading amongst themselves; that is, prices reflect the opportunity costs of the allocation. These constraints are weaker than those imposed by competitive equilibrium. See for example Milgrom (2004), Day and Cramton (2012), Day and Raghavan (2007), Day and Milgrom (2008), Goeree and Lien (2016) and Bichler and Goeree (2017a).

³ Milgrom (2004) and Gul and Stacchetti (1999) prove that when the items for sale are substitutes and bidders bid on subsets of items that provide the highest possible profit (i.e. straightforward or myopic bidding), prices will be competitive and the allocation will be efficient if the bid increment between rounds is sufficiently small.

bidders to the risk of winning a subset of their desired package.⁴ Goeree and Lien (2014) show that, in equilibrium, this *exposure problem* forces bidders to reduce their demands too early in the auction resulting in low prices and low efficiency. This raises the broader question of how to design a practical vehicle for implementing efficient and competitive outcomes when complementarities cannot be ruled out and one cannot rely on prices to guide participants to efficient and stable outcomes.

Some attention has been directed towards comparing alternative dynamic combinatorial auction formats. For instance, Kagel et al. (2010, 2014) compare the SMRA with a combinatorial clock auction with complementarities between items. They identify environments under which each auction format can be more efficient. Munro and Rassenti (2019) suggest that a *descending* price auction can sometimes solve the fitting problem better than an ascending format.

There has been less scrutiny of *sealed-bid* combinatorial auction formats. In principle, the existence of complementarities does not hinder combinatorial sealed-bid auctions from delivering efficient outcomes. For instance, the generalization of the second-price sealed-bid auction, i.e. the Vickrey–Clarke–Groves (VCG) mechanism (Vickrey, 1961), is one such option. Alas, while efficient, it suffers from significant shortcomings for practical applications (Ausubel & Milgrom, 2006; Rothkopf, 2007). We therefore focus on *first-price sealed-bid* (FPSB) combinatorial formats. While the first-price rule is ubiquitous in single-unit auctions, little is known about its performance in the presence of value synergies.⁵ In this paper we take a step towards filling this gap by examining the FPSB in an environment with strong complementarities.

Another reason for studying these formats is that the Australian Communication and Media Authority (ACMA) asked for advice regarding this format [see Bichler and Goeree (2017b) and Goeree and Louis (2019)]. To compare the performances of the SMRA and FPSB formats under realistic scenarios, our experimental framework mimics the environment of the actual Australian 900 MHz auction as close as possible. Three bidders, who put most value on packages of two or three items, vie for five indivisible items, one of which is of lower value.⁶ The combination of item complementarity and heterogeneity creates a potentially demanding bidding environment for buyers in either format. In the SMRA, buyers must manage the

⁴ This problem would be less severe if resale would be allowed, but in many jurisdictions this is not the case. In all of our theoretical analysis and the experiment we do not allow for the possibility of resale. See Filiz-Ozbay et al. (2015) for some theoretical results when resale is possible.

⁵ See, however, Kim (1996), Levin (1997) and Gentry et al. (2019).

⁶ Several features of the recent Australian 900 MHz auction are commonly encountered in other spectrum auctions. There are typically few buyers, a small number of items (i.e. blocks of spectrum) and high complementarities on either the second or third *contiguous* item obtained by the buyer. For example, auctions are being planned or have taken place in similar environments in Australia, Canada, Denmark, Italy, Austria, Switzerland, Belgium, Greece, the Netherlands and the UK (Klemperer, 2002; Earle & Sosa, 2013). Contiguity is important in these auctions in terms of the frequency of the blocks for sale. In addition, in many such auctions regulators designate a block of spectrum to be a “guard band” and impose technical restrictions on its use; this tends to lower its value relative to the remaining, otherwise identical, blocks. This heterogeneity requires that an allocation specify which particular items are won, rather than simply the number of items won.

risks of acquiring too few items and of acquiring fragmented packages (thereby losing complementarities in either case). The FPSB auction solves these problems for buyers through its package bidding format but requires bidders to arrive at an optimal strategy introspectively.

Overall, we find that the FPSB performs better than the SMRA across a range of bidding environments, in terms of efficiency, revenue, price discovery and protecting bidders from losses.⁷ Surprisingly, despite its “static nature”, the FPSB exhibits excellent price discovery properties, in that it almost always results in competitive (“core”) prices.

We also test two-stage variations of the SMRA and FPSB; these formats first establish generic quantities of items won, then determine the assignment of specific items in a second round of bidding. Such variations are thought of as potential remedies for some of the fragmentation problems present in the SMRA in practice.⁸ The two-stage FPSB still outperforms the two-stage SMRA. In both formats, however, buyers tend to bid too aggressively in the first stage, apparently underestimating the cost of avoiding the lower value item in the second stage.

1.1 Organization

The next section presents a simple theoretical model, more amenable to game theoretic analysis than the environment used in our experiments, to gain some qualitative insights into the different forces at play in each of the auction formats we consider. Section 3 contains the experimental design and procedures. Section 4 compares the FPSB and SMRA in terms of efficiency, revenue, and bidders’ profits across several environments. Price discovery is analyzed in detail in Sect. 5. The final Sect. 6 offers some conclusions and suggestions for future research. The appendices provide details about the Bayes–Nash equilibrium analysis and the experimental instructions.

2 A theoretical analysis

In this section we develop a stylized model of our experimental environment. In particular, we abstract from item heterogeneity and make strong assumptions about buyer valuations. Our intent is to generate a broad sense of the incentives at play and how these affect equilibrium outcomes. Since the analysis does not precisely match our experimental conditions, the theoretical findings are qualitative, not quantitative, in nature and (large) deviations between predicted and observed bidder behavior should be expected. Nonetheless, as we demonstrate in later sections, the model’s qualitative predictions regarding the performance of the FPSB auction relative to the SMRA are broadly consistent with the experimental evidence.

⁷ Bidders’ post-auction financial viability is essential to the efficient use of spectrum. For example, albeit not related to exposure risk, bankruptcy proceedings of a successful bidder in the 1994 American spectrum auctions precluded the use of valuable spectrum for nearly 10 years (Cramton et al., 2011).

⁸ See for example the UK 2.3 GHz and 3.4–3.6 GHz spectrum auction in 2018.

2.1 Bidding environment and value model

Three bidders compete for five items. We assume there are strong complementarities between items but allow those synergies to be strongest either when going from one to two items or from two to three items. Formally, there are two types of bidders: type X bidders who need exactly two items (i.e. they place zero value on a single item and zero marginal value on each item above two), and type Y bidders who need exactly three items (i.e. they place zero value on obtaining one or two items and zero marginal value on each item above three). We study multiple combinations of these bidder types: auctions with composition XXX , XXY , XYX , or YYY . An X type bidder draws a valuation for any pair of items uniformly from $[0, 1]$ and a Y type bidder draws a valuation for any three items uniformly from $[0, \alpha]$ for $\alpha > 1$. We consider only equilibria wherein the same types bid the same way, if such an equilibrium exists.

2.2 Bayesian nash equilibria

We summarize the structure of the equilibria for the various auctions here while relegating the technical details to Appendices A.1 to A.3. Before describing the first price sealed bid and the simultaneous multiple-round auctions, we first discuss a common theoretical benchmark for auction performance.

2.2.1 The Vickrey–Clark–Groves mechanism

An idealized benchmark to which auction formats are often compared is the Vickrey–Clark–Groves (VCG) mechanism (Vickrey, 1961).⁹ The VCG mechanism always allocates the items efficiently and its equilibrium outcome is always in the core in the environments we consider. Given its many practical drawbacks, see Ausubel and Milgrom (2006), this format is rarely used in practice and we likewise do not test it in our experiment. Nevertheless, it provides a useful theoretical comparison as a minimally-competitive mechanism.

In the VCG mechanism, bidders report their values to the seller and, based on these reports, the seller chooses the allocation that maximizes total surplus (i.e. the efficient allocation). Payments are designed such that it is a dominant strategy for

⁹ The VCG is the generalization of the single-item second-price auction to auctions with more than one item. It is the unique mechanism that induces truthful bidding, see Holmstrom (1979) and Green and Laffont (1979), using an externality-based pricing rule. Note that because of the combinatorial nature of the assignment one cannot simply elicit bids for all packages and then determine the price of each package as the highest losing bid for that package. Suppose, for instance, that there are five units and three bidders and each bidder can win at most three units. Each bidder submits a triplet (b_1, b_2, b_3) where b_k is the bid for $k = 1, 2, 3$ units. Suppose bidders 1 and 2 submit the same triplet $(10, 10, 10)$ and bidder 3 submits the triplet $(20, 20, 100)$. The restriction that bidders can win at most three items means the outcome is that bidder 3 gets three units and bidders 1 and 2 each get one unit. However, bidders 1 and 2 cannot be forced to pay 20 for their single unit.

bidders to report their true values to the seller. See Online Appendix A.1 for the formal description of the mechanism.

2.2.2 The first price auction

In the first price auction, bidders submit one bid for every possible package (i.e. subset) of items, which they pay if and only if they win the package. Since the X type bidders only value a pair of items, we need only consider their bids for two items; i.e. bids on one or three items are zero. Similarly, since the Y type bidders only value a package of three, we need only consider her bids for three items. An equilibrium in the first price auction in a particular environment will consist of a bidding function for each of the types present in that environment. The seller determines the feasible combination of bids that maximize revenue. Since each bidder only bids on a single package, these revenue-maximizing allocations are relatively simple to describe. In the XXX and XXY environments, at most two bids can be fulfilled; the seller allocates the requested number of items each to the top two bidders and nothing to the lowest bidder. In the XYX environment, the seller can fulfill at most one bid from a Y type (for three items) and one bid from the X type (for two items); the seller allocates three items to the highest bidding Y type and two items to the X regardless of her bid. In the YYY environment, at most one bid (for three items) can be fulfilled; the seller allocates three items to the highest bidder. With these revenue-maximizing allocations, equilibrium bid functions can be found using standard differential analysis. These calculations are described in Online Appendix A.2.

2.2.3 The simultaneous multiple-round auction

The simultaneous multi-round auction (SMRA) is modelled using five price clocks (one for each item), each of which ticks upward from zero whenever two or more bidders demand (i.e. bid on) the associated item. Bidders can only decrease the number of items they bid on after the auction starts. Given bidders preferences, each bidder will either bid on her entire demand (i.e. two items for type X , three items for type Y) or on no items; in the latter case we say the bidder is inactive or has dropped out. If only one bidder demands a particular item, its price clock is paused and this bidder is declared the provisional winner. If other bidders later demand this item, the price clock restarts and the item becomes provisionally unassigned. When demand on all items is at most one, the auction ends, items are assigned to their provisional winners and the winners pay the prices on the clocks for the items they won.

An equilibrium in the SMRA consists of a bidding function for each of the types present in the environment conditional on which types remain in the auction and at which prices others have dropped out. The equilibrium calculations are described in Online Appendix A.3. In environments XXX and XXY , as soon as any bidder drops out, the auctions ends. In the XYX environment, the auction ends after a type Y bidder drops out but continues after the X type drops out. In the YYY environment, the auction ends only after two Y type bidders drop out. In equilibrium, one bidder is randomly chosen to abstain from the auction while the remaining bidders compete.

Table 1 The table displays efficiency, revenue and bidder profit figures for the SMRA, first price and VCG auctions for $\alpha \in \{1, \frac{3}{2}, 2\}$

	Efficiency			Revenue			Profits		
	SMRA (%)	First price (%)	VCG (%)	SMRA	First price	VCG	SMRA	First price	VCG
$\alpha = 1$	86.4	100	100	0.440	0.458	0.458	0.631	0.646	0.646
$\alpha = \frac{3}{2}$	86.3	99.8	100	0.559	0.579	0.577	0.731	0.790	0.761
$\alpha = 2$	86.3	99.7	100	0.660	0.694	0.688	0.852	0.880	0.886

Data are averaged across type environments

2.3 Comparison of auctions

Table 1 displays expected efficiency values as well as expected revenue and payoffs for the bidders for the three mechanisms averaged across environments and for $\alpha = 1, \frac{3}{2}$ and 2.

2.3.1 Efficiency

Efficiency is calculated as

$$\text{Efficiency}_a = \frac{V_a - V_{random}}{V_{opt} - V_{random}} \times 100\%$$

where V_a denote the total surplus generated by mechanism $a \in \{\text{SMRA, First Price}\}$, V_{opt} the total maximum surplus (generated by the VCG mechanism), and V_{random} the value of randomly assigning all the items to the bidders. This definition has the advantage that it is invariant when bidders' values are multiplied by a common number (i.e. when they are measured in cents rather than dollars) or when a common number is added to all of them. Subtracting surplus generated by randomly assigning all items helps to isolate the added value of mechanisms being studied; it reflects the fact that the relevant alternative to the auction is not the withdrawal of the items from the market but random assignment of all items.¹⁰

The first price auction is perfectly efficient in all but the *XXY* environment, where it is at least 98.6% for $\alpha \leq 2$. Meanwhile, the efficiency of the SMRA varies widely between environments, with a low of 66.7% in the *YYY* environment to a high of 100% in the *XXX* environment. Both auctions approach perfect efficiency in the *XXY* environment as α tends to infinity; for $\alpha \leq 2$, the first price auction is more efficient than the SMRA.

¹⁰ In fact, radio spectrum was predominantly allocated via lottery prior to the introduction of auctions in 1994 by the FCC (Roth, 2002).

2.3.2 Seller revenue and bidder profit

The seller's revenue is the sum of the winning bidders' payments while bidder profit is the difference between the value of what bidders won and the payments they made.

As with efficiency, payments in the first price auction closely track those of the VCG; in all but the *XXY* environment, expected seller revenue and bidders payoffs in the first price auction are equal to those in the VCG auction. In the *XXY* environment, the seller's expected revenue is higher in the first price auction while bidders' profits are lower; the bidders thus absorb the loss of efficiency in this environment. Since the VCG mechanism is minimally competitive (i.e. generates the lowest competitive equilibrium revenue for the seller and highest competitive equilibrium payoffs for the bidders), the first price auction can be said to be reasonably competitive in all our environments. The seller's revenue in the SMRA fluctuates around her first price/VCG revenue between environments, being relatively low in the *YYY* environment and high in the *XXY* and *XYY* environments; the opposite pattern holds for bidders' profits. Thus, who bears the cost of the inefficiency in the SMRA depends on the environment; it is the seller in the *YYY* environment and the buyers in the *XXY* and *XYY* environments.

2.3.3 Price discovery

The use of a multiple-round auction is often justified by appealing to its ability to discover or reveal prices; that is, the process of competitive bidding is expected to determine a set of prices for items and packages of items that are "correct", in the sense that they are close to what would prevail in a perfectly competitive environment with no uncertainty.¹¹ The intuition is that bidders will reduce their demands gradually over the rounds of the auction as prices rise until supply equals demand, as in the classical Walrasian *tâtonnement* process. As shown in Milgrom (2000), this process leads to competitive prices if bidders bid truthfully (i.e. myopically) and all items are substitutes for all bidders.

Items are complements for our bidders by design. In the *XXX* and *YYY* environments, competitive prices are easy to determine: the price per item should be between the value (per item) of the highest losing bidding and the lowest winning bidder. When types are mixed, however, a competitive price may not exist. For example, consider the valuations in Table 2 for a *XXY* type environment where the numbers in bold indicate the best allocation for a total surplus of 1.7.

To allocate two items to bidder 2 and none to bidder 1, the price of a single item p must satisfy $0.35 \leq p \leq 0.4$. On the other hand, to allocate three items to bidder 3, p must satisfy $p \leq 0.3$.

Because competitive equilibrium prices do not always exist in the presence of complementarities, attention has turned to the *core* as a proper benchmark for

¹¹ Competitive equilibrium prices are prices for the items at which bidders want to purchase their efficient (i.e. value maximizing) allocation of items.

Table 2 Example of bidders' valuations in the XY environment

	Bidder 1 (Type X)	Bidder 2 (Type X)	Bidder 3 (Type Y)
2 items	0.7	0.8	0
3 items	0.7	0.8	0.9

“reasonably” competitive outcomes.¹² The core is defined by combinations of seller and buyers’ payoffs that satisfy certain stability constraints. The intuition is that auction payoffs are in the core when no coalition of bidders and seller can all do better than their auction payoffs. If we index the seller by $i = 0$ and the three bidders by $i = 1, 2, 3$ then the possible coalitions are the non-empty elements of the powerset of $\{0, 1, 2, 3\}$.¹³ A vector of payoffs $\{\pi_0, \pi_1, \pi_2, \pi_3\}$ is in the core if

$$\sum_{i \in S} \pi_i \geq v(S)$$

for all $S \subseteq \{0, 1, 2, 3\}$ where π_i is the auction profit for coalition member $i \in S$, and $v(S)$ is the maximum surplus that coalition S can generate. Competitive equilibrium prices, when they exist, always produce core payoffs. But while competitive equilibrium prices may not exist, the core is always non-empty in auction applications. As such it seems the right benchmark for competitive outcomes in settings with complementarities.

How good is price discovery in the SMRA in the absence of these assumptions? Figure 1 suggests that it is poor, relative to the first price auction. For each environment we drew 1000 valuations and calculated core payoffs for each draw using the Bayes–Nash bidding functions in Online Appendix A. We then ran each auction and calculated the distance from the auction payoff to the set of core payoffs for each player. Figure 1 shows that mean distance over these 1000 draws, staggered over types to show the distance of each type to their set of core payoffs. Aside from the $X!X!Y$ environment, the SMRA format generates payoffs that are further from core than does the FPSB format. In the $X!X!Y$ environment, while FPSB payoffs are further from core payoffs than SMRA, neither format deviates very far.

¹² See for example Milgrom (2004), Day and Cramton (2012), Day and Raghavan (2007) and Day and Milgrom (2008), each of which study mechanisms to achieve core outcomes in complete information environments. Goeree and Lien (2016) note that, when bidders values are private information, if the VCG outcome is not in the core, no core-selecting auction exists.

¹³ They are $\{0\}, \{1\}, \{2\}, \{3\}, \{0, 1\}, \{0, 2\}, \{0, 3\}, \{1, 2\}, \{1, 3\}, \{2, 3\}, \{0, 1, 2\}, \{0, 1, 3\}, \{0, 2, 3\}, \{1, 2, 3\}$, and the grand coalition $\bar{S} = \{0, 1, 2, 3\}$.




Fig. 1 The figure shows the mean distance to the set of core payoffs for the first-price auction and the SMRA for each environment with $\alpha = 2$. The bar graphs are staggered over types to show the distance of each type to their core payoff

3 Experimental design

3.1 Bidding environment and value model

In the experiment we introduce some heterogeneity across items. In particular, items are now labeled *A* through *E* and bidders are told that (any combination containing) item *A* is less valuable than (same-sized combinations containing) other items. This modification in the environment is motivated by the Australian 900 MHz auction where one block of spectrum at the end of a band was designated a “guard band” and its use needs to abide to additional technical restrictions. Heterogeneity is an important feature in many real-world spectrum auctions, but is left out of our theoretical analysis for tractability.

As is standard in auction experiments, bidders’ values are determined randomly in each auction. Nevertheless, the complexity of the bidding environment requires us to specify a correspondingly more involved value model. In particular, our value model allows for complementarities to vary in the same way as in our theoretical analysis. Again we consider two types of bidders: type *X* bidders whose per-item values peak at two items, and type *Y* bidders whose per-item values peak at three items.

Table 3 describes bidder values depending on draw, type, and whether item *A* is included. To generate a bidder’s values, an integer *R* is drawn uniformly between 25 and 35 (inclusive). This draw, together with the bidder’s type (*X* or *Y*) determines her values for each possible combination of contiguous items. Complementarities between items are realized only if items are contiguous. For example, if a bidder of type *X* wins items *B*, *C* and *E*, she earns $10 + 3R$ for *B* and *C* plus 10 for *E* rather than $10 + 4R$ for all three. Notice that the increase in value from winning a second item is higher for type *X* than for type *Y*. The increase in value from winning a third item is higher for type *Y* than for type *X*.

Table 3 Value specifications for bidders

# of consecutive items	Type X		Type Y	
	With A	Without A	With A	Without A
1	5	10	5	10
2	$10+1.5R$	$10+3R$	$10+0.5R$	$10+R$
3	$10+3.5R$	$10+4R$	$10+3R$	$10+5R$

R is an integer drawn in each round uniformly between 25 and 35 (inclusive)

3.2 Groups and matching

Subjects are placed in groups of three bidders and remained in the same group for a series of 15 auctions/periods. The fixed matching ensures that we have independent observations between groups. Each bidder is assigned a type, X or Y , which remains constant throughout the experiment. A new value for R is drawn in each period. Complementarity types could also be drawn randomly at the beginning of the experiment, but this might lead to biased results if a specific combination of types is over/under represented in our sample. Instead, we have no reason to believe that any combination of complementarity types is more relevant for applications, and therefore study all possible such combinations (XXX , XXY , XYX , or YYY) in equal proportions. Since in practice bidders' synergy-types are unknown (just like their valuations), we wish to compare the performance of the FPSB And SMRA formats across these four combinations. In particular, for each treatment we have two groups for each of the four combinations, for a total of *eight independent groups per treatment*.

Within each treatment, the draws of R were random and independent across bidders and periods. Across treatments, we used same draws to ensure that any observed difference are not due to differences in the random draws. For example, in any given period, both of the first X bidders in the first XXX group in each treatment will have the same value draw.

All of the above design choices ensure that we have treatment balance. They also allow us to employ paired tests (e.g. Wilcoxon signed-rank test) when comparing results across two treatments.

3.3 Treatments

Our main interest in comparing the SMRA and FPSB mechanisms. A detailed description of each mechanism is as follows:

1. In the *First-price sealed bid auction (FPSB)* mechanism, bidders place six bids: one for A , one for a single item other than A , one for pair of items including A , one for a pair of items not including A , one for a package of three items including A ,

and one for a package of three items not including *A*. At most one of the six bids placed by a bidder can become winning. A simple optimization algorithm finds the combination of bids that maximize revenue and the winning bidders pay their bids.

2. In the *Simultaneous multi-round auction (SMRA)*, bidders compete directly for items *A* through *E*. A price clock is associated to each of the five items. The price in the first round is five for each item. In each round, bidders indicate whether they demand an item at the price displayed on its clock. For each round and each item, one of the bidders who demands the item is randomly designated the item's provisional winner. If more than one bidder demands an item, its price increases by 15. If only one bidder demands a particular item, its price clock is paused. If other bidders later demand this item, the price clock restarts and the item is randomly provisionally assigned to one of the new bidders. When demand on all items is one (or less), the auction ends and items are assigned to their provisional winners who pay the price displayed on their clock.

An activity rule ensures that the auction progresses apace. The sum of items provisionally won by a bidder plus the items she is demanding is called her *activity*. Her *activity limit* in any round is her activity at the end of the previous round, or three if it is the first round. A bidder's activity cannot exceed her activity limit. Thus, for example, a bidder who fails to bid on any items in round one will be unable to bid in subsequent rounds.

The SMRA can result in fragmentation, i.e. a bidder winning non-contiguous items. Since value complementarities only apply to consecutive items, this is a potential source of inefficiency. In the FPSB this possibility is avoided by the algorithm that calculates the optimal allocation given bids. This, of course, gives the FPSB an advantage in terms of expected efficiency. Still, it remains an empirical question whether this theoretical advantage, which assumes some sophistication on the part of bidders who have to arrive at an optimal strategy introspectively. At the same time, due to its dynamic nature, the SMRA is typically expected to perform better in terms of price discovery. Whether this is true and how it may weigh against the performance of the FPSB in terms of price discovery remains to be seen.

One way to deal with fragmentation in the SMRA that has been used in practical applications, involves breaking down the auction into two stages.¹⁴ In the first stage bidders bid for contiguous blocks of items of different sizes. In the second stage, the location of these blocks is determined. Such a process also eliminated the possibility that two bidders may pay very different prices for otherwise homogeneous (combinations of) items. We also run treatments using the two-stage format. For completeness we apply this modification to both the SMRA and the FPSB. The details are as follows:

¹⁴ See for example the UK 2.3 GHz and 3.4–3.6 GHz spectrum auction in 2018.

3. In the *Two-stage FPSB (FPSB-2)*, bidders place three bids in the first stage: one for a single item, one for pair of contiguous items, and one for a package of three contiguous items. At most one of the three bids placed by a bidder can become winning. The winners pay their first-stage bids and proceed to the second-stage where they can bid for “not being assigned item A.” In this stage, the lowest bidder is assigned A (by itself or as part of a package, depending on how many items the bidder won in the first stage) and does not pay the second-stage bid. The other bidder(s) pay(s) their second-stage bid(s).
4. In the *Two-stage SMRA (SMRA-2)*, bidders first compete for a generic item. A single price clock is associated to the item. In each round, bidders indicate whether they demand zero, one, two, or three units of the item at the current round price. If the total demand in the round plus total units provisionally assigned for items at the current round price is fewer than five, all bidders are provisionally assigned the quantity they demanded. Otherwise, provisional winners are established in the following way. First, any current provisional winners are reassigned their provisional winnings if the round price has not increased since they were assigned. Second, the bidders demanding items at the current round price are declared provisional winners of the number of goods they demanded in random order until all five units are assigned. The last bidder provisionally assigned items in this process may be assigned fewer items than she demanded. If, at the end of the round, all provisional winners were assigned their items at the current round price, the clock price increases by ten for the next round. Otherwise the round price stays the same. The auction ends after any round with zero new demand, i.e. demand not including provisional winners. An activity rule ensures that the auction progresses apace, just as in the one-stage SMRA described above. The winners pay their first-stage bids and proceed to the second-stage where they can bid for “not being assigned item A.” The lowest bidder in the second stage is assigned A and does not pay her second-stage bid. The other bidder(s) pay(s) their second-stage bid(s).

In the initial treatments bidders are provided information about others' valuations. In particular, we describe to bidders how values are generated and the distribution of draws. On the bidding screen, each bidder is shown a table with values as well as the types (X or Y), but not the values, of the two other bidders in the group. In a series of follow-up treatments, bidders know their own value but not how these values are generated. We dub these treatments FPSB-U and SMRA-U, where the U stands for *uninformed*. These follow up experiments served as stress tests for the single-stage formats, which performed much better than the two-stage experiments in the initial experiments. Except for the informational environment, they were otherwise identical to the FPSB and SMRA respectively.

In total there were six treatments and eight independent observations per treatment, see Table 4.

Table 4 Experimental design

Treatment	Mechanism	Bids in 1st stage	Bids in 2nd stage	# Ind obs	Bidder information
FPSB	First price sealed bid	6: for 1, 2, and 3 cont. items with or without item A	None	8	Own value Other's value distributions Group types
SMRA	Simultaneous multiple-round ascending	5: for each item	None	8	Own value Other's value distributions Group types
FPSB-2	First price sealed bid	3: for 1, 2, and 3 cont. items	1: to not be assigned A	8	Own value Other's value distributions Group types
SMRA-2	Simultaneous multiple-round ascending	3: for 1, 2, and 3 cont. items	1: to not be assigned A	8	Own value Other's value distributions Group types
FPSB-U	First price sealed bid	6: for 1, 2, and 3 cont. items with or without item A	None	8	Own value only
SMRA-U	Simultaneous multiple-round ascending	5: for each item	None	8	Own value only

All treatments 3 bidders per group; 8 groups; 4 type compositions (XXX, XXY, XYY, or YYY), two groups of each one.

Each of the six treatments used one of the four different mechanisms: FPSB, SMRA or their two-stage variants. The third and fourth columns indicate the number of bids bidders need to submit in each stage. The final column indicates the information environment




Fig. 2 Aggregate results. The bars show the average efficiency, seller's revenue and bidders' profits (left, middle and right panel, respectively). Black bars indicate the theoretical VCG benchmark. Red bars correspond to treatments using an FPSB type mechanism. Blue bars correspond to treatments using an SMRA type mechanism. Treatments are grouped by the number of stages and whether bidders are informed or not about other's value distributions and the group's type composition. Observations are pooled over periods 6–15. The whiskers indicate 95% CI's for the average, based on the 8 group averages per treatment. (Color figure online)

3.4 Experimental procedures

A total of 144 subjects participated in the experiment. Subjects were recruited from University of Technology, Sydney using ORSEE (Greiner, 2015). The experiment was programmed and conducted with z-Tree (Fischbacher, 2007) and MATLAB.¹⁵ Subjects received instructions, answered a quiz and competed in a practice period, before participating in fifteen paid auctions. The experiments lasted from a little over an hour for the FPSB to 2.5 h for the SMRA-2. Participants were paid the earnings that accumulated over the 15 periods of the experiment if these were positive plus a 10 AUD show-up fee. If their cumulative earnings were negative at the end of the experiment, they were only paid the 10 AUD show-up fee. The conversion rate used in the experiment was 1 Australian dollar (AUD) for every 4 experimental points. The average earnings were 39.95 AUD including the 10 AUD show-up fee.

4 Experimental results

Figure 2 displays efficiency, seller revenue and bidder profits for all treatments, pooled over periods 6 to 15.¹⁶ In each panel, the group of bars on the left display results for the one-stage FPSB and SMRA when bidders are told the types in their groups and how values are drawn. The second group of bars displays results for the one-stage mechanisms when bidders are told only their values: FPSB-U and

¹⁵ Calculating optimal allocations dynamically is beyond the capabilities of zTree. Instead, we used zTree to design the user interface and collect subjects' bids. In the background, zTree interfaced with MATLAB automatically, the latter calculated optimal allocations and returned those to zTree to present to the subjects.

¹⁶ We do not observe any substantial change in behavior across periods in treatments using an FPSB mechanism. In treatments using SMRA, there is some learning taking place during the first few periods of the experiment, presumably as subjects get familiar with the dynamic and therefore more complex, SMRA mechanism. We therefore present our main results based on periods 6 to 15. Including the first five periods in the analysis does not change the results in a substantial way. See also Online Appendix B.

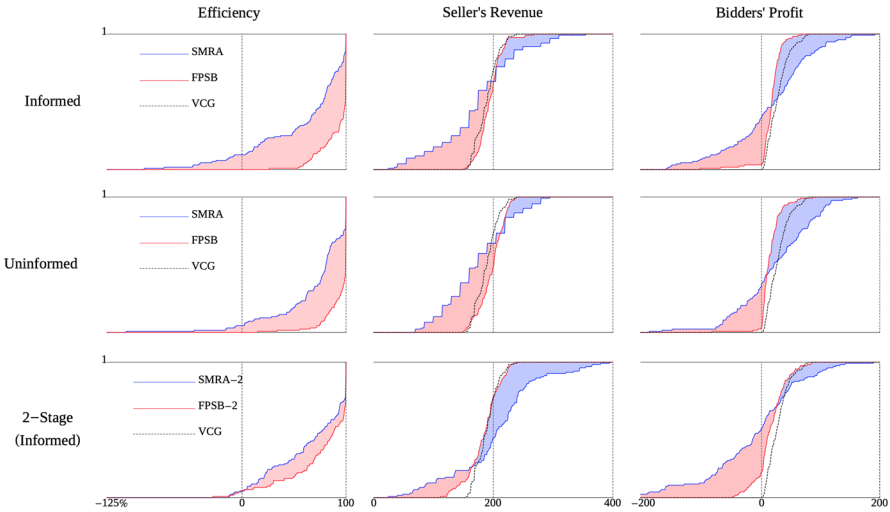


Fig. 3 Cumulative distributions of key variables. Observations are pooled over all environments for periods 6–15. The first row displays results for the treatments where bidders are told the types in their groups and how values are drawn. The second row displays results for the treatments where bidders are told only their values. All graphs include the corresponding distributions for the VCG mechanism as a benchmark

SMRA-U. The third group of bars in each panel displays results for FPSB-2 and SMRA-2. Figure 3 gives a more detailed look at the distributions of efficiency, seller revenue and bidder profits for all mechanisms, pooled over all environments.

4.1 Comparing single-stage auctions

From the first bars of left panel of Fig. 2 it is clear that the one-stage FPSB delivers substantially higher efficiency than the corresponding SMRA. In fact, efficiency in FPSB comes very close to the theoretical maximum achieved by the VCG and remained on average above 90% across all type combinations. In SMRA efficiency is substantially lower and remains on average below two thirds of the theoretical maximum achieved by the VCG mechanism. In the other two graphs one sees that the FPSB also achieves, on average, higher seller revenue and lower bidder profits than the SMRA. This can be attributed to demand reduction on the part of bidders in the SMRA. Note however that in our experiment bidders in this mechanism frequently make losses. This is a consequence of the inability to protect themselves from the exposure problem. In the SMRA, bidders competing aggressively for a package of two or three items may end up winning only one. In addition, there may be fragmentation, i.e. a bidder winning non-contiguous items.

The top panels of Fig. 4 illuminate the shortcomings of the SMRA: this format often leads to allocations where one or two bidders get a single unit. The bottom panel of Fig. 4 displays the degree to which items are sold in non-consecutive packages or remain unsold in the standard SMRA; evidently, bidders had difficulty

coordinating their bids effectively to form packages of consecutive items in the single-stage SMRA.

Note that FPSB yields higher revenue on average than VCG despite its efficiency being less than VCG's 100%. This comes at a cost to the bidders who make less than under VCG. Importantly, bidders' profits are always positive under FPSB. The reason is that FPSB fully protects bidders from the exposure problem: they can specify a separate bid for each of the (combination of) items they might win and, by submitting bids that are less than values, never risk a loss.

Result 1 Compared to the SMRA the FPSB is more efficient, yields more revenue and lower bidder profit. Bidders frequently incur losses in the SMRA due to exposure problems.

Support. The result is supported by the data shown in the graphs in Fig. 2. Further support is given in Fig. 3, where the top row graphs display the distributions for efficiency, revenue, and bidder profits for the SMRA and FPSB. This is confirmed by non-parametric tests. For instance, the Wilcoxon Signed-Rank test comparing mean efficiency in FPSB and SMRA gives a p value of below 0.008. This is further supported by regressions controlling for the groups' value type combinations. While the simple Wilcoxon Signed-Rank tests comparing revenues and bidder profits in the two treatments do not give significant differences (a Mann–Whitney U test rejects equality for revenues), the regressions controlling for value type combinations indicate that the treatment effect is indeed significant for at least some value type combinations. In fact, this is the case for revenues for all type combinations, once the data from all one-stage treatments is pooled (see Sect. 4.4). For bidder profits, the treatment effect is significant for all but one value type combination (XY). For all significance tests and the regressions controlling for value type combinations, see Tables 5 to 13 in Online Appendix B.

4.2 Comparing two-stage auctions

The two-stage SMRA is used in practice to help bidders overcome some of the problems with its single-stage counterpart. First, it avoids fragmentation, as items won are contiguous by design. It also requires bidders to focus on a specific attribute in each stage: first the number of items, then on their location. Unfortunately, it has a double exposure problem: bidders who compete aggressively for a package may end up winning only a subset (as in the single-stage SMRA) and when competing for the number of items, bidders do not know whether item A will be included or not. In fact, this second exposure problem is independent of the underlying mechanism and is inherent to the two-stage process. It can therefore also affect results in the two-stage FPSB.

The bars grouped on the right of each panel in Fig. 2 allow us to compare efficiency, revenue and bidder profits between FPSB-2 and SMRA-2. In both cases we observe a substantial efficiency loss compared to their two-stage counterparts, but on average efficiency is higher in the FPSB-2. The comparison of revenues

Fig. 4 Observed outcomes in the different mechanisms (top panel) and fragmentation in the SMRA (bottom panel)

and bidder profits seems to indicate that SMRA-2 suffers from a larger exposure problem. Bidders make high losses that translate to high seller revenues. In FPSB-2, while bidder profits are significantly lower than the ones in the VCG theoretical benchmark, losses are not very common. In fact, seller revenues are not significantly different than what they are in the VCG benchmark.

Result 2 Compared to the SMRA-2 the FPSB-2 is more efficient. But SMRA-2 yields higher revenue and lower bidder profit mainly because bidders incur losses due to exposure problems.



Support. The result is supported by the data shown in the graphs in the bottom row of Fig. 2. Further support is given in Fig. 3, where the bottom row graphs display the distributions for efficiency, revenue, and bidder profits for the SMRA-2 and FPSB-2. Overall, the differences between the two two-stage formats are not statistically significant: the p values are 0.25 for efficiency, 0.148 for revenue and 0.055 for bidder profits. The regressions controlling for value type combinations reveal that the treatment effect is significant for efficiency when the type combination is XY , while for revenue and bidder profits it is the case with type combinations XXY and YYY .

4.3 Comparing single-stage and two-stage auctions

Looking at Fig. 2 it becomes clear that the two-stage process did not help bidders in our experiment. For the SMRA efficiency did not improve significantly moving from one stage to two. At the same time, it seems that the exposure problem intensified, leading to much higher seller revenues and losses for bidders in the two-stage mechanism compared to the one-stage SMRA. As an illustration, consider a case where in the first stage, a type Y bidder might compete fiercely to win three items but finally give in (at high prices) and settle for two items. In the second stage, the value of what was won may depreciate further if the type Y bidder places the lowest bid. As a result, bidder losses in the two-stage SMRA are common and substantial. In terms of protecting bidders' from exposure risk, this format is least desirable. This is an important finding as regulators have started using such a two-stage format in spectrum applications. For example, Ofcom in the UK used two stages in the 2.3 GHz and 3.4–3.6 GHz auction in 2018.¹⁷

For the FPSB, efficiency decreased with the addition of the second stage. Apparently, bidders are able to deal well with the relatively high number of bids required in the one-stage FPSB (six bids in total). Breaking down the process

¹⁷ <https://www.ofcom.org.uk/spectrum/spectrum-management/spectrum-awards/awards-archive/2-3-and-3-4-ghz-auction>.



in two stages does not bring additional benefits, but introduces a new exposure problem. Unlike in SMRA-2, bidders in FPSB-2 do appear to protect themselves against this problem, as their profits are not significantly different than those in FPSB. Still, seller revenues are lower.

Result 3 The two-stage mechanisms result in lower efficiency (FPSB) or an exacerbated exposure problem (SMRA).

Support. The result is supported by the data shown in the graphs of Fig. 2. For the Wilcoxon Signed-Rank test comparing mean efficiency in SMRA and SMRA-2 we get p value 0.383 and for mean efficiency in FPSB versus FPSB-2 we get p value below 0.023. For the same test comparing revenues in SMRA versus SMRA-2 we get p value below 0.078 and for FPSB versus FPSB-2 we get p value 0.195. For the test comparing bidder profits in the SMRA versus SMRA-2 we get p value 0.195 and for FPSB versus FPSB-2 we get p value 0.742. Using the Mann–Whitney U test yields very similar results.

4.4 Robustness to the informational environment

In the baseline SMRA and FPSB treatments subjects knew not only their own type and valuations, but also the type of the other bidders in their group. In real applications it is not unreasonable to think that telecom companies may have some information about their competitors preferences and the degree of complementarity they face. Nevertheless, it is a valid concern that the problematic performance of the SMRA in our experiment, as stated in Result 1, may be driven by this design feature combined with the particular choice of valuation distributions used. To test the robustness of our main result with respect to the informational environment, we conducted additional treatments of the one-stage mechanisms in which bidders know their valuations but are entirely uninformed about the distribution of other bidders' valuations. These treatments are dubbed FPSB-U and SMRA-U respectively.

The second row in Fig. 2 presents the results for these additional treatments. We find that the comparison between FPSB-U and SMRA-U yields very similar results as that between FPSB and SMRA. The signs of the differences remain unchanged for all three measures: efficiency, revenue and bidders' profit. In terms of magnitudes, the difference in efficiency is reduced but remains substantial: the FPSB-U is approximately 20% more efficient than the SMRA-U. There is also a reduction in the difference in revenue between the two mechanisms, but it remains significant. Bidders' profit is again higher in the SMRA, although now the difference is not statistically significant.

Result 4 The FPSB's better performance compared to the SMRA is robust to changes in the information available to bidders about others' distribution of valuations.

Support. The result is supported by the data shown in the corresponding group of bars in the graphs of Fig. 2. Further support is given in Fig. 3, where the middle row graphs display the distributions for efficiency, revenue, and bidder profits for the SMRA-U and FPSB-U. For the Wilcoxon Signed-Rank test comparing mean efficiency in SMRA-U and FPSB-U we get p value below 0.008. For the same test comparing revenues in SMRA-U and FPSB-U we get p value below 0.023. For the test comparing bidder profits in the SMRA-U and FPSB-U we get p value 0.383. Using the Mann-Whitney U test yields very similar results. The regressions controlling for value type combination provide further support.

Overall, the change in the information available to bidders does not seem to have any effect on the average outcomes of the FPSB or the SMRA mechanism.¹⁸ Based on this, and to facilitate the presentation of results regarding price discovery, in the analysis in the following section we pool the data for each mechanism across the two informational environments. If anything, this would work in favour of the SMRA.

5 Price discovery

One justification for the use of auctions is that they are *price discovery* mechanisms. Ideally, auction prices are competitive equilibrium prices that clear the market (i.e. prices such that auction losers are happy not to be assigned any items and auction winners are happy with their assignment). Notice that in our experimental environment, the prices for items B through E must be identical. Therefore, competitive prices consist of a set of two prices p_A and $p_{-A} = p_\ell$ for $\ell \in \{B, C, D, E\}$. As we discussed in Sect. 2, such prices do not always exist, leading to the notion of core equilibrium payoffs. Since the core is always non-empty, these near-competitive payoffs will always exist.

Simply because the core is non-empty does not mean that it is easy for a particular auction format to discover prices that lead to core payoffs. For the environments considered in our experiment, the VCG auction produces core outcomes.¹⁹ In fact, the VCG outcome corresponds to the point in the core that assigns the lowest revenue to the seller and the highest profits to the bidders. In this format, truthful bidding is a (weakly) dominant strategy and the outcomes are fully efficient.

Figure 5 shows core payoffs for each of the four environments: XXX , XXY , XYY , and YYY . To produce a two-dimensional graph, the sum of bidders' profits is shown on the horizontal axis and the seller's revenue is shown on the vertical axis. All payoffs are normalized by the maximum surplus, $v(\bar{S})$, and the upper dashed line corresponds to all possible divisions of the maximum surplus among the bidders and the seller. The lower dashed line corresponds to all possible divisions of surplus from a random allocation. The subset of core constraints that dictate individual rationality (i.e. $\pi_i \geq 0$ for $i = 0, 1, 2, 3$) imply that the core is

¹⁸ The tables in Online Appendix B provide the results of formal tests to support these statements.

¹⁹ This is not necessarily the case in the presence of complementarities. In Sect. 2, VCG prices are often below core prices.




Fig. 5 Each panel shows the seller's average revenue (y-axis) and average buyer payoff (x-axis) normalized by total surplus and averaged over the last ten periods. The upper dashed line corresponds to efficient outcomes with the solid segment indicating core outcomes. The lower dashed lines correspond to random allocations of the items. These figures use pooled observations from both information treatments for the one-stage mechanisms

part of the positive orthant. The other core constraints set a minimum revenue for the seller, here given by R^{VCG} . So in each of the four panels of Fig. 5, the core corresponds to the solid segment that runs from the VCG payoff point to (0,1).

The grey triangles in each of the panels reflect alternatives to the VCG outcome that might interest a seller. These alternatives are not all fully efficient but do yield higher seller revenue than the VCG auction and generate positive profits for the bidders. As such they reflect a trade-off between efficiency and revenue that sellers typically face (e.g. in the use of reserve prices). The markers for the first price auctions are red and for the SMRA are blue.

Note that the FPSB does a remarkable job at price discovery: the red points are always close to fully efficient outcomes while providing more than VCG revenues for the seller. The SMRA formats on the other hand consistently either underperform the VCG from the seller's perspective or generate losses for the bidders (i.e. are outside the grey triangles in the figures).

Result 5 The FPSB results in closer-to-core prices than the SMRA.

Support. See Fig. 6, which parallels the theoretical Fig. 1, and demonstrates that average deviations from core prices are smaller for both X and Y type bidders




Fig. 6 The figure shows the mean distance to the set of core payoffs for the (one-stage) FPSB and SMRA treatments for each environment. Data is pooled over periods 6 to 15 and over both information treatments. The bar graphs are staggered over types to show the distance of each type to their core payoff

as well as the seller. This is true for all environments including *XXY* for which the SMRA is theoretically predicted to yield closer-to-core prices.

6 Conclusion

Our aim is to fill what we consider to be an important gap in the literature, by comparing the relative merits of ascending and sealed-bid combinatorial formats when complementarities are strong. It is posited that, in practice, the same forces in the SMRA that generate competitive prices for substitutable goods will at least mitigate any problems caused by complementarities as well as provide the seller with sufficiently competitive revenues (Cramton, 2006). Indeed, there have been notable spectrum auctions involving complements that appear to have performed quite well, such as the US regional narrowband auction in 1994 (Milgrom, 2000). At the same time, theoretical analysis shows that bidders in the SMRA are highly susceptible to the exposure problem. For a bidder whose per-item value increases in the number of items she wins, bidding up to her value for two items, for example, exposes her to the risk of having to pay a large amount for a single item that she places little value on. In equilibrium, bidders reduce their demands too early in the auction resulting in low prices and low efficiency (Goeree & Lien, 2014). Our results confirm these theoretical predictions. Price discovery in the SMRA is poor, and efficiency and revenue are low.

On the other hand, there was no systematic study in the literature of the FPSB in a similar environment, a natural candidate for an alternative. Our experimental results provide support for the potential of the SMRA to deliver poor results when complementarities are strong, as predicted by theory. In contrast, the FPSB does not

suffer from similar issues, since it protects bidders from the exposure problem by letting them place bids for every possible package they may win.

Our results present a puzzle: how can we reconcile the good performance of the FPSB compared to the SMRA with its limited usage in real applications, where the SMRA is often encountered?²⁰ A first reason may be complexity. The number of possible packages grows exponentially with the number of items on sale. This means that in large auctions bidders in the FPSB would be required to submit an unmanageable number of bids. In this case the SMRA has the obvious advantage. Still, we find it prevailing also in many occasions where the number of items on sale is relatively small, so complexity alone cannot explain the puzzle.

Perhaps the answer is also related to bidder preferences. In public consultations for spectrum auctions telecoms often favor ascending formats arguing that such auctions allow them to maintain control over their destiny. However, taking such arguments at face value might be naive. Auctions for permits, such as those for the use of spectrum, are auctions with allocative externalities. For bidders it does not simply matter what they win but also what others win, as they are your competitors in the market for telecommunication or other services where these permits apply. Incumbents in those markets want to avoid outcomes where new entrants gain significant positions through the auction. They may therefore prefer ascending formats, because at any stage they can react to make sure undesired—to them—outcomes are avoided. However, a regulator's point of view should be the exact opposite: if a smaller bidder sees more value in getting more spectrum than the major bidders, the regulator should want to raise the likelihood of such outcomes as they will obviously bolster competition and consumer welfare. If, as we suspect, bidders preferences are driven by such motivations and pressure from their side is partly responsible for the infrequent use of the FPSB in practice, our results may help strengthen the opposite view in real world lobbying contests.

Regardless of why regulators are reluctant to embrace the FPSB auction, they have recognized the need for faster auctions that avoid fragmentation and other allocative inefficiencies of the SMRA. Many regulators first moved from the SMRA to a combinatorial clock auction (see Ausubel et al., 2006). The latter combines an initial clock auction phase for gradual information revelation with a sealed package-bidding phase. When this format presented problems of its own (see Levin and Skrzypacz, 2016), regulators settled on the two-stage SMRA format with the hope that it speeds up the bidding process and avoids fragmentation, whether incidental or with the aim of splitting a competitor's spectrum holdings (Ofcom, 2014).

Further work is needed explore other apparent advantages of the FPSB format. Given its simple payment rule and its winner determination algorithm, it readily accommodates complex allocation constraints. For example, many bidders have a

²⁰ It should be noted that in recent years, there are examples of regulators that have considered the use of FPSB auction formats for the allocation of spectrum, at least at the consultation phase. These include: the FCC's Auction 108 in the US, the 3.5 GHz band auction in the Netherlands, and the 850/900 MHz band auction in Australia. Recently, Moldova's telecoms regulator ANRCETI announced the use of a FPSB auction to allocate spectrum in the 450 MHz, 900 MHz and 2.6 GHz bands.

fixed budget. In the FPSB a bidder could submit mutually exclusive (so called XOR) bids and indicate their total budget. The assignment algorithm could then pick the best allocation respecting each bidder's budget.

Another argument in favor of the FPSB in practice is that it allows new entrants to bid competitively for large packages. It is worthwhile to design future experiments to specifically test the ability of new entrants to compete in the presence of incumbents.

Of course, as with any experimental study, results can be influenced by the series of design choices we made and the characteristics of our subject pool. Also, our effort to test robustness across different bidding environments comes at the expense of sample size. For these reasons we believe that more replications are needed to ensure that the good performance of the FPSB that we find in our specific setting is robust.

More broadly, while the FPSB auction performed well in our experiments, it remains to determine how to design a practical allocation mechanism for general settings when complementarities cannot be ruled out and one cannot rely on prices to guide participants to efficient and stable outcomes.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10683-023-09805-x>.

Funding Open access funding provided by the Cyprus Libraries Consortium (CLC).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ausubel, L. M., & Cramton, P. (2004). Auctioning many divisible goods. *Journal of the European Economic Association*, 2(2–3), 480–493.
- Ausubel, L. M., Cramton, P., & Milgrom, P. (2006). The clock-proxy auction: A practical combinatorial auction design. In Cramton, P., Shoham, Y., & Steinberg, R. (Eds.), (p. 115–138). MIT Press.
- Ausubel, L. M., & Milgrom, P. (2006). The lovely but lonely vickrey auction. In *Combinatorial auctions* (pp. 57–95). MIT Press.
- Bichler, M., & Goeree, J. K. (2017). *Handbook of spectrum auction design*. Cambridge: Cambridge University Press.
- Bichler, M., & Goeree, J. K. (2017b). Review of spectrum auction design rules for the ACMA. Report prepared for the Australian Communications and Media Authority.
- Cramton, P. (1997). The FCC spectrum auctions: An early assessment. *Journal of Economics and Management Strategy*, 6(3), 431–495.
- Cramton, P. (2006). Simultaneous ascending auctions. In *Combinatorial auctions* (pp. 99–114). MIT Press.
- Cramton, P., Kwerel, E., Gregory, R., & Skrzypacz, A. (2011). Using spectrum auctions to enhance competition in wireless services. *Journal of Law and Economics*, 54, S167–S188.

- Day, R. W., & Cramton, P. (2012). The quadratic core-selecting payment rule for combinatorial auctions. *Operations Research*, *60*, 588–603.
- Day, R. W., & Milgrom, P. (2008). Core-selecting package auctions. *International Journal of Game Theory*, *36*, 393–407.
- Day, R. W., & Raghavan, S. (2007). Fair payments for efficient allocations public sector combinatorial auctions. *Management Science*, *53*, 1389–1406.
- Earle, R., & Sosa, D. (2013). Spectrum auctions around the world: an assessment of international experiences with auction restrictions (Tech. Rep.). Analysis Group.
- Filiz-Ozbay, E., Lopez-Vargas, K., & Ozbay, E. Y. (2015). Multi-object auctions with resale: Theory and experiment. *Games and Economic Behavior*, *89*, 1–16.
- Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, *10*, 171–178.
- Gentry, M., Komarova, T., Schiraldi, P., & Shin, W. (2019). On monotone strategy equilibria in simultaneous auctions for complementary goods. *Journal of Mathematical Economics*, *85*, 109–128.
- Goeree, J. K., & Lien, Y. (2014). An equilibrium analysis of the simultaneous ascending auction. *Journal of Economic Theory*, *153*, 506–533.
- Goeree, J. K., & Lien, Y. (2016). On the impossibility of core-selecting auctions. *Theoretical Economics*, *11*, 41–52.
- Goeree, J. K., & Louis, P. (2019). An evaluation of the first-price combinatorial auction for the sales of 850-900 MHz spectrum. Report prepared for the Australian Communications and Media Authority.
- Green, J., & Laffont, J.-J. (1979). *Incentives in public decision making*. Amsterdam: North Holland.
- Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with ORSEE. *Journal of the Economic Science Association*, *1*, 114–125.
- Gul, F., & Stacchetti, E. (1999). Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, *87*, 95–124.
- Holmstrom, B. (1979). Groves schemes on restricted domains. *Econometrica*, *47*, 1137–1144.
- Kagel, J. H., Lien, Y., & Milgrom, P. (2010). Ascending prices and package bidding: A theoretical and experimental analysis. *American Economic Journal Microeconomics*, *2*(3), 160–185.
- Kagel, J. H., Lien, Y., & Milgrom, P. (2014). Ascending prices and package bidding: Further experimental analysis. *Games and Economic Behavior*, *85*, 210–231.
- Kim, H.-S. (1996). Equilibrium and efficiency in auctions of complementary goods without bundling. *Economics Letters*, *52*(1), 49–54.
- Klemperer, P. (2002). How (not) to run auctions: The European 3g telecom auctions. *European Economic Review*, *46*, 829–845.
- Levin, J. (1997). An optimal auction for complements. *Games and Economic Behavior*, *18*(2), 176–192.
- Levin, J., & Skrzypacz, A. (2016). Properties of the combinatorial clock auction. *American Economic Review*, *106*(9), 2528–2551.
- Milgrom, P. (2000). Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, *108*(2), 245–272.
- Milgrom, P. (2004). *Putting auction theory to work*. Cambridge: Cambridge University Press.
- Munro, D. R., & Rassenti, S. J. (2019). Combinatorial clock auctions: Price direction and performance. *Games and Economic Behavior*, forthcoming.
- Ofcom. (2014). Public sector spectrum release: Award of the 2.3 GHz and 3.4 GHz bands (Tech. Rep.). Author.
- Roth, A. E. (2002). The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, *70*(4), 1341–1378.
- Rothkopf, M. H. (2007). Thirteen reasons why the Vickrey–Clarke–Groves process is not practical. *Operations Research*, *55*(2), 191–197.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, *16*(1), 8–37.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.