Auctions For Complements - An Experimental Analysis*

Daniel Marszalec†

This draft: February 2015

Abstract: I evaluate the performance of four static sealed-bid package auctions in an experimental setting with complementarities. The underlying valuation model comprises two items, and three bidders; two 'individual' bidders demand one item only, while the third bidder wants to win both. The rules being compared include the Vickrey and first-price auctions, Vickrey Nearest Rule and the Reference Rule. Auction-level tests find the first-price auction revenue dominant over the other three rules, while the Vickrey auction performs worst; the other two rules rank intermediate. Bidder-level tests of the experimental data reject the competitive equilibrium bidding functions that can be derived in this setting, and I find that overbidding is widespread in all four auctions, as is aversion to submitting boundary bids. I also observe behaviour consistent with collusive bidding in the Vickrey auction. Contrary to theoretical predictions, the Vickrey auction performs worst on efficiency, primarily for this reason.

The growth in popularity of auctions has seen them applied to an ever wider range of markets, including markets with multiple packages and complementarities. A stylized example of such a situation is an auctioneer selling a jacket and a pair of trousers: some buyers may only want the jacket, others may only need the trousers, but there may also be buyers who wish to buy both together to form a complete suit. The fact that the two garments match creates additional value for the buyer who wants both - this is the complementarity. More complex demand patters of a similar kind are present in the auctions for mobile telephony spectrum, contracts for serving bus routes or airport take-off and landing slots.[3] To deal with this increased complexity, a new class of mechanisms, called core-selecting auctions, have been developed and implemented, though our understanding of their incentive properties is not yet complete. I conduct a bidding experiment to evaluate the performance of two static core-selecting auctions (the Vickrey Nearest and the Reference Rule) against two older alternatives (the Vickrey and first-price auctions).

The motivation for picking the Vickrey and first-price auctions is that they cover two extremes in terms of bidder incentives. In the Vickrey auction truthful bidding is a

†Graduate School of Economics, University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo, Japan. Email: dmarszalec@e.u-tokyo.ac.jp

*I would like to thank Vince Crawford and Paul Klemperer, as well as Ian Jewitt,Thees Spreckelsen and Alex Teytelboym for their comments and criticisms. Financial support from the ESRC and RES is also gratefully acknowledged.

[3]On mobile spectrum, see Danish Business Authority (2012), ComReg (2012) and Ofcom (2012). The auction of London bus routes is discussed in Cantillon and Pesendorfer (2006). An auction solution to allocating landing slots is discussed in Federal Aviation Administration (2008).

dominant strategy, while the first-price auction gives strong incentives for bidding below value. Both auctions also embody well-known weaknesses, which have limited their use in practice: the Vickrey auction may generate very low revenue, and the first-price auction can be inefficient. The intuitive argument in favour of the two core-selecting rules is that because they make bidders' payments close-to-independent of their bids, they should induce close-to-truthful value revelation. This in turn should generate an outcome in which both revenue and efficiency are high.[4]

The main finding of this paper is the strong performance of the simplest of the four rules, the first price auction: it is revenue-dominant without losing efficiency. I cannot reject revenue equivalence between the remaining three auctions. The Vickrey auction is least efficient, and no significant efficiency difference emerges between the first-price and the core-selecting rules.

At the bidder level, I test my experimental data against the Bayesian Nash equilibrium bidding functions for all four rules, as derived by Ausubel and Baranov (2010). The theory is not supported by my experiment, and overbidding is frequent in each auction. In the core-selecting auctions, when bidders' behaviour diverges from equilibrium, they do not revert to a truth-telling rule-of-thumb; instead they attempt to game the rule to their advantage, albeit unsuccessfully. I also find evidence of attempted collusion in the Vickrey auction, which is the main explanation behind the low revenue and efficiency of this auction in my experiment. In the first-price auction when bidders deviate from theoretical equilibrium, they do so in predictable ways that do not undermine efficiency or revenue. The simplest auction is thus most robust in my experiment, and the attractive properties of the core-selective rules are not fully borne out when bidders' behaviour does not conform to expectation.

Recent experimental auction literature has focused on dynamic auctions, such as the combinatorial-clock, and simultaneous ascending auctions.[5] This strand of research has been primarily concerned about efficiency properties of those auctions, and how bidders select packages in settings with complex valuation patterns. However, many practical implementations of these dynamic designs feature a one-shot static auction as their final phase: for example, the Danish, Irish and UK spectrum auctions in 2012, all used a Vickrey-Nearest type rule to determine the final prices and allocations of licences, after a dynamic auction had been used to determine the relevant packages.[6] My work is naturally seen as investigating how these static rules perform, given that a selection of packages has already been set. At the time of writing, there had been no prior experimental work in this area.

The rest of the paper is structured as follows. The auction rules valuation model are introduced in Section 1, and the precise formulation of hypotheses which I test are discussed in Section 2. The experimental setup is presented in Section 3, and Section 4 performs a quality check of the data. Experimental results and tests of the hypotheses are performed in Sections 5 and 6. Section 7 discusses the interpretation of the results, and Section 8 concludes.

1. Auction Setup and Rule Descriptions

My model uses a setting with three bidders and two items to model the context of bidding on goods with complementarities. I label the two items as '1' and '2', and assume

[4] Sun and Yang (2006, 2009) have also proved that in the setting of my paper, there exists a dynamic incentive-compatible mechanism which finds the competitive equilibrium. In the present experiment, I only considered one-shot sealed-bid auctions, thus I could not include this mechanism in our comparison.

[5] Kagel, Lien and Milgrom (2010 and 2014), and Kazumori (2010) are good examples of this.

[6] See ComReg (2012), Danish Business Authority (2012), and Ofcom (2012).

that two of the bidders have a positive valuation on one item only. I call these the 'individual' bidders, and label them as I1 and I2, corresponding to whether they value item 1 or item 2 positively. The third bidder, J - the 'joint' bidder - has a positive value only on the bundle of items 1 and 2 together, and zero value on items 1 and 2 individually. Each bidder is permitted to bid only on the bundle they value positively, so the auctioneer always receives three bids.

To model complementarities, I assume that the individual bidders' values are drawn from a uniform distribution on [0,100], while the joint bidder's value is drawn from a uniform distribution on [0,200]. I will use b_{i1} to denote the bid of bidder I1, b_{i2} for the bid of bidder I2, and b_j for the bid of joint bidder J. The auction rule itself is described by $P(b_{i1}, b_{i2}, b_j)$, a payment vector conditional on the bid-triplet (b_{i1}, b_{i2}, b_j). Individual payments assigned by an auction mechanism to the three bidder types are labelled as p_{i1}, p_{i2} and p_j respectively, such that $P(b_{i1}, b_{i2}, b_j) = (p_{i1}, p_{i2}, p_j)$.

Prior to calculating the bidders' payments, the auctioneer solves a winner-determination problem: he picks a feasible bid-maximising allocation such that each item gets assigned to at most one bidder. In the present setting there are only two sensible allocations.[7] If the sum of individual bids is higher, the I-types win one item each; otherwise the J-type wins both.[8] The winner-determination procedure is common to all the rules I analyse.

1.1. The Vickrey Auction

The multi-unit Vickrey Auction, an extension of the standard Vickrey-Clark-Groves mechanism to the auction context, has the main aim of inducing truthful value revelation amongst the bidders. This, in turn, enables the implementation of an efficient value-maximising allocation. Irrespective of bidder type, in the Vickrey auction the price paid by each winning bidder is determined solely by the bids of the other two bidders. This price is calculated such that each bidder receives a payoff equal to the incremental surplus they bring to the auction.

For a numerical example, let $(b_{i1}, b_{i2}, b_j) = (48, 40, 60)$. Bidders I1 and I2 win the item, as the sum of their bids exceeds J's bid. The surplus that bidder I1 brings to the system is 28: without I1's bid, the auctioneer only faces the bids of $b_j = 60$ and $b_{i2} = 40$, whereby J would win both items, and the surplus - evaluated at the bidders' bids - would be 60. With I1's bid of 48, I1 and I2 win instead, and the total surplus is 88 - an increase of 28. To give I1 a surplus of 28, the payment must solve the equation $b_{i1} - p_{i1} = 28 \implies p_{i1} = 48 - 28 = 20$. By similar calculations, I2's payment is $b_{i2} = 12$.

To generalise the above reasoning after imposing a non-negativity constraint on prices, the Vickrey auction payments can be written as:

$$P^{VA}(b_{i1}, b_{i2}, b_j) = \begin{cases} (VP_{i1}, VP_{i2}, 0) & if \quad b_{i1} + b_{i2} \geq b_j \\ (0, 0, b_{i1} + b_{i2}) & if \quad b_{i1} + b_{i2} < b_j \end{cases} \tag{1}$$

$$where: \quad \begin{array}{l} VP_{i1} = \max[(b_j - b_{i2}), 0)] \\ VP_{i2} = \max[(b_j - b_{i1}), 0)] \end{array}$$

There are two well-known issues with the Vickrey auction, which limit its practical usefulness: the possibility of low revenue, and susceptibility to collusion. From equation (1) we see that in the case when $b_{i1} + b_{i2} > b_j$ with $0 < b_{i1} < b_j$ and $0 < b_{i2} < b_j$,[9] the Vickrey

[7] More allocations are feasible, but not really 'sensible': for example, only selling one item is feasible, but not sensible. Aggregate revenue could be increased by offering the unsold item at a price $\varepsilon > 0$. If a bidder's value on this item is positive, we have a Pareto improvement.

[8] Ties are broken randomly.

[9] This case corresponds to the situation where I1 and I2 together out-bid J, but neither of the individual bids, on their own, would be sufficient to out-bid the joint bidder.

auction 'leaves money on the table', in that $p_{i1}+p_{i2}<b_j$: the seller has a seen a bid that exceeds the sum of payments he receives from the winning bidders. This is equivalent to saying that the Vickrey auction outcomes frequently lie outside the core.

The core is defined as a set of allocations for which there exists no blocking coalition, such that no (sub)group of members of the system can jointly deviate to a different allocation which gives all those members a higher surplus. In the present example, the group consisting of bidder J and the auctioneer together constitutes a blocking coalition: J could offer the auctioneer a payment of $\tilde{p}_j = p_{i1}+p_{i2}+\varepsilon < b_j$, with $\varepsilon > 0$. This increases the seller's revenue, and gives bidder J a non-zero profit - so the allocation that assigns the items to I1 and I2 is not a core allocation, and the price-triplet $(p_{i1}, p_{i2}, 0)$ does not lie in the core.[10]

When $b_{i1}+b_{i2}>b_j$, the set of core payments can be defined as:

$$(p_{i1}, p_{i2}) \in \{(x,y)\,|x+y \geq b_j, x \in [0,b_{i1}], y \in [0,b_{i2}]\}\,.$$

This is the set of payments such that neither I1 or I2 pays more than their bid, but the sum of their payments weakly exceeds the bid of J. This set, along with the bids and Vickrey payments are shown in Figure 1 - the core is shaded in gray. The dotted diagonal line denotes the 'minimum revenue line', which contains all the points where the payments of I1 and I2 equal the payment of J exactly. The bold segment of this diagonal line depicts the 'minimum revenue core' (MRC),[11] which contains the points that are simultaneously in the core, and on the minimum-revenue line. The MRC depicts the combination of the lowest amounts that each of the I-types can bid, subject to them jointly out-bidding the J-type. From the seller's viewpoint, this is analogous to a 'second-price' in a single-unit auction: this is the highest observed bid after the actual winning bids have been removed.

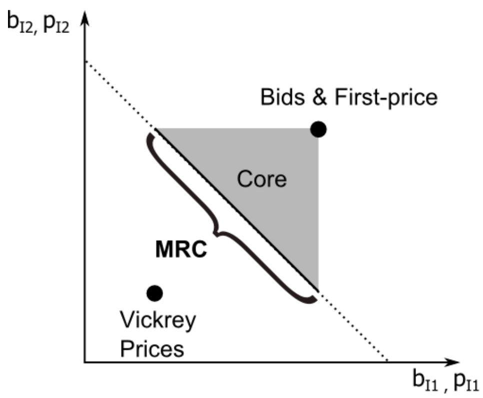

Figure 1. Vickrey prices, First–price payments and the MRC

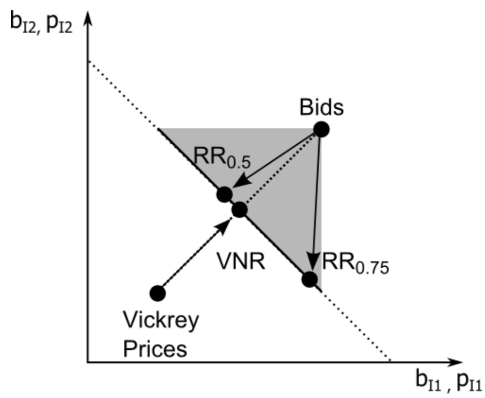

Figure 2. Vickrey Nearest, and Reference Rule with $\alpha = 0.5$ and $\alpha = 0.75$

The second weakness of the Vickrey auction is its susceptibility to collusion. We see from equation (1) that when I1 and I2 win, the payment of one is decreasing in the bid of the other.[12] If I1 and I2 want to behave cooperatively, they can both bid aggressively, which will reduce their joint payments. To collude perfectly I1 and I2 can both bid $b_{i1} = b_{i2} = 200$, which is the highest possible value that J can have. Such bids makes

[10] In the case when J wins the Vickrey payment is in the core, as then $b_j > b_{i1} + b_{i2}$.

[11] For a further detailed discussion of the MRC, see Day and Milgrom (2008).

[12] Consequently the Vickrey auction revenue is not always monotonic in bids: it is possible that an auciton with higher (individual) bids can lead to lower revenue.

sure that I1 and I2 always win, and both pay a price of 0. In less extreme cases, so long as both bidders overbid, they can still induce payments lower their Vickrey prices under truthful bidding. If only one of the two individual bidders attempts to collude and his co-bidder does not reciprocate, it is likely that the collusive bidder makes a loss.

1.2. The First Price Auction

The first-price auction, usually used for the sale of a single item, can be naturally extended to cover the case of package bidding. After the winner-determination problem has been solved, each winning bidder pays their bid in full for the item that they are allocated. The payments in the first-price auction can be summarised as:

$$P^{FP}(b_{i1}, b_{i2}, b_j) = \begin{cases} (b_{i1}, b_{i2}, 0) & if \quad b_{i1} + b_{i2} \geq b_j \\ (0, 0, b_j) & if \quad b_{i1} + b_{i2} < b_j \end{cases}.$$

Unlike the payments in the Vickrey auction, the first-price auction the winners' payments are always in the core, as shown in Figure 1. In the case when I1 and I2 win, the first-price payments will also always lie (weakly) above the minimum-revenue line. Despite its simplicity, the first-price auction with package bidding has been used in practice numerous times, including the auctioning of bus routes in London (see Cantillon and Pesendorfer, 2006) and mobile telephony spectrum in Norway in 2013.[13]

1.3. The Vickrey Nearest Rule

The Vickrey Nearest Rule (VNR) is one of many recent implementations of core-selecting auctions. One motivation behind these payment rules is to increase the revenue from Vickrey-type auctions while retaining most of their efficiency and truth-telling properties. Such a trade-off can achieved by making the winners' payments less dependent on their own bids, but still require that the payment vector lies in the core.[14] The VNR auction, as introduced by Day and Cramton (2012), first uses the submitted bids to calculate Vickrey payments, and then picks a price vector that minimises the Euclidian distance to the Vickrey payments subject to the prices being in the core.

In the case when bidder J wins, the Vickrey payment is in the core already, and the VNR implements that payment. If I1 and I2 win, the VNR will select the point on the MRC which is closest to the Vickrey payment vector, as shown in Figure 2.

Mathematically, finding the point on the MRC that is closest to the Vickrey payments involves taking an orthogonal projection of the bid vector onto the MRC. I label the outcome of such a projection as the 'preliminary shares' of bidders I1 and I2, and denote them as s_{i1} and s_{i2}. The VNR payments can then be summarised as:

$$P^{VNR}(b_{i1}, b_{i2}, b_j) = \begin{cases} (s_{i1}, s_{i2}, 0) & if & \begin{matrix} b_{i1} + b_{i2} \geq b_j, \ and \\ s_{i1}, s_{i2} > 0 \end{matrix} \\ (b_j, 0, 0) & if & b_{i1} \geq b_j + b_{i2} \\ (0, b_j, 0) & if & b_{i2} \geq b_j + b_{i1}, \\ (0, 0, b_{i1} + b_{i2}) & if & b_{i1} + b_{i2} < b_j \end{cases} \tag{2}$$

$$where \quad : \quad \begin{matrix} s_{i1} = \frac{1}{2}(b_{i1} + b_j - b_{i2}) \\ s_{i2} = \frac{1}{2}(b_{i2} + b_j - b_{i1}) \end{matrix} \tag{3}$$

The payments of individual bidders in the VNR are broken down into cases, depending

[13] Information taken from the Norwegian Post and Telecommunications Authority document "800, 900 and 1800 MHz auction - Auction Rules" (2013).

[14] The intuiton is that if incentives to deviate from truth-telling are small, bidders will bid in a near-truthful way, which would mitigate efficiency losses due to misallocation.

on the asymmetry of the bids. If, say, $b_{i1} > b_j + b_{i2}$, so that I1 on his own out-bids J by a large margin, then $s_{i2} < 0$, which implies a negative price for I2. By the non-negativity constraint on prices, we then truncate $p_{i2} = 0$, and $p_{i1} = b_j$ to remain on the MRC. The converse case applies if $b_{i2} > b_j + b_{i1}$. When the asymmetry moderate and $s_{i1}, s_{i2} > 0$, both bidders pay their preliminary share.

1.4. The Reference Rule Auction

The Reference Rule, introduced by Erdil and Klemperer (2010), is another payment rule for core-selecting package auctions. The motivation behind the rule is to make it more robust to small local deviation incentives than the VNR by further de-coupling individual payments from bids. In the VNR, individual bidders can influence their payment share by influencing the Vickrey prices, which depend on their own bid, as shown in equation (3). The innovation behind the Reference Rule is to define the bidders' payment shares in a way that further reduces the dependence on their own bids, while maintaining the core-selecting property. This is achieved defining a 'reference point' which is independent of the I-types' bids, and then selecting the final payments that are closest in Euclidian distance to that point.

I will define each individual bidder's reference price based on the bid of the joint bidder J and a sharing parameter α. I denote the corresponding reference rule by RR(α). The reference price of bidder I1 is $r_{i1} = \alpha \cdot b_j$, and the reference price for bidder I2 is $r_{i2} = (1-\alpha) \cdot b_j$, with $\alpha \in [0,1]$. Using this parametrisation, by varying α the reference point can be moved smoothly along the minimum-revenue line, with higher α setting the reference point closer I1's axis. The bidder payments in the Reference Rule then are:

$$P^{RR(\alpha)}(b_{i1}, b_{i2}, b_j) = \begin{cases} (r_{i1}, r_{i2}, 0) & if \quad \begin{array}{c} b_{i1} + b_{i2} \geq b_j,\ and \\ r_{i1} < b_{i1}, r_{i2} < b_{i2} \end{array} \\ (b_j - b_{i2}, b_{i2}, 0) & if \quad \begin{array}{c} b_{i1} + b_{i2} \geq b_j,\ and \\ r_{i1} < b_{i1}, r_{i2} > b_{i2} \end{array} \\ (b_{i1}, b_j - b_{i1}, 0) & if \quad \begin{array}{c} b_{i1} + b_{i2} \geq b_j,\ and \\ r_{i1} > b_{i1}, r_{i2} < b_{i2} \end{array} \\ (0, 0, b_{i1} + b_{i2}) & if \quad b_{i1} + b_{i2} < b_j \end{cases} \tag{4}$$

$$where : \begin{array}{l} r_{i1} = \alpha \cdot b_j \\ r_{i2} = (1-\alpha) \cdot b_j \end{array}$$

Since reference prices are only required to lie on the minimum-revenue line, and not on the MRC, it is possible that the reference point will lie outside the core. Then the point on the MRC that is closest to the reference point is a payment vector where one individual bidder (say, I1) pays his bid in full, while the other individual bidder's payment makes up the difference (between J's and I1's bid), such that the sum of payments ends up on the MRC.

In the VNR, each individual bidder's payment share always depends in part on his own bid. In the Reference Rule, so long as the *realised* reference point is on the MRC, the payment for each individual bidder is completely *in*sensitive to their own bid. The *only* case in which an individual bidder's payment depends on his bid is in the situation when the realised reference point is outside the MRC *and* he is the bidder that has to pay his bid in full. This sensitivity occurs only under certain realisation of bids, and hence has limited impact on average.[15]

[15] Erdil and Klemperer (2010) show that under plausible conditions the Reference Rule has a lower sum of 'local deviation incentives' than VNR, while the sum of 'maximum deviation incentives' is unchanged. The proof proceeds by trading off the cases where bidders have zero incentives with those where incentives

In general, as Figure 2 shows, the reference rule with $\alpha = 0.50$ generates payments different from VNR.[16] However, with $\alpha = 0.50$, the reference payments are the same as they would be in the Proxy Rule auction of Ausubel and Milgrom (2002). Hence to make the Reference Rule look significantly different from the VNR and Proxy Rule auctions, I chose to use $\alpha = 0.75$ in the main experiment. Supplementary data for the Reference Rule with $\alpha = 0.50$ was obtained from an additional experiment, detailed in the Appendix.

1.5. Comparison of the four Auction Rules

To give a concrete comparison of the four auction rules, Figure 3 summarizes the outcome from applying each rule to the bid-triplet $(b_{i1}, b_{i2}, b_j) = (48, 40, 60)$. The I-types win, and the J-type pays zero in every auction. To show the influence of varying α on the behaviour of the Reference Rule, I have calculated the payments for three values of α, denoted by RR(α). For RR(0.25) the reference prices will be $r_{i1} = 15$ and $r_{i2} = 45$, which is outside the core, so the Reference Rule payments will be truncated to lie on the boundary of the MRC. This is not the case for $RR\,(0.75)$, and the payments in that case are not in the corner of the core.

	p_{I1}	p_{I2}	Revenue
First-price	48	40	88
Vickrey	20	12	32
VNR	34	26	60
RR(0.75)	45	15	60
RR(0.50)	30	30	60
RR(0.25)	20	40	60

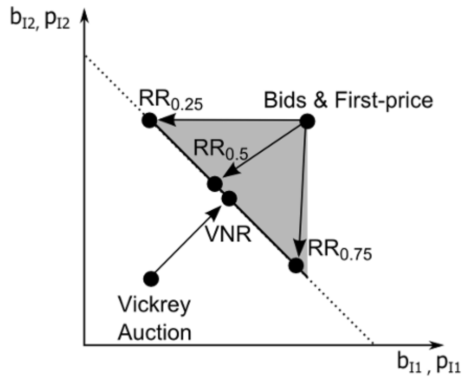

Figure 3. A numerical example of the four auction rules, with $(b_{i1}, b_{i2}, b_j) = (48, 40, 60)$

1.6. Bidding Restrictions and Collusion

None of the auctions I analyse require bidding above value in a competitive equilibrium, so in theory, a restriction prohibiting such bids should have little bite. Investigating the impact of such restrictions is nonetheless worthwhile for two reasons. Firstly, even in simpler single-item auction contexts many experimental papers, such as Kagel and Levin (1995), find that overbidding is a frequent phenomenon. Bidders bid more than theory would predict, sometimes even above their value, which can lead to negative payoffs.[17] It is useful to gauge whether such overbidding influences the performance of the rules examined here, and whether it could be the driving force behind any revenue or efficiency rankings.

The second reason for investigating bidding restrictions is that it allows me to look for collusion in the Vickrey auction. Here both individual profits as well as auction revenue can be very sensitive to the presence of overbidding. For the other three auctions no col-

---

are maximal, and comparing these with the VNR, which has moderate incentives everywhere.

[16] The Reference Rule with $\alpha = 0.50$ generates reference payments on the mid-point of the minimum-revenue line, while the VNR selects payment shares at the mid-point of the MRC. Unless $b_{i1} = b_{i2}$, these two points will differ.

[17] For a good summary of this literature and further references, see Section 1.4 of Kagel & Levin (2008), and Section I.b2 in Kagel (1995).

lusive strategies have been found.[18] Running a set of sessions with the same instructions, but with adding or removing bidding restrictions, allows for a clean and direct assessment of this effect.

2. Hypotheses

The most direct application of this kind of experiment is to test existing economic theory on the underlying auctions - this is surveyed in Section 2.1. Yet even in simpler settings and when complementarities are absent, the experimental auction literature frequently rejects theoretical predictions.[19] In addition, the standard results do not take into account the possibility of collusion, which may be a significant factor affecting the practical performance. Thus I propose some additional intuitively plausible hypotheses in Section 2.2, which can also be tested in the current experiment.

2.1. Related Theory and Literature

For my paper, the most relevant experimental work on auctions is Kagel, Lien and Milgrom (2010, 2014) and Kazumori (2010, 2014). Kagel et al. compare the performance of a combinatorial clock-auction with that of a simultaneous ascending auction for a variety of value and complementarity settings. Their interest is assessing how well the auctions perform when bidders bid only on a subset of profitable packages in each round, rather than bidding on all packages. They find that straightforward bidding - submitting bids on the most profitable package only - leads to efficient outcomes (Kagel et al. 2010), though bidders sometimes diverge from such bidding patterns to push up prices for their competitors (Kagel et al. 2014). Kazumori (2010) investigates generalized Vickrey auctions, in addition to clock-proxy and simultaneous-ascending auctions. He finds that clock-proxy auctions out-perform the generalized Vickrey auction, and also outperform the simultaneous-ascending auction when the value structure mirrored exposure. However, both these papers have looked at dynamic auctions, with complicated value and complementarity structures, and their focus has been on efficiency and package-selection.

My work, in contrast, looks at static one-shot auctions, with a fixed package structure , and allows me to check whether in a simpler context the bidding will diverge from predictions once the package-selection aspect is removed.[20] In practice, in many high-value package auctions a hybrid design is used, where a clock phase is followed by a single supplementary bidding round which determines final prices and package allocation.[21] My research is thus a complement to, rather than a substitute for, the dynamic experimental auction literature.

The papers of as Ausubel and Baranov (2010), Goeree and Lien (2009) and Sano (2010) provide theoretical foundations for optimal bidding in the auctions I analyse, under an analogous valuation model. The authors find that in the VNR the optimal strategy involves individual bidders with values below a certain cutoff to all pool into submitting a bid of zero, while all bidders with values above this cutoff should shade their bid by

[18] As of yet, there is no clear analysis as to the collusion incentives in VNR and the Reference Rule. The presumption is that being core-selecting auction rules, they should be robust to attempted collusion.

[19] Kagel (1996 and 2008) are a good overview of this literature.

[20] Kazumori (2014) has also conducted an experiment on one-shot package auctions, in a setting similar to mine, but his analysis only compares the Vickrey and Ausubel-Milgrom (2002) proxy auctions. He finds that proxy auctions revenue-superior, which is congruent with the results of this paper.

[21] The dynamic phase thus determines which packages are relevant, but does not necessarily fix the final allocation of packages to bidders.

a constant amount.[22] The optimal strategy for bidder type J is to bid his value. To obtain optimal bidding functions for the case of the first-price auction, Baranov (2010) uses numerical methods, since a solution cannot be found analytically. In the equilibrium of the Proxy Rule auction, which is equivalent to RR(0.50), low-value individual bidders pool in bidding zero, and higher value bidders will shade by a positive amount. The degree of shading decreases with v; for a low value of v, individual bidders will shade more in the Proxy Rule auction than in the VNR auction, and the ordering is reversed at higher values. At the extreme, when $v_i = 100$, there is no shading in the Proxy Rule. At the auction level, Ausubel and Baranov (2010) find that the Vickrey auction gives highest revenue, followed by the first-price auction, with VNR and Proxy Rule giving almost identical revenues, below the other two auctions. The efficiency ranking follows the same pattern as revenue.

Combining the findings of Ausubel and Baranov (2010) with the well-known prediction of truthful bidding in the Vickrey auction, I test the following set of theory-based hypotheses:

- Hypothesis HT: Bidders follow the competitive equilibrium bidding strategies.
- Hypothesis HR: The revenue ranking has Vickrey auction first, followed by first-price, with VNR and the RR(0.50) last.
- Hypothesis HE: The ranking for efficiency is the same as in HR.

2.2. Intuition-based Hypotheses

Even if bidders do not follow equilibrium strategies, they may still respond to auction incentives to some extent. It is thus worthwhile to assess the broader intuitions that could influence behaviour under the different rules. In the Vickrey auction, a bidder's price conditional on winning is independent of his bid, while there is a partial dependence in the core-selecting rules. We should hence expect to see more aggressive bidding in the Vickrey than in the core-selecting auctions. In the first-price auction, conditional on winning the price equals the bid exactly, which we would expect to invite more cautious bidding. This ranking of incentives does not apply to the J-type bidders, who face the same payment rule under all auctions except first-price. Testing whether such bidders bid truthfully is contained in the hypothesis HT, but even if that hypothesis fails, it is possible that the J-types follow a similar non-truthful bidding pattern. I propose the following intuition-based hypotheses:

- Hypothesis HB: Individual bidder types will bid most aggressively (shade the least) in the Vickrey auction, and shade most in the First Price auction.[23] The Reference Rule and VNR rank as intermediate.
- Hypothesis HJ: Bidder J bids similarly in all auctions other than first-price.

In the discussions of Day and Cramton (2012) and Erdil and Klemperer (2010), part of the motivation for core-selecting auctions is that bidders may in fact not use full equilibrium strategies, but rather follow 'rules of thumb' which appear suitable for the auction in question. The VNR and the Reference Rule were developed to minimise incentives for deviation from truthful bidding. The intuition is that because payments are 'close to independent of own bids' then bidders could find it 'close to optimal' to bid truthfully. This intuition naturally generates another hypothesis:

[22] In auction terminology, shading is defined as the difference between the bidder's true value, and the bid he submits for a given item or package.

[23] Shading is standardly defined as the difference between a bidder's value and their bid.

- Hypothesis HA: Individual bidders bid truthfully in the VNR and Reference Rule.

The final set of hypotheses I test in this experiment relate to collusion in the Vickrey auction. Collusion can be defined as behaviour that deviates from an individually optimal competitive strategy towards one that aims to maximise joint profits of the colluding parties.[24] The general tendency in the collusion literature is to provide bidders in rich bidding contexts with many opportunities to collude, and look for periods of play when collusion is successfully sustained. Examples of this approach include Goswami, Noe and Rebello (1996) and Sade, Schnitzlein and Zender (2005), who look at collusion in discriminatory and uniform-price auctions with communication. Kwasnica and Sherstyuk (2007) similarly investigate Simultaneous Ascending Auctions with repeated play (within the same bidder group), but no communication. The survey of Kagel and Levin (2008) finds that repeated play with the same opponents, and communication, tend to facilitate collusion, though this survey doesn't cover any experiments on multi-unit Vickrey auctions.

In light of the above papers, the setup of my experiment is not inherently conducive to collusion: the matching is random across periods, and communication is prohibited. My experiment was the first auction study ever run at the laboratory I used, hence few of the participants are likely to have prior auction experience.[25] The valuation setup, however, is very simple and the Vickrey auction rules are straightforward, so the collusive strategies are easy to deduce: under perfect collusion, the I-types should bid exactly 200. Even if bidders do not notice this corner solution, it is possible that the I-types realise that they can mutually benefit by bidding significantly above value. None of the other auctions in the experiment give obvious incentives for bidding in excess of value, so I would not expect bidding behaviour to change much irrespective of whether a bidding restriction is in place or not. If I do observe significant change of bidding patterns in the Vickrey auction across these two treatments, together with numerous bids in excess of value, these findings would be consistent with attempted collusion. I will thus test the following two hypotheses:

- Hypothesis HS: In auctions other than the Vickrey auction, the presence of bidding restrictions does not significantly affect bidder behaviour.

- Hypothesis HC: Removal of bidding restrictions in the Vickrey auction influences bidding behaviour. Without bidding restrictions, the I-types bid more aggressively and in excess of their value.

3. Experimental Design

The experiment was run over four sessions, and the participants were recruited from the population of Oxford graduate and undergraduate students via the mailing list at the Centre for Experimental Social Sciences (CESS) laboratory at the University of Oxford. Only students from science and social science subjects were included in the recruitment mailshot, and no participant was allowed to play in more than one session. The experiment itself was programmed using the zTree software of Fischbacher (2007), and run at the CESS laboratory. Sessions lasted up to two and a half hours, with average earnings of around £35 ($\approx$ \$55).[26]

[24] Playing a collusive strategy in itself is not necessarily non-equilibrium behaviour - in games where multiple equilibria exist, a 'collusive' outcome can be one of such equilibria.

[25] I cannot exclude the possibility that they would have participated in auction experiments elsewhere.

[26] A sample of the instructions is available in the Online Appendix, on my website, at http://daniel.marszalec.com.

During each session, the same group of participants played in each of the four auctions. The attendance was between 18 to 30 participants per session. After receiving the instructions for a given auction type, the participants were allowed to ask clarifying questions, and then were presented with an understanding test. Upon passing the test they participated in two payoff-irrelevant practice rounds, followed by the ten payoff-relevant rounds of the same auction rule. This design yielded 140 auction-round observations for each rule from the sessions without bidding restrictions and 160 auction-rounds with bidding restrictions present. The matching of participants to groups and bidder types was random each round, and communication was not permitted. Once the paying rounds of a given auction type were complete, the instruction sheets for that auction were collected, and the instructions for the next auction were distributed.[27]

A sample of the understanding test that the participants were required to complete is provided in the Online Appendix. The test was administered on paper, and there were few failures.[28]

To allow for an analysis of the importance of overbidding and possible collusion in the Vickrey auction, two of the four sessions were run with the bidding restrictions in place, prohibiting the bidders from bidding above value. In the other sessions the bidding restrictions were removed, and all three bidders were allowed to bid any number in $[0, 200]$.[29] The participants were paid for each auction rule based on their profits in two randomly selected rounds (out of the ten played); if the sum from these two rounds was negative, the payoff for that auction was truncated to zero. Final payments were calculated as the sum of payoffs from all four auction types, plus a show-up fee.

4. A Comment on Data

Since the experimental design is within subjects, we need to verify that bids are independent across auctions. To assess this degree of dependence, I ran series of pairwise estimations of Kendall's τ correlation parameter and tested its significance.[30] None of the tests for I-type bidders rejected a no-correlation null, with all p-values>0.15. The tests on the J-type bidders also fail to reject the hypothesis that the bidding patterns are uncorrelated at the 95% level. These results suggest that there is little correlation between bidding pattern across auction types, and that the assumption of independence between treatments for testing purposes is acceptable.

In addition to the four sessions where bidders bid in all four auction rules, I also ran another set of experiments in an analogous setting, but focusing only on the effects of α in the Reference Rule; the details of these experiments are outlined in the Appendix.[31] Due to time-constraints (and participant fatigue), it was not feasible to run both $\alpha = 0.75$ and $\alpha = 0.50$ treatments in the main sessions. Since the data for RR(0.50) was available, I have included it in the comparisons for the present paper, though with the caveat that it is possible that participants' behaviour in RR(0.50) would be somehow influenced by

[27] The ordering of the auction rules was: [VCG,VNR,RR,FPS] in one set of sessions, and [VCG, FPS,RR,VNR] in another. These orderings were generated randomly, but for consistencly the same pair of orderings was used in both restricted and unrestricted bidding sessions.

[28] On average, between one or two out of every thirty subjects failed the test.

[29] The bidders were made aware that when bidding was unrestricted, even though they would never pay more than their bid, they could nonetheless end up with a negative payoff if they overbid and win at a price that exceeds their valuation.

[30] The purpose of this test is to check that the assumptions of the statistical test I use later are satisfied. While values are independent by design, I must check that the bidding process itself did not induce a strong pattern of dependece.

[31] The data collected in the supplementary experiment consisted of 140 auciton-rounds for each rule - the same number as in the unrestricted bidding sample of the main experiment.

their *not* playing in the other three auctions.

To allay concerns of such an effect, the supplementary experiments also contained a treatment where $\alpha = 0.75$, and this permitted for a consistency check between the two datasets. I tested the 'no difference' null using Mann-Whitney and Kolmogorov-Smirnov tests on the raw bid data, as well as direct tests of means and medians; in each case, the null could not be rejected even at the 90% level. These results suggest that the behaviour for the $\alpha = 0.75$ case is similar in both the main experiment as in the supplementary sessions, so the effects of presenting the Reference Rule in the two different setting are likely to be minor.

5. Auction-level Results

Revenue, surplus and efficiency are the three main parameters of interest for evaluating auction performance. Revenue is often of foremost importance to sellers, while bidders are primarily interested in their own surplus. From a welfare or policy point of view efficiency is also relevant, so that the items are allocated to the highest-value buyers.[32] An overview of these three parameters in the unrestricted bidding sample are provided in Table 1.[33]

Table 1: Revenue, Efficiency and Surplus Summary

	Vickrey	First Price	VNR	RR(0.50)	RR(0.75)
revenue	67.6 (56.9)	91.5 (37.1)	68.2 (41.2)	77.0 (42.3)	71.1 (46.3)
surplus	44.1 (67.6)	29.8 (28.1)	57.9 (39.1)	48.9 (49.3)	46.7 (49.6)
efficiency (%)	88.9 (22.2)	97.5 (8.4)	97.7 (9.1)	94.9 (13.8)	95.1 (12.8)

Means reported, standard deviation below. Revenue and surplus reported. as points. The calculations are based on all 140 experimental auction rounds.

One immediately visible characteristic of Table 1 is how distinct the first-price auction looks from the others: the revenue is higher, surplus is lower, and both variables have lower variance than in the other auctions.

Results from the pairwise tests and comparisons on revenue equivalence are shown in Table 2. The first-price auction revenue-dominates all other rules, while pairwise comparisons between the Vickrey, VNR and Reference Rule cannot reject revenue equivalence. Though revenue in the Vickrey auction is lower than under VNR and Reference Rule, this difference is not statistically significant. I also cannot reject equivalence between the two kinds of Reference Rules with different values of α. This revenue ranking runs contrary to hypothesis HR, which I reject. The first-price auction performs better than predicted, and the Vickrey auction underperforms.

Mirroring the results from the revenue figures above, the first-price auction generates less bidder surplus than any of the other three rules: all pairwise tests reject in this direction at a confidence level of 95% or stricter (see Table 2). All other pairings fail to reject the zero-difference null. Pairwise testing confirms the intuitive conclusion from Table 1: the first-price auction is very different from the others, giving higher revenue

[32] Efficiecny here is calculated as: $100\% \cdot \frac{\text{sum of winning bidders' values}}{\text{sum of values under value-maximising allocation}}$

[33] A parallel analysis for the restricted-bidding sample is conducted in the Online Appendix.

Table 2: Pairwise Auction Revenue and Surplus Comparisons

Revenue	Vickrey	VNR	RR(0.50)	RR(0.75)
FirstPrice	29.0***	24.0***	15.0**	23.0***
Vickrey		−3.0	−13.0	−7.0
VNR			−9.0	−1.0
RR(0.50)				8.0

Surplus	Vickrey	VNR	RR(0.50)	RR(0.75)
FirstPrice	−16.0**	−24.0***	−17.0***	−17.0***
Vickrey		−10.0	−2.0	−1.0
VNR			8.8	8.0
RR(0.50)				0.0

Reported values are for median-difference of (row - column). Rejections of zero-difference null at 90%/95%/99% level indicated by */**/***; Bonferroni-Holm corrections applied.

and lower surplus.[34]

Assessing efficiency using a direct median-comparison test, as above, is unhelpful, because in all the treatments the median efficiency is 100%. A Kruskal-Wallis test nonetheless rejects with p-value<0.005, suggesting that efficiency is not homogenous across auctions. Hence I ran a series of Mann-Whitney tests, pairwise for each combination of auctions; this allows me to check the distribution of efficiency in each pairing, and to see whether a clear dominance pattern emerges. All but one pairwise comparisons against the Vickrey auction reject at the 95% level or stricter, with Vickrey auction giving lower efficiency.[35] No other strict ranking pattern emerges from the pairwise tests. These findings provide evidence to reject hypothesis HE, according to which the Vickrey auction should be most efficient.

5.1. Bidding Constraints and Bidder Behaviour

I check the impact of bidding constraints by comparing the raw bid patterns across the two treatments, as summarised in Table 3.[36] Removing bidding constraints only significantly changes behaviour in the Vickrey auction. The bids are higher when restrictions are lifted, with a median difference of 30 for bidder I1, and 20 for I2. To put these numbers in perspective, recall that I-type values are uniform on [0,100] implying a median value of 50 - the median increase in bids is at least 40%. The median-difference test accordingly rejects for all bidder types under the Vickrey auction at the 99% confidence level,[37] but none of the other auctions register any rejections.

On this evidence, I cannot reject hypothesis HS: bidding constraints have no impact on first-price, VNR and Reference Rule auctions. In subsequent portions of the paper, I will conduct the analysis using data from the sessions with unrestricted bidding; a parallel

[34] The revenue and surplus conclusions of this section are precisely mirrored in the results from the restricted-bidding sample, and are included in the Online Appendix.

[35] The single auction that does not reject pairwise efficiency equivalence with the Vickrey auction is RR(0.50).

[36] The RR(0.50) auction is not included in this comparison, since none of the supplementary sessions were run with bidding restrictions.

[37] These are calculated using the Hodges-Lehmann method, implemented through the SomersD package in Stata (Newson, 2006).

Table 3: The influence of bidding restrictions on bids

Case		Vickrey	First-Price	VNR	RefRule(0.75)
Bidder I1	Medians	84.0 | 50.0	35.0 | 34.5	45.0 | 40.0	45.0 | 39.5
	Median Difference	30.0***	−2.0	3.0	5.0
Bidder I2	Medians	75.0 | 56.5	30.0 | 30.0	50.0 | 39.5	45.5 | 44.0
	Median Difference	20.0***	−2.0	5.0	4.0
Bidder J	Medians	136.0 | 90.0	65.0 | 79.5	100.0 | 90.0	106.5 | 91.0
	Median Difference	27.0***	−8.0	7.0	11.0

Medians reported as: Unrestricted | Restricted. Median difference tested via the Hodges-Lehmann method.
Rejections of zero-difference null at 90%/95%/99% level indicated by */**/***.

analysis for the restricted-bidding sessions is available in the Online Appendix. The large difference registered in the Vickrey auction is consistent with hypothesis HC on collusion, and this finding will be further analysed in Section 6.3.

6. Bidder-level Results

With the exception of the Reference Rule with $\alpha = 0.75$, all other auction settings analysed in this paper offer symmetric incentives for both I-type bidders, and the data from these two sub-cases could be pooled for analysis. This intuition is confirmed by the data: in the symmetric auctions, Mann-Whitney tests for the zero-difference null fail to reject on both the bid and shading variables (all p-values >0.15). For the purpose of further analysis in this section, the data for I1 and I2 types will thus be pooled in all auctions except RR(0.75), where I will consider both types separately.

To give an overview of individual bidding and assess hypothesis HB, Table 4 shows a set of pairwise median-difference tests across auctions for the bid variable. I-types bid the most in the Vickrey auction, and the least in first-price. The three core-selecting auctions rank as intermediate, and show no significant difference from each other. The intuition of hypothesis HB cannot be rejected - the data shows that indeed Vickrey auction induces aggressive bidding, while first-price discourages it.

Table 4: Pairwise Comparison of I-types' Bidding Behaviour

Bids	Vickrey	VNR	RR(0.50)	RR(0.75)[I1]	RR(0.75)[I2]
FirstPrice	−44.0***	−14.0***	−16.0***	−13.0***	−13.5***
Vickrey		30.0***	26.0***	30.0***	27.0***
VNR			−2.0	0.0	0.0
RR(0.50)				3.0	1.0

Reported values are for median-difference of ("row" - "column").
Rejections of zero-difference null at 90%/95%/99% level indicated by */**/***;
Bonferroni-Holm corrections applied.

When assessing the validity of Hypothesis HJ - that the J-type bidders bid similarly in all auctions except first-price - the Kruskal-Wallis tests for equality of populations rejects (p-value=0.005), suggesting that there are differences in bidding behaviour across auction types. On this evidence, the data reject hypothesis HJ.

6.1. Bidder-level Tests of the Theory

The theory results being tested in this section base on the equilibrium bidding functions derived for the first-price, VNR, and RR(0.50) auctions by Baranov and Ausubel (2010). As no analytical results are available for RR(0.75) due to the asymmetry between I1 and I2, I obtained the equilibrium bidding functions numerically.[38] In all of the core-selecting auctions equilibrium bidding requires the I-types to bid exactly zero when their values are sufficiently low, and attempt to free-ride on the other I-type out-bidding the J-type on their own. Table 5 shows that experimental results diverge significantly from theory, and Figure 4 provides an illustration of how experimental bidding functions for I-types compare to their theoretical counterparts.[39,40]

Table 5: Bidder-level Tests of the Theory

I-types	Vickrey	FirstPrice	VNR	RR(0.50)	RR(0.75)-I1	RR(0.75)-I2
Bid	80.0(48.0)***	31.5(18.2)***	48.5(35.8)***	45.0(3.1)***	50.0(32.7)***	45.5(48.5)
Win%	67.1(52.1)***	47.1(45.0)	47.9(36.4)***	39.3(32.9)**	52.9(35.7)***	52.9(35.7)***
Surplus	31.0(39.0)*	14.3(35.3)***	26.5(30.5)*	21.0(32.6)**	14.9(41.0)***	25.8(29.8)
J-type						
Bid	136.0(92.0)***	65.0(47.1)***	100.0(98.5)	122.5(112.0)**	106.5(94.5)**	
Win%	32.9(47.9)***	52.9(55.0)	52.1(63.6)***	60.7(67.1)	47.1(64.3)***	
Surplus	31.0(48.0)*	25.0(70.3)***	55.0(70.3)	45.0(63.7)***	47.0(61.7)	

For bid and surplus, experimental medians reported; theory-based medians in parentheses.
Sign-test used for testing bid and win% variables, median-based permutation test used for surplus.
Rejections of zero-difference null at 90%/95%/99% level indicated by */**/***.

For I-types, the bidding variable rejects in all sub-cases, with the exception of the I2-bidder in the RR(0.75) auction; the general pattern indicates that I-type bidders bid more than predicted by theory. Furthermore, the I-types bid exactly 'zero' much too rarely: theory would predict a total of 531 bids at zero in my data, whereas only 83 were submitted. Beyond the misunderstanding of bidding incentives, it is likely that 'boundary effects' - the aversion to bid exactly at the boundary of the bidding support - may contribute to the scarcity of zero-bids.[41]

The J-types also overbid relative to theory in all auctions except VNR. However, in the core-selecting auctions and the Vickrey auction, the overbidding of the I-types dominates, which results in them winning more often than expected. Consequently the I-types also receive lower surplus, conditional on winning, in all cases except the I2-bidder in RR(0.75).

[38] The method used is similar to that which Baranov (2010) uses to obtain the equilibrium bid functions in the first-price auction.

[39] Analogous graphs for the J-types are provided in the Online Appendix.

[40] In Table 5, I use standard non-parametric tests for all variables except the 'surplus'. The surplus is calculated 'conditional on winning' which introduces a complex pattern of dependence across the 'experimental' and 'theoretical' samples: there are situations where an actual bid won in the experiment, whereas the corresponding theoretically predicted bid would not have won (and vice versa). Thus the samples are neither independent, nor matched-pairs. Given this dependence, I cannot use bootstrapping and use permutation-based tests instead. For further discussion of permutation tests, see Good (1994).

[41] A good analysis of this effect is Palfrey and Prisbrey (1997) in the context of public-goods contributions. In the present experiment, there is no way to test for this effect directly.

The variable for winning probability does not reject in the first-price auction, suggesting that though both I and J types overbid considerably, this does not affect their relative winning chances. Conditional on winning, both types make less profit in the first-price auction than theory predicts.

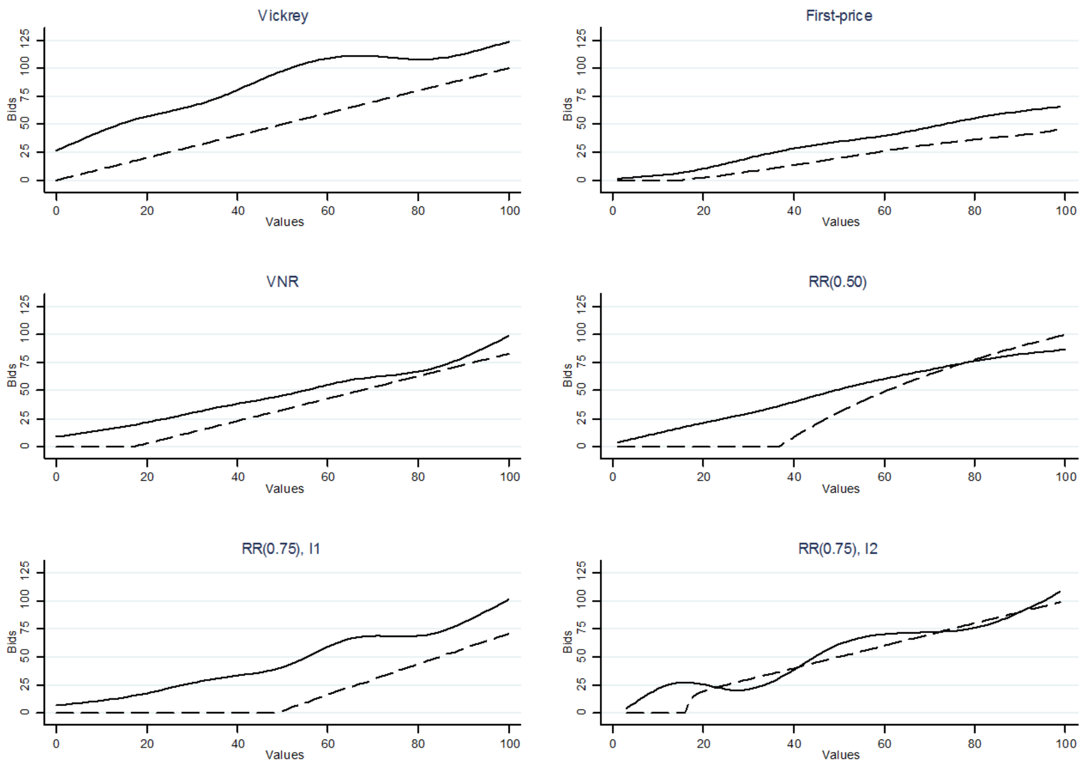

Figure 4. Bidding functions for I-type bidders: experimental (solid) and theoretical (dashed). Curves fitted are cubic splines with 7 knots and no monotonicity constraints imposed.

The broad conclusions from Table 5 and Figure 4 suggest that in all auctions the I-types overbid significantly relative to theory, therefore winning too often, but making lower profits than predicted. Correspondingly, in all auctions except first-price, the J-type wins too rarely, and when he does win he makes little profit. Jointly, these findings lead me to reject hypothesis HT - competitive equilibrium bidding theory is not supported by my data.

Hypothesis HA, on truthful bidding in core-selecting auctions, similarly finds no support in my experiment. The highest p-value generated by a sign-test for truthful bidding is for the I2-type in the Reference Rule, and here p-value =0.039, which is still a rejection at the 95% level. For all other cases, the sign-test generates p-values <0.001. When deviating from theory, the bidders do not use a truth-telling rule-of-thumb.

6.2. Evaluating Bidder Sophistication

The standard theoretical benchmark assumes that all bidders follow their Bayesian Nash equilibrium strategies; this is also the case that my experimental results were tested against. But this benchmark may be invalid: perhaps bidders in the experiment *expect* that their opponents deviate from the equilibria derived by Ausubel and Baranov (2010). According to a 'sophisticated behaviour' hypothesis of Costa-Gomes, Crawford and Broseta (2001), the bidders may be trying to best-respond to the actual play of their opponents, rather than to theoretical predictions. If this is the case, then the fact that Bayesian Nash equilibrium bidding was rejected should be surprising: such a strategy may not be a best response to actual play.

To assess whether sophisticated bidding could explain the divergence from theory, I calculate profits and winning probabilities for all bidder types under the additional scenario where each of the three bidder types unilaterally plays the equilibrium strategy, while the other two bidders play as they did in the experiment. If profits from actual bidding are higher than they would be if that bidder type unilaterally engaged in equilibrium play, then the observed bids may indeed be a best response to actual behaviour of the opponents. The results from this comparison are shown in Table 6.

Table 6: Sophisticated Bidding: Actual bids vs. unilateral deviation to Bayesian Nash bidding

I-types	Vickrey	FirstPrice	VNR	RR(0.50)	RR(0.75)-I1	RR(0.75)-I2
Win%	67.1(55.7)***	47.1(38.6)***	47.9(44.6)**	39.3(34.3)***	52.9(35.7)***	52.9(49.3)
Surplus	31.0(40.5)*	14.3(31.8)***	26.5(31.6)**	21.0(31.3)**	14.9(40.3)***	25.8(29.0)
J-type						
Win%	32.9(26.4)**	52.9(27.9)***	52.1(50.7)	60.7(55.7)	47.1(42.9)	
Surplus	31.0(39.0)	25.0(73.3)***	55.0(58.0)	45.0(48.5)	47.0(57.5)	

For surplus, experimental medians reported; ‘sophisticated bidding’ medians in parentheses.
Sign-test used for testing the win% variable, median-based permutationtest used for surplus.
Rejections of zero-difference null at 90%/95%/99% level indicated by */**/***.

For individual bidders, the winning probability and conditional profit variables reject the zero-difference null in all cases except for the I2-type in the RR(0.75) auction. In all these cases, the unilateral deviation towards equilibrium bidding would lead to a (slightly, but significantly) lower winning probability, but a much higher surplus conditional on winning.[42] Since in Table 5 the I2-type's bidding in RR(0.75) was not significantly different from theory, it is unsurprising that a unilateral deviation towards theory does not lead to higher conditional profit. The results suggest, however, that the vast majority of I-type bidders are not engaging in sophisticated bidding.

The results for the J-type are more varied. In the first-price auction a unilateral deviation is profitable for the J-type for the same reason as it is for the I-types: the payment conditional on winning is then much lower. A similar deviation does not significantly improve profits in any of the other auctions, nor does it much affect winning probabilities in VNR and Reference Rule. In these auctions, the I-types' bids influence their payments in addition to the winning probability, but since J's payment depends only on I-types' bids, the foremost effect of equilibrium bidding is to reduce the probability of winning. The only way in which such a change in strategy would increase the profit, conditional on winning, is by excluding some of the cases where J-type wins after overbidding (and making a negative profit). Table 6 does reflect that this effect is present, since benefits from deviation towards theory are positive, but not sufficiently to be significant.

Since the sophistication hypothesis is rejected in seven of eleven sub-cases, it does not offer a plausible explanation for bidders' deviation from the theory. Following the

[42] If instead of ‘surplus conditional on winning’ I used ‘unconditional surplus’ instead, a sign-test on this variable rejects even more strongly. It would also reject in the additional case of the I2 bidder in RR(0.75).

Bayesian Nash equilibrium bidding functions would leave each bidder type no worse off, even if their opponents did not follow suit.

6.3. Collusion in the Vickrey Auction

In Section 5.1 I found that without bidding restrictions, the I-types bid significantly more aggressively in the Vickrey auction, which is consistent with the collusive hypothesis HC. The next step is to evaluate whether such a bidding pattern can plausibly be attributed to collusion, or whether other explanations are more plausible.

The most direct method for checking whether collusion is present is to look for instances of perfect collusion, where both I1 and I2 bid 200. This criterion is very stringent and of limited use if mis-coordination is frequent. Perfect collusion occurs in only 5 out of 140 rounds of play; in these 5 instances, the joint profit of the I-types is 110, while the average for the whole sample is 54.

To move beyond checking for perfect collusion, another plausible benchmark is needed. Looking for overbidding in excess of value alone is insufficient, because such bidding is frequently found even in simpler single-item auctions where no collusive motive is present.[43] Furthermore, modest overbidding is sometimes attributed to a 'desire to win' effect: if bidders enjoy the phenomenon of winning in itself, they will bid more aggressively, even if this reduces their profit.[44] The significance of this effect is higher in rules where the influence of the bidder's own bid on their price is lower: the increased likelihood of winning looks evident, while the payoff-consequences look less obvious.

The experimental setup allows me to construct a benchmark that measures this 'desire to win' effect, and use that to deflate the data from the Vickrey auction. The I-type payments in VNR and RR(0.75) auctions are designed so as to mitigate the effect of own bids on the payment. While this isolation is not perfect, it does nonetheless provide the bidders with an opportunity to bid more aggressively without expecting large payoff-consequences. Looking at the differences in shading in these two auctions with, and without, bidding restrictions allows me to construct a proxy for the 'desire to win' effect. I use this measure as my non-collusive benchmark: if the change in bidder behaviour in the Vickrey auction is significantly greater than the benchmark, then collusion may be present.

To gauge the extent of the collusion attempts, I will use the amount of overbidding (in excess of the benchmark) and the frequency with which such bids are submitted. If I only observe moderate and occasional overbidding, collusion is not very plausible, and such deviations could easily be attributed to miscalculation. If a significant portion of the data feature overbidding by a considerable amount, it is unlikely that such behaviour is purely accidental.

From Table 3, the largest median difference between restricted and unrestricted bidding treatments occurs in the Reference Rule for the I1-type, and the difference is -2. As expected, when bidding restrictions are lifted, this bidder type bids more aggressively.[45] I then use a sign-test to check whether the shading by I-types in the Vickrey auction is more than this amount; the test rejects with a one-sided p-value ≈ 0.006. The amount of overbidding in the Vickrey auction exceeds the 'desire to win' benchmark, and triggers suspicions of collusion. A summary of the numbers of overbidding I-types, as well as their median surplus, is show in Table 7, below.

[43] In second-price auctions, overbidding is found by Kagel and Levin (1995) and more recently Cooper and Fang (2008).

[44] For an overview, see Kagel and Levin (1995).

[45] Note that this is the median decrease in shading, and though the median amount of shading is still positive, 25% of the bids of this bidder type involve overbidding relative to value.

Table 7: Numbers of Overbidding I-types

Overbid by more than:	Vickrey	First-price	VNR	RR(0.75)
0	166 (15.8)	7 (-6.4)	67 (12.5)	77 (4.3)
5	151 (13.7)	5 (-8.8)	52 (7.8)	59 (2.3)
10	136 (12.5)	4 (-11)	34 (2.3)	42 (-1.1)
20	116 (9.8)	1 (-26)	19 (-6.1)	23 (-8.5)
30	101 (6.7)	0 (NA)	12 (-15.0)	16 (-21.5)
50	79 (3.7)	0 (NA)	5 (-32.4)	6 (-53.7)
75	55 (-0.1)	0 (NA)	3 (-61.3)	5 (-67.2)

Mean surplus in brackets. Total number of I-type bids is 280 under all rules.

The number of overbidding I-types is much higher in the Vickrey auction at all overbidding levels than in any other auction. As the expected value of an I-type bidder is 50, overbidding by 30 is already 60% above the expected value, and over 40% of bids are in this group. Furthermore, almost 20% of all submitted bids are 75 points or above value; this magnitude of overbidding is unlikely to be accidental, especially given how rarely similar deviations occur in the other auctions.

Bidders in the Vickrey auction still make more profit than they would by behaving similarly in any of the other auctions. By overbidding as much as 50 points, the I-types in the Vickrey auction still make a positive surplus (with a mean of 3.7), whereas in other auction types by this point the surplus is negative. Since overbidding is both most prevalent and most profitable in the Vickrey auction, it is likely that this pattern can be attributed to attempts of collusive bidding.[46]

Despite its prevalence, overbidding is not overall profitable for the bidders involved. The rejection of the 'sophisticated bidding' hypothesis showed that I-types in the Vickrey auction would do better by unilaterally deviating towards truthful bidding. The data describes a pattern where I-type bidders attempt to collude, despite frequent mis-coordination. As a result, the Vickrey auction underperforms doubly: even though in Section 5 it gave low revenue to the seller, at the individual level this has not translated into higher bidder surplus. Both the seller, and the bidders, end up significantly worse off than theory predicts.

7. Discussion

Table 8 summarizes the outcome of the hypotheses that I have tested in this paper. At the auction level, the theory-based hypothesis HR on revenue was rejected due to the superior performance of the first-price auction, and the equally poor outcomes form the Vickrey auction. The data did not support the hypothesis of full efficiency in the Vickrey auction either: instead, this auction ranked as least efficient. No significant differences among the other rules emerged, so overall hypothesis HE was also rejected.

The acceptance of hypothesis HB shows that bidders were broadly responding to auction incentives in the ways we would intuitively expect. However, the data rejects more precise hypotheses on bidding behaviour. For the first-price auction, this finding is similar to results on overbidding in single-unit contexts. In the core-selecting auctions, the VNR and Reference Rule, the picture is more complex. Participants with low values did not

[46] The findings of Table 7 would not significantly change if I looked at the amoung of 'bidding in excess of equilibrium prediction' rather than looking at overbidding relative to true values.

Table 8: Outome of the hypothesis tests

Hypothesis	Outcome
HR: The revenue ranking is Vickrey>First-price>VNR≈RR(0.50)	Rejected
HE: The efficiency ranking is the same as in HR	Rejected
HB: Bidding is most aggressive in the Vickrey auction, least in first-price	Accepted
HT: Bidders follow competitive equilibriums strategies	Rejected
HA: I-types bid truthfully in VNR and Reference Rule	Rejected
'Sophistication hypothesis'	Rejected
HJ: J-types bid similarly in all auctions except first-price	Rejected
HS: Bidding constraints have no effect in first-price, VNR and RR	Accepted
HC: Bidding behaviour in Vickrey Auction is consistent with collusion	Accepted

submit zero bids often enough, and more generally all types bid more aggressively than predicted. This leads to the rejection of hypothesis HT. Furthermore, the participants did not bid truthfully in any of the core-selecting auctions, whereby I rejected hypothesis HA. Neither theory, nor 'rule of thumb' behaviour is a good explanation of bidding in my experiment.

The rejection of the 'sophistication hypothesis' showed that unilateral deviations towards equilibrium bidding would be profitable for I-type bidders in five out of six cases, which suggests that participants were also not best-responding to each other's actual bidding behaviour either. The current experimental design does not offer a fuller explanation as to what is the cause of the emerging pattern. Future work in this area will look at the influence of expectations to evaluate whether the divergence from theory is due to incorrect expectation formation, or sub-optimal bidding in response to correct expectations.

The behaviour of I-type bidders in the Vickrey auction is consistent with attempted collusion, even if full collusion rarely manifests. In all other auctions the presence of bidding constraints had no impact, as shown by the acceptance of hypothesis HS, while in the Vickrey auction extensive overbidding was observed when constraints were removed. The extent of the overbidding was far above what I could attribute to a 'desire to win' effect, and the number of extremely high bids was much higher than in any of the other auctions.

A natural interpretation of finding collusion in the setting of my paper is to relate it to practical one-shot auctions, in contrast to the collusion literature which looks at repeated play. An example of this would be a one-off sale of government assets with a pure efficiency objective, and no concern for revenue. My results suggest that even if revenue in itself is unimportant, the potential for attempts at collusive bidding in a Vickrey auction is high, and that would be sufficient to undermine its efficiency properties. Conversely, a policy with a pure efficiency objective may be counterproductive.

8. Conclusions

My main finding is the surprisingly good performance of the first-price auction: it generated most revenue, without any corresponding efficiency loss. Conversely, the performance of the Vickrey auction was unexpectedly poor: contrary to the expectation of 100% efficiency, it ranked last on this criterion. Given that efficiency concerns are frequently used to argue against the use of first-price mechanisms in high value auctions, these experimental results provide evidence to allay such concerns. The core-selecting auctions tied with the first-price auction on efficiency, and were revenue-equivalent with the Vickrey auction; they were not "the best of both worlds", but also never ranked last, contrary to theoretical predictions.

At the individual level, I found that actual bidding diverged significantly from Bayesian Nash equilibrium predictions. Bidders frequently bid in excess of the theoretical benchmark, and occasionally even above their own valuation. This behaviour could not be attributed to sophistication, as the observed bids never resulted in higher profits compared to a unilateral deviation towards equilibrium bidding. In the core-selecting auctions, bidders also did not resort to using a truth-telling rule-of-thumb: I found no evidence for the intuition that payments close-to-independent of own bids would induce close-to-truthful bidding. The behaviour I observed in the Vickrey auction was consistent with attempts at playing collusively, even though such attempts were rarely successful. The Vickrey auction generated neither high revenue, nor high bidder surplus. Even if the auction context is complex, my results propose, simple and easy to understand rules may still perform better and more predictably than complicated alternatives, whose actual performance does not match their theoretical predictions.

9. Appendix A: The Variable-α Experiment

In the proofs that Erdil and Klemperer (2010) use to analyse the incentive properties of the Reference Rule, the reference point itself does not change the relevant deviation incentives on aggregate. However, it can significantly affect the relative amount that each bidder will have to pay, conditional on winning, and this may have non-trivial behavioural implications. Numerical calculations have shown that as α changes so do the optimal bidding functions, resulting in extremely disparate optimal bidding functions for the two types as α tends to either 0 or 1.[47] This additional experiment set out to examine whether such variation would also emerge in the laboratory.

When investigating asymmetries, it is useful to introduce some additional notation. Let K denote the upper end of the support of the value distribution of the I1-type. Then asymmetries in the valuations of the two I-types can be modelled as follows: set $v_{i1} \sim U\left[0, K\right]$ and $v_{i2} \sim U\left[0, 200 - K\right]$. This keeps the sum of supports (and hence the expected total value) of the two I-type bidders the same as that of the J-type bidder, but when $K \neq 100$ the I-types are no longer symmetric. The nature of asymmetry in my experiment can then be summarised by two parameters: α and K. I consider four cases:

- Setting 1: $\alpha = 0.50$ and K=100 (i.e. $v_{i1}, v_{i2} \sim U\left[0, 100\right]$)
- Setting 2: $\alpha = 0.75$ and K=150 (i.e. $v_{i1}, \sim U\left[0, 150\right], v_{i2} \sim U\left[0, 50\right]$)
- Setting 3: $\alpha = 0.75$ and K=100 (i.e. $v_{i1}, v_{i2} \sim U\left[0, 100\right]$)
- Setting 4: $\alpha = 0.50$ and K=150 (i.e. $v_{i1}, \sim U\left[0, 150\right], v_{i2} \sim U\left[0, 50\right]$)

This particular combination of α and K allows me to investigate two main issues. Firstly, I can check whether it is the asymmetry of the α parameter itself that influences behaviour; for this comparison, I look at the cases where the support of the two bidders' valuations stays constant, and the α varies.[48] Secondly, I can assess whether it is the magnitude of α *relative to* the 'expected valuation' of the bidders that matters; here I compare the cases where the ratio of $\frac{E(v_{i1})}{E(v_{i2})} = \frac{\alpha}{1-\alpha}$, to those where it is not.[49]

[47] In the limit, as $\alpha \to 0$ or $\alpha \to 1$ an analytical solution is possible. The solution entails the I-type bidder with the infinitesmal 'reference share' bidding truthfully, while the other I-type shades by a large amount.

[48] Here the relevant comparisons are: Setting 1 v.s. Setting 3 , and Setting 2 v.s. Setting 4.

[49] Here the relevant comparisons are: Setting 1 v.s. Setting 4 , and Setting 2 v.s. Setting 3.

The experimental setup of these session was analogous to the main experiment in this paper, with the exception that here only one set of instructions was given out at the beginning of the experiment. These instructions outlined how variations in the α parameter influenced reference payments in the Reference Rule.[50] The participants were allowed to ask questions whereafter they proceeded to complete an understanding test.[51] Upon successful completion of the test, the participants were informed which α parameter and which valuation model would apply in the given section of the experiment. They subsequently played two practice rounds, followed by ten payment-relevant rounds in each setting.[52] The duration of the sessions in the Alpha-experiments was two hours on average, generating mean earnings of £27 (~\$43).

9.1. Results of the Variable-α Experiment

Comparing bidder-level results in the asymmetry experiment poses complications that are not present in the main experiment. Direct tests of bidding variables cannot be conducted across settings where K varies, because these tests will reject by default due to the bidding support being different across the compared cases.

This problem does not arise, however, when performing tests while holding K fixed. When I test for the effects of varying α only, holding K fixed, none of the four test-pairing for the I-type bidders reject a zero-difference null even at the 90% level. Hence α on its own does not significantly influence individual bidding.

An alternative to using direct bid data is to look at bid ratios,[53] but this approach will artificially inflate differences in the cases where $K \neq 100$. Here the two I-types have a different value support, and the I2-bidder with a narrower support is more likely to exhibit large variation in the bid ratios. The tests are hence likely to over-reject a zero-difference null, though using non-parametric tests reduces the likelihood of this mistake. With that caveat in mind, I ran battery of median-difference tests for both I-types on their bid ratios, and still found only one statistically significant difference. The I2-type's bid-ratios in Setting 4 ($\alpha = 0.50,\ K = 150$) test as significantly lower than in all other cases. This is an intuitive finding, as in this case the I2-type can be seen to be in a particularly weak position: they have a bidding support of only [0,50], but their 'preliminary share' of the payments is a disproportionately higher 50%. As a result, in this setting the I2 type bids more cautiously. No other ranking emerges from the pairwise tests.

A final hypothesis that I test on the individual bidder data is to check whether setting the α proportionately to the ratio of expected values of the two I-types affects bidding. It is, for example, possible that bidders would have a preference for equality or some notion of fairness, as found by Battalio, Van Huyck and Gillette (1992) in the context of two-person coordination games. To test for this effect, I pooled the data from settings 1 and 2, where α is set 'proportionately', and tested it against the pooled data from settings 3 and 4. Median-difference tests for both I1's and I2's bidding ratios failed to reject the zero-difference null (p-values>0.22 in both cases). Thus I could not find any influence of proportionality on bidding at the individual level.

From the J-types' perspective, all four settings are identical, so we should expect them to bid similarly in all four cases. A Kruskal-Wallis test for this hypothesis marginally rejects with a p-value=0.046, indicating that the J-types do not bid the same way across the four settings. In pairwise tests for bidding and shading, various individual pairings reject, but no coherent pattern emerges. It appears that the J-type bidders are trying

[50] The instructions are available from the author on request.

[51] The rate of failures was three out of 45 participants in this phase of the experiment.

[52] The order of the Cases in the experimental sessions was from 1 to 4 in the first session. The ordering was reversed for the other session.

[53] These are calculated as the ratios of bid relative to the value of the bidder.

to best respond differently to the I-types' actual bidding across the different settings, ignoring the prediction that truthful bidding should be optimal every time.

At the auction level, the main variables of interest are again revenue, surplus and efficiency. A summary of these parameters across the four settings is shown in Table 9. Setting 1 immediately stands out: revenue is almost 10 points higher than in the other three settings, while surplus is lower by a similar amount. Efficiency is high in all four settings, and the differences are small.

Table 9: Revenue, Surplus and Efficiency Summary from alpha experiment

	K=100|α=0.50	K=150|α=0.75	K=100|α=0.75	K=150|α=0.50
revenue	77.0 (42.3)	65.5 (41.0)	62.6 (38.4)	64.2 (40.9)
surplus	48.9 (49.3)	61.1 (51.4)	58.2 (44.1)	63.8 (49.1)
efficiency	94.9 (13.8)	95.3 (15.0)	96.9 (12.0)	96.0 (15.1)

Means reported, standard deviations below.

A series of pairwise median-difference tests for revenue is summarised in Table 10. The results hence confirm that the symmetric setting with K=100 and $\alpha = 0.50$ is revenue-superior to the other three cases, with the tests rejecting the zero-difference null with 90% confidence or stricter. No significant revenue differences emerge amongst the other pairings. Correspondingly, Setting 1 also yields significantly lower surplus than Setting 4 (p-value=0.009). Finally, a Mann-Whitney test for differences in efficiency fails to reject between Settings 1 and 2, but it does reject the zero-difference null between Setting 1 and Settings 3 and 4 with p-value=0.015 and p-value=0.002; after applying the Bonferroni-Holm corrections, these rejections remain significant at the 90% and 95% levels, respectively. This implies that Setting 1 is less efficient, and no other pairings yield a rejection of the zero-difference null. Using the RR(0.50), or the Proxy Rule, in a symmetric setting yields superior revenue, but lower efficiency.

Table 10: Pairwise Revenue-difference Tests for variable-α experiment

	K=150|α=0.75	K=100|α=0.75	K=150|α=50
K=100|α=0.50	12.5*	14.0**	13.0*
K=150|α=0.75		2.0	0.0
K=100|α=0.75			−1.0

Reported values are for median-difference of (row - column).
Rejections of zero-difference null at 90%/95%/99% level
indicated by */**/***; Bonferroni-Holm corrections applied.

The final test of interest at the auction level checks whether revenue and efficiency are sensitive to setting the α proportionately to the bidders' expected values. This comparison is particularly significant for its policy implications relating to the relevance of 'reference points'. If the proportional cases where $\frac{E(v_{i1})}{E(v_{i2})} = \frac{\alpha}{1-\alpha}$ perform significantly better, this would be supporting evidence in favour of the flexibility inherent in the Reference Rule. A median-difference test for revenue rejects with a p-value=0.037; the median-difference is 7 points in favour of the proportional settings. A corresponding Mann-Whitney test for efficiency rejects with a p-value<0.001, and indicates that there is only a 46% chance of drawing a higher efficiency value from the 'proportional' sub-sample. In practice the differences in efficiency are low - on average around 1.3 points - so the statistical significance here may not have much economic importance. This pair of findings gives some support to the view that selecting a reference point appropriately in relation to the relative values of the assets for sale may yield superior revenue results.[54]

Overall, the findings of the sessions on asymmetries did not offer many conclusive answers as to the influence of α. While I found some significant auction-level results in favour of setting α appropriately, the bidder-level data showed little sensitivity to α.

REFERENCES

Ausubel, L.M. & Baranov, O.V., 'Core-Selecting Auctions with Incomplete Information', *University of Maryland Working Paper* (2010).

Ausubel, L.M. & Cramton, P., 'Auctioning many divisible goods', *Journal of the European Economic Association*, vol. 2, no. 2-3, 480–493 (2004).

Ausubel, L.M., Cramton, P. & Milgrom, P., 'The clock-proxy auction: A practical combinatorial auction design', *Combinatorial Auctions*, pp. 115–138 (2006).

Baranov, O., 'Exposure vs. Free-Riding in Auctions with Incomplete Information', *Working Paper* (2010).

Cantillon, E. & Pesendorfer, M., 'Auctioning bus routes: the London experience', in Ausubel L., Cramton P. & R., Steinberg eds., *Combinatorial Auctions* (Cambridge, Mass., USA: MIT Press, 2006).

Cooper, D.J. & Fang, H., 'Understanding overbidding in second price auctions: An experimental study', *The Economic Journal*, vol. 118, no. 532, 1572–1595 (2008).

Costa-Gomes, M., Crawford, V.P. & Broseta, B., 'Cognition and Behavior in Normal-Form Games: An Experimental Study', *Econometrica*, vol. 69, no. 5, 1193–1235 (2001).

Day, R. & Milgrom, P., 'Core-selecting package auctions', *International Journal of Game Theory*, vol. 36, no. 3, 393–407 (2008).

Day, Robert W & Cramton, Peter, 'Quadratic core-selecting payment rules for combinatorial auctions', *Operations Research*, vol. 60, no. 3, 588–603 (2012).

Erdil, A. & Klemperer, P., 'A new payment rule for core-selecting package auctions', *Journal of the European Economic Association*, vol. 8, no. 2-3, 537–547 (2010).

Fischbacher, U., 'z-Tree: Zurich toolbox for ready-made economic experiments', *Experimental Economics*, vol. 10, no. 2, 171–178 (2007).

Goeree, J.K. & Lien, Y., 'On the impossibility of core-selecting auctions', *Institute for Empirical Research in Economics University of Zurich Working Paper No.* (2009).

Good, P., *Permutation Tests: a practical guide to resampling methods for testing hypotheses* (Springer, New York, 1994).

Goswami, G., Noe, T.H. & Rebello, M.J., 'Collusion in uniform-price auctions: experimental evidence and implication for treasury auctions', *Review of Financial Studies*, vol. 9, no. 3, 757 (1996).

Kagel, J.H., 'Auctions: A Survey of Experimental Research', *The Handbook of Experimental Economics*, pp. 501–585 (1995).

[54] This conclusion, however, is heavily influenced by the revenue-superior performance of the $RR\,(0.50)$ auction in the symmetric case.

Kagel, J.H. & Levin, D., 'Auctions: A survey of experimental research, 1995-2008', *Handbook of Experimental Economics Vol 2.*, vol. 2 (2008).

Kagel, J.H., Lien, Y. & Milgrom, P., 'Ascending Prices and Package Bidding: A Theoretical and Experimental Analysis', *American Economic Journal: Microeconomics*, vol. 2, no. 3, 160–185 (2010).

Kazumori, E., 'Auctions with package bidding: An experimental study', *Working Paper* (2010).

Kazumori, E., 'Vickrey Auctions Versus Proxy Auctions: An Experimental Study', *working paper* (2014).

Kwasnica, A.M. & Sherstyuk, K., 'Collusion and Equilibrium Selection in Auctions', *The Economic Journal*, vol. 117, no. 516, 120–145 (2007).

Newson, Roger, 'Confidence intervals for rank statistics: Somers' D and extensions', *Stata Journal*, vol. 6, no. 3, 309 (2006).

Palfrey, T.R. & Prisbrey, J.E., 'Anomalous behavior in public goods experiments: How much and why?' *The American Economic Review*, vol. 87, no. 5, 829–846 (1997).

Sano, R., 'An Equilibrium Analysis of a Package Auction with Single-Minded Bidders', *Working Paper* (2010).

Sherstyuk, K., 'Collusion in private value ascending price auctions', *Journal of economic behavior & organization*, vol. 48, no. 2, 177–195 (2002).

Sun, N. & Yang, Z., 'Equilibria and indivisibilities: gross substitutes and complements', *Econometrica*, vol. 74, no. 5, 1385–1402 (2006).

Sun, N. & Yang, Z., 'A Double-Track Adjustment Process for Discrete Markets With Substitutes and Complements', *Econometrica*, vol. 77, no. 3, 933–952 (2009).

ComReg, 'Multi-band Spectrum Release: Information Memorandum', **URL:** `http://www.comreg.ie/_fileupload/publications/ComReg1252.pdf` (2012).

Danish Business Authority, '800 MHz Auction: Information Memorandum', **URL:** `http://erhvervsstyrelsen.dk/file/251159/information-memorandum-800mhz-auction.pdf` (2012).

Federal Aviation Administration, 'Congestion Management Rule for LaGuardia Airport; Final Rule', *Federal Register*, vol. 73, no. 198, 60573–60601 (2008).

Norwegian Post and Telecommunications Authority, '800, 900 and 1800 MHz auction - Auction Rules', **URL:** `http://eng.npt.no/portal/page/portal/PG_NPT_NO_EN/PAG_NPT_EN_HOME/PAG_RESOURCE_TEXT?p_d_i=-121&p_d_c=&p_d_v=128156` (2013).

Ofcom, 'The award of 800 MHz and 2.6 GHz spectrum Information Memorandum', **URL:** `http://stakeholders.ofcom.org.uk/binaries/consultations/award-800mhz/statement/IM.pdf` (2012).

Teytelboym, A., 'Essays on Networks and Market Design', *D.Phil Thesis, University of Oxford* (2013).