Available online at www.sciencedirect.com
SCIENCE DIRECT®
International Journal of Industrial Organization
24 (2006) 555–581
International Journal of Industrial Organization
www.elsevier.com/locate/econbase

Using first-price auctions to sell heterogeneous licenses

Jacob K. Goeree [a], Theo Offerman [b], Arthur Schram [b,*]

[a] *Division of the Humanities and Social Sciences, Mail code 228-77, California Institute of Technology, Pasadena, CA 91125, USA*

[b] *CREED, Faculteit der Economische Wetenschappen en Econometrie, Universiteit van Amsterdam, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands*

Received 16 January 2004; received in revised form 4 February 2005; accepted 28 July 2005
Available online 20 October 2005

Abstract

This paper considers three alternative ways to sell heterogeneous licenses via a first-price format when there is single unit demand. It has been suggested that incorporating a first-price element may bolster competition in this case [Klemperer, Paul D., 2002. What really matters in auction design. Journal of Economic Perspectives, 16, 169–189]. In a controlled laboratory setting, we compare the performance of the simultaneous first-price auction, the sequential first-price auction and the simultaneous descending auction with that of the simultaneous ascending auction. The experiments involve several bidding environments of varying complexity. We find that the simultaneous ascending auction achieves the highest levels of efficiency but also has drawbacks: (i) its revenues are low and variable, (ii) per-license profits vary, and (iii) the incidence of winner's curse outcomes is high. Seller's revenues are highest when the licenses are sold in a sequential first-price auction, in decreasing order of quality.

JEL classification: C72; D44

Keywords: Multiple-unit auctions; Experiments; Winner's curse; Entry

1. Introduction

Economists have long advocated the use of auctions to assign scarce resources. Ronald Coase (1959) first proposed auctioning the spectrum and William Vickrey later pushed auctions for many years. Spectrum auctions became a reality in the US in 1994 and since then

* Corresponding author.
E-mail address: schram@uva.nl (A. Schram).

0167-7187/$ - see front matter
doi:10.1016/j.ijindorg.2005.07.011

thousands of licenses have been assigned in this way, raising over $40 billion in the process (Cramton, 2002). In all but two instances, licenses were assigned using the simultaneous ascending auction, which is the natural extension of the English auction to multiple related goods. More recently, the US spectrum auctions have been successfully copied throughout the world. In Europe, for example, many countries employed the simultaneous ascending auction to award third-generation mobile phone (UMTS) licenses, raising a staggering amount of nearly $100 billion.

Many governments now consider auctions as an allocation device for far more than just the mobile phone spectrum. The Dutch government, for example, has auctioned off spectrum for both DCS1800 and UMTS networks but is also considering (or has considered) using auctions for, i.a., the allocation of slots for highway gas stations, wireless local loop spectrum and telephone numbers. In 2002 it was planning a 'zero-base' auction to assign FM-radio licenses for commercial applications. In fact, the same bureau within the Dutch government that organized the UMTS auction was assigned responsibility for planning the zero-base auction.[1] In doing so, it realizes that it cannot simply copy the previous auction design. Every auction should be tailored to the specific context in which it will take place (e.g., Klemperer, 2002). In designing the zero-base auction, the Dutch government hoped to learn from the experiences in previous auctions in the Netherlands, and elsewhere.

One important lesson is related to the auction revenue. In a few cases, like the Netherlands, revenues of the UMTS auctions were disappointing. This led to speculation that the simultaneous ascending auction is ineffective in raising revenue in 'uncompetitive' situations, in particular when the number of licenses equals the number of incumbents who have clear advantages over a few potential entrants (Jehiel and Moldovanu, 2003; Klemperer, 2002). The intuition is that strong bidders can simply 'trail' weaker ones and keep overbidding them by the minimum bid increment until they drop out. As a result, weak bidders anticipate they have no chance of winning and quit the auction early (or do not enter at all), which adversely affects revenues.

Paul Klemperer (2002) first pointed out that in uncompetitive situations, auctions with a first-price element may be more attractive to entrants since they offer weak bidders a positive chance of winning (in equilibrium). Stated differently, first-price formats contain an element of surprise as strong bidders may forego a license when they shade their bids too much. This uncertainty forces strong bidders to bid more competitively, with a positive effect on revenues. Klemperer (2002) suggests that the hybrid 'Anglo-Dutch' auction might have worked better for the sale of third-generation mobile phone licenses in the Netherlands. For the case of K *homogeneous* licenses, this auction works as follows. In the first stage, the license price rises until $K+1$ active bidders are left. In the second stage, the $K+1$ remaining bidders make a single sealed bid no less than the final price level of the first stage. The K highest bidders in the second stage each obtain a license and pay their own (sealed) bids.

In the zero-base auction, incumbent firms would again be advantaged. As a consequence, only a few entrants were expected to participate. Hence, DGTP was considering the possibility of adding a first-price element to the auction design. There is a complicating factor that makes the introduction of a first-price element a non-trivial problem, however. As in the UMTS

[1] This is the 'Directoraat-Generaal Telecomminicatie en Post' (DGTP). DGTP is currently part of the Dutch Ministry of Economic Affairs. In 2003, the Ministry responded to strong pressure from existing stations and dropped the idea of an auction. Instead, it decided to use an administrative process to assign stations.

auction, the licenses in the zero-base auction are heterogeneous in 'quality'. The FM-licenses differ substantially in their demographic coverage. This heterogeneity makes it unclear how the *first-price* element should be added to the auction format.[2] In particular, it is not obvious how to adapt the rules of the Anglo-Dutch auction to the *heterogeneous* case. One issue is whether bidders are allowed to 'switch' licenses between the first and second stage (and, if so, how the minimum prices in the second stage are determined).

More generally, little is known about the effect of 'pure' first-price elements in uncompetitive situations with heterogeneous goods. We address this issue in this paper. Though the interest in multi-unit auctions of heterogeneous goods has grown tremendously over the past few years (Klemperer, 2002), first-price formats have largely been neglected in this literature. An important reason is that first-price auctions with heterogeneous objects are intractable from a theoretical point of view. Therefore, we decided to employ controlled laboratory experiments to compare the performance of the first-price formats vis-à-vis each other and the simultaneous ascending auction.

In order to single out the effects of using a first-price format in these situations, we decided not to investigate the two-stage Anglo-Dutch auction. Instead, we study various first-price formats per se and compare these to the standard simultaneous ascending auction. This research project was undertaken at the request of DGTP. At this stage, the government was more interested in gaining general insights with respect to the issue described than to obtain specific advice on the design of the auction at hand. Consequently, as agreed by DGTP, we decided to ignore specific details of the planned zero-base auction because they are not thought to be essential for a test of the robustness of the proposed first-price formats. The insights to be obtained can then be used in the design of future auctions in which heterogeneous goods are allocated in uncompetitive environments.

The remainder of this paper is organized as follows. In Section 2 we give an overview of the auction formats considered. Section 3 describes the experimental design. In Section 4 we present the experimental results. Section 5 concludes.

2. Auction formats considered

We will compare three first-price formats to each other and to the more traditional (second price) simultaneous ascending auction. The first design was suggested by DGTP. It is the 'simultaneous first-price' auction where all bidders simultaneously submit bids for every license for sale. The licenses are then assigned such that the sum of the winning bids (revenue) is maximized and each bidder gets at most one license. Bidders who win a license pay a price equal to their own bid. Although this design resembles the standard discriminative auction for homogeneous licenses, it is not at all clear how one should bid in this auction. Nor is it clear that the manner in which a first-price element is incorporated into the design is optimal.

An alternative format is the sequential first-price auction, where licenses are put up for sale one by one. When heterogeneous licenses are sold via a sequential first-price auction, a decision has to be made regarding the order of sale. If the best licenses are sold first, the most intense competition is directed at the most valuable licenses, which could enhance revenue. On the other

[2] In contrast, it is straightforward to design an auction with a first-price element when licenses are homogeneous. Examples are the Anglo-Dutch auction described above and the discriminative auction (Harris and Raviv, 1981; Cox et al., 1984).

hand, with a reverse order, information about common value aspects of the licenses may be revealed before the best licenses are sold, which also stimulates competitive bidding.[3] In the experiment, we vary the order in which products are offered for sale to study such order effects.

The final first-price format we study is the ‘simultaneous descending auction’, where license prices are determined by a number of clocks (one for each license). All clocks start at the same high price level and fall at equal speeds. Once a bidder stops a clock, she buys the particular license at the price where she stopped the clock. The clocks of the remaining licenses continue to fall until they are stopped by remaining bidders. This auction format was first suggested by Bulow and Klemperer (1994), but has not been studied in detail.[4]

The simultaneous descending auction resembles a sequential auction where the best license is sold first. One difference, however, is that in the simultaneous descending auction, bidders possess more information than in a sequential auction. When the clocks have reached a certain price level, bidders know the remaining licenses will sell for less than that price. In contrast, in a sequential auction with best licenses sold first, bidders can never rule out that licenses will be sold at higher prices later in the auction. When best licenses are sold last in the sequential auction, information is not perfect either: bidders cannot rule out a lower price for later licenses. Since more information warrants more aggressive bidding, this suggests that the simultaneous descending auction may generate higher revenues than the sequential auction. On the other hand, the simultaneous descending auction may facilitate collusion, especially in asymmetric private value settings. If one bidder is known to have a very high value for a particular license, others may be inclined to wait until this bidder stops the clock of the preferred license. If the high-value bidder anticipates others’ behavior, she may decide to let the price fall to very low levels before stopping the clock, which would result in low revenues.

Finally, we compare the performance of the three first-price formats with that of the simultaneous ascending auction. The popularity of the latter format is due to the flexibility and transparency it provides to bidders who can choose to bid on a license knowing its current price and those of other licenses. In addition, for the case where each bidder can acquire at most one license (and the private values paradigm applies), theory predicts that the simultaneous ascending yields an efficient allocation (Demange et al., 1986). An alternative possibility is to employ a Vickrey auction, in which the efficient outcome is also supported in equilibrium. Leonard (1983) proposed a simple mechanism to implement the Vickrey auction. In an experiment, Olson and Porter (1994) compare the performance of the simultaneous ascending auction to this Vickrey–Leonard format. The simultaneous ascending auction resulted in higher revenues and efficiency than Vickrey–Leonard, although the difference in efficiency was not significant.

Recent literature, however, has also identified some drawbacks of the simultaneous ascending auction. First, it facilitates collusion as bidders have the opportunity to retaliate when others deviate from a collusive agreement. Such retaliatory strategies are not possible in

[3] For example, when the final license is put up for sale, the identities of future competitors are known, which makes it is easier to estimate the value of a license.

[4] A design that is closely related to the simultaneous descending auction is the so-called ‘right-to-choose auction’. In this auction, the highest bidder obtains the right to choose any one of the objects. Then, the right to choose from the remaining objects is auctioned sequentially until all objects are allocated. Notice the resemblance to the simultaneous descending auction. The latter implies a decreasing price from one choice to the next, however. Though the right-to-choose auction certainly deserves investigation, it is not included in our set-up because of its similarity to the simultaneous descending auction and because we wanted to limit the number of alternatives studied.

the first-price formats we consider. Second, (as discussed above) the simultaneous ascending auction may discourage entry when incumbents have ex ante advantages. In this case, entrants may feel that showing up is useless because they will always be outbid in an ascending auction (Klemperer, 2002).[5] In contrast, in first-price auctions, entrants have a strictly positive chance of winning. Third, ex ante asymmetries may lead to low revenues since incumbents can simply trail and overbid an entrant by the minimum amount until the entrant drops out. In auctions with a first-price element, incumbents are forced to bid aggressively as they only get a single chance and may not obtain a license if they take too much risk by bidding too low (Maskin and Riley, 2000).

In our evaluation of the distinct auction formats, we use several criteria. We consider standard properties such as efficiency and revenue, and the auction's capability to prevent collusion. An aspect of practical importance is the variance of the auction's revenue. A format that generates high revenues on average, but shows a lot of variation in outcomes, is not robust. Another important question is "how attractive is the auction for entrants?" Finally, we explore the likelihood that a particular auction format gives rise to subsequent litigation. In particular, if the allocation resulting from the auction is perceived to be unfair, or when a winner of a license bids too high and actually incurs a loss, costly lawsuits become more likely.

In our comparison of laboratory results, we are not guided by theoretical results pertaining to the first-price formats. Hence, the external validity of our experiments requires that we incorporate in our design the most important characteristics of the setting in which the government will run these auctions. Given the reasons depicted above for considering first-price auctions in the first place, we will focus on settings where the number of bidders is relatively low compared to the number of licenses.[6] In addition, we consider several bidding environments in the experiment: (i) we start with a symmetric situation where no a priori differences between bidders exist and where bidders do not face uncertainty regarding their private values for the licenses. (ii) Next, we introduce value asymmetries between incumbents and newcomers. In this setting, there are as many licenses as strong incumbents, and there is one weak entrant. (iii) Then we introduce some common value uncertainty, by adding a random component to bidders valuations that affects all bidders and all licenses the same way.[7] (iv) Finally, we give entrants the choice to participate in the auction or profit from some outside opportunity. In other words, entrants face an opportunity cost when taking part in the auction. This final part of the experiment allows us to determine the attractiveness of different auction formats for entrants.

We test the four auction formats under each of these environments. Each situation is held constant across formats, enabling us to compare the auctions in various settings under ceteris paribus conditions. This paper illustrates a powerful application of the experimental method: it can be used to shed light on problems that are deemed intractable from a theoretical perspective.

---

[5] Another way to model ex ante advantages is by introducing allocative payoff externalities in the incumbents' preferences (Hoppe et al., 2004). The idea is that an incumbent is better off when another incumbent acquires a license than when a newcomer buys a license. In an experiment, Goeree et al. (2004) investigate the effects of allocative externalities on entry, revenue and efficiency.

[6] As pointed out by Klemperer (2002), most formats perform well with a large number of bidders.

[7] For example, prior to the European UMTS spectrum auctions, there was considerable uncertainty about future consumer demand for third-generation mobile phone services. This uncertainty affected the values of the licenses for all bidders in the same manner.

3. Experimental design

We applied a between-subject design with respect to the auction formats and a within-subject design for the bidding environment. Both are described in more detail below. The experiment was completely computerized and consists of four parts. The Dutch instructions are more elaborate than the description offered here and use common and neutral terms. A translation of the instructions for one of the treatments is contained in Appendix A. In the experiment, subjects earn points, which are exchanged into guilders at the end of the experiment (we used a rate of 1 point = 0.25 guilders ≈ 0.11 Euro). At the start of the experiment all subjects receive a starting capital of 50 points, which they did not have to pay back.[8] They play two practice periods to get acquainted with the computer software. Then the experiment of 28 periods starts. In each period, three heterogeneous licenses are sold and bidders can acquire at most one of them. Each of the four treatments employs a different auction mechanism to allocate the three licenses.

3.1. Auction formats

In the *simultaneous first-price auction* (SimFPA), bidders submit at most three bids at the same time, one for each license. The licenses are then allocated such that the sum of the winning bids is maximized under the restriction that at most one license is assigned to each bidder. A winning bidder pays a price equal to her own bid (for the license she is assigned). Bidders can refrain from bidding on a license by not submitting a bid for it.

In the *sequential first-price auction* (SeqFPA), licenses are sold one after another. The order in which the three licenses are sold varies from period to period. When the first license is put up for sale, every bidder can submit a bid. The highest bidder obtains the license at a price equal to her own bid and is excluded from bidding on subsequent licenses. The remaining bidders receive information about the price paid after which they can bid on the next license. This process is repeated until all three licenses are sold. A bidder can indicate that she refrains from bidding on a license by not filling out a bid for it.

In the *simultaneous descending auction* (SDA), license prices start at some common price level and then decrease point by point. If bidders want to buy a license, they can do so by clicking the button for the license. The price clock for the license then stops falling and other bidders no longer have the possibility to select this license. The prices of other licenses continue to decline until they are selected by one of the other bidders. Each buyer pays a price equal to the level at which she selects the particular license. The price of a license cannot fall below zero; if a license is not selected at price of zero, it will not be sold. A period ends when all licenses have been sold or when it is clear that they will not be sold.

In the *simultaneous ascending auction* (SAA), bidders select one of three licenses at the start of a period for an initial price of 0. If at most one bidder selects a particular license then the price of that license does not increase. If two or more bidders select the same license, its price starts to rise point by point. When a bidder decides the price for a license is too high, she can switch to another license or indicate that she does not want to buy any license that period by clicking the button “I do not want to buy”. If a bidder switches to a license that has already been selected by

[8] In the experiment, subjects were rarely bothered by budget constraints. In only 12 out of 3584 cases (0.3%) a subject’s balance went below 25 points (the minimum of 9 points materialized once). All these cases occurred in the first 5 periods. The average budget was 177.07 points (with a standard deviation of 88.75).

Table 1
Design

Part	Bidders		Private values		Common value	Entry
(1) Periods 1–7	4 symmetric		L1 $U[0,50]$		1	No
			L2 $U[0,40]$			
			L3 $U[0,30]$			
			Weak	Strong		
(2) Periods 8–14	3 strong, 1 weak	L1	$U[0,50]$	$U[20,70]$	1	No
		L2	$U[0,40]$	$U[15,55]$		
		L3	$U[0,30]$	$U[10,40]$		
(3) Periods 15–21	3 strong, 1 weak	L1	$U[0,50]$	$U[20,70]$	$U[0.5,1.5]$	No
		L2	$U[0,40]$	$U[15,55]$		
		L3	$U[0,30]$	$U[10,40]$		
(4) Periods 22–28	3 strong, 5 weak	L1	$U[0,50]$	$U[20,70]$	$U[0.5,1.5]$	Yes
		L2	$U[0,40]$	$U[15,55]$		
		L3	$U[0,30]$	$U[10,40]$		

one other bidder, the price for that license starts to rise. When a bidder is the only one who has selected a particular license, she cannot switch to another license nor can she decide not to buy a license. Such a bidder can change only after another bidder selects the same license. Once a bidder drops out of the auction, she cannot re-enter that period. A period ends when each license has been selected by at most one bidder, at which time no further price increases occur. The active bidders then purchase their selected licenses at prices equal to the levels at which the last rival for the license switched.

In all treatments, winning bidders are informed about their profit (or loss) at the end of a period. Then a new period starts and again three licenses are sold.

3.2. Parts of the experiment

The experiment consists of four parts of seven periods each. See Table 1 for details. Only after one part of the experiment has finished do subjects receive the instructions for the next part. When subjects have finished reading the instructions on the computer screen, they are given a written summary of the instructions for that part.[9]

In each of 8 sessions, 16 subjects participated, so that we have data on 4 groups of 8 players in each treatment. In the first three parts, four bidders compete for three licenses. The composition of a group of bidders is kept constant for statistical reasons.[10] In part 1, bidders are symmetric, i.e. their values are drawn from the same distributions. In parts 2 and 3, three strong bidders and one weak bidder compete for the three licenses. This setup captures the situation where strong incumbents fight against a relatively weak newcomer and the number of licenses equals the number of incumbents.

In part 4, two groups of four bidders are combined into one group consisting of three strong and five weak bidders (this group remains the same for the final seven periods of the

[9] At the start of parts 2, 3, and 4, the earnings of subjects are increased by 10 points, which they do not have to pay back after the experiment.

[10] In order to avoid repeated game effects, subjects are not informed about this choice. Subjects did not ask the question whether the composition of a group would remain constant across periods or not. If they would have asked this question, we would have told them that this information would not be revealed in this experiment.

experiment).[11] Prior to the auction, weak bidders sequentially decide whether or not to participate in the auction. To avoid over-entry or under-entry due to miscoordination among weaker players, we let them choose sequentially whether or not to participate in the auction. When making the entry decision, a weak bidder knows her values for the licenses and the common value signal (discussed below). If a weak bidder chooses not to participate, she takes part in a lottery based on a single throw of a (computer animated) six-sided die. Since the auction payoff is uncertain we decided to make the payoff of the outside option uncertain as well, to avoid a possible bias in entry behavior. With probability one-half (outcomes 1, 2, and 3) the lottery pays 0 points and with probability one-half (outcomes 4, 5, and 6) it pays 6 points. The outcomes of die throws are independent across subjects and periods. If a weak bidder decides to participate in the auction, she loses the opportunity to play the lottery. Weak bidders make entry decisions one by one, knowing the entry choices other weak bidders have made before them. The three strong bidders make no entry choice; they have to participate in the auction. Before the auction starts, bidders are told how many weak bidders will participate with the three strong bidders.[12]

The reasons for using only a limited number of seven rounds per environment are twofold. First, our primary focus is not on behavioral convergence, but rather we are interested in creating an environment that is relevant for the policies under consideration. Bidders rarely participate more than a few times in this kind of government-run auctions, and, hence, we believe that behavior in early rounds is most relevant for our purposes. Second, even if we were interested in behavioral convergence, we would likely need to run at least 60–80 repetitions within the same environment (based on the ‘individual information’ treatment in Armantier, 2004). This is impossible in practice due to the amount of time per round that our environments require.[13]

3.3. Bidders' value and information

In each period, bidders are assigned new values. Throughout the experiment, bidder i's valuation $V_{i,j}$ for license j is the product of her private value $\mathrm{PV}_{i,j}$ and a common value CV:[14]

$$V_{i,j} = PV_{i,j} * CV. \tag{1}$$

At the start of a period, bidders are told their private values, which differ across bidders and licenses. The distributions of the private values and the common value signal change between the distinct parts of the experiment. It is common knowledge that in part 1 each bidder's private value for license 1 is a draw from a uniform distribution on [0,50], which we denote by U[0,50]. Bidders' private value for licenses 2 and 3 are drawn from a U[0,40] and U[0,30] distribution

[11] In parts 2, 3, and 4, the roles of weak and strong bidders rotates among the subjects. At the start of each period subjects receive information about their roles.

[12] When a license is sold in parts 2, 3, and 4, subjects immediately receive information about whether the buyer is a weak or a strong bidder.

[13] A possible danger in our setup is that the results are affected by order effects. As a consequence, spill-over effects from one part to the other may occur. In theory, this could add noise to our data, leading to convergence of results. Given that we measure significant treatment effects, with a clear and sensible interpretation, however, we expect that order effects are small at most. We chose for a within treatment design to be able to compare the auction formats in a variety of plausible environments.

[14] An interpretation of the multiplicative structure in (1) in terms of the Dutch zero-base auction is that the value of a FM-license can roughly be modeled as the revenue per listener (the private value part) times the number of potential listeners that can be reached (the common value part).

respectively. Hence, on average, license 1 is the best license and license 3 the worst. In parts 2 and 3, we introduce asymmetries. One weak bidder's private value is drawn from the same distributions as in part 1. Three strong bidders, however, receive private values drawn from U[20,70], U[15,55], and U[10,40] distributions for licenses 1, 2, and 3 respectively. Notice that private values of strong bidders are, on average, higher than those of weak bidders. Private values are independent across licenses, subjects, and periods.

The common value is equal for each bidder and each license. In parts 1 and 2 there is no uncertainty about the common value, which is set equal to 1. In parts 3 and 4, uncertainty is introduced by drawing the common value from a U[0.5,1.5] distribution, with each penny amount being equally likely. The common value is not revealed to subjects who, instead, receive signals, CVS_i. These signals are equal to the common value plus some 'noise':

$$CVS_i = CV + \varepsilon_i \tag{2}$$

where ε_i represents an individual specific error term drawn from a $U[-0.5,0.5]$ distribution. The common value is independent of private values and independent across periods. The error terms are independent across subjects, independent of the common value, and independent across periods. The common value is revealed to all subjects at the end of a period.

Finally, in every auction, bids are restricted to lie between the lowest and highest possible value of the license. Hence, the lowest bid allowed is 0. Moreover, bids cannot exceed 50, 70, 105, and 105 in parts 1, 2, 3, and 4 respectively. These numbers coincide with the levels at which prices start in the simultaneous decreasing auction.

3.4. Subjects

A total of 128 subjects were recruited at the University of Amsterdam. The experiment lasted between 2.5 and 3.5 hours. Subjects' earnings ranged from 38.75 to 119.50 guilders with an average of 77.10 guilders and a standard deviation of 17.13 guilders (77.10 guilders equals about 35 Euro). Theoretically, there is the possibility of bankruptcy. A subject goes bankrupt when her cash balance becomes negative, in which case she would have to leave the experiment and the computer would make choices for her. Subjects are completely informed about the bankruptcy procedure before the first period. In the sessions run, no subject went bankrupt.

4. Results

We present the results in five parts: in Section 4.1 we compare efficiency levels and revenues. Section 4.2 studies the impact of the auction format on entry and in Section 4.3 we discuss potential causes for costly litigation. Section 4.4 determines which order is most profitable when licenses are sold via a sequential first-price auction. Finally, in Section 4.5, we will have a closer look at bidding behavior in the simultaneous first-price auction.

4.1. Efficiency and revenues

The top panel of Table 2 shows the efficiency levels realized in the different parts of the experiment. Efficiency is defined as follows: we first add up the values of the bidders that won a license and denote the result by V_{realized}. Next, we determine the (optimal) allocation that would maximize the sum of the bidders' values and write the corresponding total value as V_{optimal}. To avoid unrealistically high efficiency levels when the lower bound of bidders' values is different

Table 2
Some statistics on efficiencies (in %), revenues, and profits*

		SAA	SimFPA	SeqFPA	SDA
Efficiencies	Part 1	95.2	87.5	79.7	84.7
	Part 2	98.3	92.0	82.9	88.6
	Part 3	88.8	82.0	73.4	74.1
	Part 4[a]	88.7	88.2	86.3	77.4
		78.5	74.8	78.8	69.0
Revenues[b]	Part 1	54.6	56.6	61.9	57.7
		16.5	*9.5*	*10.4*	*14.3*
	Part 2	77.8	83.0	*87.1*	84.4
		27.4	*13.9*	*12.0*	*14.3*
	Part 3	73.0	78.6	83.9	81.7
		27.5	*16.6*	*17.5*	*18.5*
	Part 4	88.4	89.0	91.7	88.4
		28.1	*18.9*	*21.0*	*24.9*
	Part 1	8.78	7.78	5.72	7.27
		8.53	*6.36*	*6.35*	*7.96*
Profits[c]	Part 2	12.03	10.21	8.42	9.63
		12.38	*8.29*	*7.35*	*8.08*
	Part 3	13.64	11.56	9.32	10.03
		15.31	*12.60*	*12.06*	*11.93*
	Part 4	3.57	3.42	2.89	2.82
		8.82	*7.52*	*7.33*	*7.32*

*SAA=simultaneous ascending auction; SimFPA=simultaneous first-price auction; SeqFPA=sequential first-price auction; SDA=simultaneous descending auction. Part 1=symmetric private values, certainty about common value; Part 2=3 strong bidders, 1 weak bidder, certainty about common value; Part 3=3 strong bidders, 1 weak bidder, uncertainty about common value; Part 4=3 strong bidders, 5 weak bidders with outside option; uncertainty about common value.

[a] The (top) bottom entry lists efficiency (without) taking into account the opportunity costs of participating in the auction.

[b] The number in italics is the standard deviation of the revenue.

[c] The number in italics is the standard deviation of the bidders' profits.

from zero, we subtract from V_{realized} and V_{optimal} the sum of bidders' values that would result if the licenses were assigned randomly:

$$\text{Efficiency} = \frac{V_{\text{realized}} - V_{\text{random}}}{V_{\text{optimal}} - V_{\text{random}}} \times 100\%. \tag{3}$$

From Table 2 it is obvious that the simultaneous ascending auction is most efficient in almost all parts of the experiment, and the simultaneous first-price sealed-bid auction is generally second-best.

Notice that in all formats, efficiency levels in part 2 are higher than in part 1. This is intuitive, as it is more difficult for the 'weak bidder' to win a license when asymmetries are introduced. A drop in efficiency levels occurs when uncertainty about the common value component is introduced in part 3. The reason is that a bidder with a high private value but a low common-value signal may be outbid by a bidder with a lower private value but more optimistic common-value information. Such an outcome lowers efficiency, for which only the ranking of private values matters (see also Goeree and Offerman, 2002).

In the final part with endogenous entry, efficiency levels are roughly the same in all treatments except for the simultaneous descending auction. This is true both when the opportunity costs of entry are ignored (top line of part 4) and when they are taken into account (bottom line of part 4). By opportunity costs of entry we mean the outside (lottery) payoff that bidders forego if they enter the auction. In most treatments, too many bidders entered the auction, which is why the number in the bottom line of part 4 is less than that in the top line.[15] The sequential first-price auction yields the least over-entry and the highest level of efficiency (78.8%) when entry is endogenous and its costs are taken into account.

The middle panel of Table 2 lists the revenues generated by the various auctions. Revenues are simply defined as the sum of the three winning bids, i.e. the payments winning bidders had to make. While the simultaneous ascending auction is generally most efficient, it also generates the lowest revenue. The sequential first-price auction is most profitable in all parts of the experiment and its revenues are roughly 10% higher than that of a simultaneous ascending auction.

It is instructive to compare the revenues of the different auction formats in part 3. First note that, in all auction formats, revenues decrease when uncertainty is introduced in part 3. The intuition for this result is that uncertainty increases bidders' information rents, which positively affects their profits (shown in the third panel). With lower levels of efficiency and higher profits for the bidders, the auction's revenue is necessarily lower. Further, it is often argued that the simultaneous ascending auction raises more money when there is uncertainty about the common value because bidders learn from others' bidding behavior, which allows them to bid more aggressively. This intuition is based on Milgrom and Weber (1982) who derive the revenue superiority of the English auction in a symmetric setting. Klemperer (1998) shows that this intuition does not generalize to auctions with small asymmetries between bidders, in which case a first-price element can be expected to raise more money. In our experiment, this is confirmed. The simultaneous ascending auction generates lower revenue in part 3, showing that asymmetries can have a large impact on revenues.

The importance of information generation can also be assessed by comparing the simultaneous descending auction and the sequential first-price auction. As we argued in the Introduction, these two formats are closely related. The main difference is that bidders in the simultaneous descending auction possess more information. At the same time, the simultaneous descending auction may be more vulnerable to collusion. Table 2 shows that the sequential first-price auction even raises slightly more revenue than the simultaneous descending auction and does so at a lower variance.[16] Again the information-generating properties do not seem very important.

Surprisingly, the differences in revenues across auction formats tend to disappear when entry is endogenous (part 4). One would expect that auctions with a first-price element would generate even more revenues relative to the simultaneous ascending auction in this part because they are believed to stimulate entry. In Section 4.2 we will give an explanation for this unexpected result.[17]

[15] Cox et al. (2001) report *less* entry than predicted by Nash in a single-object first-price common value auction. This difference could be the result of a more severe winner' curse in their experiment, where only one winning bidder is selected from six, eight, or twelve bidders. In contrast, in our setup three winning bidders are selected from eight bidders, which means that winning is less informative and there is less danger of a winner's curse.

[16] As one of the referees rightly pointed out, when there is a common value component, ex-ante bidder asymmetries heighten winner's curse considerations for weak bidders, thus leading to passive bidding by weak bidders and low revenues. Therefore, low revenues should also be expected in the absence of collusion. However, the fact that the sequential descending auction raises (slightly) less revenue than the sequential first-price auction even in the environments where only private value components exist (parts 1 and 2) cannot be explained by winner's curse and suggests that collusion is facilitated in the sequential descending auction.

[17] Appendix B presents two tables that list significance levels of tests comparing efficiency and revenue across treatments. The differences in efficiency between the simultaneous ascending auction and the first-price formats, as well as the differences in revenue between the simultaneous ascending auction and the sequential first-price auction are significant in the first three parts of the experiment.

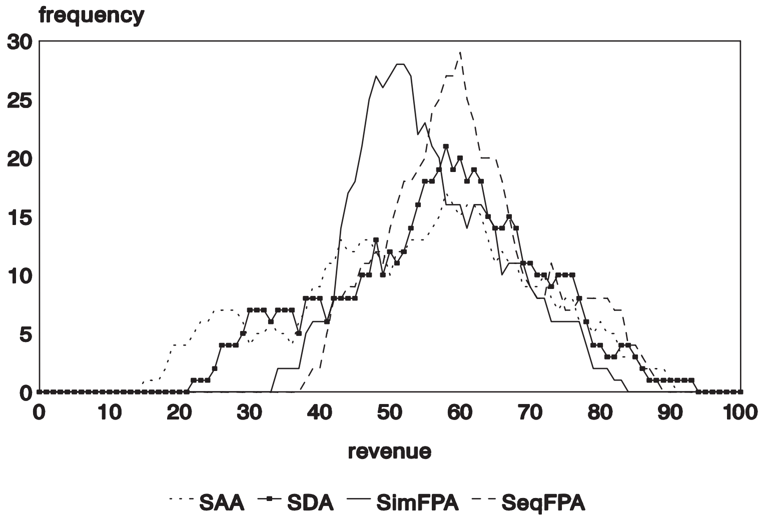

Fig. 1. Frequency distribution of revenues in part 1. For each revenue-level on the horizontal axis, the vertical axis reports the number of auctions generating a revenue within the interval [revenue−5, revenue+5].

The bottom panel of Table 2 shows that bidders' profits are highest in the simultaneous ascending auction. Not only the level of profits is important, however. In view of the European UMTS auctions, politicians (and the public) judged that winners should pay 'reasonable' prices for the licenses. In other words, the auction design should not result in undeserved windfall profits nor should it lead to bankruptcies caused by overbidding. In this perspective, the variability of bidders' profits is important. Note that this variability is substantially higher in the simultaneous ascending auction than in the first-price auctions, which thus have a better chance of yielding 'reasonable' prices.

Likewise, it is risky for government decision-makers if the auction results in very high or very low revenues. Table 2 includes a measure for the variance of revenues. Note that the simultaneous ascending auction not only yields the lowest revenues, but the variability is also the largest. In other words, the revenue in the simultaneous ascending auction is less predictable and varies more than that of other formats. This result is in accordance with a theoretical result for auctions with symmetric private values. Waehrer et al. (1998) show that in the class of standard auctions, the first-price auction is the superior choice for risk-averse sellers who are interested in raising revenue at low variability. The effect can also be observed in Figs. 1–3, which show revenue histograms of the four auctions. One noticeable feature of these figures is that the simultaneous ascending auction (and to a lesser extent the simultaneous descending auction) yields a histogram that is much more spread out than that of other formats, especially when asymmetries exist. Consider, for instance, Fig. 2 where the simultaneous ascending auction results in revenues of less than 20 in quite a few cases. And revenues in the third part of the experiment, where a common value component is added to bidders' valuations, are almost uniformly distributed between 20 and 110, see Fig. 3.[18]

To summarize, the simultaneous ascending auction, and to a lesser extent the simultaneous descending auction, generate low and variable revenues in uncompetitive situations. It is well known that such outcomes yield suspicion of collusion.

[18] We did not add a figure for part 4 since we only had half as many data points for this part.

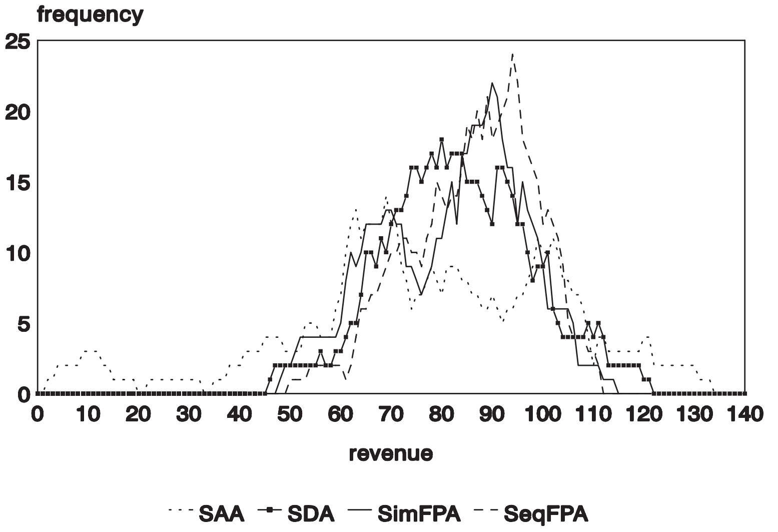

Fig. 2. Frequency distribution of revenues in part 2. For each revenue-level on the horizontal axis, the vertical axis reports the number of auctions generating a revenue within the interval [revenue − 5, revenue + 5].

4.2. Performance of weak and strong bidders and entry

One alleged problem with the simultaneous ascending auction is that weak bidders perceive their chances of winning as low and feel they may as well drop out early. In contrast, weak bidders have better chances in a first-price format, where their probability of winning depends on how much risk strong bidders are willing to take.

This intuition is borne out by Table 3, which shows the fraction of licenses won by weak and strong bidders across treatments. For those cells in Table 3 that correspond to parts 2 and 3 of the

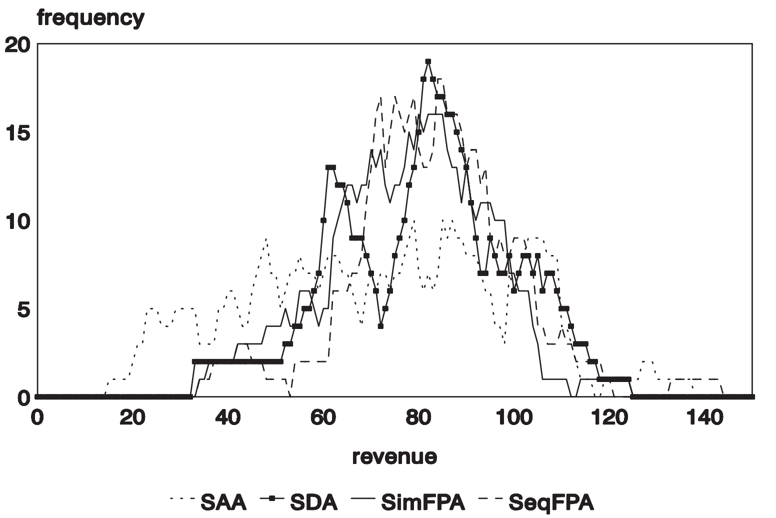

Fig. 3. Frequency distribution of revenues in part 3. For each revenue-level on the horizontal axis, the vertical axis reports the number of auctions generating a revenue within the interval [revenue − 5, revenue + 5].

Table 3
Probability of winning for weak and strong bidders

	SAA	SimFPA	SeqFPA	SDA
Part 2[a]	88.7	85.1	85.7	86.3
	33.9	44.6	42.9	41.1
Part 3[a]	88.1	83.3	86.3	81.2
	35.7	50.0	41.1	53.6
Part 4[a]	75.0	65.5	72.6	66.7
	15.0	20.7	16.4	20.0

[a] The top (bottom) entry lists the probability that the strong (weak) bidder buys.

experiment, three times the top number (3 strong bidders) plus one times the bottom number (one weak bidder) add up to three (the number of licenses). For the cells corresponding to part 4, the equivalent rule is 3 times the top number plus five times the bottom number add up to three. Notice that the weak bidder's probability of winning is 20% to 40% lower in the simultaneous ascending auction compared to the first-price formats.

However, a weak bidder's entry decision is not only determined by her probability of winning but also by her profit if a license is won. Table 4 shows average profits of weak and strong bidders across treatments. These are averaged across all participating bidders and, therefore, represent the expected profits in the various treatments. Note that in 7 out of 9 cases, the expected profit for a weak bidder is higher for the first-price auction than for the corresponding simultaneous ascending auction. One of the exceptions is part 3 of the experiment, where weak bidders earn *less* in the sequential first-price auction than in the simultaneous ascending auction. This may explain why entry is lower in the former format (as shown in the bottom panel of Table 4). Furthermore, weak bidders generally make higher profits in the simultaneous descending auction and the simultaneous first-price auction. Hence, in part 4 where entry is endogenous, they tend to enter more in these formats. In this way, endogenous entry provides a force towards revenue equivalence of different auction formats because weak bidders enter more frequently in auctions where they expect to make higher profits.

4.3. Causes of litigation

Next, we discuss possible causes of litigation. The simultaneous first-price auction may create problems when bidders discover they had the highest bid for a certain license but

Table 4
Profits of weak and strong bidders[a]

		SAA	SimFPA	SeqFPA	SDA
Profit	Part 2	15.46	12.40	10.39	11.88
		1.73	3.66	2.50	2.86
	Part 3	17.13	13.68	11.54	11.65
		3.16	5.20	2.66	5.18
	Part 4	8.93	8.06	6.73	7.35
		0.35	0.63	0.59	0.10
Entry part 4	Realized	2.18	2.36	1.89	2.54
	Efficient	0.43	0.43	0.43	0.43

[a] In the upper part of the table the top (bottom) entry lists the expected profit for strong (weak) bidders.

Table 5
Occurrences of regretful bidders in the SimFPA[a]

		# Licenses not awarded to highest bidder			
		0	1	2	
# Unhappy winners	0	79	46	8	67.9%
	1		47	12	30.1%
	2			4	2.0%
		40.3%	47.4%	12.2%	

[a] Each cell shows the number of simultaneous first-price auctions that resulted in a particular outcome. A 'license not awarded to the highest bidder' is one which was assigned to a subject that bid less for the license than at least one other subject. An 'unhappy winner' is a subject that (i) was not assigned a license for which she was the highest bidder but instead received a different license, and (ii) profited less from the license she was assigned than she would have from the license for which she was the highest bidder.

were awarded a different one. This is not necessarily bad for those bidders, however, since the profit margin earned on the license assigned may exceed that of the license for which they had the highest bid.[19] Table 5 lists the number of times that bidders in the experiment were assigned a particular license while they had the highest bid and a higher profit margin for a different license. In roughly one-third (30.1%+2.0%) of the simultaneous first-price auctions we ran, there is at least one winner who would have preferred another license for which she had the highest bid. It may happen that such 'unhappy' bidders file a lawsuit after the auction in an attempt to change the assignment of licenses.[20]

There are other potential reasons to start lawsuits, which could occur in all formats. For example, when the profit margins across licenses vary significantly, the final allocation may be perceived as unfair and the bidder that is worst off may consider a lawsuit. By aggregating subjects' choices in the experiment we determine in which format the profits derived from the three licenses are most homogeneous. Table 6 shows that the simultaneous first-price auction results in the *lowest* variance of profits. So, from this perspective, the simultaneous first-price auction provides the *least* temptation to go to court.

Finally, we consider the probability that a bidder pays more for a license than it is worth. Such 'cursed' winners may try to renegotiate the conditions of the sale, resulting in

[19] This is illustrated in the following example. Suppose:

- Bidder 1 bids 20 for license 1, 15 for license 2, and 0 for license 3.
- Bidder 2 bids 18 for license 1, 10 for license 2, and 0 for license 3.
- Bidder 3 bids 0 for license 1, 0 for license 2, and 2 for license 3.
- Bidder 4 bids 0 for license 1, 0 for license 2, and 0 for license 3.

The simultaneous first-price auction assigns license 1 to bidder 2, license 2 to bidder 1, and license 3 to bidder 3, with a maximum revenue of 35. Bidder 1 has the highest bids on licenses 1 and 2, and it seems likely that this bidder will compare the profit margins on these licenses. If bidder 1 has a value of 22 for license 1 and a value of 20 for license 2, she will be happy with the final assignment: she makes a profit of $20-15=5$ when assigned license 2, while she would have earned $22-20=2$ if she would have been assigned license 1. On the other hand, if bidder 1's value is 25 for license 1 and 17 for license 2, she will be unhappy with the assignment, as her profit is 2 under the current assignment while it would have been 5 otherwise.

[20] Notice that we take a short-cut at the important issue of litigation. Ideally, one would like to analyze litigation issues in a supergame where bidders may try to change the auction outcome in a litigation phase after the auction. Such an extensive analysis is beyond the scope of this paper, however.

Table 6
Fair division of profits[a]

	SAA	SimFPA	SeqFPA	SDA
Part 1	347.9	138.7	255.5	299.2
Part 2	520.6	224.4	241.3	261.2
Part 3	719.4	476.2	480.3	468.4
Part 4	527.0	434.5	464.6	474.4
All parts	529.0	301.9	345.5	361.7

[a] Each entry displays the 'variance' in profits that result from the three different licenses, L1, L2, and L3, where variance is defined as: $(\text{profit L1} - \text{profit L2})^2 + (\text{profit L1} - \text{profit L3})^2 + (\text{profit L2} - \text{profit L3})^2$.

costly lawsuits and delays.[21] Table 7 lists the observed percentages of losses in the four treatments. It shows that the simultaneous first-price auction protects bidders better from a winner's curse than the simultaneous ascending auction. Notice that because there are three licenses for four bidders, rational bidders will let their strategy only depend on their own signal(s). If one bidder drops out, the auction is finished, so there is no possibility to condition the strategy on the drop-out level of others. Naïve bidders may argue differently, however. The fact that all other bidders are still active may convince a bidder that the others have more favorable information about the common value aspect than expected and may encourage her to stay a bit longer in the auction than originally planned. Such a possibility of irrational updating was first proposed to explain low revenues in Dutch single-unit private value auctions where some subjects argue that the fact that others have not yet pushed the button makes it more likely that they have low private values and thus they decide to delay their decision to stop the button. A similar loss-enhancing factor does not exist in the simultaneous first-price auction, which may explain that we observe larger winner's curse problems in the ascending auction than in the simultaneous first-price auction. Notice further that the winner's curse is less strong than in single-unit auctions because there is only a single loser so winning is not as informative.

4.4. The sequential auction: order matters

In all but one of the formats studied here, all three licenses are offered for sale simultaneously. The exception is the sequential first-price auction where licenses are sold one by one. Since some licenses are (on average) better than others, the obvious question becomes "which license should be put up for sale first?" We compare two possible ordering strategies for the seller: 'best foot forward' and 'best for last'. In the former, license 1 is offered first, then license 2, and finally license 3. The latter strategy employs the reverse order.[22]

[21] For example, in the C-block broadband PCS auction organized by the FCC, overly attractive installment payments led to speculative bidding and caused all major bidders to default and declare bankruptcy. Due to subsequent bankruptcy litigation, large part of the C-block spectrum was left idle for years (Cramton, 2002).

[22] In a theoretical analysis, Chakraborty et al. (2001) take a different angle at the question whether best goods should be sold first or last in a sequential auction. In their model, the seller has private information about two stochastically equivalent, independently distributed goods. Unlike in our setup, two different sets of bidders participate in the two auctions. Bidders of the second auction observe the price realized in the first auction. The authors find the result that best for last may be the revenue-maximizing strategy (as long as knowledge of the first period price reduces buyer information rents more for the better good than for the worse good) but best foot forward may be the unique pure strategy equilibrium.

Table 7
Probability of a loss

	SAA	SimFPA	SeqFPA	SDA
Part 1	0.0	0.0	1.8	1.2
Part 2[a]	1.8	0.6	1.8	0.0
	5.3, 1.3	0.0, 0.7	4.2, 1.4	0.0, 0.0
Part 3[a]	9.5	6.6	13.1	9.5
	20.0, 8.1	6.0, 7.9	26.1, 11.0	13.3, 8.7
Part 4[a]	21.4	15.5	21.4	25.0
	38.1, 15.9	34.5, 5.5	39.1, 14.8	46.4, 14.3

[a] The top entry lists the aggregate probability and the two bottom entry list the probability for the weak and strong bidder respectively.

The effect of sequencing on revenues is shown in Table 8 that focuses on the subset of auctions where the licenses were sold in strictly decreasing or strictly increasing order of value. Notice that in all parts, revenues are higher with the best foot forward strategy, i.e. when the best license is sold first. This is an intuitive result, since the best foot forward strategy assigns the best licenses when competition is most fierce. Remarkably, however, the result also holds in parts 3 and 4, where uncertainty about the common value could result in a countervailing effect: in the best for last strategy the information revealed in the sale of less valuable licenses stimulates aggressive bidding on the better licenses, which are sold at the end. The best foot forward strategy leads to significantly higher revenues than the best for last strategy.[23]

4.5. Bidding in the simultaneous first-price auction

One exciting reason for using experiments is to study the performance of new auction designs. To the best of our knowledge, the simultaneous first-price auction has not yet been employed in practice. In order to get a better insight in how it works, we take a closer look at bidding behavior in this auction design.[24] Although the equilibrium strategies are not known, we have a number of intuitive conjectures with respect to bidding in the simultaneous first-price auction. We consider individual bid functions of the format

$$b_{ij} = f_j(v_{i1}, v_{i2}, v_{i3}), \tag{4}$$

where b_{ij} denotes individual i's bid, $i=1,\ldots, 4$, and v_{ij} her value for license $j=1,\ldots, 3$.[25] We formulate the following conjectures.

C1. *$\partial f_j/\partial v_{ij}>0$. In words, i's bid for license j is positively correlated with i's value for license j.*

C2. *$\partial f_j/\partial v_k \leq 0$, $k \neq j$. In words, i's bid for license j is negatively correlated with i's value for license $k \neq j$. One way to see this is the following. Ceteris paribus, an increase in v_{ik} will increase the bid for license k (through C1) and, therefore, the probability of winning k. Because i can obtain at*

[23] A Wilcoxon rank test that uses for each group of bidders the average revenues per block of 7 periods as data points yields a significance level of $p=0.05$ ($n=28$).

[24] We restrict the analysis to the simultaneous first-price auction, because of the lack of theory for this format. In addition, the number of observations (bids) available for the analysis is much lower for the other formats.

[25] For notational convenience, we do not include a subscript for the round in the experiment.

Table 8
Effect of the order of licenses on revenue in SeqFPA[a]

Order of licenses for sale	Part 1	Part 2	Part 3	Part 4	All parts
3–2–1 (increasing in value)	59.8	86.1	81.1	84.5	76.9
	10.0	*12.5*	*24.4*	*21.0*	*20.3*
1–2–3 (decreasing in value)	61.0	90.8	85.6	96.0	81.5
	10.5	*12.2*	*14.5*	*21.6*	*19.2*

[a] Standard deviation in italics.

most one license, this implies an overall increase in the probability of obtaining a license. Hence, i can afford to take more risk and increase the profit margin on $j \neq k$ by lowering her bid.

C3. *b_{ij} (common value) $<$ b_{ij} (private value). Because the common value element forces the bidder to correct for the winner's curse, a bid will be lower, ceteris paribus, in part 3 of the experiment than in part 2.*

In order to test these conjectures, we estimated bid functions (4) in a simultaneous model (using OLS) with fixed effects to correct for multiple observations per individual. We did so separately for parts 1, 2, and 3 of the experiment and separately for strong and weak bidders. In part three, we corrected subjects' private values using the individual signal of the common value. First, we used a quadratic specification of f_j, but the results did not reject a null-hypothesis of linearity.[26] Therefore, from here onward, we will assume that the bid functions are linear.

Table 9 gives the results for estimated linear bid functions of the form $b_{ijt} = \beta_{0j} + \beta_{j1} v_{i1t} + \beta_{j2} v_{i2t} + \beta_{j3} v_{i3t} + \mu_i + \varepsilon_{ijt}, j = 1, \ldots, 3$, where t denotes the round (auction), μ_i is the fixed effect and ε_{ij} is a white noise error term. Table 9 shows strong support for C1 and C2. In all 15 cases, $\partial f_j / \partial v_{ij}$ ($= \beta_{jj}$) is positive with high levels of statistical significance. Moreover, $\partial f_j / \partial v_{ij} = \beta_{jk}$, $k \neq j$, is negative 28 out of 30 times (the two exceptions being small and statistically insignificant). Though small, 11 of these 28 coefficients are statistically significant.

A closer look at the estimated bid functions reveals that they are quite similar across environments. The coefficients concerning the direct effect do support C3, however: in all 6 cases (3 for strong bidders and 3 for weak bidders) the coefficient is lower in part 3 than in part 2. The direct effect is roughly 2/3 for private values and 1/2 for common values. This seems quite high, given the high probability of obtaining a license (3 licenses for 4 bidders). Note, however, that the bid is reduced by the negative effect of the values of the other licenses, even though the direct effect strongly dominates. Finally, note that the intercepts are not significantly different from 0 when there are only private values. When common values are introduced in part 3, we find that the mitigated reaction to value is compensated by an increase in the intercept.

All in all, bidding in the simultaneous first-price auction appears to follow intuitive rules. This and the robustness across environments indicate that bidders are responding to the mechanism in a consistent way, even if we cannot establish whether their behavior

[26] The regressions included 90 second-order terms (6 quadratic terms per bid function, 1 function for each of the 3 licenses, 3 estimations for strong bidders in parts 1, 2, and 3 and 2 for weak bidders in parts 2 and 3). After estimation, only 5 out of 90 coefficients turned out to be significant at the 5% level, which is what one would expect to observe in absence of an effect. The 5 'significant' coefficients were spread across different equations.

Table 9
Linear bid functions in SimFPA[a]

		Intercept	Value license 1	Value license 2	Value license 3
Part 1 (only strong bidders)	Bid license 1	1.015 (0.51)	**0.689** (31.8)	−0.082 (3.10)	−0.090 (2.54)
	Bid license 2	2.860 (1.67)	−0.034 (1.85)	**0.620** (27.5)	−0.136 (4.51)
	Bid license 3	1.076 (0.85)	−0.022 (1.62)	−0.036 (2.16)	**0.568** (25.50)
Part 2, strong bidders	Bid license 1	1.369 (0.48)	**0.661** (23.7)	−0.080 (2.33)	−0.066 (1.46)
	Bid license 2	0.107 (0.04)	−0.051 (2.00)	**0.681** (21.9)	−0.077 (1.87)
	Bid license 3	0.665 (0.33)	−0.090 (4.52)	0.003 (0.12)	**0.666** (20.5)
Part 2, weak bidders	Bid license 1	3.498 (0.59)	**0.669** (7.87)	−0.172 (1.84)	−0.034 (0.33)
	Bid license 2	−1.905 (0.44)	−0.070 (1.13)	**0.671** (9.89)	0.018 (0.24)
	Bid license 3	0.878 (0.29)	−0.005 (0.11)	−0.074 (1.53)	**0.660** (12.4)
Part 3, strong bidders	Bid license 1	15.46 (5.16)	**0.463** (13.7)	−0.083 (1.82)	−0.063 (1.10)
	Bid license 2	13.42 (5.60)	−0.151 (5.60)	**0.536** (14.7)	−0.097 (2.12)
	Bid license 3	5.532 (2.99)	−0.084 (4.03)	−0.074 (2.65)	**0.539** (15.2)
Part 3, weak bidders	Bid license 1	5.670 (1.11)	**0.594** (11.8)	−0.090 (1.18)	−0.211 (2.97)
	Bid license 2	8.742 (1.83)	−0.070 (1.05)	**0.590** (8.25)	−0.070 (1.05)
	Bid license 3	2.516 (0.73)	−0.094 (2.76)	−0.023 (0.44)	**0.424** (8.84)

[a] Absolute t-values in parentheses; coefficient for effect of value v_j on b_j in bold.

corresponds to equilibrium bidding. The estimated bid functions do help us understand some of the results presented in the previous subsection, however. For example, bidders avoid the winners' curse because they adjust their bid functions in a sensible way.

5. Conclusions

We used experiments to compare the performance of three first-price formats to the more conventional simultaneous ascending auction. We considered the familiar efficiency and revenue criteria, but also other factors such as variance in revenues, entry, winner's curse outcomes, and reasons for subsequent litigation. To test the robustness of each auction, we changed the environment during the experiment. In the first part, four bidders knew their values for the three licenses with certainty and there were no a priori differences between bidders. In the second part, asymmetries were introduced: three strong bidders (incumbents) competed with one weak bidder (newcomer). The third part introduced (common-value) uncertainty into bidders valuations, which affected all bidders and licenses the same way. In the final part, five weak bidders made sequential decisions to enter the auction or profit from an outside opportunity; in the auction, weak bidders had to compete with three strong bidders.

In the first three parts of the experiment, the simultaneous ascending auction was most efficient. Hence, if efficiency is the sole criteria used to judge the auction's performance, there is no reason to switch to a first-price format. In practice, however, allocative efficiency is not observable and discussion will likely focus on the auction's revenue. In the first three parts of the experiment, auctions with a first-price element yielded significantly higher revenues than the simultaneous ascending auction.[27] In particular, the highest revenues occur with a sequential format where the best licenses are sold first.

[27] In addition, revenues of the simultaneous ascending auction are more variable and this format is therefore most likely to result in really low revenues, which bear the suspicion of collusion.

As expected, auctions with a first-price element offer weak bidders a better chance to acquire a license than the simultaneous ascending auction. Thus, the idea that first-price elements enhance competition in this environment finds some support in our results. These results do point to an interesting refinement of this idea, however. The motivation for entry lies not so much in the probability of winning as in the expected profit. In some of our first-price designs we observed more aggressive bidding, causing lower expected profits and less entry than in the SAA. In particular, aggressive bidding in the sequential first-price auction yielded the *lowest* profits for weak bidders in part 3 resulting in the *least* entry in part 4. The simultaneous descending auction and the simultaneous first-price auction offer better prospects to weak bidders and resulted in the highest level of entry in part 4. Comparing revenues in part 4 shows that endogenous entry provides a force towards revenue equivalence of different auction formats.

The allocations of licenses that result from first-price auctions show less variation in profit margins than the allocations produced by the simultaneous ascending auction. Variation in per license profits may cause a feeling of injustice and result in lawsuits after the auction. When valuations are uncertain, the simultaneous first-price auction shows the lowest incidence of winner's curse outcomes. In other words, the chance that the buyer of a license incurs a loss is smallest using this format. Winner's curse outcomes are often followed by costly renegotiation processes in or out of court. It should be noted, however, that a bidder in the simultaneous first-price auction may not get a license for which she had the highest bid. This may cause problems when a bidder profits less from the license she was assigned. In fact, in roughly one-third of all the auctions conducted, at least one of the buyers is an 'unhappy' winner who may challenge the allocation of licenses. These results are a consequence of the way in which bidders behave in the simultaneous first-price auction. We observe consistent linear bid functions that follow intuitive rules and respond to changes in the environment in a sensible way.

The main motivation for our study lies in the notion that in asymmetric, uncompetitive environments, the simultaneous ascending auction may lose some of its appeal and auctions with first-price elements might be an attractive alternative. We created an extreme environment of this type to provide a test bed for this notion. Our results support the idea that the SAA performs less well in these circumstances than is usually the case, though it still outperforms the first-price alternatives on some criteria, such as efficiency. The bottom line of our results is that in uncompetitive environments, the choice of an auction format depends crucially on the evaluation criteria considered to be important.

Acknowledgements

This research was sponsored by the Directoraat-Generaal Telecommunicatie en Post (DGTP) of the Dutch Ministry of Transport, Public Works, and Water management (V&W). We are very grateful to CREED programmer Jos Theelen for programming the experiment. We would like to thank three referees, John Bernard, Charles Noussair and seminar participants at the Economic Science Association Meetings (Boston, 2002; Strasbourg, 2002) and University of Arizona economic theory seminar (2003) for useful comments. We gratefully acknowledge financial support from the National Science Foundation (SBR 0094800), the Dutch National Science Foundation (VICI 453.03.606), the Alfred P. Sloan Foundation, and the Royal Netherlands Academy of Arts and Sciences.

Appendix A. Instructions for simultaneous first-price auction (translated from Dutch)

Welcome to this experiment on decision making! You can make money in this experiment. Your choices and the choices of the other participants will determine how much money you will make. Read the instructions carefully. There is a calculator, paper, and a pen on your table, which you can use during the experiment. Before the experiment starts, we will hand out a summary of the instructions and there will be two practice rounds.

The experiment: You will earn points in the experiment. At the end of the experiment your points will be exchanged into guilders at a rate of 25 cents per point. You will start with 50 points, which you do not have to pay back. The experiment consists of 4 parts each consisting of 7 rounds. Only when a part is completely finished will you receive the instructions for the next part. Each round you will be part of a group of 4 persons and three products will be sold in each group in each round.

Values of the products: Before a round starts you will be assigned a value for each product. Product 1 is usually the best (most valuable) product and product 3 is usually the worst. For each participant:

The value of *product 1* lies between 0 and 50 points, where each number is equally likely.
The value of *product 2* lies between 0 and 40 points, where each number is equally likely.
The value of *product 3* lies between 0 and 30 points, where each number is equally likely.

For each participant, the value of one product is independent of the value of any other product. For each product, the value for one participant is independent of the value for other participants. Therefore, your values will (very) likely differ from those of others. At the start of a round you will get to know your own values. You will not know the values of others and other participants will not know your values.

Sale of the products: In each round, each participant can buy at most one product. Three products are sold simultaneously. Each participant submits a bid for each of the three products. The computer assigns the products such as the sum of the winning bids is as large as possible and each participant gets at most one product. (When there is more than one way to maximize the sum of the winning bids, the computer selects one of these ways at random.) A participant who is assigned a product will pay a price equal to her own bid for the product.

A bid has to be larger than or equal to 0 points and smaller than or equal to 50 points. A bidder can also decide not to buy a product by not submitting a bid for the particular product. If no one submits a bid for a particular product, the product will not be sold.

Example: Here follows an example to illustrate the sale procedure. The numbers in the example are chosen arbitrarily and do not indicate what will happen in the experiment.

Bidder 1 bids 30 for product 1.
Bidder 1 bids 29 for product 2.
Bidder 1 bids 1 for product 3.
Bidder 2 bids 24 for product 1.
Bidder 2 bids 16 for product 2.
Bidder 2 bids 1 for product 3.

Bidder 3 bids 1 for product 1.
Bidder 3 bids 1 for product 2.
Bidder 3 bids 1 for product 3.

Bidder 4 bids 1 for product 1.
Bidder 4 bids 1 for product 2.
Bidder 4 bids 2 for product 3.

The outcome is then:

Product 1 is sold to the second bidder at a price of 24.
Product 2 is sold to the first bidder at a price of 29.
Product 3 is sold to the fourth bidder at a price of 2.

The sum of the winning bids is equal to 55 (24+29+2=55). Note that bidder 1 is assigned product 2, while (s)he also had the highest bid for product 1. However, if product 1 had been assigned to the bidder 1, product 2 to bidder 2, and product 3 to bidder 4, the sum of the winning bids would have been 48 (30+16+2=48), which is less than 55.

Profit and loss: The buyer will not literally receive a product. Instead the buyer will receive an amount equal to the value of the product minus its price (in points). Note that the buyer makes a loss when (s)he pays a price above (her) his value for the product. A participant who does not buy a product earns 0 points. To summarize:

If you buy a product in a round : Profit = your value minus the price you pay

If you do not buy a product in a round : Profit = 0

As mentioned above, a buyer makes a loss if (s)he pays more than (her) his value for the product. Just as a profit is automatically added to the amount earned up to that round, a loss will automatically be subtracted. It is conceivable that at some point your earnings will become negative. Since we do not want you to owe money to us, you will have to leave the experiment without earning any money in that case. A participant with negative earnings will be replaced by the computer. Participants who are matched with the computer will know this in advance.

Results of a round: At the end of a round, results will be communicated. You will be told the prices at which the products are sold, and in case you buy a product, you will also be told your profit. Then a new round will start and again three products will be sold. Each participant receives three new values for the products. Your values for the products in one round will be independent of your values in any other round.

End: You have reached the end of the instructions. If you want to you can read (parts of) the instructions again. When you are ready, please push the button “READY”. When all participants have pushed “READY”, the experiment will start. Once the experiment has started, you will NOT be able to return to these instructions. Before the experiment starts, a summary of the instructions will be handed out and two practice rounds will be carried out. Your gains or losses during these practice rounds will NOT be added to or subtracted from your earnings.

If you still have questions, please raise your hand!

Instructions part 2: The second part of the experiment lasts for 7 rounds. At the start of the first round of the second part your earnings will be increased by 10 points (you do not have to pay back this amount). The rules in the second part of the experiment are by and large the same as the rules of the first part. The only difference between parts 1 and 2 concerns the way in which the values of the products are determined.

Values of the products: There will be two types of buyers. In each round, there will be 3 large bidders and 1 small bidder in each group.

For a large bidder:

The value of *product 1* lies between 20 and 70 points, where each number is equally likely.
The value of *product 2* lies between 15 and 55 points, where each number is equally likely.
The value of *product 3* lies between 10 and 40 points, where each number is equally likely.

For the small bidder:

The value of *product 1* lies between 0 and 50 points, where each number is equally likely.
The value of *product 2* lies between 0 and 40 points, where each number is equally likely.
The value of *product 3* lies between 0 and 30 points, where each number is equally likely.

For both types of buyers the value of product 1 is on average higher than that of product 2, and the value of product 2 is on average higher valuation than that of product 3. Large bidders have on average higher values than the small bidder. Nevertheless, it can happen that the small bidder values a product more than one (or more) large bidder(s).

For each participant, the value of one product is independent of the values of other products. For each product, the value for one participant is independent of the values for others. Hence, your values will (very) likely differ from those of others. At the start of a round you will get to know your values. You will also be told whether you are a small or a large bidder (in some rounds you will be a small bidder, in other rounds you will be a large bidder). You will not know the values of other participants, and other participants will not know your values.

Information: Bids are restricted to lie between 0 and 70 points. After the results of a round have been communicated, a new round starts in which three new products will be sold. Again it will be determined for each bidder whether (s)he is a small or large bidder. Each participant receives new values for the products. Your values for the products in one round are independent of your values in any other round.

End: You have reached the end of the instructions. If you want to you can read (parts of) the instructions again. When you are ready, please push the button "READY". When all participants have pushed "READY", the second part will start. Before the second part starts, a summary of the instructions will be handed out. If you still have questions, please raise your hand!

Instructions part 3: The third part of the experiment lasts for 7 rounds. At the start of the first round of the third part your earnings will be increased by 10 points (you do not have to pay back this amount). The rules of the third part are by and large the same as the rules of the second part. The only difference between parts 2 and 3 concerns the way in which the values of the products are determined. This will be more complicated than in parts 1 and 2. So, please read these instructions carefully. In part 3 you will not precisely know the value of a product.

Preliminary values and common percentage: Until now you were told your values for the products at the start of each round. In the remainder of the experiment, you will only know your preliminary values at the start of each round. Your "final value" for a product equals a percentage of your preliminary value. This "common percentage" is the same for all products and for all participants. The common percentage can be larger or smaller than 100%, so your final value can be larger or smaller than your preliminary value. Preliminary values are determined in the same way as the values in previous parts:

For a large bidder the preliminary value of

product 1 lies between 20 and 70 points, where each number is equally likely.
product 2 lies between 15 and 55 points, where each number is equally likely.
product 3 lies between 10 and 40 points, where each number is equally likely.

For the small bidder the preliminary value of

product 1 lies between 0 and 50 points, where each number is equally likely.
product 2 lies between 0 and 40 points, where each number is equally likely.
product 3 lies between 0 and 30 points, where each number is equally likely.

At the start of a round, participants only get to know their own preliminary values.

Common percentage and signal: In each round, the common percentage is the same for all licenses and products. The common percentage lies between 50% and 150%, where each percentage is equally likely.

No participant knows the common percentage with certainty. Instead, each participant receives a signal about the common percentage. A participant's signal is equal to sum of the common percentage plus an "error". This error is a percentage between −50% and 50%, where each percentage is equally likely. For example, if the common percentage is 90% and the error is −10%, then the signal is 80%.

The error for one participant is independent of the errors for others. As a result, participants will (very) likely receive different signals. So even though the common percentage is the same for all participants, they receive different signals. The error for one participant is also independent of the common percentage and the preliminary values.

At the start of each round, participants will only get to know their own signal.

Example: Here follows an example to illustrate the procedure to determine the values. The numbers are chosen arbitrarily and do not indicate what will happen in the experiment.

Suppose you are a large bidder and:

Your preliminary value for product 1 is 40.
Your preliminary value for product 2 is 28.
Your preliminary value for product 3 is 20.
Your signal is 125%.

Suppose you buy product 2 at a price of 27. After the round has finished you are informed that the common percentage is 110% (hence your error was +15%). Your final values for the products are then:

Your final value for product 1 is 110% of 40 = 44.
Your final value for product 2 is 110% of 28 = 31.
Your final value for product 3 is 110% of 20 = 22.

In this case, your profit equals your value for product 2 minus the price you paid: $31-27=4$.

Information: Bids are restricted to lie between 0 and 105 points. At the end of a round you will be told the value of the common percentage. After the results of a round have been communicated, a new round starts in which three new products will be sold. Again it will be determined for each bidder whether (s)he is a small or large bidder. Each participant receives new preliminary values for the products. A new common percentage will be determined and each participant will receive a new signal. Your preliminary values and signals in one round are independent of your preliminary values and signals in any other round.

End: You have reached the end of the instructions. If you want to you can read (parts of) the instructions again. When you are ready, please push the button "READY". When all participants have pushed "READY", the third part will start. Before the third part starts, a summary of the instructions will be handed out. If you still have questions, please raise your hand!

Instructions part 4: The fourth part of the experiment lasts for 7 rounds. At the start of the first round of the fourth part your earnings will be increased by 10 points (you do not have to pay back this amount). The rules of the fourth part are by and large the same as the rules of the third part.

In each round, you will be part of a group of eight participants. Three participants play the role of large bidders and five participants will be small bidders. Again three products are sold in each round. Before bidding starts, the small bidders can decide whether they want to participate in the auction or take part in a lottery.

Lottery: In the lottery each participant has a 50% chance to earn 0 points and a 50% chance to earn 6 points. The computer throws a die, and if the die throw is 4, 5, or 6, the participant receives 6 points and if the die throw is 1, 2, or 3, the participant earns 0 points. The die throws for one participant are independent of those for others. And the die throws in one round are independent of those in other rounds.

If a small bidder decides to participate in the auction, she forgoes the chance to play the lottery. As before, the small bidder can earn money in the auction by buying a product.

A large bidder does not have the possibility to play the lottery. They have to participate in the auction. In the auction, the three large bidders and the small bidders that decided to participate can bid for the products in the same way as in part 3.

Information: Small bidders decide sequentially whether they participate in the auction or lottery. The order in which small bidders decide is determined randomly by the computer. Before a small bidder decides (s)he will be informed about (her) his preliminary values and signals and about the decisions that other small bidders made before (her) him. After all small bidders have decided to participate in the auction or lottery, all participants will be told how many bidders participate in the auction.

All other rules are the same as in part 3.

End: You have reached the end of the instructions. If you want to you can read (parts of) the instructions again. When you are ready, please push the button "READY". When all participants have pushed "READY", the fourth part will start. Before the fourth part starts, a summary of the instructions will be handed out. If you still have questions, please raise your hand!

Appendix B. Statistical tests

Table B1. Test results efficiency

	SDA	SimFPA	SeqFPA
SAA	1 > 0.02	1 > 0.02	1 > 0.02
	2 > 0.01	2 > 0.01	2 > 0.01
	3 > 0.01	3 > 0.09	3 > 0.01
	4a > 0.07	4a > 0.72	4a > 0.72
	4b > 0.07	4b > 0.14	4b < 1.00
SDA	X	1 < 0.58	1 > 0.12
		2 < 0.26	2 > 0.48
		3 < 0.04	3 > 1.00
		4a < 0.07	4a < 0.07
		4b < 0.14	4b < 0.07
SimFPA	X	X	1 > 0.26
			2 > 0.05
			3 > 0.09
			4a > 0.47
			4b < 0.27

In each cell 1 refers to periods 1–7; 2 to periods 8–14; 3 to periods 15–21, 4a to realized efficiency in periods 22–28 and 4b to the efficiency levels corrected for opportunity costs in periods 22–28; > (<) indicates that the efficiency of the treatment in the row is greater (smaller) than the efficiency of the treatment in the column. After the > (<) sign the p-value of the Wilcoxon rank test is displayed. The tests use independent average data per group as observations (for comparisons 1, 2 and 3 there are 8 pairwise data; for 4a and 4b there are 4 pairwise data).

Table B2. Test results revenue

	SDA	SimFPA	SeqFPA
SAA	1 < 0.21	1 < 0.33	1 < 0.01
	2 < 0.16	2 < 0.33	2 < 0.05
	3 < 0.12	3 < 0.26	3 < 0.07
	4 = 1.00	4 < 0.72	4 < 0.47
SDA	X	1 > 0.78	1 < 0.44
		2 > 1.00	2 < 0.78
		3 > 0.48	3 < 0.58
		4 < 0.72	4 < 0.14
SimFPA	X	X	1 < 0.12
			2 < 0.40
			3 < 0.33
			4 < 1.00

In each cell 1 refers to periods 1–7; 2 to periods 8–14; 3 to periods 15–21 and 4 to periods 22–28; > (<) indicates that the revenue of the treatment in the row is greater (smaller) than the revenue of the treatment in the column. After the > (<) sign the p-value of the Wilcoxon rank test is displayed. The tests use independent average data per group as observations (for comparisons 1, 2 and 3 there are 8 pairwise data; for 4 there are 4 pairwise data).

References

Armantier, Olivier, 2004. Does observation influence learning. Games and Economic Behavior 46, 221–239.

Bulow, Jeremy, Klemperer, Paul D., 1994. Rational frenzies and crashes. Journal of Political Economy 102 (1), 1–23.

Chakraborty, Archishman, Gupta, Nandini, Harbaugh, Rick, 2001. Best Foot Forward or Best for Last in a Sequential Auction, mimeo, Baruch College, CUNY.

Coase, Ronald, 1959. The federal communications commission. Journal of Law and Economics 2, 1–40.
Cox, James, Smith, Vernon L., Walker, James L., 1984. Theory and behavior of multiple unit discriminative auctions. Journal of Finance 39, 983–1010.
Cox, James C., Dinkin, Sam, Swarhout, James T., 2001. Endogenous entry and exit in common value auctions. Experimental Economics 4, 163–181.
Cramton, Peter, 2002. Spectrum Auctions, Handbook of Telecommunications Economics. In: Cave, Martin, Majumdar, Sumit, Vogelsang, IngoElsevier Science B.V., Amsterdam, pp. 605–639.
Demange, Gabrielle, Gale, David, Sotomayor, Marilda, 1986. Multi-item auctions. Journal of Political Economy 94, 863–872.
Goeree, Jacob K., Offerman, Theo, 2002. Efficiency in auctions with private and common values: an experimental study. American Economic Review 92 (3).
Goeree, Jacob K., Offerman, Theo, Sloof, Randolph, 2004. Demand Reduction and Preemptive Bidding in License Auctions, working paper, University of Amsterdam.
Harris, Milton, Raviv, Artur, 1981. Allocation mechanisms and the design of auctions. Econometrica 49, 1477–1499.
Hoppe, Heidrun, Jehiel, Philippe, Moldovanu, Benny, 2004. License Auctions and Market Structure, working paper, University of Bonn.
Jehiel, Philippe, Moldovanu, Benny, 2003. An Economic Perspective on Auctions, Working Paper, Mannheim University 36, 271–308.
Klemperer, Paul D., 1998. Auctions with almost common values. European Economic Review 42, 757–769.
Klemperer, Paul D., 2002. What really matters in auction design. Journal of Economic Perspectives 16, 169–189.
Leonard, Herman B., 1983. Elicitation of honest preferences for the assignment of individuals to positions. Journal of Political Economy 91, 461–479.
Maskin, Eric S., Riley, John G., 2000. Asymmetric auctions. Review of Economic Studies 67, 413–438.
Milgrom, Paul R., Weber, Robert J., 1982. A theory of auctions and competitive bidding. Econometrica 50, 1089–1121.
Olson, Mark A., Porter, David, 1994. An experimental examination into the design of centralized methods to solve the assignment problem with and without money. Economic Theory 4, 11–40.
Waehrer, Keith, Harstad, Ronald M., Rothkopf, Michael H., 1998. Auction form preferences of risk-averse bid takers. Rand Journal of Economics 29, 179–192.