ELSEVIER

Journal of Economic Behavior & Organization
Vol. 42 (2000) 323–346

JOURNAL OF
Economic Behavior
& Organization

www.elsevier.com/locate/econbase

The choice of cooperative strategies in sealed bid auctions

Anthony M. Kwasnica*

Department of Management Science & Information Systems, Pennsylvania State University, 339 Beam Business Administration Building, University Park, PA 16802, USA

Received 7 May 1998; received in revised form 2 August 1999; accepted 30 August 1999

Abstract

Isaac and Walker (1985) demonstrated that stable cooperative agreements are often formed in sealed bid auctions. By extending the experimental environment to multiple, simultaneous sealed bid auctions, it is possible to not only examine whether bidders collude in experimental auctions, but also to determine the strategies they choose in order to implement their collusive agreements. Bidders have a variety of strategies, as opposed to the single auction, that they may select from. In the experiments reported, bidders often choose cooperative strategies that are consistent with truthful revelation of their values. However, bidders occasionally select strategies that ignore individual incentive compatibility. An analysis of the choice of strategies in different informational and distributional environments as well as of the auction outcomes provides some insights into these unusual choices.

JEL classification: C92; D44

Keywords: Collusion; Auctions; Cooperation

1. Introduction

Experimental research has repeatedly shown that non-binding communication between participants can lead to choices which Pareto dominate those which are dictated by individual incentives. In prisoner's dilemma games, communication reduces the number of defect choices despite the fact that such a choice is a dominant strategy (Dawes et al., 1977). A similar result holds for the voluntary provision of public goods. The existence of non-binding preplay communication consistently leads to increased provision of the public good (Isaac

* Tel.: +1-814-863-0542; fax: +1-814-863-2381.
E-mail address: kwasnica@psu.edu (A.M. Kwasnica)

PII: S0167-2681(00)00092-5

et al., 1985; Isaac and Walker, 1988). In the context of oligopoly, firms are better able to coordinate choices and, thus, increase prices through communication (Friedman, 1967; Daughety and Forsythe, 1987; Binger et al., 1992; Brown-Kruse et al., 1993). Isaac et al. (1984), Isaac and Plott (1981), and Clauser and Plott (1993) demonstrate that institutional structure drastically effects the prospects for profitable communication in their comparison of the ability of conspirators to raise prices in posted-offer markets as opposed to their relative ineptitude in double auctions.

The story of bidder collusion in auctions is similar. While bidders can increase their total surplus by agreeing to limit bidding, they face an individual incentive to deviate from any such agreement in order to reap a one period windfall profit. In addition, even if we assume that bidders can overcome this incentive problem and compel all bidders to cooperate in each auction, the bidders will face an additional incentive problem: the presence of asymmetric information. In auctions, bidders must formulate a cooperative agreement when each agent only knows his own valuation(s) for the object(s) being sold. Therefore, it is reasonable to expect that any agreement (or decision rule) that bidders reach will be constrained by incentive compatibility. A decision rule is incentive compatible if it is in the interest of each bidder to truthfully report his valuation(s).[1] McAfee and McMillan (1992) demonstrate that if bidders are not allowed to make side payments, the only incentive compatible agreement is the random determination of the winning bidder. Any decision rule, other than competitive bidding, that uses information on bidders' valuations will not be incentive compatible since bidders will want to lie in order to increase their chance of winning the object. Isaac and Walker (1985) provide relatively good support for this approach in their experimental results. They show that, despite the individual incentive to defect from cooperative agreements, bidders usually manage to keep their bids near zero with the winner being chosen randomly.[2] However, Isaac and Walker suggest that bidders may be attempting to increase surplus by picking bidders with higher valuations with greater probability as evidenced by a higher than expected efficiency in some of the auctions.

If we assume that multiple units are being sold in a set of simultaneous first-price auctions, the incentive compatibility constraints on the choice of cooperative strategies are reduced.[3] Pesendorfer (1996) suggests a cooperative strategy, the ranking mechanism, which is incentive compatible and allows bidders to select a higher valued bidder with greater probability. Kwasnica (1997) extends this approach by showing that the ranking mechanism, along with other incentive compatible strategies, interim dominates the randomization procedure which is the only incentive compatible strategy in the single unit case. When multiple units are for sale, bidders are willing to tradeoff the chance of being selected as the single bidder in a market where they have a low valuation in return for a greater chance in a market where they have a higher valuation allowing their agreements to use more information than in the

[1] The Revelation Principle states that any Bayes–Nash equilibrium outcome can be replicated by an incentive compatible decision rule (Gibbard, 1973). Therefore, the incentive compatibility requirement is equivalent to assuming that, while the bidders have decided to cooperate, they formulate their cooperative strategy through some non-cooperative game.

[2] In their experiments, stable cooperative agreements were formed in 7 out of 10 experimental sessions.

[3] The use of simultaneous auctions is not entirely unfamiliar. The FCC auction for bandwidth is of a similar nature. Milgrom (1996) suggests using this sort of auction for the determination of Carrier of Last Resort privileges. Pesendorfer (1996) describes school milk markets which function in this manner.

single unit case. Therefore, simultaneous multiple unit auctions are an ideal environment for the closer examination of the choice of cooperative strategies. Whereas, in the single unit auction, any deviation from randomization would be viewed as a violation of incentive constraints, here a variety of incentive compatible strategies are available which provide differing payoffs. There are, of course, still strategies which violate incentive compatibility.[4] The increased number of strategies (bids) that each bidder takes in each experimental period lends itself well to a more systematic examination of exactly which strategies are being chosen by bidders.

With these advantages in mind, a series of simultaneous first-price auctions were conducted in order to shed light on three primary questions:

1. Do bidders form cooperative agreements in simultaneous first-price auctions?
2. If so, what types of strategies do they utilize?
3. What effect do these strategies have on the outcome of the auction?

The first question investigates the first incentive issue. Can bidders cooperate in auctions? As the preponderance of previous experimental results suggest, we find that the answer is an emphatic yes. Bidders can look beyond the instantaneous incentives to defect in order to reap long term gains just as they can in prisoner's dilemma and public good experiments. The next two questions look specifically at the types of agreements formulated. Do these agreements look like the predictions suggested by the imposition of an incentive compatibility constraint? Here the answers are less clear. In most experiments (7 out of 10), bidders find agreements that are very close to those predicted by the theory of collusion in auctions. However, at times, the choices are not consistent with the incentive compatibility constraint. Bidders choose strategies that are such that they *should* lie about their private information. However, given that they do not lie about their valuations, bidders actually fair better under these agreements.

2. Experimental design

In each experiment, five bidders participated in five simultaneous single unit first-price auctions, in which five objects were sold. An experimental design with the same number of bidders as objects was chosen for two reasons. First, we expect cooperative agreements to be more successful here (as opposed to a setting with fewer objects). Since our interest is primarily in the observation of cooperative strategies, such a design should maximize the number of observations. Second, when the number of bidders and objects is the same, bidders may utilize a relatively simple but less profitable strategy of assigning (ex ante) one bidder to each market.

In the first five periods of each session no communication was allowed. Then, in the next 12–17 periods subjects were allowed to communicate between each period.[5] In general, subjects were undergraduate students at the California Institute of Technology. However,

[4] It would be foolish to a priori rule out strategies which violate incentive compatibility given the preponderance of experimental evidence suggesting that subjects who communicate can, at times, overcome these limitations.

[5] The number of periods varied due to treatments completed at the end of some experiments. We do not report those results here but the interested reader can find them in Kwasnica (1997).

a few graduate students and staff members were participants. Each subject participated in only one experiment.[6]

Bidders were required to place a bid of at least one unit of experimental currency (francs) in each market.[7] This restriction ensured that subjects were unable to monitor adherence to collusive agreements via the sound of computer keys being hit indicating submission of bids. There were no other restrictions placed on bidding. If ties occurred in the highest bids, the computer software randomized between the high bidders to determine the winner.

2.1. Symmetry

In the symmetric environment, valuations for all five markets and bidders were drawn from the same distribution. Integer values between 1 and 1000 were drawn using the discrete uniform distribution. Under the assumption that bidders are risk neutral, the unique, symmetric Bayes–Nash equilibrium bid function is

$$b_{ij}(v_{ij}) = 0.8v_{ij} \quad \text{for all } i \text{ and } j.$$

Since the bid functions are symmetric and strictly monotonic, under the risk neutrality assumption, the auction is ex post efficient.[8] All bidders received new valuation draws after each auction period.

2.2. Asymmetry

In each market, one bidder's values were drawn from the stochastic dominant distribution, $F(v) = v^2/1000^2$, taking values between 1 and 1000 as well.[9] In that market, the other four bidders values were drawn from the discrete uniform distribution with values between 1 and 1000. More specifically, bidder 1 had his values drawn from $F(v)$ in market A, 2 in B, 3 in C, 4 in D, and 5 in E. The identity of that bidder was announced to all participants and fixed for all periods.

The asymmetric environment was used in order to give bidders a stronger incentive to use an inefficient but simple cooperative strategy: assign sole bidding rights to the bidder with the preferred distribution in each market. The expected efficiency of such a strategy is 80% (as opposed to 60% under symmetry) but there exist more sophisticated strategies which dominate it.[10]

When bidders are behaving non-cooperatively, the Bayes–Nash equilibrium bid function can be estimated numerically. Fig. 1 is a plot of the estimated bid functions for each market

---

[6] Instructions can be found at http://www.hss.caltech.edu/∼akwas/collusion_instructions.html. The simultaneous first-price auctions were implemented using auction software designed by Wes Boudeville and Dave Porter.

[7] The conversion rate of francs to US$ was either 250 or 500. Thus, a minimum bid of 1 franc was either US$ 0.004 or 0.002.

[8] An auction is ex post efficient if the winning bidder has the highest valuation for the object. If bidders have different levels of risk aversion, the auction may no longer be ex post efficient.

[9] The discrete analog to this distribution was actually used.

[10] We say one strategy dominates another if for all possible valuation draws all agents prefer that strategy.

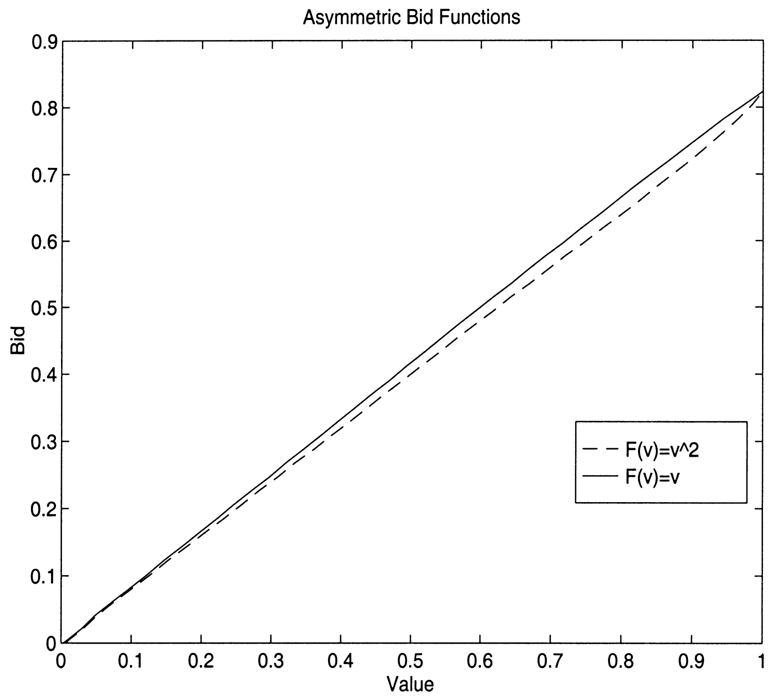

Fig. 1. Asymmetric bid functions.

when bidders have values drawn from the above distributions.[11] If bidder 1 has values drawn from the stochastically dominant distribution, $b_1(v) \leq b_i(v)$ for all other i and for all valuations. For all values, bidder 1 will think he is more likely to have the highest value. Therefore, he will shade his bid down slightly more than those with values drawn from the uniform distribution. This means that the bidder with the highest value will not necessarily be the highest bidder, and competitive bidding will not be ex post efficient. However, in this case, the expected efficiency of competitive bidding is extremely close to 100% (at 99.983%).

2.3. Communication

After the end of the fifth period it was announced that communication between bidders was to be allowed. In the following periods, bidders were allowed to communicate after having learned their valuations for that period. The following statement was handed out and read to subjects, who were then allowed to ask questions.

> 'Sometimes in previous experiments, participants have found it useful, when the opportunity arose, to communicate with one another. You are going to be allowed this opportunity while the computers are reset between periods.
>
> There will be some restrictions.

[11] BIDCOMP2, a program developed by John Riley, was used to estimate these bid functions.

You are free to discuss any aspects of the experiment (or the market) that you wish, except that:

- You may not discuss any quantitative aspects of the private information on your value sheets.
- You are not allowed to discuss side payments or to use physical threats.

Since there are still some restrictions on your communications with one another, an experimenter will monitor your discussion between periods. To make this easier, all discussions will be at this site.

Remember, after the computers have been reset between periods (and the next period has begun) there will be no discussion until after the end of the next period.

We will allow a maximum of 4 min in any one discussion session.'

In addition to these instructions, in every experiment except the first, subjects were also told that the number of rounds had been fixed. This announcement assured participants that lengthy communication would not reduce profits via a reduced number of periods. In most experiments, subjects had no problem understanding the limitations of their communication and only occasional reminders (or clarifications about the form of acceptable information) were required.

2.4. Information conditions

The *limited information* environment was the most restrictive information condition utilized by Isaac and Walker. The only information available to participants was the identity of the winning bidders and the prices they paid. A second, more limited, information condition not used by Isaac and Walker was the *zero information* condition which reported only the winning bids to the bidders. The identity of the winning bidder in each auction was unknown to everyone except the winner. Under the zero information condition, the participants could only determine who had placed winning bids through voluntary discussion. We would expect the increased difficulty in identifying and punishing deviant bidders to have made the zero information condition less conducive to cooperative behavior.

Ten experiments were completed with six experiments utilizing the symmetric environment and the other four using asymmetric valuation draws. Six of the experiments were conducted under the limited information setting. The final four experiments used the more limited zero information condition. A general summary of the experiments can be found in Table 1.

3. Experiment results

Subject earnings averaged US$ 33.75 across all experiments. No experimental session lasted longer than 2 h with the average length closer to 1.5 h. There did not appear to be significant variance of subject profits between and within periods. [12]

The behavior of the bidders in the first five periods of the auction, when communication was not allowed, was similar to previously observed results. Cox et al. (1988) found that

[12] However, there was significant variation of profits across experiments due to the choice of cooperative strategy.

Table 1
Experimental design

Experiment	Number of periods	Environment	Information condition
1	20	Symmetric	Limited
2	22	Symmetric	Limited
3	22	Symmetric	Limited
4	22	Symmetric	Limited
5	20	Asymmetric	Limited
6	20	Asymmetric	Limited
7	20	Asymmetric	Zero
8	18	Asymmetric	Zero
9	18	Symmetric	Zero
10	17	Symmetric	Zero

bidders often place bids above the risk neutral Nash equilibrium prediction. However, for extremely low valuations, where bidders have little chance of winning, bidders typically place extremely low bids (often 0). The estimation of a simple linear regression on the bids placed in the auctions with symmetric valuations demonstrates these results. Estimating the linear regression of $b_i = \beta_1 + \beta_2 v_i + \epsilon$ for each bidder should lead to estimates of $\hat{\beta}_1 = 0$ and $\hat{\beta}_2 = 0.8$ if bidders are playing the risk neutral Nash equilibrium. Cox, Smith, and Walker found that for many bidders $\hat{\beta}_1 < 0$ and $\hat{\beta}_2 > 0.8$. We found similar behavior in our experiments: 17 out of 30 (57%) subjects exhibited $\hat{\beta}_1 < 0$ and 16 out of 30 (53%) exhibited $\hat{\beta}_2 > 0.8$.

3.1. *Do bidders form cooperative agreements?*

The results of Isaac and Walker suggest that successful cooperative behavior should be expected here. One indication of the presence of collusive behavior would be a significant drop in bidding prices. The average bid in periods with communication drops to near zero. While the average bid in no communication periods was 428 francs, it was only 9.6 francs in communication periods. A reduction in bid levels is not necessarily an indicator of profitable collusive behavior. Isaac and Walker use an index of monopoly effectiveness (M) which is the proportion of maximum total possible surplus captured by the bidders:

$$M = \frac{\sum_{j=1}^{5} v_j^* - b_j^*}{\sum_{j=1}^{5} \max_i v_{ij}},$$

where v_j^* is the valuation of the winning bidder in market j and b_j^* is his bid. In our experiments, M increases from an average of 0.265 in no communication periods to 0.912 when communication is allowed. Bidders are capturing a larger proportion of the total surplus available. Perhaps the strongest evidence of successful cooperative behavior is that, despite a change in the conversion rate from 250 to 500 francs per US$, average bidder per period profits rose from US$ 0.93 to 1.51. In comparison, if the bidders were placing bids consistent with the risk neutral Nash equilibrium, they would have earned US$ 0.33 on average in the communication periods.

Remark 1. *When communication is allowed, under both environments and information conditions, collusive agreements are formed and are stable.*

In all 10 experiments, very few deviations from collusive agreements were evident. In early periods, bidders occasionally placed bids that were obviously not consistent with the collusive agreement. After the first two periods of communication, there were only three out of 129 periods in which bidders made notable deviations from the cooperative agreements. Notable deviations are classified as bids that are easily distinguished, by the other bidders, as not being consistent with the agreed upon strategy. For example, in an auction where bidders decided to choose a particular bidder as the single bidder in a particular market, a bid greater than one by another participant in that market would be considered a notable deviation. However, this does not mean that bidders were not making minor, less noticeable deviations from agreed upon strategies (see Section 3.2.1). In contrast to Isaac and Walker, where collusion occasionally broke down, there is no evidence of sustained deviations in our experiments. [13]

3.2. What types of strategies do bidders utilize?

Closer examination of the periods where communication was allowed reveals heterogeneity in the choice of cooperative strategies between experimental sessions. Two distinct strategies can be discerned from the data and observation of preplay communication. The first, and most common strategy, was the utilization of bid rotation. Bid rotation strategies are characterized by the selection of a bidder as the sole bidder in each auction. This bidder placed a low bid greater than 1 franc (typically 2–5 francs) while all other bidders bid 1 franc in the auction. A second strategy that was observed in the data was a linear bid reduction agreement. This strategy entails the agreement by all bidders to place bids which are linear transformations of their actual valuations. Since bidders are required to submit whole franc bids of at least 1 franc, linear bid reduction will, in general, lead to higher average bids than bid rotation. [14]

Remark 2. *In 7 out of 10 experiments, bidders used a bid rotation strategy. In experiments where bid rotation was not used, bidders used a linear bid reduction strategy.*

The easiest method for discerning these two different strategies was observation of preplay communication. In the seven experiments where bid rotation was used, bidders attempted to reach some resolution of who would bid in each market. However, in the three bid reduction schemes, they determined a level of bidding. The difference between these experiments can also be seen in the level of bids placed. In the seven rotation experiments, the average bid placed was 2.8 francs. In the linear bid reduction experiments, the average bid was 23 francs.

[13] The graphs of bidders' surplus in Appendix A demonstrate the consistency of the cooperative agreements.

[14] Bid rotation strategies lead to average bids that are close to 1 franc since all bidders except one are bidding 1 franc.

Table 2
Mean deviations from reduced bidding agreement

	Experiment		
	1	3	4
Mean deviation	−0.06481	−0.00774	0.167219
Standard error	0.131955	0.076929	0.286164
Observations	375	425	425

3.2.1. Linear bid reduction

In two of the linear bid reduction experiments, the bidders agreed to place bids that were 1% of redemption values.[15] In the other, bidders agreed to place bids that were 10% of valuations. Kwasnica (1997) shows that such agreements violate individual incentive compatibility. Only an agreement to bid 80% of their valuations would be incentive compatible.[16] Since their values are not ex post verifiable, bidders can increase their bids beyond either the 1 or 10% level without detection, and increase their probability of winning the object. Therefore, bidders would be expected to bid higher than their particular bid reduction agreement dictates. Fig. 2 shows the deviations from the agreed upon strategy. In all experiments, the vast majority of the bids are within 1 franc of that predicted by their particular collusive agreement. In all three experiments the null hypothesis that the mean deviation is equal to zero can be rejected at the 95% confidence level. Surprisingly, however, in two of the experiments, mean deviations are significantly below zero implying bidders were actually bidding slightly below the agreement. Only in one experiment were deviations significantly above zero (see Table 2). As shown in Fig. 2, most of these deviations are small. The Bayes–Nash equilibrium of these linear bid reduction agreements predicts that bidders would misreport their values much more than was observed.[17]

Remark 3. *Bidders choose linear bid reduction strategies that are not incentive compatible. However, bidders rarely significantly misreport their redemption values.*

The fact that these linear bid reduction agreements are replicated and appear to be relatively stable creates problems for the theory. Why are these bidders not shading their bids up in two experiments? Bidders seem to ignore their individual incentives despite the fact that detection of placing higher bids is very difficult. Section 3.3 provides one possible explanation for the choice of this strategy.

3.2.2. Bid rotation

The majority of the experiment sessions (7 out of 10) lead to bidding strategies that were classified as bid rotation agreements. There are many different mechanisms that are incentive compatible and look like bid rotation outcomes. The choice of mechanism by the group has significant implications for efficiency and thus the percentage of maximum total

[15] Bidders in Experiment 4 quickly switched from a 10 to a 1% rule after two periods.
[16] Kwasnica (1997) shows that all reduced bidding strategies which are incentive compatible must yield the same expected utility as the outcome of the (non-collusive) Bayes–Nash equilibrium of the auction.
[17] Many of the deviations can be accounted for as rounding errors by the bidders.

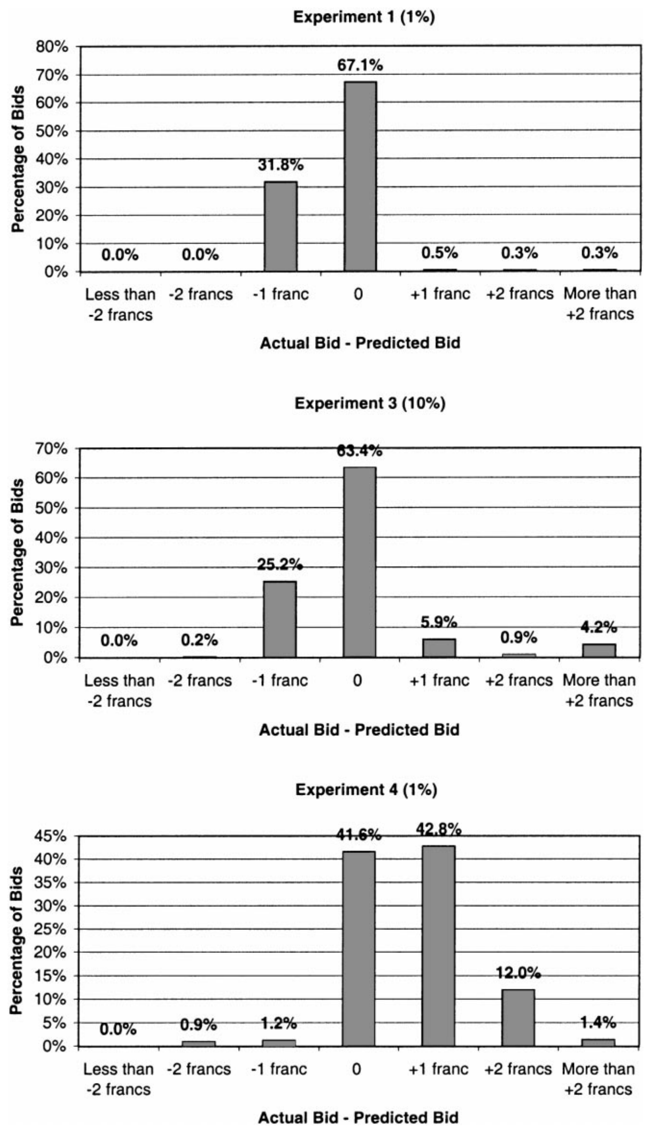

Fig. 2. Deviations from linear bid reduction.

surplus captured by the bidders. Four behavioral strategies that can lead to outcomes similar to those observed in these seven experiments are: [18]

1. Ranking mechanism (R),
2. Serial dictator mechanism (SD),

[18] This list of strategies is by no means exhaustive.

		Bidder				
		1	**2**	**3**	**4**	**5**
Market	**A**	1	3	3	5	5
	B	2	5	2	3	2
	C	3	2	1	1	1
	D	4	1	4	2	3
	E	5	4	5	4	4

Fig. 3. An example.

3. Random assignment mechanism (A), and
4. Perfect information (P).

The ranking (R), serial dictator (SD), and random assignment (A) mechanisms are incentive compatible mechanisms discussed in Kwasnica (1997). The ranking mechanism was originally suggested by Pesendorfer. If bidders were using this strategy they would report the relative ranking of their valuations. Then, the bidder with the highest rank would be selected as the sole bidder in each auction. If multiple bidders have the same rank, the sole bidder would be determined randomly between them. Under the serial dictator mechanism, an order of bidders would be selected (presumably randomly) and then the bidders would proceed in turn, each choosing an auction in which he wanted to be the sole bidder. The different choices that these two strategies lead to are best illustrated by an example.

In the example described by Fig. 3,[19] if the random draw of dictators yields the order $(1, 2, 3, 4, 5)$, then bidder 1 would select first and choose market A, 2 would select market D, 3 would select C, 4 would select B, and 5 would have no choice but to select market E. However, the ranking mechanism would select bidder 1 as the sole bidder in market A, either 1, 3, or 5 in market B, either 3, 4, or 5 in market C, 2 in market D, and 2, 4, or 5 in market E. The random assignment mechanism (A) is the extension of random choice of the winner in the single unit auction to multiple unit auctions. The sole bidder is chosen randomly in each market, or, equivalently, each bidder is assigned a single auction in which he is the sole bidder throughout the experiment. Kwasnica (1997) demonstrates that this sort of mechanism is dominated (in the interim stage) by both the ranking and serial dictator strategies. Finally, the perfect information (P) strategy describes the possibility that bidders may perfectly collude by somehow determining the bidder with the highest valuation in each market.[20] The objective is to determine which of these possible mechanisms was most likely utilized in each of these experiments. Three techniques that shed light on the choice of a strategy by bidders are:

1. Observation of preplay discussion,
2. Comparison of expected efficiencies with observed efficiencies, and
3. Comparison of predicted market division with observed choices.

[19] The numbers in the figure represent each bidder's rank of the markets with one being the highest.

[20] This is a highly unexpected outcome given the limitations on bidder communication. In fact, Kwasnica (1997) demonstrates that, when side payments are not allowed, there is no incentive compatible strategy that can obtain this outcome. However, it is still possible that this may be the *best* predictor of group behavior.

Table 3
Expected efficiencies

Behavioral strategy	Expected efficiency	
	Symmetry (%)	Asymmetry (%)
R	90.60	92.12
SD	85.20	86.63
A	60.00	80.00
P	100.00	100.00

Discussion: While observing bidder discussion is not a rigorous test for the predominance of one model over the other, simply listening to the conversations of the bidders can provide a great deal of insight into the intentions of the bidders. Bidder discussion was typically closer to the ranking mechanism than to the serial dictator mechanism. In most cases, bidders would begin their discussion by naming what they wanted first (their highest rank). If there was no conflict, discussion ended. If there was disagreement, those who had chosen conflicting markets would attempt to reach a compromise by naming their next best market. It is easy to see that such an iterative procedure leads to outcomes predicted by the ranking mechanism under the restriction that no bidder be chosen more than once. If the group discussion was consistent with the serial dictator mechanism, once a bidder had named a market in which he wished to bid, no other bidder could pick that market. Typically, conversation between bidders did not take this form. [21]

Efficiencies: An auction is (ex post) efficient if for each object the high bidder is the bidder with the highest valuation. Efficiency is calculated as the percentage of maximum possible surplus obtained in the auction:

$$\text{Efficiency} = \frac{\sum_{j=1}^{5} v_j^*}{\sum_{j=1}^{5} \max_i v_{ij}}.$$

Expected efficiency is calculated as the average ex post efficiency achieved over all possible outcomes. The expected efficiencies for each of the behavioral strategies in the symmetric and asymmetric environments are given in Table 3. If a group is utilizing a particular mechanism, the average of the observed efficiencies should converge to these efficiencies. The null hypothesis that the mean efficiency for each experiment was not different than the 90.60%, for the symmetric environment, and 92.12%, for the asymmetric environment, predicted by the ranking mechanism cannot be rejected at a 95% level of confidence for all seven experiments (see Table 4). However, in five of the seven experiments, the null hypothesis that the mean efficiency was equal to that predicted by the serial dictator mechanism (85.20 and 86.63%) can be rejected at a 95% level of confidence. [22] Obviously, the observed efficiencies are also significantly different from the 60 and 80% predicted

[21] However, it is possible that there may have been some first-mover advantage; the bidder who made his announcement of preferred markets first got his favored market more often. Since the order of discussion was not recorded, this factor cannot be analyzed for these experiments.
[22] Comparison of the mean bidder surplus yields similar results since bids placed are close to zero.

Table 4
Mean efficiencies — rotation

	Experiment						
	2	5	6	7	8	9	10
Mean Efficiency (%)	90.52	92.70	90.48	92.77	87.94	89.94	90.69
Standard error	0.0466	0.0545	0.0985	0.0755	0.1100	0.0729	0.0841
Observations	17	15	15	15	13	13	12

by an assignment mechanism. The perfect information model can also be rejected under this test in all seven experiments. A comparison of observed outcomes favors the ranking mechanism as the best description of behavior in each of the seven experiments.

Comparing the choices: While the first two methods provide some support for the ranking mechanism, they are somewhat unsatisfying in their approach. Analysis of discussion is purely ad hoc and relies upon the judgement of the experimenter who observed the experimental session. Comparison of mean observations utilizes outcomes rather than choices.

A more rigorous test would actually involve comparing the choices of the bidders with the choices predicted by each model. Initial examination of choices in each particular experiment indicates that the ranking mechanism is a good predictor of choices; 87% of all observed choices are consistent with the ranking mechanism. However, other mechanisms also correlate well with the observed choices. Thus, the likelihood based classification procedure of El-Gamal and Grether (1995) is used to provide a more rigorous statistical comparison of all the proposed models. Let $C_t = \{(c_1, c_2, c_3, c_4, c_5) | c_i \in \mathbb{Z}, 1 \le c_i \le 5,\ \ i = 1, 2, 3, 4, 5\}$ be the class of behavioral rules for each period such that each bidder is selected as the sole bidder in a particular market. For example, $c_1 = 2$ indicates that bidder 2 was selected as the sole bidder in market A. Each model predicts a subset $B_t \subset C_t$ and $B = B_1 \times B_2 \times \cdots \times B_{p_s}$ where p_s is the number of periods completed in an experiment. Each experimental session is treated as a single subject, s, and it is assumed that each 'subject' chooses exactly one behavioral strategy. The error probability, ϵ, is assumed to be the same for all individuals, experimental sessions' and choices. The choice by individual i in period t for a particular experimental session is denoted by a_{ti}. Then, for all B, let

$$x^s_{B,ti} = \begin{cases} 1 & \text{if } a_{ti} \in B_t \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

Let

$$X^s_B = \sum_{i=1}^{n} \sum_{t=1}^{p_s} x^s_{B,ti} \tag{2}$$

be the total number of choices predicted correctly for a particular session. The likelihood can be found to be

$$f^{B,\epsilon}(x^s) = \left(1 - \frac{\epsilon}{3}\right)^{X_B^s} \left(\frac{2\epsilon}{3}\right)^{p_s n - X_B^s} \tag{3}$$

for each behavioral strategy.[23] Under the assumption that participants in all S experiments are using the same mechanism, the maximum likelihood estimate is given by

$$(\hat{B}, \hat{\epsilon}) = \arg\max_{B,\epsilon} \prod_{s=1}^{S} f^{B,\epsilon}(x^s). \tag{4}$$

The algorithm suggested by El-Gamal and Grether is used in order to estimate the maximum likelihood estimate for any set of k behavioral strategies. Then, using a penalty function given by

$$g(k) = k\ln(4) + k\ln(3) + S\ln(k), \tag{5}$$

k is chosen to maximize the information criterion,

$$\mathrm{IC}(k) = \ln\left(\prod_{s=1}^{S} \max_{h\in\{1,\ldots,k\}} f^{\hat{B},\hat{\epsilon}}(x^s)\right) - g(k). \tag{6}$$

Using this technique, we can test the ability of the four possible mechanisms to explain the observed choices by each experimental session. The choices of the ranking mechanism (R) are easily characterized by saying that an error was made in a particular market if the bidder chosen was not the individual with the highest rank in that market. Unfortunately, the serial dictator mechanism (SD) cannot be characterized as easily. Each possible permutation of the five bidders can potentially lead to a different choice of market assignment predicted by the mechanism. Nearly any observed choice can be predicted by the mechanism. For any particular experiment the number of possible combinations of choices across periods is 120^{p_s} (which is 8.92×10^{24} in the experiment with the fewest periods). The choices predicted by the SD mechanism can be limited to a smaller set by assuming that each experimental group agrees to rotate the order of selection in each period. Thus, if the order of choosing was $1, 2, 3, 4, 5$ in period t then it would be $2, 3, 4, 5, 1$ in period $t + 1$. This limits the number of combinations predicted by the serial dictator mechanism to a more manageable 120 combinations. While limiting the serial dictator mechanism in this manner does, a priori, make it less likely that it will be classified as the best fitting model, it is reasonable to assume that no individual bidder would approve of any combination that did not evenly spread out the right to pick early in the order since early picking leads to higher individual surplus. The assignment mechanism (A) assumes that each bidder is selected as the sole bidder in his favored market when distributions are not symmetric. Thus, bidder 1 is assumed to always be the sole bidder in market A, bidder 2 in B, bidder 3 in C, bidder 4 in D,

[23] It is assumed here that, if a bidder made an error, he chose the correct strategy with probability one third and another strategy with probability two thirds. In reality, a bidder could have a choice of between five (if he happens to be choosing first or if there is little conflict) and one (if he is choosing last or there is a great deal of conflict) markets. Since, on average, he will have a choice of three markets, $(1/3, 2/3)$ is selected as an approximation.

Table 5
Percentage of choices explained by models — individual experiments

	Experiment						
	2 (%)	5 (%)	6 (%)	7 (%)	8 (%)	9 (%)	10 (%)
R	85.88	90.67	85.33	93.33	78.46	89.23	85.00
SD	84.71	65.33	62.67	64.00	86.15	86.15	86.67
A	29.41	49.33	40.00	42.67	33.85	23.08	30.00
P	56.47	58.67	61.33	60.00	50.77	53.85	60.00

Table 6
Estimated models

No. of models	Rule(s) chosen	No. classified	$\hat{\epsilon}$	$g(k)$	IC
1	R	435	0.39	2.485	−203.101
2	R, SD	333, 108	0.354	9.822	−197.893
3	R, SD, *	333, 108, 0	0.354	15.145	−203.216
4	R, SD,*,*	333, 108, 0, 0	0.354	19.644	−207.715

and bidder 5 in E. [24] Finally, the perfect information model (P) represents the choices that would be made if the bidders were able to actually aggregate their information perfectly. The bidder with the highest value is picked in each market.

Table 5 presents the data for each experiment. In all experiments, the ranking and serial dictator mechanisms better explain the data than either assignment or perfect information. Table 6 reports the results of the maximization of the information criterion in order to determine the optimal number of rules to choose. Using two rules best explains the choices observed in the seven experiments. In experiments 2, 5, 6, 7 and 9, the ranking mechanism is the behavioral strategy which best fits the experimental data. However, the serial dictator mechanism significantly adds to the explanatory power of the model in experiments 8 and 10. Using this classification procedure, it is possible to rule out the random assignment model of collusive behavior. Also, bidders were apparently unable to perfectly aggregate information.

Remark 4. *The ranking mechanism is the best description of behavior in the rotation scheme experiments. However, the serial dictator mechanism cannot be ruled out in some experiments.*

The combination of these three methods of determining which bid rotation scheme was used gives strong evidence in favor of the ranking mechanism. The serial dictator mechanism, however, still appears to be a strategy that is used occasionally by groups in this

[24] While this is an assumption based upon the optimal, ex ante strategy in the asymmetric environment (assign bidding rights to the bidder with the stochastically dominant distribution), any other assignment of bidders to the same single market throughout the experiment performs just as poorly.

Table 7
The effect of treatments

Values	Information	
	Limited	Zero
Symmetric	3 Reduced, 1 Rotation	2 Rotation
Asymmetric	2 Rotation	2 Rotation

setting, especially in Experiment 8, in which both the observed efficiency and the choices of markets correlate well with the serial dictator mechanism.

Remark 5. *Linear bid reduction mechanisms are only observed under the limited information and symmetric environments.*

All three instances of utilization of bid reduction strategies were in experiments where bidders had uniform valuation draws in all five markets and were informed of the identity of the winning bidders (Table 7). While Isaac and Walker found no significant patterns between collusive agreements and their two information conditions of full information and limited information, this result demonstrates that information matters. [25] While it may not be significant in determining whether bidders collude, it does appear to affect their choice of strategy. Bidders seem to be less willing to select a strategy which violates incentive constraints when they have less ex post information. While this result is consistent with numerous previous experimental studies which have shown that the ability to monitor helps sustain cooperation (Kagel and Roth, 1995, pp. 403–411), it is unique in that it demonstrates a switch from one cooperative strategy to another. The utilization of a less cooperative strategy in the asymmetric environment also has some precedence. Isaac and Walker (1988) found that asymmetries in public goods experiments tended to decrease the level of voluntary contributions. While a complete breakdown of cooperation is never evident here, this result suggests that bidders' choice of strategies is subtlely effected by the environment.

3.3. What effect do different strategies have on the outcome of the auction?

The choice of cooperative strategies can drastically affect the results of the auction. The differences between mechanisms can best be seen by examining the efficiency of the auction and the amount of surplus accruing to the bidders.

3.3.1. Efficiency

Despite the apparent problems with enforceability, linear bid reduction agreements have advantages from a social welfare standpoint. In the three experiments that exhibited these collusive agreements, average efficiencies were 99.26, 99.38, and 98.36%. A rank sum test shows that the mean efficiency for these experiments is greater than the mean efficiency of experiments where bidders used rotation schemes. [26]

[25] The full information condition was a less restrictive environment that reported all the bids placed in the auction.
[26] The Wilcoxon–Mann–Whitney test with correction for ties yielded $z = 8.267$ which is greater than any reasonable critical value of the standard normal distribution.

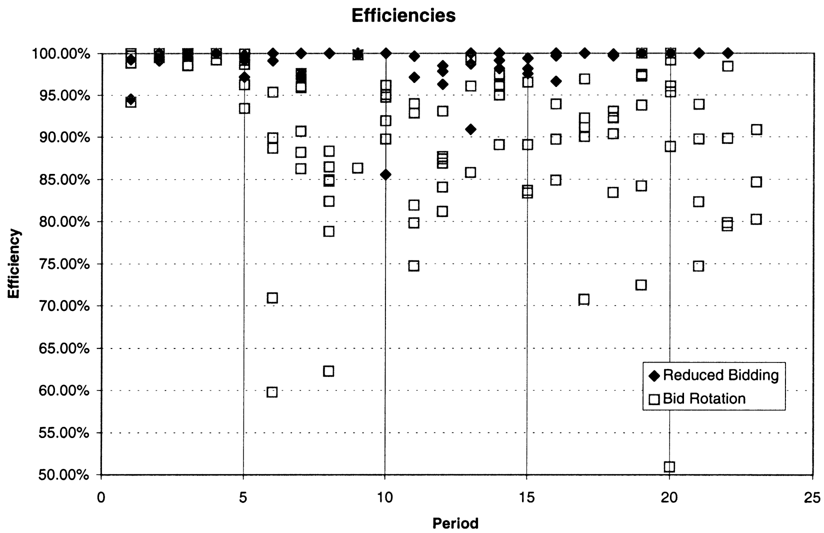

Fig. 4. Reduced bidding vs. rotation efficiencies.

Remark 6. *Linear bid reduction yields higher average efficiency than bid rotation.*

This result is due to the ability of bid reduction to select the highest valued bidder (assuming people do not deviate from the agreement). Fig. 4 shows the efficiencies for the experiments where reduced bidding was observed as opposed to the other seven experiments. Efficiency is also fairly stable under the linear bid reduction agreements. When bidders are using rotation schemes, efficiency varies significantly due to the imprecision of the ranks. However, the bid reduction agreements consistently yield efficiencies near 100%. The variance of the observed efficiencies for the seven non-reduced bidding experiments was always higher than the variance for the three linear bid reduction experiments.

3.3.2. Bidder surplus

The overall level of profitability for the bidders is best described by the index of monopoly effectiveness which reports the proportion of total possible surplus captured by the bidders. In experiments where bidders used rotation schemes, the average M was 0.898. The two 1% bid reduction experiments yielded an average effectiveness of 0.970, and the 10% bid reduction experiment averaged $M = 0.885$.

Remark 7. *The index of monopoly effectiveness is highest for the bidders under the 1% bid reduction rule.*

The 1% bid reduction agreement was the most successful (profitable) collusive agreement. This result highlights the apparent tradeoffs between these strategies. If bidders do not lie about their values, linear bid reduction yields a much higher efficiency than rotation schemes. This increase in the size of the pie more than accounts for the increased level of bids required by a 1% agreement. The 10% agreement, on the other hand, entails too high a level of bidding to actually increase profitability.

These two conclusions are the best argument in favor of a linear bid reduction mechanism. Bidders select reduced bidding because it is more profitable than rotation mechanisms, despite the fact that it is not consistent with individual incentives. It appears that something in the nature of communication in the group decision making process allowed the bidders to ignore this problem.

4. Conclusion

Given the previous experimental results, the fact that bidders colluded in these auctions is not very surprising. A more interesting result is how the bidders decided to implement cooperation. In most experiments (7 out of 10), bidders chose to follow strategies that were remarkably similar to those predicted by the imposition of incentive compatibility constraints. However, the results of three experiments suggest that cooperative behavior is not always bound by these constraints. In these three experiments, bidders selected a mechanism, linear bid reduction, that violated incentive compatibility. The current state of the art in theoretical research does not explain these results. What leads these bidders to ignore their individual incentives? The experimental results provide a suggestion. These strategies lead to outcomes that Pareto dominate the outcomes of incentive compatible mechanisms. When bidders are bidding truthfully, under the reduced bidding agreement, the auction is nearly efficient which translates to higher levels of bidder surplus. This choice is consistent with previous results that found that participants select behavior which although not consistent with individual incentives, leads to a higher return.

Why didn't we observe linear bid reduction in all 10 experiments? It appears that changes in informational and distributional conditions play a vital role in determining whether participants will feel bound by individual incentive compatibility. In an environment with symmetry and limited information, bidders select strategies that are not incentive compatible. However, increased asymmetry or decreased information leads bidders to look toward incentive compatibility in the choice of a mechanism. The open question, therefore, is whether there is a theory that would allow for the inclusion of these factors in determining what sort of choice a group will make.

Acknowledgements

I thank John Ledyard, Simon Wilkie, Dave Porter, Julie Scott and an anonymous referee for their comments and patience.

Appendix A. Bidder and seller surplus

Figs. 5–14 demonstrate the consistency of the cooperative agreements.

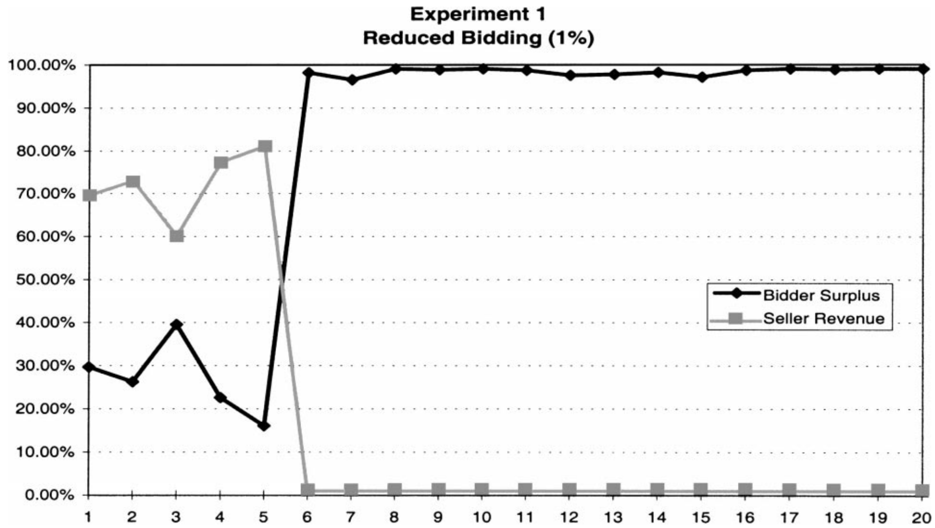

Fig. 5. Experiment 1.

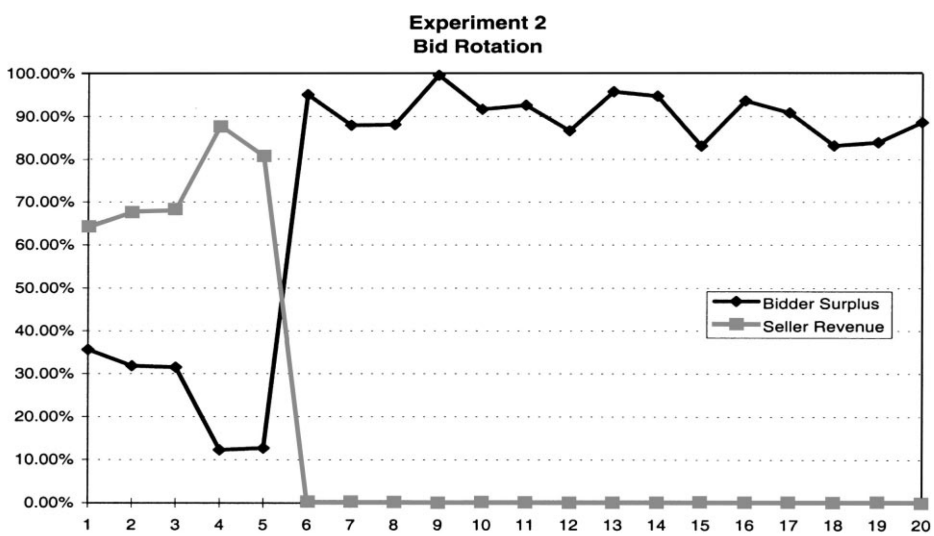

Fig. 6. Experiment 2.

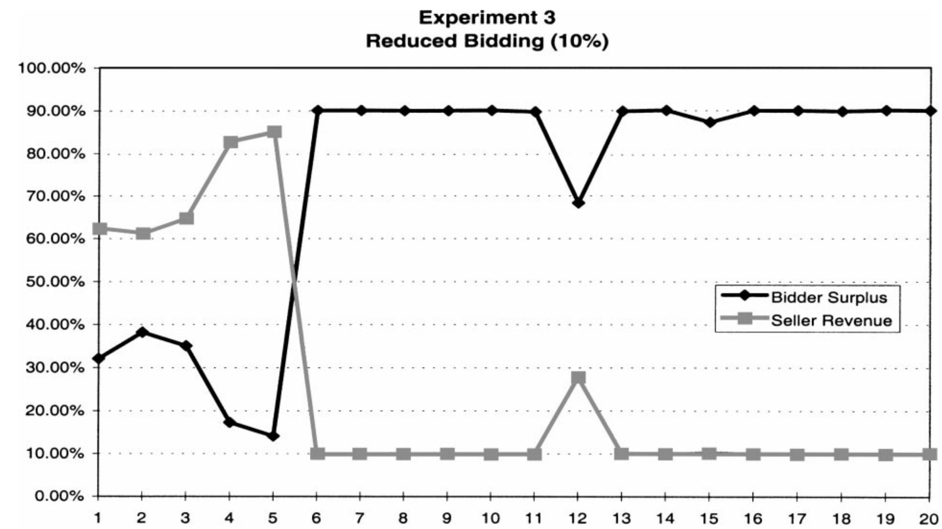

Fig. 7. Experiment 3.

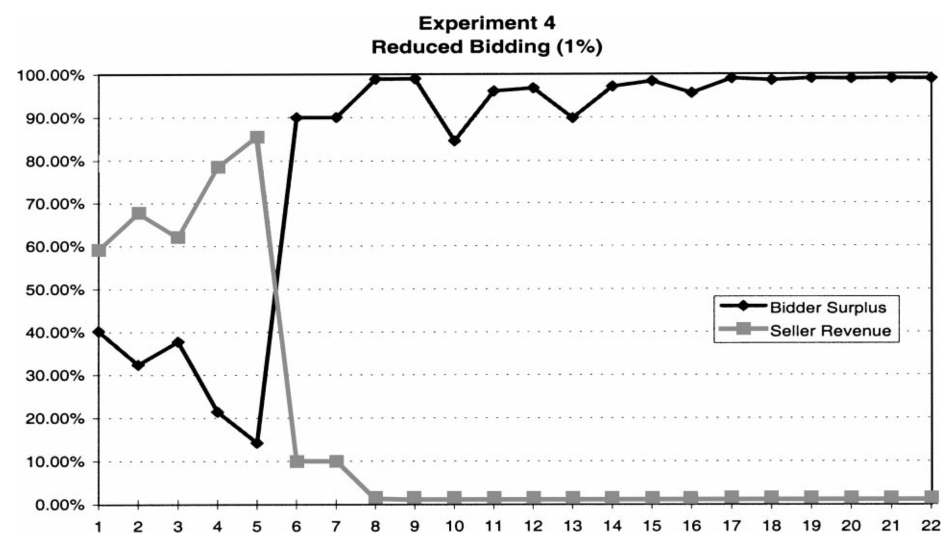

Fig. 8. Experiment 4.

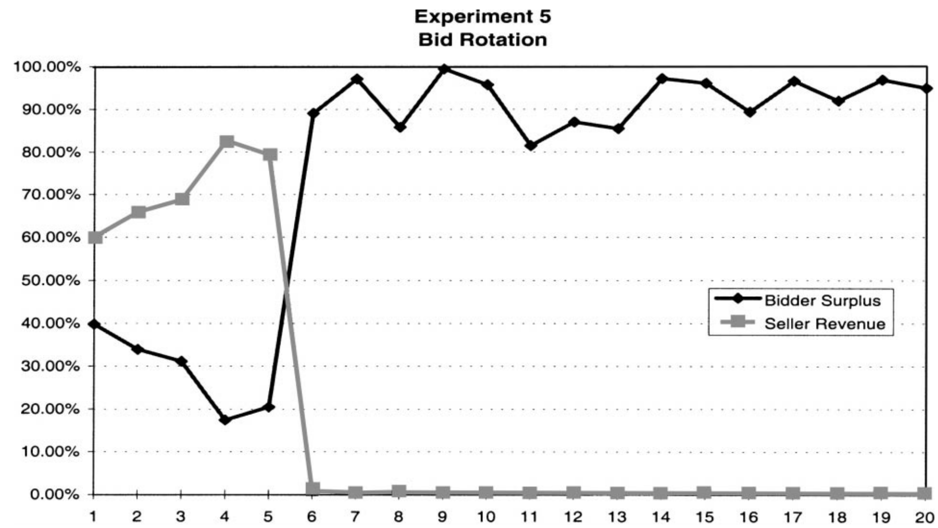

Fig. 9. Experiment 5.

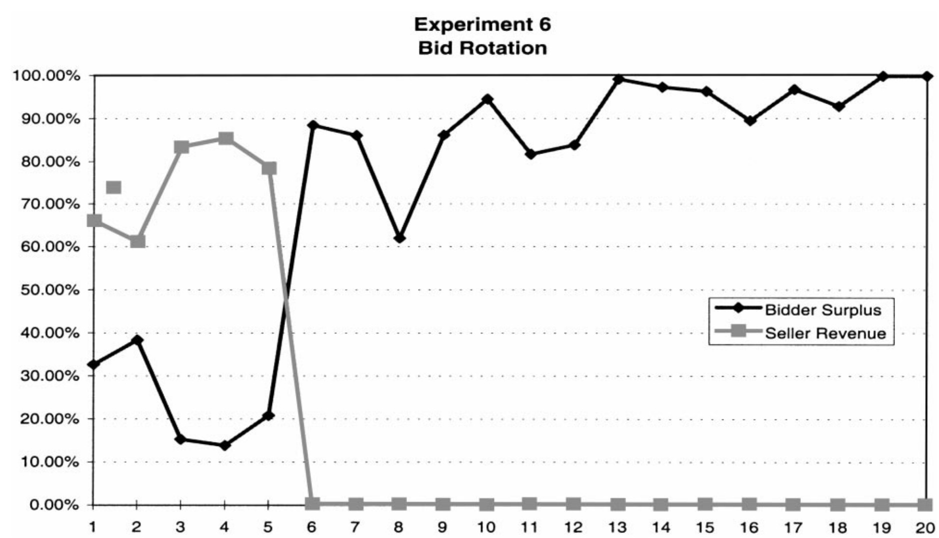

Fig. 10. Experiment 6.

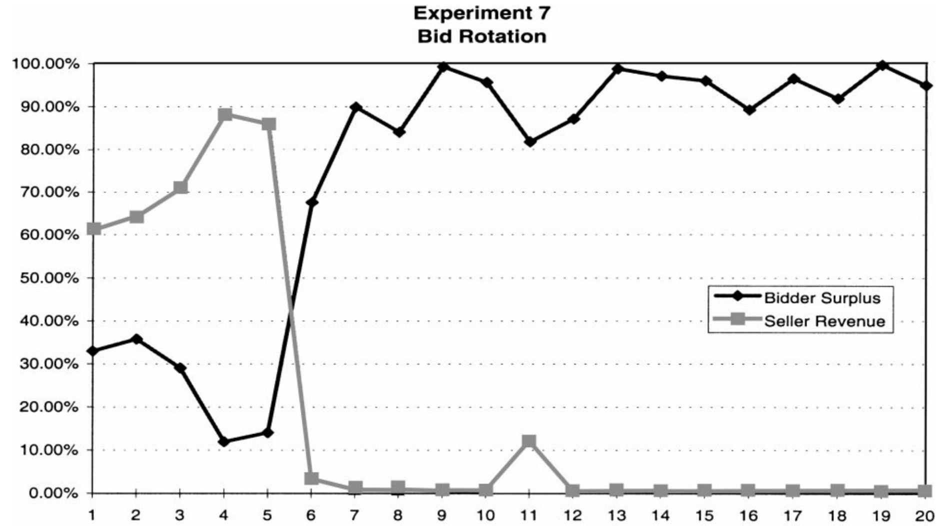

Fig. 11. Experiment 7.

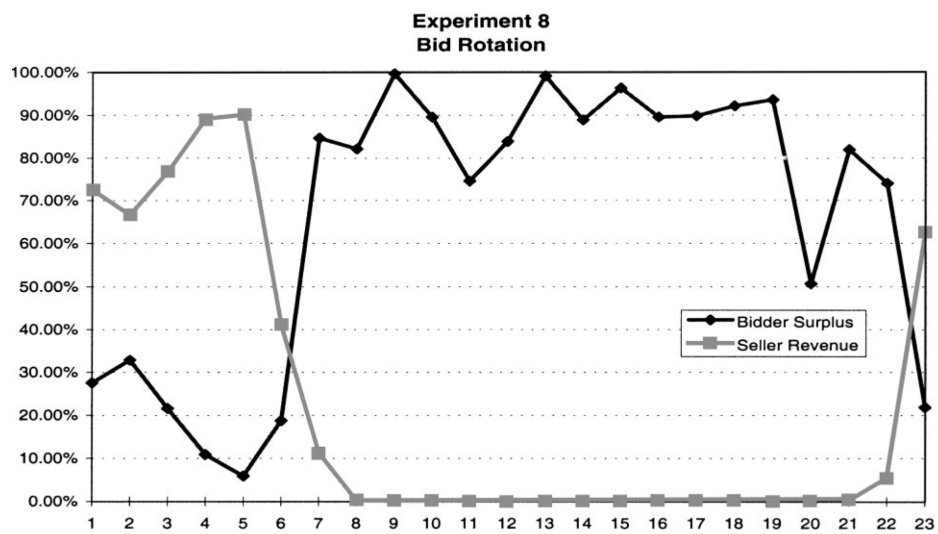

Fig. 12. Experiment 8.

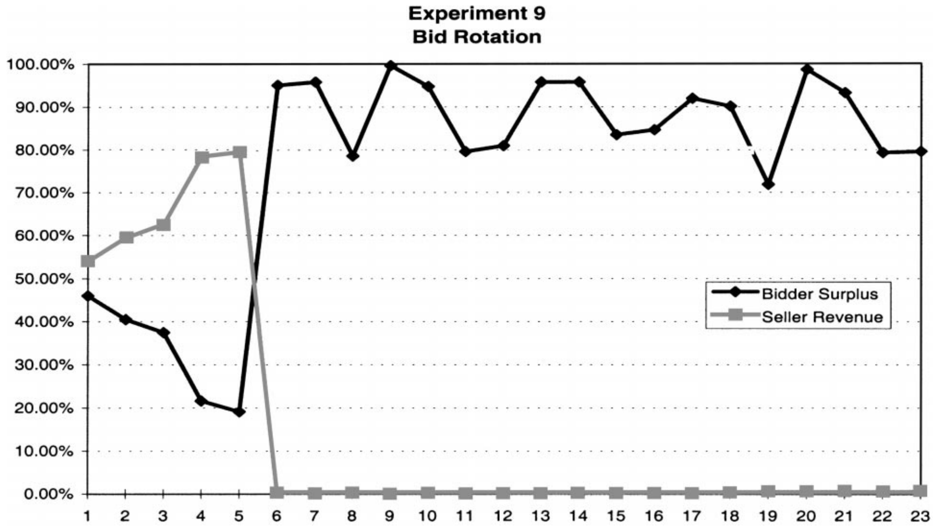

Fig. 13. Experiment 9.

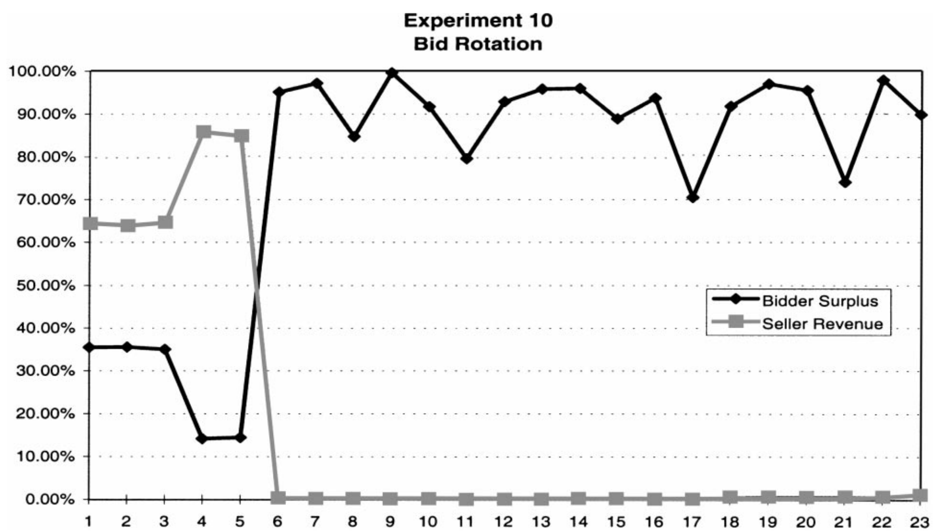

Fig. 14. Experiment 10.

References

Binger, B.R., Hoffman, E., Libecap, G.D., Shachat, K.M., 1992. An Experimetric Study of the Cournot Model. Discussion Paper 92-13, University of Arizona.

Brown-Kruse, J., Cronshaw, M.B., Schenk, D.J., 1993. Theory and experiments on spatial competition. Economic Inquiry 31, 139–165.

Clauser, L., Plott, C., 1993. On the anatomy of the nonfacilitating features of the double auction institution in conspiratorial markets. In: Friedman, D., Rust, J. (Eds.), The Double Auction Market. SFI Studies in the Sciences of Complexity, Addison-Wesley, Reading, MA (Chapter 12).

Cox, J.C., Smith, V.L., Walker, J.M., 1988. Theory and individual behavior of first-price auctions. Journal of Risk and Uncertainty 1, 61–99.

Daughety, A.F., Forsythe, R., 1987. Industrywide regulation and the formation of reputations: a laboratory analysis. In: Bailey, E. (Ed.), Public Regulation: New Perspectives on Institutions and Policies. MIT Press, Cambridge, MA.

Dawes, R.M., McTavish, J., Shaklee, H., 1977. Behavior, communication and assumptions about other people's behavior in a commons dilemma situation. Journal of Personality and Social Psychology 35, 1–11.

El-Gamal, M.A., Grether, D., 1995. Are people Bayesian? Uncovering behavioral strategies. Journal of the American Statistical Association 90 (432), 1137–1145.

Friedman, J.W., 1967. An experimental study of cooperative duopoly. Econometrica 35 (3), 379–397.

Gibbard, A., 1973. Manipulation of voting schemes: a general result. Econometrica 41, 587–602.

Isaac, R.M., McCue, K.F., Plott, C.R., 1985. Public goods provision in an experimental environment. Journal of Public Economics 26, 51–74.

Isaac, R.M., Plott, C.R., 1981. The oppurtunity for conspiracy in restraint of trade. Journal of Economic Behavior and Organization 3, 1–30.

Isaac, R.M., Ramey, V., Williams, A.W., 1984. The effects of market organization on conspiracies in the restraint of trade. Journal of Economic Behavior and Organization 5, 191–222.

Isaac, R.M., Walker, J.M., 1985. Information and conspiracy in sealed bid auctions. Journal of Economic Behavior and Organization 6, 139–159.

Isaac, R.M., Walker, J.M., 1988. Communication and free-riding behavior: the voluntary contribution mechanism. Economic Inquiry 26, 585–608.

Kagel, J.H., Roth, A.E. (Eds.), 1995. The Handbook of Experimental Economics. Princeton University Press, Princeton, NJ.

Kwasnica, A.M., 1997. Collusion in Multiple Object Simultaneous Auctions: Theory and Experiments. Social Science Working Paper 1010, California Institute of Technology, Pasadena, CA.

McAfee, R.P., McMillan, J., 1992. Bidding rings. American Economic Review 92 (3), 579–599.

Milgrom, P.R., 1996. Statement of Paul R. Milgrom: Attached to GTE's Comments in Response to Questions. CC Docket 96-45.

Pesendorfer, M., 1996. A Study of Collusion in First-Price Auctions. Submitted for publication.