

Combinatorial Auctions in the Information Age: An Experimental Study*

John Morgan[†]
University of California at Berkeley

June 2002

Abstract

In private values settings, the Vickrey-Clarke-Groves (VCG) mechanism leads to efficient auction outcomes, while the theoretical properties of the Simultaneous Ascending (SA) auction are not well understood. This leads us to compare the properties of an SA and a VCG auction in an experimental setting with private values for multiple objects having complementarities. Statistically, we find little to distinguish the two auctions with both auction forms achieving more than 98% efficiency and extracting roughly 95% of the available surplus. Finally, in contrast to experimental results in single object VCG settings, the theoretical prediction of demand revelation in the multiple object VCG auction is largely supported in our experiments.

JEL Classification Nos.: D44, L50

Keywords: Auctions, Experiments

*This paper has benefited greatly from the comments provided by Rachel Croson, Vijay Krishna, David Lucking-Reiley, Martin Sefton, and Tim Van Zandt, and Robert Willig. The author thanks Tanjim Hossain for his superlative research assistance. He thanks Princeton University for their financial support of this research. The author also gratefully acknowledges support from the National Science Foundation provided under grant SBR 9618648.

[†]All correspondence should be addressed to: John Morgan, Haas School of Business, 685 Faculty Building #1900, Berkeley, CA 94720. e-mail: morgan@haas.berkeley.edu

1 Introduction

The rapid decline in the cost of and gains in the efficiency of information technology over the last two decades continue to change the landscape in which items are bought and sold. Auctions, in particular, have gained much greater prominence as a means of determining prices at which goods are bought and sold, largely as a result of improved communication and information processing capabilities of the personal computer and the Internet. At a consumer level, the auction site eBay has transformed the market for buying and selling collectibles as well as a host of other products. EBay, with a market capitalization of over \$17 billion and sales of \$0.75 billion per year continues to grow rapidly (132% increase in net income year on year in 2001) despite the general slowdown in online markets. As Lucking-Reiley (2000) points out, the ubiquity of eBay has fundamentally changed the scope for using auctions as a means of allocating items by dramatically expanding the potential market for an item. On the consumer side, for instance, the collectibles market has expanded enormously in scale and scope as a result of the online auctions. As Lucking-Reiley puts it, “an item which might have been relegated to the trash heap in Shreveport, for lack of local interest, can now find its way to an enthusiastic collector in Boise.” Nor is eBay even close to being the largest

source of auction transactions—it is dwarfed in sales volume by business to business and business to government transactions conducted via auction.

It is not just in the expansion of the set of potential buyers and sellers that information technology has changed the nature of auction. The complexity of feasible auction mechanisms has also changed considerably owing to huge gains in processing power. Nowhere is this better illustrated than in the construction of the auction mechanism for the sale of bandwidth by the US and other governments in recent years. Governments would hire auction theorists for advice about how to auction licenses and, remarkably, the somewhat complicated schemes that the theorists came up with were incorporated in practice.

In view of these changes in the landscape in which auctions are conducted, it seems appropriate to re-evaluate the practicality of certain auction mechanisms which were earlier ruled out or discarded owing to their perceived complexity. First among these is the theoretically desirable auction mechanism first proposed by Vickrey. The focus of this paper is to study the efficiency properties of the auction mechanism used by the US and other countries in auctioning telecoms licenses—the simultaneous ascending auction—as compared to Vickrey’s mechanism. In particular, we focus on

efficiency and revenues of these competing auction forms in a complicated bidding environment characterized by demand for multiple objects and synergies for bidders depending on the combination of objects acquired. The idea behind this structure is to mimic some of the key strategic elements present in the auctions of spectrum rights held in the US and elsewhere. Before proceeding, it is useful to present a bit of background on some of the earliest of these auctions.

In August of 1993, Congress granted the FCC the authority to auction off thousands of personal communications services (PCS) licenses; rights to use the electromagnetic spectrum defined by both wavelength and geographic coverage. In granting the FCC this authority, Congress charged them with the task of allocating the licenses in a manner encouraging the “*efficient and intensive use of the electromagnetic spectrum*” [italics added]. The FCC settled on using a simultaneous ascending (SA) auction to allocate spectrum rights.¹ McAfee and McMillan (1996) have argued that the principal reasons for choosing the SA auction over sealed bid or sequential auction forms arise largely because of interdependencies over the objects being offered. That

¹This mechanism was proposed independently by Milgrom and Wilson as well as McAfee.

is, some licenses may represent close substitutes for one another while other represent complements.²

While the theoretical properties of the simultaneous ascending auction have not been characterized, a competing alternative with desirable theoretical properties is a multiple object version of the second-price sealed-bid mechanism proposed by Vickrey (1961) (which is a special case of the Groves (1973)-Clarke (1971) class of mechanisms). This mechanism is known to implement efficient allocations in dominant strategies; hence its use would seem to be consistent in achieving the efficiency objective set forth by Congress in allocating spectrum rights.

In the parlance of auction theory, the Vickrey-Clarke-Groves mechanism (hereafter ‘VCG auction’) is a combinatorial auction. That is, bidders reveal values, not just for each of the licenses being allocated separately, but also for each combination of licenses being allocated. However, the precise allocation procedure in the VCG auction differs from previously considered (and rejected) combinatorial bid mechanisms surveyed in McMillan (1994). Specifically, the combinatorial bid mechanism examined in Palfrey (1983) has

²Licenses covering a similar geographic region at certain bandwidths are substitutes for one another; whereas those covering contiguous geographic regions may be viewed as complements in enabling their owner to obtain a larger geographic “footprint”.

the auctioneer deciding strategically which objects to bundle together and which to sell separately. In the VCG auction, the decision to bundle is determined by the revealed demands of the bidders rather than through strategic considerations of the auctioneer. The VCG auction is closer in spirit to the iterative Vickrey Groves (IVG) mechanism considered in Banks, Ledyard, and Porter (1989). However the VCG auction differs along two lines: first, its allocation procedure is determined after a single round rather than through iterations; and second, its rules for awarding the objects are somewhat more complicated than the adding up rules used in Banks et al. (1989). Thus, to our knowledge, this paper represents the first experimental examination of the properties of the VCG auction relative to the SA auction.

While revenue maximization was not an explicit goal of the FCC in allocating spectrum rights, it seems sensible to think that, given a choice of efficient auction forms, the FCC would likely choose a form yielding more revenues rather than fewer. The VCG auction has the desirable theoretical property that, in the class of efficient Bayesian mechanisms, it maximizes expected revenue. This result, due to Krishna and Perry (1997), extends the more familiar Green and Laffont (1979) result to the case of vector valued private information implementation problems. The theoretical revenue

properties of the SA auction are not known.

Although demonstrably efficient in theory, the VCG auction has been criticized in practice along several grounds. First, even in the simple single object case, the theoretical prediction of demand revelation often does not occur. For instance, it has been frequently observed that in single object second price sealed bid auctions with independent private values, subjects persistently bid above the dominant strategy prediction (see Kagel (1995) for a survey).³ Such violations of the dominant strategy prediction often result in inefficient allocations. Second, with interdependencies among the objects, the vector of valuations which bidders are required to submit in the VCG auction increases geometrically in the number of objects. The complexity of administering a VCG auction grows rapidly with the number of objects being auctioned. This complexity has been an important objection to its practical use. However, advances in processing speed and power in recent years suggest that even at the scale contemplated by the FCC, computing power is sufficient to handle the complexity of the mechanism for most auctions currently being contemplated. A further objection to the sealed-bid version of the VCG auction is that it does not allow for bidders to incorpo-

³In the open outcry version of this auction, such overbidding is not present.

rate new information from others bids into their own bids. This objection, however, may be overcome by running a dynamic version of the VCG auction along the lines suggested by Ausubel (2000). Finally, with multiple object demand, the VCG auction can have undesirable equity properties—the bidder bidding the highest for the object or objects may end up paying a lower price than some other bidder. We return to this issue in section 2.

In this paper, we compare the efficiency, revenue and bidding properties of the VCG and the SA auctions in the presence of complementarities among the objects being allocated. Specifically, we conduct laboratory experiments allocating three objects with private values and complementarities in a VCG auction and in a stylized version of the simultaneous ascending auction. Statistically, we find little to distinguish the two auctions, with both auction forms achieving more than 98% efficiency and extracting roughly 95% of the available surplus. In addition, unlike most of the findings in the survey by Kagel (1995), the theoretical prediction of demand revelation in the VCG auction is largely borne out in our experiments.

The remainder of the paper proceeds as follows: Section two presents the information and valuation structure used in the experiment and establishes the theoretical properties of the VCG auction as well as the main features of

the SA auction. Section three gives an overview of the experimental design. Section four reports the main results from the experiments. Section five concludes.

2 Theory

We consider a model in which fifteen bidders $\{1, 2, \dots, 15\}$ compete for three objects $\{A, B, C\}$ under some auction form α . Bidder i receives a vector of private signals, \bar{v}_i , about her value for all possible combinations of the objects. Let \bar{V}_i denote the set of all possible values of \bar{v}_i .

Let $v = \{\bar{v}_1, \bar{v}_2, \dots, \bar{v}_{15}\}$ denote the matrix of values for all bidders, and let $v_{-i} = \{\bar{v}_1, \bar{v}_2, \dots, \bar{v}_{i-1}, \bar{v}_{i+1}, \dots, \bar{v}_{15}\}$ denote the matrix of values of all bidders except i . Likewise, let $V = \times_{i=1}^{15} \bar{V}_i$ and $V_{-i} = \times_{j \neq i} \bar{V}_j$.

In this set of experiments, \bar{v}_i consists of $\{v_i^A, v_i^B, v_i^C, v_i^{AB}, v_i^{BC}, v_i^{AC}, v_i^{ABC}\}$ where

$$\begin{aligned} v_i^j &\sim \text{Uniform}(10, 20) && \text{for } j = \{A, B, C\} \\ v_i^{jk} &= v_i^j + v_i^k + \text{Uniform}(0, 5) && \text{for } j, k = \{A, B, C\}, j \neq k \\ v_i^{ABC} &= \max_{j \neq k \neq l} (v_i^{jk} + v_i^l) + \text{Uniform}(0, 2) \end{aligned}$$

for all $i = \{1, 2, \dots, 15\}$. Information about the distribution of valuations is common knowledge among the bidders; however, the particular realizations

of the random variables are private information to each bidder. It is useful to notice that this structure of valuations leads to complementarities between and among the objects. The model also includes an auctioneer who is commonly known to value the objects at zero.

An auction form α is efficient if it allocates the three objects in a way which maximizes the aggregate value of the objects to the bidders. Let K denote the set of possible allocations of the objects among the bidders. A particular allocation $k = \{\bar{k}_1, \bar{k}_2, \dots, \bar{k}_{15}\}$ is a 15×7 matrix where element $k_i^j = 1$ if object $j \in \{A, B, C, AB, BC, AC, ABC\}$ is assigned to bidder i and zero otherwise. An allocation rule is a mapping $\kappa : V \rightarrow K$.

Definition 1 *An allocation rule κ^* is ex post efficient if for all $v \in V$*

$$\kappa^*(v) \in \arg \max_{k \in K} \sum_{i=1}^{15} \bar{v}_i \bar{k}_i$$

2.1 VCG Auction

In the VCG auction, bidder i submits a bid vector \bar{b}_i for all combinations of the objects. Define b to be the concatenation of \bar{b}_i for all $i = 1, 2, \dots, 15$, and b_{-i} to be the matrix b excluding the i th vector. The objects are awarded to

bidders such that the sum of the accepted bids is maximized. That is

$$\kappa^*(b) \in \arg \max_{k \in K} \sum_{i=1}^{15} \bar{b}_i \bar{k}_i.$$

It is useful to define the allocation rule when bidder i is not present at the auction. This is given by

$$\kappa^*(b_{-i}) \in \arg \max_{k \in K} \sum_{j \neq i} \bar{b}_j \bar{k}_j;$$

that is, the allocation maximizing aggregate surplus under the constraint that no objects be allocated to bidder i .

Winning bidders pay the highest value that would have been bid for the set of objects awarded to the winner, had the winning bidder not been present at the auction. That is, the payment rule $t_i(b)$ for bidder i is given by

$$t_i(b) = \sum_{j \neq i} \bar{b}_j \cdot \kappa_j^*(b_{-i}) - \sum_{j \neq i} \bar{b}_j \cdot \kappa_j^*(b)$$

Proposition 2 *The dominant strategy truth-telling equilibrium of the VCG auction is ex post efficient.*

Proof. Bidder i 's optimization is to choose a bid vector \bar{b}_i to maximize

$$\bar{v}_i \cdot \kappa_i^*(\bar{b}_i, b_{-i}) - \sum_{j \neq i} \bar{b}_j \cdot \kappa_j^*(b_{-i}) + \sum_{j \neq i} \bar{b}_j \cdot \kappa_j^*(\bar{b}_i, b_{-i}).$$

Since, for a given allocation, the payment by bidder i is independent of his bid, bidder i is only affected by his bid through changes in the allocation. It

then follows by standard arguments that truth-telling is a dominant strategy, that is, $\bar{b}_i = \bar{v}_i$. By construction, under truth-telling κ^* leads to the ex post efficient allocation. ■

Finally, it is a simple matter to verify that participation constraints are likewise satisfied since bidding zero for all objects guarantees a non-negative payoff.

Despite its efficiency properties, the VCG auction, by not tying the price paid directly to a bidder's bid amount, can have the undesirable property that a high bidder ends up paying less for the object than some lower bidder. To see this, consider the following simple example involving only two objects: Bidder 1 values object A at 100 and object B at 0, but values the pair of objects at 106. Bidder 2 values object A at 0, object B at 8, and the pair at 10. Under the VCG auction, each bidder truthfully reveals his value and the objects are assigned efficiently. In this case, that means object A is assigned to bidder 1 and object B to bidder 2. Bidder 1 pays the loss to bidder 2 of his participation in the auction, which in this example is 2. Likewise, bidder 2 pays the loss to bidder 1 of his participation, which in this example is 6. Thus, bidder 1 pays only 2 for the object that he values at 100, whereas bidder 2 pays 6 for an object whose value is only 8. From a practical perspective,

explaining why the bidder 1 got the obviously very valuable object A for a pittance might be difficult.

2.2 Simultaneous Ascending Auction

In the simultaneous ascending auction, we attempted to capture a stylized version of the main elements of the final rounds of the true FCC implementation. Since the rules for the true FCC auction procedure comprise more than 150 pages, simplifications were required for experimental implementation. Below, we briefly explain the auction procedures used in the experiment.

The auction consists of an unknown number of *rounds*, where, in each round, eligible bidders make separate bids on any or all of the objects (but not combinations of the objects). When only one active bidder remains for each of the objects being allocated, the auction ends and the high bidders for each of the objects receive the objects and pay their amounts bid.

In the first round, consistent with FCC practices, there was a minimum opening bid of \$10. This reservation price was chosen to be consistent with the stipulated goal of the FCC that the “opening bid...provide bidders with an incentive to start bidding at a substantial portion of license value, thus ensuring a rapid conclusion of the auction.” (FCC, 1994). Our reserve price

lies (weakly) below the lowest value any bidder might place on any object; thus, we avoid any strategic trade-off between efficiency and revenues as in Riley and Samuelson (1981). In subsequent rounds, a *bid increment rule* was instituted for each object requiring that the minimum bid must be 50 cents higher than the previous round's high bid.

Finally, an *activity rule* stated that once a bidder chose not to bid for an object in a given round, subject to opening bid and bid increment requirements, she was no longer eligible to bid for the object in any subsequent rounds. Our activity rule is a considerable simplification over FCC procedures which use an up-front payment to determine initial eligibility and then use activity rules to adjust these eligibility levels. Our simplified design is intended to model the final stage of an FCC auction, which has the effect of making a bidder ineligible to bid on an object if they fail to bid on that object in any round.

To derive results on efficiency, we model the SA auction as a simultaneous English auction for all objects. To be concrete, as long as the bid increment is small enough, we may think of the SA auction as consisting of a clock starting with a low price and slowly ascending. The clock shows the current price for all objects. Bidders then decide at what point to drop out of the

auction for each object. When only one bidder has not dropped out, the price for that object is set at the price of the last person to drop out, and so on.

The main result here is to establish that allocations are not, in general, efficient in this auction. To establish this result most starkly, consider the simplest case of 2 bidders and 2 objects. Suppose that the valuations are as follows: Bidder 1 values object A at 7, B at 2, and the pair at 9 (i.e., there is no synergy for bidder 1). Bidder 2 values A at 5, B at 0, and the pair at 8 (i.e., object B is only useful to bidder 2 in conjunction with object A). Notice that the efficient allocation in this case is for bidder 1 to receive both objects. Suppose that bidder 2's strategy calls for him to drop out of the market for B at price p . If $p > 2$, then 2 will inefficiently win object B. Thus, for an efficient equilibrium to exist, it must be the case that 2 drops out of the B market at a price below 2.

But now suppose that bidder 1's value for object A is 4 instead of 7 while all of the other valuations remain the same. In particular, since bidder 2's valuations are unchanged, his bidding strategy prior to any bidder exiting the auction must likewise be unchanged.

In this case, the efficient allocation calls for bidder 2 to receive both

objects. However, since in the putative equilibrium bidder 2 is dropping out at a price below 2, this means that bidder 1 will inefficiently receive object B. Since all of the inequalities in the example are strict, it occurs with positive probability. Thus, we have shown:

Proposition 3 *The SA auction is not ex post efficient.*

The degree of inefficiency of course depends on the magnitude of the synergy terms. When there are no synergies, both auction forms are efficient whereas positive synergies lead to the possibility of inefficiency. The magnitude of the synergy terms is considerable in our experiment thus leading to the possibility of inefficiency.

3 Experimental Procedures

The experiment consisted of two sessions conducted at Princeton University in the Spring semester of 1997. Subjects were recruited from the undergraduate population using posters advertising payoffs between \$10 and \$25. In the first session subjects participated in six auctions consisting of three VCG auctions and three simultaneous ascending auctions. The order in which the auction forms were performed was determined randomly prior to the start of

the first session. In the second session, the order was reversed to alleviate any presentation effects in the design. To ensure that both sessions used the same procedures, we adopted a written protocol. In all sessions, the subjects were seated in a large room, read a set of instructions, and given an opportunity to ask questions.

Throughout each session, the only permitted communication between subjects was via their formal decisions. These decisions involved indicating a bid for each object (or combination of objects in the VCG auction only) on a bid submission form. Monitors waited until all subjects had made their decisions before collecting these forms—this preserved the simultaneity of moves required by the theoretical model. The monitors then entered the bids in a spreadsheet and then executed a macro to calculate: a) the winning bidders and the prices paid in the VCG auction; or b) the bid increment rule and the set of active or winning bidders in the simultaneous ascending auction.⁴ In the event of a tie, a winner was determined randomly from among high bidders.

Before being asked to bid, subjects received a handout listing their val-

⁴Detailed descriptions of the precise allocation and payment determination algorithms are available upon request from the author.

uations for all possible combinations of the objects in all six auctions. To ensure comparability, identical realizations of the random valuations (permuted by bidder number) were used in the second session and across auction treatments.

The feedback received by the subjects naturally differed between the auction forms. In the VCG auction, subjects were informed of the objects they received and the prices paid. They then recorded this information. In the simultaneous ascending auction, subjects received feedback at the end of each round regarding the high bidder for each object. This information was also recorded. Throughout, monitors verified that the calculations entered were correct.

At the conclusion of the experiment, the subjects were paid in private and in cash according to their final balance sheet cash balance. This consisted of a \$5 show-up fee, a \$10 starting balance, plus (minus) the proceeds from the six auctions. To aid in retaining experimental control, subjects were informed in the instructions that in the event of losses in excess of their initial \$10, they would be required to leave the session. This bankruptcy event did not occur in any of the sessions.

4 Results

Throughout we present the pooled results of both sessions in terms of efficiency, revenues, and bidding behavior. Since we observed few differences in behavior across sessions, pooling seems justified.

4.1 Efficiency and Revenue

We begin by considering how each auction fared in achieving the efficient allocation of the objects. Efficiency is calculated by dividing the value of the objects actually realized by the bidders by the theoretical maximum obtainable. The associated standard deviation (σ) is given below each of the mean figures.

(TABLE 1 ABOUT HERE)

A visual inspection of the efficiency column of Table 1 highlights the fact that the experimental efficiency results of the two mechanisms are quite close to one another. We can make this more precise by using a Wilcoxon matched-pairs signed-ranks test by matching the six pairs in our sample for each treatment. In testing efficiency in this manner, we obtain a z statistic of .342; thus we fail to reject the null hypothesis that mean efficiency is the same under the two treatments against both one and two-sided alternative

hypotheses. Likewise, we fail to reject the null hypothesis that the efficiency attained by the VCG conforms to theoretical predictions.

Revenue, calculated as a fraction of what would be obtained under full extraction, is given in the last column of Table 1. Again, visual inspection reveals little difference in the revenue properties of the two mechanisms.

Analogous to our efficiency calculations, we can test the equality of median revenues across the two auction forms by using a Wilcoxon matched-pairs signed-ranks test. Testing the equality of revenues yields a z statistic of $-.943$ which is also not significant at conventional levels against one and two-sided alternative hypotheses.

(TABLE 2 ABOUT HERE)

Table 2 presents the efficiency and revenue achievement measures for all 12 auctions. There seem to be no correlation between the sequence of the auction and efficiency or revenue achievement. Regressing revenue on efficiency with all 12 auctions we get a very small but negative ($-.076$) coefficient. However, in SA auction the coefficient is positive ($.276$) and in VCG auctions it is negative ($-.742$). Nevertheless, they are insignificant at 95% confidence level in all cases.

4.2 Bidding

The VCG has the theoretical prediction that it is a dominant strategy for bidders to submit bids equal to their values for the objects. For each of the objects, the dominant strategy prediction was the modal bidding response. Table 3 shows the frequency that subjects' bids were exactly equal to their underlying valuations for the VCG auction.

(TABLE 3 ABOUT HERE)

To see whether there is a pattern of systematic underbidding or overbidding, it is useful to consider bids as a percentage of underlying values in each of the auction forms. To compute bids for AB, AC, etc., we simply summed the final bids for each bidder for each of the individual objects. This is given in Table 4.

(TABLE 4 ABOUT HERE)

In contrast to single object second price sealed bid auctions in which subjects persistently overbid; here, we see underbidding for the object. Likewise, the first-price nature of the SA auction also leads to downward shading of bids relative to underlying valuations. Figure 1 presents scatterplots of bids against values for each of the object combinations $\{A, B, \dots, ABC\}$ for the SA and VCG auctions respectively. Visual inspection of the VCG plots suggests

that underbidding in the VCG auction is mainly due to the submission of a few “non-serious” bids close to zero. Apart from these outliers, the theoretical prediction of full demand revelation appears to be borne out in the data. In the SA auction, there is far more dispersion of bids relative to underlying valuations.

(FIGURE 1 ABOUT HERE)

We can use a signed-ranks to test the theoretical prediction that median bids are equal to underlying valuations in the VCG auction. Specifically, for each bidder in each VCG auction, we compute the difference between bids and underlying values for each object or group of objects. We then compute the median of these differences for each bidder and compare this to the theoretical prediction. Table 5 presents the signed-ranks test results for both VCG and SA auctions.⁵ Here, p-values are from $\Pr(stat \leq |z|)$.

(TABLE 5 ABOUT HERE)

Only the bid for object A is less than its value at the 5% significance level. For all other combinations of objects we fail to reject the null hypothesis that bids are equal to the underlying values. This is in contrast to Banks et al. (1989) who obtained overbidding in a variant of the VCG auction. A similar

⁵Sign test results are qualitatively similar.

test may be conducted by comparing the mean bid to the mean valuation for each bidder. However, as Figure 1 shows, there are a number of “non-serious” outlier bids near zero which have the potential to significantly affect the results of the test. Thus, we eliminate all bids less than or equal to one dollar from the sample. Using this procedure, the sign test results are not qualitatively different from results in Table 5.

One might speculate that full revelation, rather than reflecting a strategic decision on the part of the bidders, in fact, reflects a heuristic bidding strategy which happens to coincide with the equilibrium. Were this the case, one might expect to see something similar in the SA auction.

The results for the SA auction are quite different. Here, for all combinations of objects, we reject the hypothesis of demand revelation. The results for the SA auction suggest the need for caution in attributing demand revelation observed in the VCG auction purely to heuristics such as bidding one’s value. Kagel, Harstad, and Levin (1987) speculate that the overbidding observed in single object VCG auctions results from defective reasoning by subject along the following lines: by bidding above value, subjects correctly believe that their chances of winning the auction are increased; however, subjects place no probability on the event that by overbidding and

winning the auction, losses might occur. In contrast, the complexity of the multiple object VCG payment scheme perhaps forces subject to recognize the possibility that overbidding may lead to losses and hence offsetting the perceived benefit of the increased probability of winning the auction.

Notice that despite the differences in the bidding functions of the two auctions, the monotonicity of bids relative to values and the differences in the payment rules are such that both auctions lead to essentially the same allocations and generate the same revenue. This is broadly consistent with implications of the *revenue equivalence theorem* (Myerson, 1981).

5 Conclusion

Our experimental results show that the revenue and efficiency properties of the VCG auction are comparable to a stylized version of the SA auction currently being employed by the FCC. Moreover, in contrast to findings for single object Vickrey auctions, the VCG auction leads bidders to truthfully reveal their values for combinations of objects with complementarities. Despite its administrative complexity, the VCG auction is implementable in a laboratory setting (at least for three objects) and resolves the auction out-

comes more quickly than the multiple round SA auction (taking roughly half as long as an SA auction to conduct). The sealed bid nature of the VCG auction might also serve to circumvent signaling and other attempts at what might be viewed as collusive behavior in the SA auction. In short, our results (combined with recent allegations of collusion in the existing auction framework) suggest a reconsideration of the merits of the VCG auction.

Several important extensions are still needed in evaluating the merits of the VCG versus the SA auction. First, robustness checks in the magnitude and direction of the externality terms as well as the underlying information structure would seem useful. Since our experimental auctions were conducted under an independent private values with complementarities framework, we have, in some respects, minimized problems associated with the winner's curse. As McAfee and McMillan (1996) point out, the ascending nature of the SA auction is designed precisely to mitigate winner's curse problems often associated with sealed bid auctions, such as the VCG. Thus, it would be useful to compare the SA and VCG auctions in a common (or affiliated) values setting where winner's curse effects are present. Such an extension, however, will dramatically affect the equilibrium properties of the VCG auction in that demand revelation will no longer be a weakly dominant strategy.

References

- [1] Ausubel, L. “System and Method for an Efficient Dynamic Auction for Multiple Objects,” U.S. Patent Number 6,026,383, issued 15 Feb 2000.
- [2] Banks, J., J. Ledyard, and D. Porter, 1989. “Allocating uncertain and unresponsive resources: an experimental approach,” *RAND Journal of Economics*, 20, 1-25.
- [3] Clarke, E., 1971. “Multipart pricing of public goods,” *Public Choice*, 2, 19-33.
- [4] Green, J. and J.-J. Laffont, 1979. *Incentives in Public Decision Making*, Amsterdam: North-Holland.
- [5] Groves, T., 1973. “Incentives in teams,” *Econometrica*, 41, 617-31.
- [6] Kagel, J., 1995. “Auctions: A survey of experimental research” in Kagel and Roth eds. *Handbook of Experimental Economics*, Princeton: Princeton University Press.
- [7] Kagel, J., R. Harstad, and D. Levin, “Information impact and allocation rules in auctions with affiliated private values,” *Econometrica*, 55 (1987), 1275-1304.

- [8] Krishna, V. and M. Perry, 1997. "Efficient mechanism design," *mimeo*.
- [9] Lucking-Reiley, D. "Auctions on the Internet: What's Being Auctioned, and How?" *Journal of Industrial Economics*, September 2000, vol. 48, no. 3, pp. 227-252
- [10] McAfee, P. and J. McMillan, 1996 "Analyzing the airwaves auction," *Journal of Economic Perspectives*, 10, 159-75.
- [11] McMillan, J., 1994. "Selling spectrum rights," *Journal of Economic Perspectives*, 8, 145-62.
- [12] Myerson, R., 1981. "Optimal auction design," *Mathematics of Operations Research*, 6, 58-73.
- [13] Palfrey, T., 1983. "Bundling decisions by a multiproduct monopolist with incomplete information," *Econometrica*, 51, 463-483.
- [14] Riley, J. and W. Samuelson, 1981. "Optimal auctions," *American Economic Review*, 71, 381-92.
- [15] Vickrey, W., 1961. "Counterspeculation, auctions, and competitive sealed tenders," *Journal of Finance*, 16, 8-37.

Table 1: Comparisons of Efficiency and Revenue

Auction Form	Efficiency	Revenue
SA (σ_{SA})	98.5% (.036)	94.7% (.029)
VCG (σ_{VCG})	98.4% (.026)	95.8% (.032)

Table 2: Summary of Efficiency and Revenue in Each Auction

	Session 1						Session 2					
Auction	1	2	3	4	5	6	1	2	3	4	5	6
Type	SA	SA	VCG	VCG	VCG	SA	SA	VCG	VCG	VCG	SA	SA
Efficiency	100%	100%	97.9%	100%	93.3%	100%	91.1%	99.2%	100%	100%	100%	100%
Revenue	95.0%	90.2%	97.8%	92.7%	99.4%	95.4%	92.7%	93.8%	92.5%	99.1%	97.5%	97.6%

Table 3: Frequency of Dominant Strategy Bidding in the VCG Auction

Object	A	B	C	AB	AC	BC	ABC	All
	42%	39%	43%	26%	38%	37%	40%	39%

Table 4: Bids as a Percentage of Values

Object	A	B	C	AB	AC	BC	ABC
SA	87.2%	86.7%	89.7%	81.0%	81.9%	81.1%	79.6%
VCG	92.5%	91.7%	92.0%	94.1%	92.9%	93.6%	96.3%

Figure 1

Bids in the SA auction are shown in the left column, and bids in the VCG auction are shown in the right column.

