Experiments Testing Multiobject Allocation Mechanisms

John O. Ledyard, David Porter, and Antonio Rangel

California Institute of Technology
Pasadena, CA 91125
jledyard@hss.caltech.edu
dave@hss.caltech.edu
arangel@hss.caltech.edu

This paper reports the results of over 130 auctions conducted under controlled conditions to examine the robustness of several auction mechanisms to allocate multiple objects. The simultaneous discrete auction process used by the Federal Communications Commission to allocate Personal Communications licenses was contrasted with a sequential auction and a combinatorial auction over a variety of demand conditions. In test environments created to check only the minimum competency of the procedures, the simultaneous discrete auction process produces highly efficient allocations, approaching levels similar to those found with a continuous form of the auction, and it outperforms a sequential auction. However, in environments created to stress test the procedures, a combinatorial auction outperforms the simultaneous discrete auction.

1. Introduction

During the discussion and evaluation of proposals for the design of the Federal Communications Commission (FCC) mechanism to sell the spectrum, over 130 auctions were run under controlled conditions at Caltech for the National Telecommunications and Information Administration (NTIA), the FCC, and others.[1] While these data were used in those debates, we do not intend to relive that process here. Instead, in this paper, we reexamine these data and try to extract some useful information for those who may, in the future, be involved in the difficult task of creating mechanisms to auction multiple items.

1. Some of the trials and the data generated are described in a report to the FCC. See Ledyard et al. (1994). For a discussion of the role of experimentation in the FCC design process, see Plott (1996). We would like to thank Robin Hanson for his design of the spatial environments.

Journal of Economics & Management Strategy, Volume 6, Number 3, Fall 1997, 639–675

The two major design questions we can say something about are (1) should the items be auctioned off sequentially or simultaneously? and (2) should package bidding be allowed? Our main conclusion is that, over a very wide range of environments, package bidding mechanisms (weakly) dominate simultaneous mechanisms, which in turn (weakly) dominate sequential mechanisms. This conclusion is based on three observations derived from a close look at the data.

First, in environments with multiple items to be allocated, if those items are homogeneous and substitutes, then little coordination between buyers is needed and the only role of the mechanism is to sort bidders with high values from bidders with low values. Both the sequential and simultaneous mechanisms seem to work very well at finding efficient allocations in these ''easy'' environments.

Second, in environments with multiple items to be allocated, if those items are heterogeneous, then some coordination among bidders is necessary to achieve high-value allocations even if there are only low synergy values. Simultaneous auctions provide a first step at this coordination that sequential auctions might have difficulty in providing.

Third, in environments with heterogeneous goods exhibiting complementarities, significant coordination is required for an auction or allocation mechanism to perform well with respect to efficiency or revenue. Sequential auctions perform poorly. Simultaneity is clearly necessary but not sufficient to attain high efficiencies. The simultaneous, one-price-per-item auction tends to produce outcomes that are either high in efficiency, revenue, *and* losses or low in efficiency, revenue, and losses. Package bidding seems to help a lot in systematically attaining high efficiency, high revenue, and no losses.

In the rest of this paper we explain and detail the collection of experiments. In Section 2, we provide information on the environments covered, the mechanisms tested, and the performance measures used. In Section 3, we describe the data and our findings. In Section 4, we address some of the questions and issues left unanswered by these data.

Because we expect that many readers may not have a background in the methodology of applied mechanism design, we have included a brief introduction in an appendix. Policy makers and theorists interested in applying the results in this paper should read Section A.1 carefully.

2. What Did We Do?

During the actual FCC design process, a wide range of questions were continuously thrown at the experimentalists who were trying to provide insights and data, as fast as possible, about situations for which theory had virtually nothing to say. Whenever the experimentalists found a problem with a current manifestation of the proposed designs, new proposed solutions were immediately put forward. No careful theoretical analysis or experimental design was followed, nor could one be, given the urgency of the situation. Nevertheless, we think that the experiments that were done can be organized in a reasonably coherent fashion and that, while they do not cover the entire territory one may wish they had, some fairly straightforward conclusions can be drawn for future designs.

2.1 Questions and Method

There were two major design questions with respect to the auction rules about which the data we have reveal some information:

1. Should the items be auctioned off sequentially or simultaneously?[2]
2. Should package bidding be allowed?

There were other issues that achieved some relative importance at various times during the design process but for which there still is neither any convincing theory nor enough experimental evidence on which to base a judgment. Should there be a withdrawal rule or not, and if so, in what form?[3] What should be the appropriate stopping rule? Should activity rules be required, and if so, what should they be? How many waivers should be allowed? While there are some data that might provide light, we feel that more experiments and theory are needed before anything conclusive can be said, and so we will not address these secondary questions in this paper.

Experimental methods in economics provide a type of "wind tun-

2. A hybrid design was also considered, which involved comparing the results of a simultaneous sealed bid of all items and a sequential open outcry auction of each item. See Plott (1997) for a description of the process and data on its comparative performance. Milgrom (1995) also has a description of this proposal. We do not cover that design here.

3. There were data and theory on one proposed rule, to allow withdrawal at any time for free. These suggested that such a rule would destabilize the auction and produce low efficiencies in the allocation and low revenue. [See, for example, Banks et al. (1989) and Milgrom (1995).] The rule was eventually eliminated from further consideration. Porter (1996) provides an experimental analysis of the withdrawal rule currently used in the FCC auction and finds that there is a positive relationship between individual losses and allocative efficiency when the rule is imposed.

nel'' within which to test mechanism designs. These tests can be a valuable source of scientific information that one can use to determine the likely performance of new mechanisms in new environments. The process is simple and very similar to the testing of airfoils in wind tunnels or the testing of hull shapes in towing tanks. One first simulates the environment, in our case by inducing the constellation of participants' valuations and the information they each have about these valuations. Then a mechanism is provided and allowed to operate within the testbed environment. Performance is measured. With enough variation in the environments and enough variation in the mechanisms, one can begin to reach some conclusions about details in design that affect performance. Hunches and arguments loosely based on inappropriate theory can be replaced by facts.

Testbed experiments can be a valuable source of data about the performance of newly designed mechanisms for which there are no examples in operation. As an illustration, see Ledyard et al. (1994b) for the research that led to the Cassini trading mechanism—a bulletin-board trading system now in use as a project management device in the design and construction of the Cassini spacecraft for a mission to Saturn.[4] It may well have been the very first active mechanism for trading worldwide over the Internet. For other illustrations, see Plott (1994). As with any evidence, including theory, testbed data must be weighed carefully. But if used intelligently they can eliminate bad designs, provide comparative performance data, and actually help a decisionmaker come to good conclusions during the design process.

2.2 The Economic Environments We Used[5]

All of the *economic* environments reported in this paper are derived from the following generic setup. A set of n objects, labeled $x_1,...,x_n$, are to be allocated to m agents. Agent i's profit function is $V_i\,(x_{1i},....,x_{ni})$ where $x_{ki} = 1$ if and only if agent i is awarded item k. Thus, an agent knows what they will be paid if they successfully acquire any particular subset of the items. In some cases below, agents will be assumed to have common knowledge about aspects of others' values. In other cases agents will know nothing *a priori* except their own valuations. In creat-

4. Another mechanism that is successfully running in practice and that was developed with the aid of laboratory testbeds is the ACE market. ACE is now operating in Los Angeles at least four times per year as a call market for trading emissions credits—a very complex process. Those interested in the details can go to the WorldWideWeb page at www.ace-mkt.com, which includes a link to download the client software.

5. More details can be found in the Appendix, Section A.3.

ing this generic setup, we realized we were abstracting from correlated and asymmetric information. We did so, not because we thought such information was unimportant, but because we wanted to concentrate the limited resources we had on the performance of the proposed auctions when problems such as the winner's curse were absent. In the presence of those problems some fundamental performance properties might have been either obscured or exacerbated. This abstraction is, of course, a defect of the existing research that can and should be corrected in future work.

In this class of environments the most efficient allocations solve the problem

$$\max_{x} \sum_{i} V_i(x_{1i}, \ldots, x_{ni})$$

subject to

$$\mathrm{x}_{ji} = 0 \text{ or } 1,$$

$$\sum_{j} x_{ji} = 1 \qquad \text{for each } i.$$

For purposes of reporting the test results, we split the environments we used into three somewhat arbitrary classes. The first, which we refer to as *easy*, involve constellations of values for which no reasonable mechanism should have any problems achieving efficient allocations. If a proposed mechanism had failed to perform well in these situations, one would have been fairly sure that it would also fail to perform in more complicated environments. One can think of these as minimal competency tests. These environments have at least two features that make the allocation problem easy for a mechanism: there are no significant coordination issues that require the mechanism to fit complementary demands together, and there is a competitive equilibrium (CE) price vector with one price per item that is sufficient to support the 100%-efficient allocation.

2.2.1 Easy Environments. The first easy class of environments used is one in which values are additive, i.e., the value function is of the form $V_i(x_{1i}, \ldots, x_{ni}) = \sum_j V_i(x_{ji})$. In the experiments, there were six items for sale to six demanders. Each demander knew that his value and those of other participants were to be drawn uniformly, with replacement, from a fixed list of ten value sheets.[6]

A second easy class of environments used had items that were

6. More details can be found in the Appendix, Section A.3.

homogenous and had decreasing marginal values, i.e., $V_i(x_{1i}, \ldots, x_{ni}) = V_i(\sum_j x_{ji})$ with $V_i'(\sum_j x_{ji}) < 0$. In the experiments, there were eight participants and ten units to be allocated. Each participant had decreasing demands for up to four units. In addition, subjects knew the number of units being auctioned, the number of participants, and who bid on what items; however, they did not know the distribution over which values were drawn.

The final easy class of environments we used was one in which items were homogenous but had increasing marginal values (superadditive values), i.e. $V_i'(\sum_j x_{ji}) > 0$. We also constructed the value functions so that a competitive equilibrium price existed. In the experiment, there were eight participants and ten units to be allocated. Each participant had increasing demands for up to as many as four units. Subjects knew the number of units being auctioned, the number of participants, and who bid on what items. However, subjects did not know the distribution over which values were drawn or even if a single price would clear the market.

2.2.2 Moderate Environments. The second sets of values we consider are cases in which the degree of difficulty for mechanisms is raised a bit. We consider this a move towards the actual possibilities. We introduce heterogeneity into the environment. We believe, and the data support, that heterogeneity can significantly increase the difficulty any auction design has producing efficient allocations. By effectively increasing the dimension of the commodity space from one to n, prices must now not only separate high-value bidders from low-value bidders; prices must also coordinate the demands of bidders across commodities. Prices in one market affect the demands in another, and general equilibrium phenomena become important. Theory and data and intuition from environments with homogenous objects are not sufficient background for analyzing environments with heterogeneity.

The first moderate environment we used was similar to the easy superadditive case described above, except that values were selected so that there did not exist a single market clearing price for the items. So unless the mechanism can produce nonlinear pricing, either the outcome must result in losses to at least one bidder, or some participant must forgo the pursuit of potentially profitable opportunities. As in the easy superadditive case, subjects knew the number of units being auctioned, the number of participants, and who bid on what items. However, subjects did not know the distribution over which values were drawn or if a single price could clear the market.

The second moderate environment we used was the assignment problem that incorporates heterogeneity but allows buyers *to redeem* *one and only one item*. The mechanism is presented with a coordination problem that has features of what might happen if bidders had budgets. The individual payoff in an assignment environment is given by

$$V_i(x_{1i}, \ldots, x_{ni}) = \sum_{j}^{n} V_i(x_{ji})\cdot\theta_{ji},$$

$$\theta_{ji} \in \{0, 1\} \quad \text{and} \quad \sum_{j}^{n} \theta_{ji} = 1.$$

The only problem facing the mechanism is the coordination of the demanders. In the experiments, there were six items for sale to six demanders. Each demander knew that his value and those of other participants were to be drawn uniformly, with replacement, from a fixed list of ten value sheets (see Appendix, Sec. A.3, for more details).

The final moderate environment we used was one in which individuals value the heterogeneous items offered more in groups than singly. That is, for some items and some agents, preferences may have the property that $V(\{a,b\}) > V(\{a\}) + V(\{b\})$. This structure of preference is often characterized as possessing "complements" or "synergies." The theory that guided the creation of these testbed environments can be found in Bykowsky et al. (1995). These experiments cover a variety of cases in which the mechanism must coordinate the bidders and guide them to best fit together. In some of the cases those trying to assemble packages can be exposed to losses if they try to build a particular package but are eventually outbid for a piece of it.

A constellation of values with this potential risk can be seen in the environment provided in Table I. There, three subjects are competing for three heterogeneous items (a, b, and c) with the values listed. Subject 1 has the highest value for each of the items. However, subjects

TABLE I.
SIMPLE FITTING ENVIRONMENT

	Values		
Packages	Subject 1	Subject 2	Subject 3
a	2	2	4
b	2	4	2
c	4	2	2
ab	23	24	27
bc	24	27	24
ac	27	23	23
abc	42	32	32

2 and 3 have high values for packages that overlap at item *b*. It is easy to see that there are no competitive equilibrium prices for this case. We also included some constellations of values with synergies and for which a simple competitive equilibrium existed.[7]

Subjects knew the number of units being auctioned and the number of participants and who bid on what items. However, subjects did not know the distribution over which package values were drawn.

2.2.3 Hard Environments. Finally we turn to three classes of environments that were intentionally constructed to test the limits and robustness of each of the mechanisms. In these environments value synergies over specific packages of items (spatial demands) are predominant. This forces bidders to coordinate their bids to find the highest-value fit among packages of items.

The first hard environment we used had three items and three demanders. Each demander had values for single items and also a synergy value for all three items. The values of bidders were determined as follows:

1. The integer values for the single items were drawn independently from a triangular distribution with support [0,98].
2. The value for the three-item package was determined by adding a number randomly selected from the interval [0,149] to the highest value for item *a*, *b*, or *c* drawn in step 1.

The efficient outcome can have either the entire set of items going to one demander or one item going to each demander. This class of environments was designed to contrast the interests of a user who wants a major package with the interests of some single-item bidders. We created this environment after hearing many suggestions in the policy analysis of the FCC auction design that package bidding would bias the results in favor of those wanting many items. We thought it important to study that unsubstantiated claim and, in doing so, to give that argument its best chance. Subjects in this experiment knew the distribution from which the values were drawn.

The second hard environment we used had five bidders and six

7. At the time we thought that the performance of some of the proposed auctions would differ according to whether a competitive equilibrium existed or not. As it turns out, that conjecture was shown to be wrong by the data, and so we have reported the data from both situations together.

TABLE II.
VALUES IN A SPATIAL FITTING EXAMPLE

Bidder 1	Packages:	*f*	*cd*	*bcf*	*bde*	*abe*
	Values:	9	22	128	130	120
Bidder 2	Packages:	*b*	*df*	*ae*	*af*	*abd* [a]
	Values:	8	28	24	27	130
Bidder 3	Packages:	*c*	*a*	*d*	*bd*	*abf*
	Values:	2	3	8	20	119
Bidder 4	Packages:	*e*	*abc*	*adf*	*bdf*	*aef*
	Values:	10	117	112	128	125
Bidder 5	Packages:	*cf*	*de*	*cef* [a]	*bef*	*abcdef*
	Values:	29	25	117	125	142

[a] Optimal fit.

items (a, b, c, d, e, and f) to be allocated. Table II shows the nature of the problem faced by an allocation mechanism. Bidders 2 and 5 have three-item packages that exactly fit together and have the highest possible value of any combination of packages. The problem is that every one has high-value, three-item packages, all of which overlap. The task of the mechanism is to guide the owners of the components of the optimal allocation to find each other. This environment was created to give package bidding its best chance. The distribution over which the values were drawn was given as common information to subjects in this experiment.

The third hard environment used was designed to investigate the boundary case described in Banks et al. (1989) in which multidimensional demands of lumpy sizes must fit into a box with fixed dimensions (a *network* problem). This environment involves values in which bidders receive payoffs only for packages that are highly interrelated and must fit with the packages of other similar demanders. In particular, there were two resources with fixed supply of 20 units each. Each subject had nine two-dimensional packages of the fixed resources they could select from, for which they would receive value.

2.3 Allocation Mechanisms Tested

Into these multifaceted environments we threw three mechanisms: a sequential ascending-bid auction, a simultaneous ascending-bid batch auction, and a continuous package bidding auction. The mechanism designs were taken from the early debates in the winter of 1993 over which mechanism the FCC could or should use to allocate PCS spectrum. The debate dealt primarily with the allocative efficiency of the mechanisms and their revenue-generating properties. We describe below each of the mechanisms as we implemented them in the various testbeds.

2.3.1 The Sequential Ascending-Bid Auction. As its name implies, the sequential auction mechanism allocates one unit at a time in some sequential order. The method used in the laboratory testbed was to auction them off in random order. All participants knew, before the bidding began, the order in which the units would be auctioned off. Each unit was allocated using an ascending-oral-bid auction. Of all the mechanisms we tested, this was clearly the easiest to implement.

2.3.2 The Simultaneous Ascending-Bid Batch Auction. This mechanism operates across of a series of rounds. Once each round, individuals submit a sealed bid on each of as many items as they wish. After a round closes, the highest bid submitted for each item, the *standing bid*, is identified and displayed along with all other bids submitted. An allocation is made when bidding stops.

2.3.2.1 Activity and Update Rules. In order to be able to submit a bid in the round, a participant must have been *active* in the previous round. To be active a participant must have submitted an acceptable bid in the previous round or have had the standing bid two rounds back.[8] In order for a bid to be *acceptable* in a round, it had to be at least 10% higher than the standing bid for the item.

A *second-stage activity rule* was imposed if the auction did not close before round 8.[9] This rule restricted the number of items for which a participant could bid in a round. The restriction was that a participant could bid for at most a total number of items equal to (1) the number of acceptable bids placed in the previous round for items for which the participant did not have the standing bid for the item, plus (2) the number of items for which the participant had the standing bid two rounds back but no longer had the standing bid. In addition, the participant could always bid on those items for which he currently had the standing bid.

8. In some experiments participants were provided with two waivers that they could use to stay active. The original purpose of the waivers was to ensure that a bidder with logistical problems in entering a bid was not penalized. In our experiments logistical problems were not an issue. Still, over half of the allotted waivers were used by participants, presumably for strategic reasons.

9. In the spatial environment the second-stage activity rule was not used.

These activity rules are exactly those that were employed in the FCC's auction for nationwide narrowband PCS licenses.

2.3.2.2 Withdrawal Rule. A withdrawal rule allowed participants to delete any of their standing bids before a round began. After such a withdrawal, the price of that item was dropped to zero and that bid became the standing bid of the experimenter. An individual who withdrew his bid paid a *penalty* equal to the maximum of the difference between the amount of the bid he withdrew and the highest bid submitted after his withdrawal, and zero.[10]

2.3.2.3 Stopping, Pricing, and Allocation Rule. The simultaneous ascending-bid batch auction stopped if no acceptable bids were submitted in a round or if the process reached some round after round 13. In the latter case, the actual round the process was to be stopped at was announced two rounds ahead of time.[11] When the process stopped, the items were awarded to the participants with the standing bids in that round. Withdrawal penalties were also calculated at this time.

2.3.3 The Package Bidding Mechanism. This is the AUSM mechanism described in Banks et al. (1989).[12] It is similar to the continuous ascending-bid auction[13] but with two special features. First, participants are allowed, *but not required*, to submit bids for packages of items as well as for individual items.[14] That is, they can say ''I am willing to

10. In the spatial environments, a slightly different withdrawal rule was used. In those experiments, one could withdraw all of one's bids when the auction stopped. The withdrawn items were then offered in a random but sequential order to the next highest bidder on that item. This was an early idea put forward by some for the FCC auction design.

11. In the spatial environments a hard stop rule was not used. Instead, the mechanism was allowed to run its course. That is, the auction stopped and all markets closed simultaneously if and only if no new bids were entered in a round. This was the stopping rule eventually chosen by the FCC.

12. Other references with details about AUSM include Ledyard et al. (1996) and Bykowsky et al. (1995).

13. AUSM can and has been run as a batch process. This requires that an optimization routine be run after each round, but with modern computers and software and with economic incentives driving the structure of the bids, there have been no computational problems in practice.

14. There seems to be a widely held misperception that AUSM and related package bidding mechanisms require that each bidder submit a bid for every possible package or 2^n bids. This is wrong. Just as in the simultaneous ascending-bid auction, bidders need only bid on those items they truly want and think they have a chance of winning. In fact, if package bidding is allowed, in equilibrium fewer bids are needed to support an efficient allocation than in the simultaneous ascending-bid auction. Data from testbeds and from real-world use suggest that package bidding generates no more serious bids per person than any other mechanism and, indeed, may actually generate fewer.

pay \$100 for the package $\{a, b, c\}$'' and not have to identify a separate bid for each item. With such a bid, they are requesting to be allocated a and charged \$100 if and only if they are also allocated b and c. This bid is accepted if and only if \$100 is more than the sum of the standing bids for the packages that contain a, b, and c. So, for example, if there is a standing bid of \$35 for a, a standing bid of \$50 for b, and a standing bid of \$5 for c, then the bid of \$100 for $\{a, b, c\}$ wins. If, however, a bid of \$75 for $\{b, c\}$ is made before the bid of \$100 for $\{a, b, c\}$ is submitted, then the standing bids are the bid for a and the bid for $\{b, c\}$. The \$100 bid for $\{a, b, c\}$ is then no longer large enough to become a standing bid. It would need to be greater than \$110. Because it can sometimes take several small package bids to displace a large package bid, a second special feature of the continuous AUSM mechanism[15] is a bulletin board on which bidders can post small bids that are not large enough to displace a current winner but that might be part of a collection of bids that would be large enough. This standby queue of bids is always available for bidders to combine with to displace a large package bid.

2.4 Summary of Procedures

All of the experiments were conducted at the California Institute of Technology using the student population as the subject pool. All of the mechanisms, other than the sequential auction, were computerized. For the AUSM and simultaneous batch auctions, experienced subjects were used who had three hours of training in the rules of the mechanisms and software. In Table III are listed the relevant information for each of the experimental sessions.

3. What Did We Find?

We present the results in three parts. First, we provide a brief summary of the measured performance, both efficiency and revenue, of the mechanisms in the various testbed environments. Second, we look more closely at some results from the hard testbeds, the ones designed to stress the limits of the mechanisms. Finally, we present some observations based on our reading of the totality of the evidence. We also answer the two design questions raised in Section 2.

15. When AUSM is run in the batch format, this feature is not necessary, because every bid is processed simultaneously. Even when it is available, it is rarely used. See Kwasnica et al. (1997) for a comparison of a batch AUSM mechanism with a continuous AUSM mechanism.

TABLE III.
EXPERIMENTAL DESIGN

Mechanism	Environment	Number of Experimental Sessions	Subjects	Comments
Sequential	Additive	1[a]	Inexperienced	
	Decreasing	1[a]	Inexperienced	Conducted at end of simultaneous sessions
	Assignment	2[a]	Inexperienced	
	Spatial fitting	2[b]	Experienced	
Simultaneous batch	Additive	1[c]	Experienced	
	Decreasing	4[c]	Experienced	
	Superadditive CE[d]	4[c]	Experienced	Fixed ending round used
	Superadditive without CE	5[c]	Experienced	
	Assignment	2[c]	Experienced	
	Fitting	3[c]	Experienced	No second stage and no fixed ending round
	Spatial	4[c]	Experienced	As above; withdrawal at end only
	Spatial fitting	3[c]	Experienced	
	Network	1[a]	Experienced	No withdrawal
AUSM	Fitting	2[a]	Experienced	
	Spatial	5[b]	Experienced	Standby queue used
	Spatial fitting	4[b]	Experienced	
	Network	2[a]	Experienced	

[a] David Porter designed and conducted these experiments.
[b] Robin Hanson and David Porter designed and conducted these experiments.
[c] Antonio Rangel, David Porter, and John Ledyard designed and/or conducted these experiments.
[d] Competitive equilibrium.

In organizing the data, we use three standard performance measures: efficiency, revenue, and bidders' surplus.[16] Since we are working in environments in which value is measured in terms of profit, we measure efficiency in the usual way as the aggregate value achieved by the mechanism as a percentage of the maximum possible. It has been correctly pointed out by some that the absolute value of this measure is not particularly illuminating.[17] However, comparing values across mechanisms in the same environments can be informative. For exam-

16. We of course also look at other dimensions of interest. Of particular interest in the debate about package bidding is the extent to which a mechanism biases the outcomes in favor of bidders who only want a small number of items as opposed to those who get high value from large packages.

17. In the Appendix, Section A.2, we expand and explain this problem.

ple, if I were to tell you that mechanism M1 produced observed efficiencies between 90% and 96% and that mechanism M2 produced efficiencies between 86% and 89%, for the same structure of payments, you would be justified in concluding that M1 outperformed M2 on that class of environments with respect to attaining efficient allocations. We provide such comparative data below.

Our second measure of performance, revenue, is simply the number of dollars collected from the participants in payment for the items. Again the absolute numbers do not necessarily provide any guidance for the designer. The usual solution is to use revenue as a percentage of the predicted market equilibrium prices. Unfortunately, in some of the environments we report on below, no such prices exist. So, in our analysis, we will use revenue as a percentage of the maximum value attainable. Again, it is the relative values of revenue collected across mechanisms in the same environments that can be informative.

Finally, because it is a measure of the user's gains from participation and because it can serve as an explanation for why some potential participants argued for particular designs, we include data on bidders' surplus. In this case, this is simply the difference between the value attained and the revenue paid. Bidders' surplus as a percentage of the maximum value attainable is thus simply the difference between the efficiency and revenue percentages.

3.1 Basic Performance—Efficiency and Revenue

3.1.1 Easy Environments. In Figures 1 and 2 are plotted the efficiency and revenue percentages achieved by the mechanisms tested in environments with additive values, with decreasing returns, and with increasing returns with competitive equilibria. These data confirm our prior intuition that all of the mechanisms would do well in such unchallenging situations. The variation observed in the performance of the batch process, relative to what would be efficient, should be expected in light of the requirement that bids increase by at least 10% in each round.

3.1.2 Moderate Environments. In Figure 3 and 4 are the data from the tests in which the mechanisms were exposed to moderately more difficult parametric conditions. These were one with homogeneous goods with increasing marginal values and no competitive equilibrium, the assignment problem, and a simple fitting problem.

The mechanisms now begin to separate themselves. The simultaneous mechanisms do a better job of finding efficient allocations and

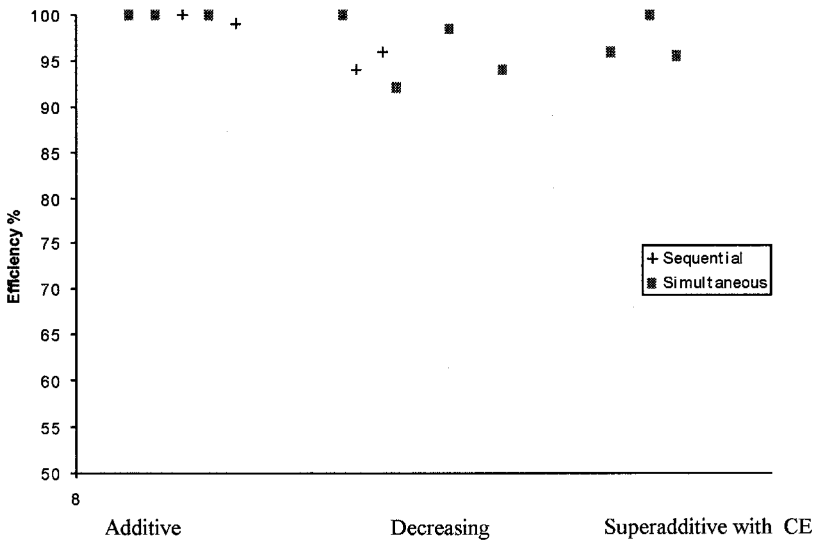

FIGURE 1. MECHANISM EFFICIENCY—EASY ENVIRONMENTS

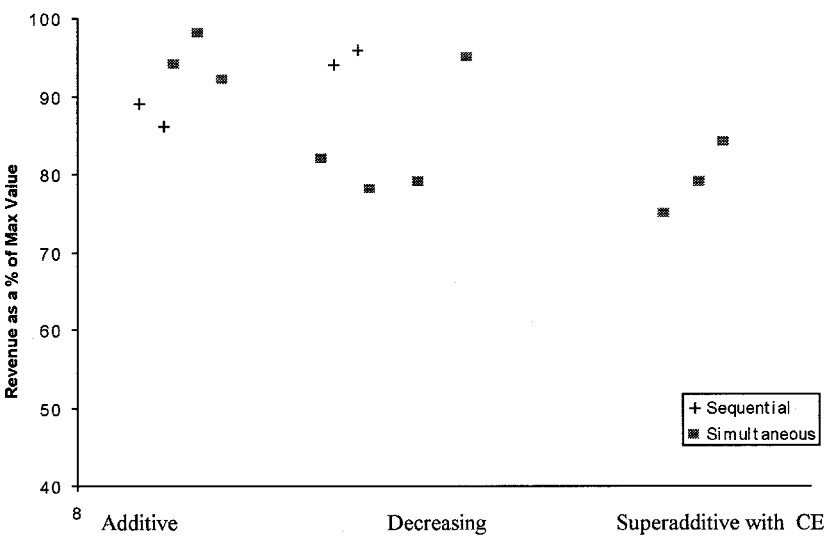

FIGURE 2. MECHANISM REVENUE—EASY ENVIRONMENTS

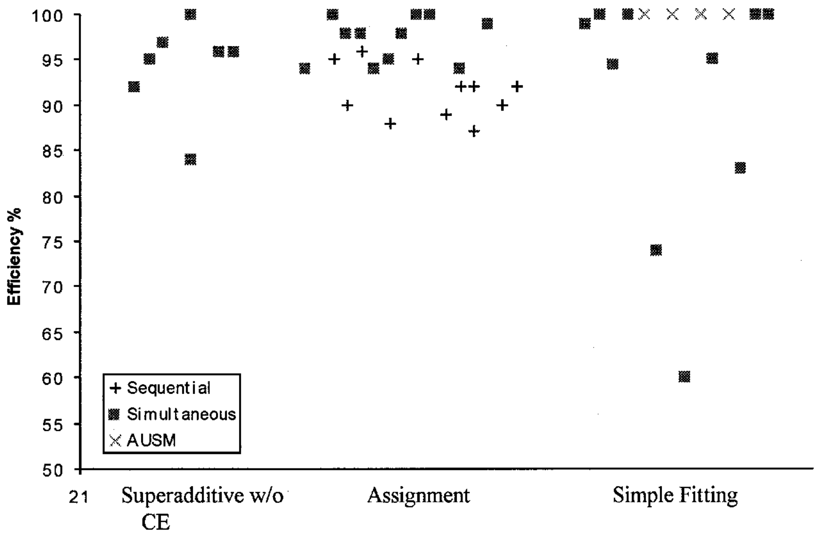

FIGURE 3. MECHANISM EFFICIENCY—MODERATE ENVIRONMENTS

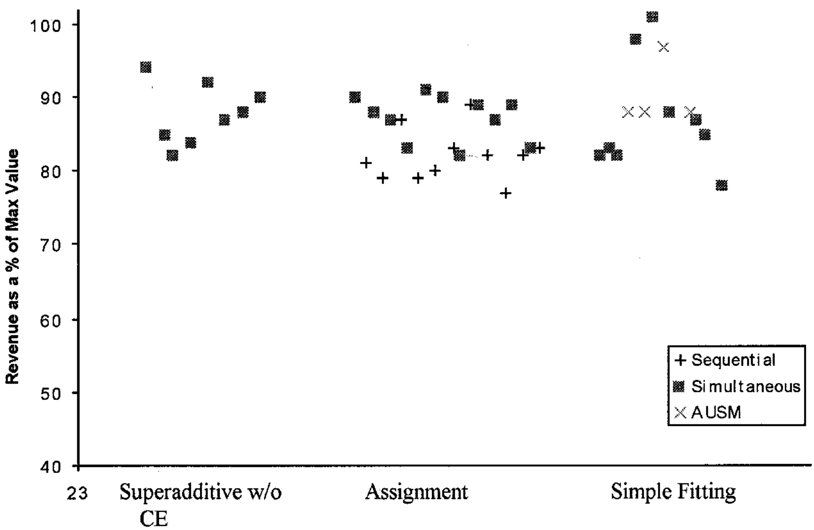

FIGURE 4. MECHANISM REVENUE—MODERATE ENVIRONMENTS

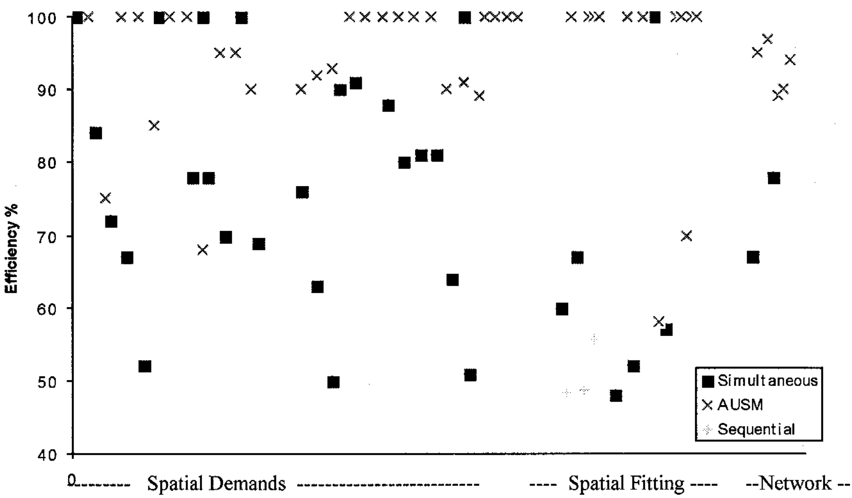

FIGURE 5. MECHANISM EFFICIENCY—HARD ENVIRONMENTS

produce more revenue than the sequential mechanism. However, relative to the easy environments, both the simultaneous and sequential auctions now exhibit some losses in efficiency. The packaging mechanism seems to work very well in the simple fitting environment. One finding of some interest is that, in the simple fitting case, the revenue produced by the simultaneous auction can actually exceed the maximum value of the final allocation. This means that one or more of the bidders has sustained significant losses—has paid more for the items won than they are actually worth to that bidder. This is not a winners'-curse phenomenon: remember that there is no correlated information. It is the result of using an inappropriate mechanism in an environment with complementarities among heterogeneous items.[18]

3.1.3 Hard Environments. In Figure 5 and 6 are the data from the testbeds that used the most difficult coordination environments. These data are perhaps the most revealing about the relative performance capabilities of each mechanism. Stress tests often highlight strengths and weaknesses missed under more normal conditions.

The results seem very clear from these figures. First, without package bidding, there are major losses in allocative efficiency. Thus, it

18. See Bykowsky et al. (1995) for how this might happen.

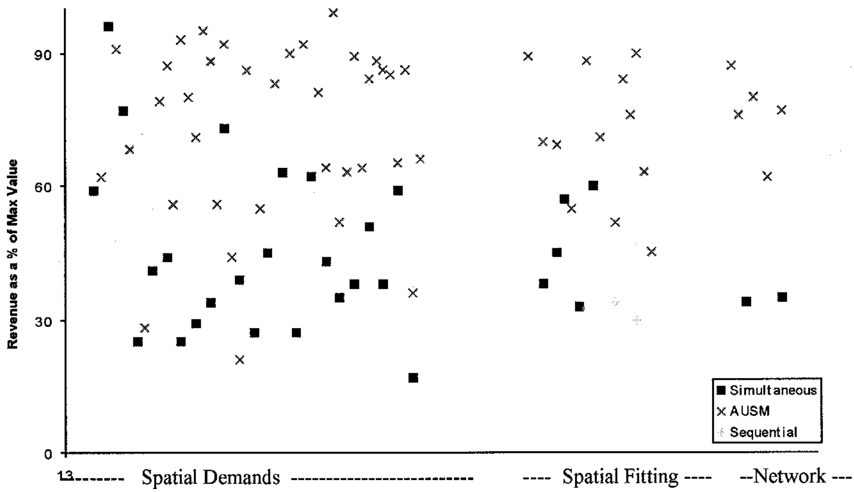

FIGURE 6. MECHANISM REVENUE—HARD ENVIRONMENTS

appears that simultaneous auction processes are necessary but not sufficient to coordinate demanders. A simultaneous auction does eliminate single-item efficiencies. But in complex environments with nonconvexities arising from heterogeneous spatial returns to scale, one price per item is simply not enough information to guide bidders to efficient allocations.[19] Opportunities for economies of scale, scope, and fit are easily missed. Allowing bids for packages leads to improvements in efficiencies and revenue, and losses are controlled. Efficiencies are improved because bidders can find and bid on those packages with significant complementarities without bearing the risk of losing part of that package. Revenue is improved for the same reason.[20] However, the gain in efficiency and revenue from allowing package bidding appears

19. The theory behind this observation can be found in Calsamiglia (1977), Jordan (1987), and Mount and Reiter (1996).

20. In ascending-bid auctions, winning prices are driven by the values of the second-best allocations, since winners must bid enough to ration the losers out. With packaging, the second-best allocations can be more easily found and bid, and, with synergies, these are worth more than the sum of the values of the single-item allocations. Thus package bidding generally yields higher revenue without losses.

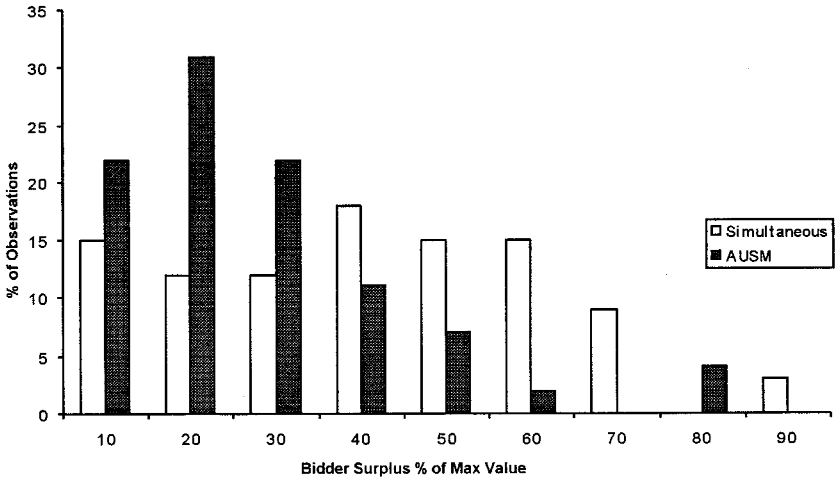

FIGURE 7. DISTRIBUTION OF BIDDER SURPLUS

to come at some expense to bidders' surplus. In Figure 7 are displayed the distribution of bidders' surplus (net profit)[21] as a percent of the maximum value.

3.2 A Slightly Deeper Look

There is a three-way tradeoff in the design of mechanisms between efficiency, revenue, and bidders' surplus. In Figures 8 and 9 we exhibit this tradeoff separately for two classes of the hard environments: the spatial and the spatial fitting. In the figures, points to the northeast represent higher efficiencies, since revenue percentage plus surplus percentage equals efficiency percentage. The distribution of points along the lines of equal efficiency represents the distribution of the surplus between the seller (revenue) and the buyers (bidders' surplus). We chose these tests because they appear to generate the starkest differentiation in performance between the simultaneous and the packaging mechanisms. Such differentiation allows some insight into the particular strengths of each mechanism.

21. Net profit is simply $V_i(x_{1i}, \ldots, x_{ni}) - \sum_j b_{ji}$. Here, b_{ji} is the amount paid by i for item j.

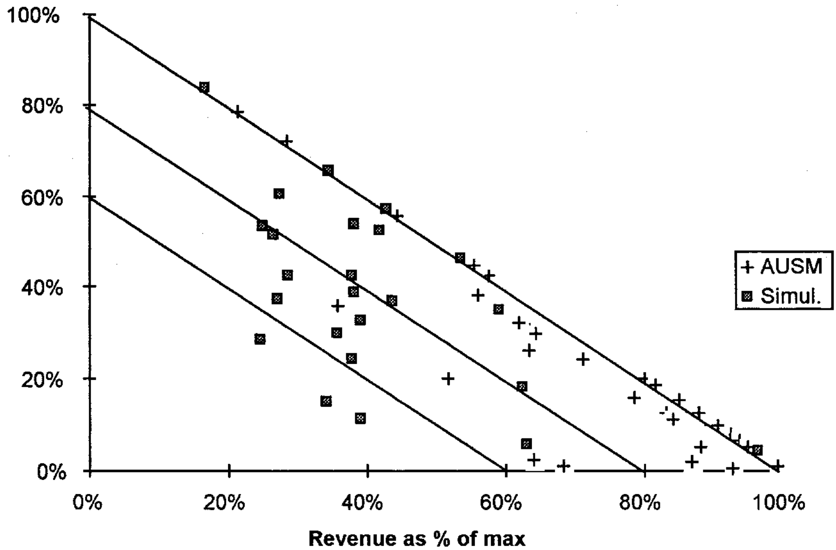

FIGURE 8. AUSM VS. SIMULTANEOUS (SPATIAL)

3.2.1 Giving the Simultaneous Auction Its Best Chance. The spatial testbed is probably the ideal example of a situation that gives the simultaneous ascending-bid auction its best shot at outperforming the package-bidding auction. It places the interests of single-item demanders in direct conflict with those of demanders who wish the entire collection of items. To make that conflict even starker, the highest-valued demander of the collection is also one of the single-item demanders. Thus, in situations in which the single-item demanders should win (i.e., when that is the efficient outcome), not only must they outbid the demander with the highest value for the whole collection, but also the demander of the whole must be part of the effort to outbid himself.[22]

22. This should certainly lead to the exposure of a threshold problem for AUSM, if one exists. It has been the speculation of some that package bidding creates a situation in which those bidders who value large packages have an advantage over bidders who only value single items. The speculation is that when it is efficient for single-item demanders to win, AUSM will let the large package bidders win instead, leading to low efficiencies. This is called the threshold problem, because the single-item bidders have to coordinate to jointly overcome the threshold provided by the large package bid. The counter bias, that the simultaneous mechanism may let single-item bidders win in tests when it is efficient for large packages to win, is generally not mentioned in these discussions.

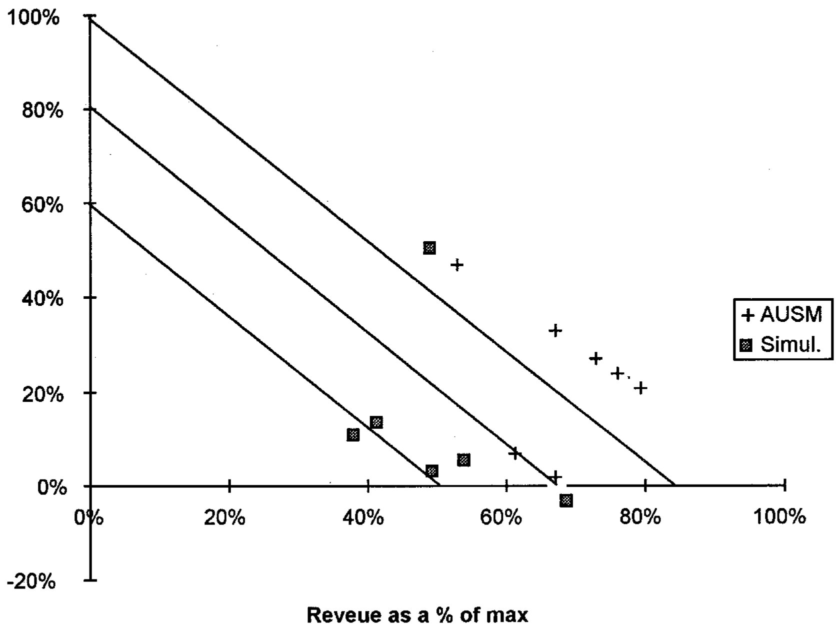

FIGURE 9. AUSM VS. SIMULTANEOUS (SPATIAL FITTING)

In these tests, AUSM tends to generate outcomes that are, relative to those of the simultaneous mechanism, high in both efficiency and revenue but low in surplus. The simultaneous ascending-bid auction does better in generating bidders' surplus, but does so in many tests at a serious loss in efficiency. The numbers are straightforward. In 90% of its tests, AUSM yielded more than 80% efficiency. In only 33% of its tests did the simultaneous mechanism yield more than 80% efficiency. On a relative basis, using efficiency as the appropriate measure, the data seem to reject the charges that AUSM has a threshold problem. A closer examination, however, reveals some evidence to the contrary. In 22 of the 35 tests of AUSM in this environment, the 100% efficient outcome was for the single-item demanders to win. AUSM produced that outcome only 45% of the time. In the other 13 tests the 100%-efficient outcome was for the demander of the whole to win. AUSM produced that outcome 100% of the time. It appears that in this extreme test for the existence of a threshold problem, there are signs that AUSM has one but that its effect on efficiency and revenue is low. The numbers for the simultaneous mechanism are almost the exact opposite of those for AUSM. In 17 tests of the simultaneous mechanism in this environment, the 100%-efficient allocation was for the single-item demanders to win. The simultaneous mechanism produced that outcome only 75% of the time.[23] In the other 7 tests, the demander of the whole should win to produce 100% efficiency. The simultaneous mechanism produced that outcome only once, or 14% of the time.[24] Clearly the data from this extreme testbed highlight the fact that each of the two mechanisms possesses an unmistakable bias. AUSM seems to be biased slightly in favor of large-package demanders, while the simultaneous mechanism seems to be biased seriously in favor of single item demanders. Nevertheless, if one is interested in generating highly efficient allocations, then AUSM clearly dominates in these tests.

Looking at the other data from the spatial tests, we note that in 90% of the tests AUSM produced revenue in excess of 50% of the maximum possible. In only 20% of its tests did the simultaneous mechanism yield more than 50% of the maximum possible revenue. With respect to bidders' surplus, the simultaneous mechanism yielded over 40% of the maximum possible surplus in over 50% of its tests, and over 30% of the maximum surplus in over 70% of its tests. AUSM, on the other hand, yielded over 40% of the maximum surplus in only 15% of its tests, and over 30% of the maximum in only 20%. A straightforward policy observation follows. If the surplus attained is all that is important to the potential participants in an auction, and if efficiency is allowed to take a back seat to self-interest, then bidders should be expected to argue for the simultaneous auction, while the seller should be expected to argue for the inclusion of package bidding.

3.2.2 Giving AUSM Its Best Chance. The spatial fitting testbed is probably the ideal example of a situation that gives the package-bidding auction its best shot at outperforming the simultaneous ascending-bid auction. It highlights situations in which the efficient allocations involve no single-item buyers and no buyers who want the entire collection. Rather, the efficient allocation is usually one in which two buyers each buy three of the six items. Also, and as importantly, the second-

23. This seemingly high failure rate of 25%, in situations for which the mechanism seems particularly suited (compare that with AUSM's 0% in its ''good'' environments), occurs because of the internal conflict faced by the bidder who is the high-value large-package demander. That bidder must choose at some point during the auction to go for one unit instead of all three. If the bidder waits too long to withdraw from his pursuit of three, then a misallocation of the single items can occur and that bidder can actually face losses. That seems to have happened quite often in these tests.

24. Compare that with AUSM's 45% in its ''bad'' environments.

best allocation (remember, that's the one that drives prices) also involves two buyers buying three items each but usually in a different configuration than that of the first-best allocation. So coordination is the key to success in these testbeds.

Turning to the data, one can see that the conflict between bidder and seller is no longer as sharp as it was in the spatial tests and, in a certain sense, can be said to be not there at all. Here, with rare exceptions, package bidding leads simultaneously to higher efficiency and higher revenue and higher bidders' surplus. The efficiency increase that occurs by including package bidding apparently creates enough surplus to allow both sides of the market to be better off. The numbers are again straightforward. In only two tests did AUSM fail to achieve 100% efficiency, while in only one of its tests did the simultaneous mechanism exceed 70% efficiency. So much coordination is needed to find both the best allocation and the next-best allocations that the single-price-per-item structure of the simultaneous auction doesn't have a chance. In the spatial fitting tests, AUSM yielded bidders' surplus less than 20% only twice. The simultaneous mechanism, on the other hand, managed to yield higher than 20% surplus only once. In fact, because bidders are exposed as they try to acquire packages in the simultaneous auction, many bidders actually lose money. In one test, the losses were so high that there was a *negative* bidders' surplus and only 70% efficiency; that is, bidders paid out more in revenue to the seller than they would make as owners of the items they bid on and won. These losses did not, however, yield the highest revenue to the seller. In most cases that was accomplished by allowing packaging.[25] In 90% of its tests, AUSM exceeded 50% of the maximum possible revenue. In only 20% of its tests did the simultaneous mechanism exceed 50% of the maximum possible revenue.

3.3 Summary of Our Observations and Their Support

In Section 1, we identified two major design choices that the data might help provide answers for. These were (1) sequential or simultaneous, and (2) package bidding or not. We have exhibited a lot of data from various testbed combinations of mechanisms and environments, many of which were created to influence those choices. Because the experiments were designed "on the fly," the data may not be as definitive as one might wish. On the other hand, we believe there is still enough

25. As a side note, the data also support the claim that allowing package bidding does not give the package of the whole any particular advantage.

evidence to allow us to make some observations that will stand up to further examination.

Observation 1: *If the items are homogeneous, then the answer to (1) is that it doesn't matter. The answer to (2) in these environments is unknown but probably unimportant.*

The support for this observation comes from the data in Section 3.1.1. In environments with multiple items to be allocated, if those items are homogeneous and substitutes, then little coordination between buyers is needed and the only role of the mechanism is to sort bidders with high values from bidders with low values. Both the sequential and simultaneous mechanisms seem to work very well at finding efficient allocations in these "easy" environments.[26]

Observation 2: *If the items are heterogeneous, then the answer to (1) is that simultaneous is better. If the extent of complementarity between items is small, the answer to (2) is that it probably doesn't matter.*

The support for this observation comes from the data in Section 3.1.2. In environments with multiple items to be allocated, if those items are heterogeneous, then some coordination among bidders is necessary to achieve high-value allocations even if there are only low synergy values. Simultaneous auctions provide a first step toward this coordination in a way that sequential auctions are unable to.[27] Packaging doesn't seem to either help or hurt relative to the simultaneous mechanism.

Observation 3: *If there are significant complementarities, then the answer to (2) is that package bidding is significantly better.*

The support comes from the data in Section 3.1.3. In environments with heterogeneous goods exhibiting complementarities, significant coordination is required for an auction or allocation mechanism to perform well with respect to efficiency or revenue. Sequential auctions perform poorly. Simultaneity is clearly necessary but not sufficient to attain high efficiencies. The simultaneous one-price-per-item auction seems to produce outcomes that are either high in efficiency, revenue, *and* losses or low in efficiency, revenue, and losses. Package bidding

26. We do not have any data on the performance of AUSM, the package-bid mechanism, in these simple environments with homogeneity, but we find it plausible that packaging could actually hurt here by guiding bidders to try to attain allocations that are inefficient. Of course, it is also plausible that AUSM, like the other mechanisms, will also perform well in these easy environments.

27. See Milgrom (1995, pp. 14, 15) for a concise discussion of why this might be so.

seems to help a lot in attaining high efficiency, high revenue, and no losses.

4. What Next?

In this section we try to point to some of the future research that we think is vital to the creation of better designs of complex auctions. There is a major gap between theory, scientific evidence, and practice in the design of these mechanisms. Until there are some serious breakthroughs in the theory of heterogeneous, multiunit auctions, it is also likely that experimental evidence will have to suffice.

4.1 Stopping, Activity and Withdrawal Rules

The simplest and most needed research is the testing of various straightforward variations in existing designs. Among these variations are withdrawal rules and stopping rules, including the various activity rules that have been proposed or used. Nothing systematic has yet been done to provide the research needed to answer questions that came up in the design of the FCC auction. For example, there are no serious theoretical discussions about whether withdrawal rules do any good at all and, if so, what are the better rules.[28] Discussions are naive. From an individual's standpoint, the possibility of withdrawal allows a bidder to be more aggressive,[29] to try risky fitting strategies at lower risk, and (maybe) to avoid losses incurred "by mistake." From a strategic point of view (i.e., when the reactions of the other players are also considered), some of these benefits may disappear. Losses occur for sure only when prices are high and the end of the auction is near, in exactly those cases in which no one is left to bail the loser out. However, it is still possible that the apparent reduction in risk will increase efficiency and revenue at the cost of increased losses. A second strategic effect is less benign. The lowering of the risk of loss could lead an opponent to try to drive the price of an item up to force you to give it up and, more importantly, because of that to release another item at a loss. This type of strategy can lower both efficiency and revenue. What will really happen remains to be carefully studied.

With respect to stopping rules, we also have virtually no system-

28. A first step in the direction of providing some scientific evidence about the effect of withdrawal rules can be found in Porter (1996).

29. It lowers the *expected* cost of not acquiring a piece of a package in a simultaneous auction.

atic data or theory that can provide guidance about what stopping rule should be used in which situation.[30] Stopping and its corollary, the encouragement of active bidding, remain very much an art both in the lab and in practice. The FCC chose to allow bidding to continue until no new bids are entered, over another serious proposal: that bidding stop item by item when no new bids are entered for that item. Their choice in turn required that some form of activity rule be designed that would force participation, since otherwise all bidders would have an incentive to wait for others to go first.[31] On the other hand, in markets for emissions permits for the Los Angeles basin we have, seemingly successfully, used a much different stopping rule.[32] Those auctions close at the end of a round if the aggregate value of the standing bids does not increase by more than 5% over the previous round.[33] This yields a much faster closure to the process and requires no activity rules other than the very natural and simple one that all high bids are binding and must be improved on to be displaced. What is not known with any certainty is whether this faster stopping creates more or less revenue or more or less efficiency than, say, the FCC rules, and whether such a finding would depend in any systematic way on the environment. With the extreme importance of these issues, it is very surprising that there is virtually no theoretical or experimental research.

4.2 Complexity

4.2.1 Computational and Strategic Complexity. Another level of open questions in the design of multiple-unit, heterogeneous goods auctions involves issues of complexity in both mechanism and environment. This is especially important when one begins to anticipate the impact of scaling up the experimental tests to something closer to the actual application. One of the issues one must face in comparing mechanisms, as in the question of whether to allow package bidding or whether to use a continuous or batch process, is how to judge the computational and strategic complexity of each approach. These concepts lie behind some of the discussions during the FCC design process but have never been satisfactorily defined and measured.[34] For exam-

30. A first step in this direction can be found in Kwasnica et al. (1997).
31. Banks et al. (1989) has some relevant observations on this phenomenon.
32. More details can be found at the World Wide Web address www.ace-mkt.com.
33. There is also a provision for a maximum number of rounds. In the LA emissions markets that we designed, this number is currently 5.
34. Some initial research has begun in this direction. For a leading example, see Rothkopf et al. (1995).

ple, batch processing gives all bidders time to think through their next response and so seemingly simplifies their problems; on the other hand, because of the sealed-bid nature of each round of batch processing, bidders have to anticipate their competitors' responses and never know for sure what a provisionally winning incremental bid is. As another example, package bidding provides bidders with a strategically easier way to coordinate their own bids and minimizes their exposure to losses; on the other hand, if used in batch mode, package bidding requires the auctioneer to use an optimization algorithm and confronts bidders with some coordination complexity when collections of small bids are needed to produce efficiency. There are solutions to many of these problems, but a complete study of the tradeoffs, including developing methods to measure the effects of the strategic complexity on bidders, is long overdue.

4.2.2 Environmental Complexity. Proposed mechanisms must be studied in more complex environments. Two variations that are obvious to consider include environments in which bidders have budget constraints and environments in which there are correlated or affiliated values across multiple, heterogeneous items. The first of these is important to study because it appears there were a number of bidding teams in the FCC auctions who were given a budget by their senior management and told to do as well as they could within that constraint.[35] One might expect this to be a common situation in large, complex auctions. It is our conjecture, and that of several others, that in an environment with budgets, package bidding is going to yield better performance than the simultaneous auctions. To the contrary, a few others have opined that with a good withdrawal rule the simultaneous auction will do as well as a packaging auction. The correct answer awaits further study.

Auctions for single items have been studied in great detail, both in theory and in experiment, in environments with correlated values. Evidence of the winner's curse has been found, and it has been shown that ascending-bid auctions allow better information aggregation than sealed-bid auctions. This was one of the compelling reasons behind the decision by the FCC to use an ascending-bid auction instead of a sealed-bid auction. One might hope that having multiple heterogeneous items would not change these results a great deal. But since sequential, simultaneous, and package-bidding auctions all provide different informa-

35. There are many reasons why such a constraint might exist, but a leading candidate would be principal-agent problems.

tion to bidders during the auction, it is possible that systematic differences in performance could appear in environments with correlated values that would reverse the findings above. Our conjecture is that this will not happen, but this needs to be studied both theoretically and experimentally before we can be sure.

4.3 Designs?

Finally, some purely speculative thoughts on what the future will bring in the design of auctions to price and allocate a large number of multiple heterogeneous items at one time. Under the rubric of moderate fixes, we think there are two that are the easiest and most productive. One would be the development of really good stopping rules. These would drive bidding activity without using complex eligibility requirements and activity rules, they would cause convergence to equilibrium reasonably rapidly, and they would not impose a lot of strategic complexity on the bidders. The second development would be user-friendly package bidding. This would reduce the seeming complexity facing the bidder while allowing the significant improvements in revenue, efficiency, and bidder surplus that such bidding creates.[36]

In the more speculative realm of really new approaches, we suggest one. In Banks et al. (1989) we considered a number of mechanisms and chose AUSM on the basis of the evidence there. It has proven to be a flexible, successful mechanism in many applications. But there was another mechanism we considered, the iterative Vickrey mechanism. In that design, we tried to capture some of the demand-revealing aspects of Vickrey's original mechanisms (see Vickrey, 1961), while introducing some of the cognitively easier aspects of simple iterative bidding found in standard English auctions.[37] We were looking for demand revelation because we believed that if there were little strategic loss from bidding one's true values, then the strategic complexity of simultaneously bidding for multiple items would be significantly reduced and good performance would be more likely. We also believed that with the appropriate iterative procedure, bidders would not need to submit bids for all packages in each round (a possibility that would destroy the imple-

36. We continue to believe that the unsupported claims by some that package bidding is ''too complex'' are exaggerated and unfounded. Our experience in the applications of large auctions with truckers (an 800-item auction for Sears) and with environmental engineers (the ACE pollution permit market) suggests the contrary.

37. Some of these ideas can be found in Rassenti et al. (1982), but they only consider a sealed-bid mechanism, which doesn't appear to do the desired job. Some of these ideas can be found in Ausubel (1996), but the mechanism there seems extremely complex and difficult to implement even in a very simple testbed.

mentability of the mechanism). Because we were unable to provide appropriate commitment rules,[38] the mechanism did not perform as well as we had hoped. However, we believe there is still an iterative Vickrey design to be found that will minimize strategic complexity, allow package bidding, and provide excellent efficiency and revenue performance. If so its performance could easily surpass that of all of the mechanisms studied in this paper.

Appendix

A.1 Some Background Methodology on Applied Mechanism Design

The FCC auction designers' problem was to create a mechanism to allocate and price a number of heterogeneous items. The goal, at least as initially stated by the FCC, was to allocate those items to the highest-value users.[39] The basic problem, common to most mechanism design efforts, was that the information needed to solve this problem (the values of the items) was best known, if at all, by the various potential users and not by the FCC. Further, none of these potential users had any incentive to precisely reveal their information to the FCC. There is a standard solution to this problem that has been developed over a number of years of basic research in economics and other disciplines: *If one can predict the performance of various mechanisms over a range of possible user values for the items, then one does not need to know the details of the specific values to achieve one's goals; one need only select the appropriate mechanism to achieve the desired outcome*.

The idea is simple. A mechanism, such as a particular auction format, works as follows. Participants bring their own information and valuations to the auction. The auction is then held, and the participants use their information to determine how they interact with each other through the mechanism. The interaction between individual behavior and the auction rules produces an allocation of the items and payments for those items. Operating the same auction on a different constellation of values will generally produce a different allocation and payment distribution. Operating a different auction format on the same constellation of values across individuals will generally also produce a different allocation and payment distribution. We refer to the relationship a

38. These would be rules that require bids to be somewhat binding so as to prevent cheap-talk uses of the bidding process, which in turn could prevent the auction from converging.

39. See, e.g., Milgrom (1995, pp. 13–14).

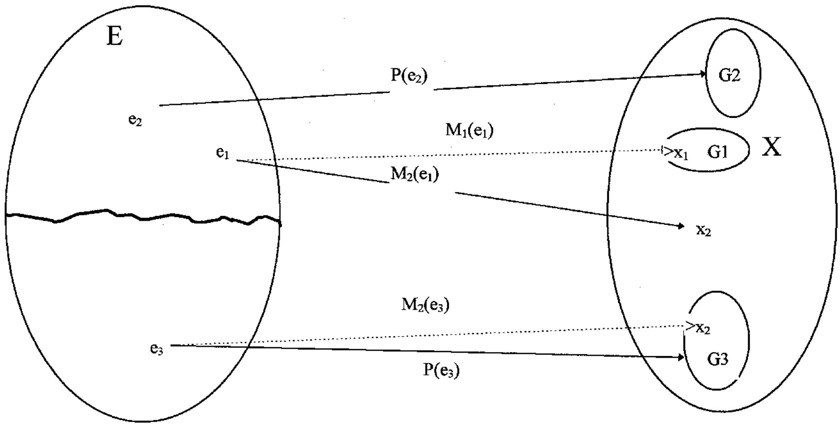

FIGURE 10.

mechanism creates between the particular constellation of values and the allocation and payments as the *performance of the mechanism.*

Consider the simple diagram in Figure 10. *E* is to be thought of as a set of possible constellations of values with a single point in *E* representing the true valuations. We call elements of *E* *environments.* *X* is to be thought of as the set of possible allocations of items and payments for those items that might result. We call elements of *X* *outcomes.* A mechanism's *performance* is then a mapping from *E* to *X*. So, for example, mechanism *M*1 produces outcome *x*1 if the environment is *e*1, while mechanism *M*2 produces outcome *x*2 in that same environment. The policy issue in mechanism design is to determine the standard of performance that is desired for the mechanism to be chosen. For example, should the mechanism try to produce an outcome that maximizes the aggregate value of the allocation? Such a standard will usually be another mapping from some part of *E* to *X* and will look like *P* in Figure 10. So the performance standard *P* asks the outcome to be something in the set *G*2 if the true environment is *e*2 and something in the set *G*3 if the environment is *e*3. If the scientific evidence can establish the performance of various mechanisms and if the performance of at least one of those mechanisms is consistent with the performance standard of the decisionmaker over the part of *E* within which the decisionmaker thinks the true valuations lie, then the design problem has been solved. In Figure 10, if the decisionmaker thinks that the true *e* is somewhere in the upper half of *E*, if that decisionmaker has the performance standard P in mind, and if $M1$ produces something in $P(e)$ for all of the e in the top half, then *even though no one knows the true constellation of values*, the mechanism design problem is solved by using $M1$.[40]

The *science* of applied mechanism design, then, is focused on providing the best evidence possible on the performance of mechanisms in a variety of environments. The *policy* of applied mechanism design is focused on using those findings to pick the most appropriate mechanism for the situation. In this paper we concentrate on the science only.

A.2 Possible Problems with Efficiency as a Performance Measure

It has been correctly pointed out by some that the absolute value of this measure is not particularly illuminating. For example, even if you knew that a mechanism produces 95% efficiency on average over a class of environments, there would still be no basis for you to know whether this is good or bad. An example easily illustrates this. Suppose there are two items to allocate, A and B. Further suppose bidder 1 is to be paid \$6 if she gets A, \$10 if she gets B, and \$20 if she gets both. Suppose bidder 2 is to be paid \$4 if he gets A, \$15 if he gets B, and \$20 if he gets both. The optimal allocation is that 1 gets A and 2 gets B, for a total profit (before payments) of \$21. If, instead, the actual outcome were that 1 gets B and 2 gets A, the profit (before payments) would be \$14, for an efficiency of 66%. Now suppose I wanted to make this look a little better. I could simply add \$300 to each possible payoff. This would not change the incentives to each agent (it is only a lump-sum payment), but it would yield a significantly better-looking efficiency measure, $(310 + 304)/(315 + 306) = 98.87\%$. So the absolute number is of little value. However, comparing values across mechanisms in the same environments can be informative. For example, if I were to tell

40. Decisionmakers involved in policy should be forewarned. Even with the best scientific evidence about the performance of mechanisms, arguments about the possible location of the true state of the world can derail good intentions. If the future participants are asked to provide advice during the design phase and if these participants know something about the true e, they may have an incentive to provide arguments intended to improve their final allocation. So in Figure 10, a potential participant may know that e is truly in the top half of E. If that participant likes $x2$ better than $x1$ and if that participant knows the policymaker's performance standard is P, then that participant may argue strongly that "e must be in the bottom half." If that argument is successful, then the decisionmaker will select mechanism $M2$, since $M2(e)$ is in $P(e)$ for e in the lower half of E. But then when the mechanism $M2$ is run in conjunction with the true world, say $e1$, the outcome $x2$ occurs. The participant is better off; the policymaker may not be.

you that mechanism $M1$ produced observed efficiencies between 90% and 96% and that mechanism $M2$ produced efficiencies between 86% and 89%, for the same structure of payments, you would be justified in concluding that $M1$ outperformed $M2$ on that class of environments with respect to attaining efficient allocations.

A.3 Details on the Environments

A.3.1. Easy Environments.

A.3.1.1 Additive Values. The first easy environment examined is one in which values are additive, i.e., the value function is of the form $V_i(x_{1i}, \ldots, x_{ni}) = \sum_j V_i(x_{ji})$ In the experiments, there were six items for sale to six demanders. Each demander knew that his value and those of other participants were to be drawn uniformly, with replacement, from a fixed list of ten value sheets, shown in Table IV. In this environment, a competitive equilibrium always exists for each item, and its competitive-equilibrium price is equal to the second-highest value for that item. For example, if a participant drew sheet 4, his profit before any payment for item c would be 900 and his gross profit for item a would be only 100. If he obtained both items a and c he would be paid a gross profit of 1000. The subjects also knew the distribution of the possible value draws. That is, they knew Table IV and the process by which values were assigned to each subject.

A.3.1.2 Decreasing Values. Items in this environment are *homogenous*, i.e., $V_i(x_1, x_2, \ldots, x_n) = V_i(\sum_j x_j)$ with $V_i'(\sum_j x_j) \leq 0$, so that we have a downward-sloping demand with a competitive equilibrium (CE). One of the demand conditions we used is given in Figure 11. In this environ-

TABLE IV.
Value-Sheet Space

Item Sheet	a	b	c	d	e	f
1	900	450	400	350	300	250
2	400	600	800	600	400	200
3	800	600	400	200	400	600
4	100	100	900	400	300	200
5	400	800	400	200	0	200
6	900	600	300	200	100	0
7	300	300	300	300	300	900
8	750	250	250	750	400	400
9	400	200	400	600	800	600
10	850	350	350	650	150	150

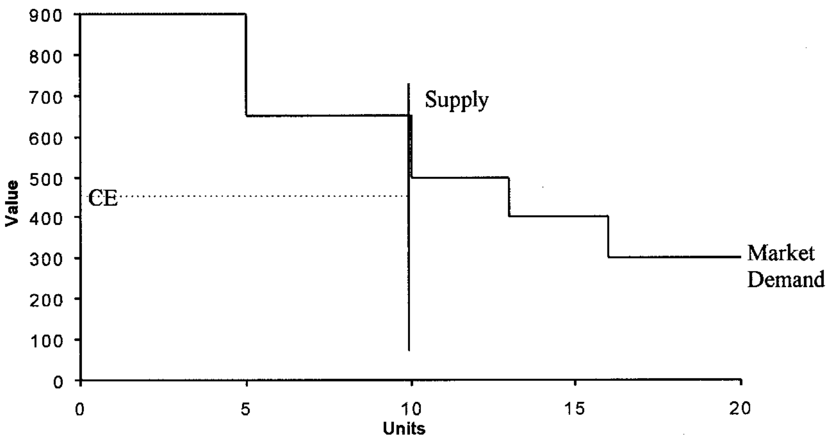

FIGURE 11. DECREASING-MARGINAL-VALUE ENVIRONMENT

ment there were six participants and ten units to be allocated. Each participant had decreasing demands for up to four units.

A.3.1.3 Superadditive Values with Competitive Equilibrium. The items are homogenous but we have $V_i'(\sum_j x_j) \geq 0$. In addition, there is a price that clears the market. An example is given in Figure 12. Thus, while

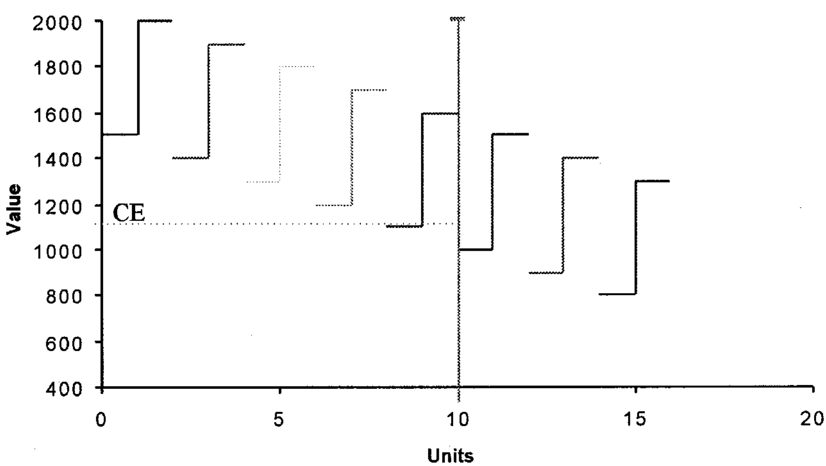

FIGURE 12. SUPERADDITIVE ENVIRONMENT WITH CE

there exists a risk of assembling the package of items, a single (CE) price exists. In the figure, each step represents a participant's marginal return function for two units. Thus, there were eight participants in the experiments.

A.3.2 Moderate Environments

A.3.2.1 Superadditive Values without Competitive Equilibrium. An example is given in Figure 13. There are eight participants. In the figure, each step represents a participant's marginal return function for the first three units. The marginal return for more than three units is zero. Notice that in this environment there is no single price equilibrium. As soon as bids go over 100, losses must occur. When price is above 100, bidder 5, for example, might then withdraw his bid and accept any price above 70. This result occurs because the mechanism does not allow nonlinear pricing. Thus, either the outcome must result in losses, or at least one bidder must forgo the pursuit of potentially profitable opportunities.

In the experiments, subjects were provided with the following common information about their environment: the number of units being auctioned, the number of participants, and who had the standing

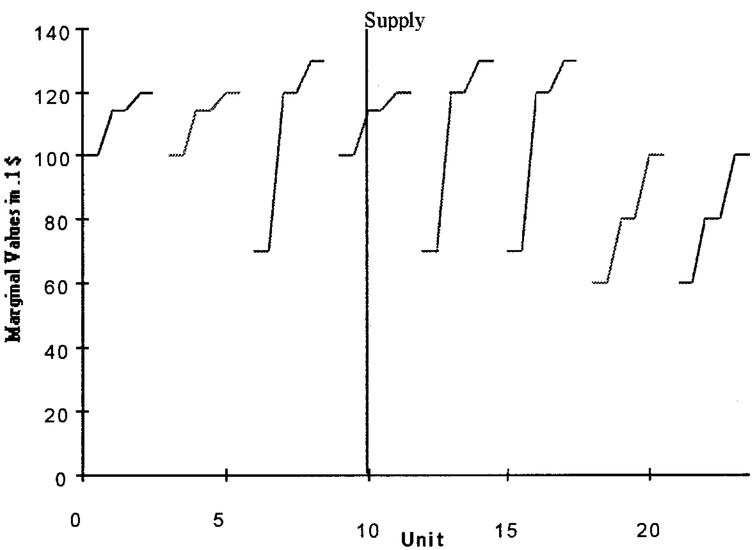

FIGURE 13. SUPERADDITIVE ENVIRONMENT WITH NO CE

bid on which items. The subjects did not know the distribution over which the values were drawn or even if a single price could clear the market.

A.3.2.2 The Assignment Problem. The environment here is exactly that described in Olson and Porter (1994). In these experiments we used the same parameters as in the additive-value environment (see Table 11 for values) with the added restriction that demanders can use one and only one item to make a profit. In addition, we used the same common knowledge structure where each demander knew that his value and those of other participants were to be drawn uniformly, with replacement, from the fixed list of ten value sheets in Table IV.

A.3.3 Hard Environments. These environments are to provide boundary cases that test the robustness of mechanisms that do not allow individuals to package demands. In general they are extensions of the case described above with more constraints and priors over values provided to participants. We detail the three specific environments we used below.

A.3.3.1 Spatial Demands. Values for three items called *a*, *b*, and *c* along with a value for the full package *abc* were drawn from a common knowledge distribution as follows:

1. The integer values for the single items are drawn independently from a triangular distribution with support [0,98] (Fig. 14).
2. The value for the package *abc* is then determined by adding a number randomly selected from the interval [0,149] to the highest value for *a*, *b*, or *c* in step 1.

This parameter set can generate values in which a competitive equilibrium price for each of the items *a*, *b*, and *c* exists or not.

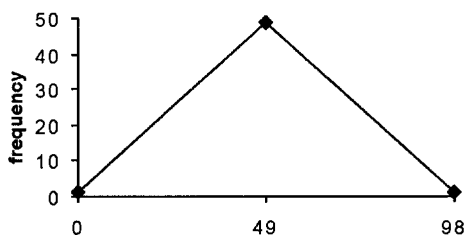

FIGURE 14.

A.3.3.2 Spatial Fitting. This environment is one in which individuals have nonadditive preferences for specific packages of items and must find how they fit together. There are five participants and six heterogeneous items to allocate. The structure of the problem is as follows:

1. The single-item packages called *a*, *b*, *c*, *d*, *e*, and *f* have their integer values drawn independently from the uniform distribution with support [0,10].
2. The two-item packages $\{a,b\}$, $\{a,c\}$, . . ., $\{e,f\}$ have their integer values drawn independently from the uniform distribution with support [20,40].
3. The three-item packages $\{a,b,c\}$, . . ., $\{d,e,f\}$ have their integer values drawn independently from the uniform distribution with support [110,140].
4. A single value is given for the six-item package $\{a,b,c,d,e,f\}$ drawn from [140,180].

A total of 25 unique packages from the total possible generated from steps 1–4 above were given to participants. The main point to note is that two three-item packages clearly form the largest total value. However, this optimal package configuration is likely to be overlapped by many other competing packages. The task of the mechanism is to guide the owners of the optimal packages to find each other.

References

Ausubel, L.M., 1996, ''An Efficient Ascending-Bid Auction for Dissimilar Objects,'' Preliminary paper, University of Maryland, January.

Banks, J., J.O. Ledyard, and D. Porter, 1989, ''Allocating Uncertain and Unresponsive Resources: An Experimental Approach,'' *RAND Journal of Economics,* 20, 1–23.

Bykowsky, M.M., R.J. Cull, and J.O. Ledyard, 1995, ''Mutually Destructive Bidding: The FCC Auction Design Problem,'' Caltech Social Science Working Paper 916, January.

Calsamiglia, X., ''Decentralized Resource Allocation and Increasing Returns,'' *Journal of Economic Theory*, 14(2), 263–283, April.

Jordan, J.S., 1987, ''The Informational Requirements of Local Stability in Decentralized Allocation Mechanisms,'' in T. Groves, R. Radner, and S. Reiter, eds, *Information, Incentives, and Economic Mechanisms*, Minneapolis: University of Minnesota Press.

Kwasnica, A.M., J.O. Ledyard, D. Porter, and J. Scott, 1997, ''The Design of Multi-round Multi-object Auctions,'' Preliminary Paper, delivered to the Public Choice Society, March 21.

Ledyard, J.O., C.R. Plott, and D. Porter, 1994a, ''A Report to the FCC and to Tradewinds International Inc.: Experimental Tests of Auction Software, Supporting Systems and Organization,'' Final Report for Subcontract Number 005 (Prime Contract FCC-94-12), September.

——, D. Porter, and A. Rangel, 1994b, ''Using Computerized Exchange Systems to Solve an Allocation Problem in Project Management,'' *Journal of Organizational Computing*, 4(3), 271–296.

——, C. Noussair, and D. Porter, 1996, "The Allocation of a Shared Resource Within an Organization," *Economic Design*, 2, 163–192.

Milgrom, P., 1995, "Auctioning the Radio Spectrum," in *Auction Theory for Privatization*, Cambridge University Press, Chapter 1.

Mount, K.R., and S. Reiter, 1996, "A Lower Bound on Computational Complexity Given by Revelation Mechanisms," *Economic Theory*, 7(2), 237–266.

Olson, M. and D. Porter, 1994, "An Experimental Examination into the Design of Decentralized Methods to Solve the Assignment Problem with and without Money," *Economic Theory*, 4, 11–40.

Plott, C.R., 1994, "Market Architectures, Institutional Landscapes and Testbed Experiments," *Economic Theory*, 4, No. 1.

——, 1997, "Laboratory Experimental Testbeds: Application to the PCS Auction," *Journal of Economics & Management Strategy*, this issue.

Porter, D., 1996, "The Effect of Bid Withdrawal in a Multi-Object Auction," Working Paper 982, California Institute of Technology.

Rassenti, S.J., V.L. Smith, and R.L. Bulfin, 1982, "A Combinatorial Auction Mechanism for Airport Time Slot Allocation," *Bell Journal of Economics* 13, 402–417.

Rothkopf, M.H., A. Pekec, and R.M. Harstad, 1995, "Computationally Manageable Combinatorial Auctions," Preliminary Paper, April.

Vickrey, W., 1961, "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, Vol. 16 (May), 8–37.