A New and Improved Design for Multi-Object Iterative Auctions

Anthony M. Kwasnica
John O. Ledyard
David Porter
Christine DeMartini*

April 25, 2002

Abstract

In this paper we present a new improved design for multi-object auctions and report on the results of experimental tests of that design. We merge the better features of two extant but very different auction processes, the Simultaneous Multiple Round (SMR) design and the Adaptive User Selection Mechanism (AUSM) of Banks, Ledyard, and Porter (1989). Then, by adding one crucial new feature, we are able to create a new design, the Resource Allocation Design (RAD) auction process, which performs better than both. We are able to demonstrate, in both simple and complex environments, that the RAD auction achieves higher efficiencies, lower bidder losses, and faster times to completion without increasing the complexity of a bidder's problem.

1 Introduction

Theory, experiment and practice suggest that, when bidder valuations for multiple objects are super-additive, combinatorial auctions are needed to increase efficiency, seller revenue, and bidder willingness to participate (Bykowsky, Cull, and Ledyard 1995, **?**, Ledyard, Olson, Porter, Swanson, and Torma forthcoming). A combinatorial auction is an auction where bidders are allowed to express bids in terms of packages of objects. The now famous FCC spectrum auctions are are good example of the relevance of these issues. In 35 auctions events before 2002, the FCC used what is known as a Simultaneous Multiple Round (SMR) auction to allocate bandwidth. This auction format does not allow package bidding. It is reasonable to expect that some bidders might receive extra benefits by obtaining larger, more contiguous portions of the spectrum. Many of the auctions divided the spectrum by geographic location. A firm might enjoy cost savings if they could purchase two adjacent locations. However, without package bidding, a bidder cannot express that preference potentially lowering the efficiency of the auction. If the bidder attempts to acquire both licenses through bidding on the licenses individually, they might be forced to expose themselves to potential losses. The high number of bidder defaults on payments might, in part, be evidence of losses caused by the lack of package bidding. In response to these difficulties, the FCC plans on allow package bidding in future auctions.

While the potential utility of combinatorial auctions is considerable, combinatorial auction have not yet been used widely in practice.[1] The successful implementation of a

*The authors thank Evan Kwerel, Paul Milgrom and Charles Plott for their helpful comments.

combinatorial auction requires one to overcome a number of hurdles. Two widely recognized and discussed issues are:

i. The computational complexity of the winner determination problem, and

ii. The complexity of the bidding environment for the bidder.

Computation complexity comes from the fact that determining a set of winning bids - those that maximize the sum of the bid prices subject to feasibility constraints - is NP-complete. While Rothkopf, Pekec, and Harstad (1998) has shown that computational issues can be limited via limitations on acceptable bids and other strategies, and others have found promising algorithms (Sandholm, Suri, Gilpin, and Levine 2001). Significant computational challenges still persist - particularly when the number of objects available for bid grows and bids are likely to express high degrees of substitutability.

A combinatorial auction is also difficult for a bidder. Bidders must formulate valuations for a large number of potential bid combinations, and then they must formulate an optimal bidding strategy given those valuations. The fact that game theorists have had little success characterizing optimal strategies in simple auctions with more than two objects is testament to the difficult of this problem for the bidders.

This paper reports on a new combinatorial auction design called Resource Allocation Design (RAD) that, we believe, demonstrates the practicality of designing an iterative combinatorial auction that could easily be scaled up to handle large allocation problems. The design realizes that an important part of the auction design is the clarity and usefulness of the information provided to the auction participants. The strong performance of RAD in the laboratory setting suggests that one can economize on the complexity of the information provided to the bidders without significant negative consequences.

One area where these computational obstacles have their greatest impact is in the determination of prices in the auction. While it might seem logical to simply let the prices equal the bids that winning bidders placed, this approach provides limited information to bidders. Ideally, one would like to come up with prices that make the bidder's problem simpler.

Game theory provides one suggestion - construct Vickrey prices for the auction. Vickrey prices are personalized prices on all bids placed that have been shown to eliminate all strategic incentives for bidders. Assuming bidders have correctly formulated their valuations, bidders should be willing to submit a full, honest report of those values to an auctioneer who has committed to Vickrey prices. So, in theory, if bidders are faced by Vickrey prices they should realize that their only strategy is to simply bid their value for each particular combination they value. This approach has a number of drawbacks. First, calculation of these prices requires the auctioneer to calculate an additional NP-complete mixed integer program for each bidder.[2] Second, it is not clear that bidders will interpret these bids correctly. Even in Vickrey auctions for one object, bidders systematically deviate from the ideal strategy of simply stating ones value. Finally, if one views bidders as 'learning' about their valuations and profitable bids in the course of the auction, Vickrey prices provide little information to the bidder about potentially profitable combinations of new bids.

[1]There are some notable exceptions including Sears Logistics Services, the Automated Credit Exchange and the upcoming FCC auctions.

[2]The calculation is essentially the winner determination problem with the bids of the bidder in questions removed from the calculations.

A second potential approach is closely related to the economic theory of competitive equilibrium. A set of prices - one for each object - are said to constitute a competitive equilibrium if, given these prices, the supply of objects equals demand (i.e., excess demand is zero). Competitive equilibrium prices could be calculated in the combinatorial auction setting. However, in order to create such prices in a combinatorial auction, one would be required to create a price for each potential combination of bids. In a large auction, the number of bids on a particular combination is likely to be sparse. This might cause nearly any price to be a competitive equilibrium price limiting the informational content. For example, if a bidder won the combination of objects AB and no other bidders placed a bid on that specific combination, then a feasible set of competitive prices would have the price anywhere between 0 and the winning bid on AB. Another difficulty with the formulation of competitive prices on all combinations is that, even if one is able to formulate meaningful prices, the cognitive burden on the bidders would be quite high. It is not clear that bidders would find it useful - or have the inclination to - examine all these prices.

The approach we take in this paper is to take the issues of computational and cognitive complexity seriously and to formulate a combinatorial auction mechanism that attempts to ease those burdens. We present an iterative combinatorial auction that presents approximately competitive equilibrium prices on each object to the bidders. These prices help ease the two practical design issues discussed earlier. First, the computation of prices occurs by completing a series of nearly instantaneous linear programs. Therefore, the auctioneer needs to conduct only one NP-complete computation, the winner determination itself. Second, the prices present information on the level of bidding for all objects in a manner that is simple and natural for the bidders. Instead of looking at prices on all subsets, the bidders are present with a price for each object. There is no reason, in theory, to expect these prices to work well. However, we demonstrate, through the use of human subject experiments, that these auctions can perform quite well under a number of reasonable performance measures.

How should one evaluate different auction designs? The theory has proven intractable. Computer simulations fail to capture the important details behind human cognition. When auctions occur in practice it is difficult if not impossible to learn much other than that they worked. Because the true values of the items to the bidders are unknowable, it is not possible to measure such important performance indicators as efficiency, maximum possible revenue, or bidder losses. So if progress is to be made in studying complex auctions in complex environments, we must turn to the laboratory. The use of the laboratory as a test bed for complex auctions in complex environments began with Ferejohn, Forsythe, and Noll (1979), Smith (1979), Grether, Isaac, and Plott (1981) and Rassenti, Smith, and Bulfin (1982). This methodology has proven to be fairly successful in providing guidance for the design of a variety of implemented auctions (Plott 1997, Ishikida, Ledyard, Olson, and Porter 1998, Ledyard, Porter, and Rangel 1997). Building on knowledge from theoretical and practical experience, one can create test bed environments in the lab which exhibit as much complexity or simplicity as one wishes. In these environments, one can test any auction. With laboratory control, one can calculate performance measures unknowable in the field. One can precisely answer questions such as: did the highest value bidders win the items, was there a bidder who wanted a particular configuration and did not get it, and were there bidders who, because of the auction design, bid more for an item than it was truly worth to them?

We evaluate the RAD design in the lab using both a complex and simple test bed. We are able to compare the performance of RAD to a version of the SMR auction used by the FCC. Since we used the same test bed as previous experiments, we are also able to compare the performance of RAD, when possible, with that of the Adaptive User Selection Mechanism (AUSM) originally proposed by Banks, Ledyard, and Porter (1989), which is widely regarded as one of the first combinatorial auction mechanisms.

In Section 2, we describe the background of our search for a high performance multi-object auction design. In section 3, we formally describe the SMR and AUSM designs and the RAD auction. In Sections 4 and 5, we describe the test bed and the performance measures we use to evaluate the design. In Section 6, our findings are offered. Finally, in Section 7, we provide our conclusions and work that remains to be done.

2 The Context

As most theorists realize, it is relatively simple to describe a demand-revealing, efficient auction for this problem. A natural extension of the famous Vickrey sealed-bid auction will award the objects to the highest valuing bidders and eliminate any incentive for them to misrepresent their preferences. But if K items are being auctioned, each agent's bid would need to be 2^K numbers – potentially creating a very large, very complex problem. Further, if there is any affiliation in the values of bidders then sealed bid auctions of this sort are thought to be less efficient than auctions that allow bidders to learn as they bid (Milgrom and Weber 1982).[3] Even when only one object is for sale, bidders in experimental sessions often do not understand the demand revealing incentives of the Vickrey auction (Kagel, Harstad, and Levin 1987). Progressive auctions, such as the English auction, usually perform quite well in the laboratory (Coppinger, Smith, and Titus 1980). The auction design problem we address here is to find a suitably defined progressive auction that performs well when multiple heterogeneous objects are for sale and bidders have complementarities in their preferences.

There are two types of progressive auctions one might consider: iterative and continuous auctions. A continuous auction is similar to the classic English auction where bids can be submitted at any time. Iterative auctions proceed in a series of rounds, which last a specified period of time. During a round, bidders have the opportunity to place bids before the auctioneer considers any of the bids placed in the round. Once a bid is submitted, in the case of a continuous auction, or at the end of a round, in the case of an interative auction, the auctioneer processes the bid(s) and identifies *provisionally winning* or standing bids. These are the bids that will win if no new bids are forthcoming. In all cases, the auctioneer then provides information back to the bidders. The process repeats until a *stopping rule* is satisfied. At that time the provisionally winning bids become winning bids. Readers familiar with mechanism design theory will recognize an iterative auction as a special case of a resource allocation process as originally described by Hurwicz (1960). Readers familiar with experimental economics will recognize an iterative auction as a special case of a microeconomic system as originally described by Smith (1982).

To understand the possibilities and choices facing the designer of multi-object auctions, we begin by recalling the key features of two vastly different designs: the Simultaneous Multiple Round (SMR) design (Milgrom 1995) and the Adaptive User Selection Mechanism (AUSM) (Banks, Ledyard, and Porter 1989).

[3] Dasgupta and Maskin (2000) show that the Vickrey auction could, in theory, be extended to this setting by allowing for bids to be functions that allow each bidder to state what his value would have been were the other bidders' information revealed.

The SMR design allows only single-item bids, is iterative, and has an eligibility based stopping rule (i.e., a *use-it-or-lose-it* feature) driven by a minimum increment requirement for new bids. The SMR design was used extensively by the FCC to run early bandwidth auctions. On the other hand, AUSM allows package bids, is continuous, and is stopped at the discretion of the auctioneer. An iterative version of AUSM was used by Sears Logistics Services to procure trucking services (Ledyard, Olson, Porter, Swanson, and Torma forthcoming). Three aspects of the design are the same for each: Each winning bidder pays what they bid, provisionally winning bids are determined by maximizing potential revenue subject to feasibility, and provisionally winning bids remain as a standing commitment until replaced by another provisional winner.

Each auction represents a compromise, which is the result of a sequence of choices of design characteristics. Each choice often leads to one side of a seeming unavoidable trade-off. Therefore each auction process has its potential weaknesses. In this paper we focus on potential failures in performance in the areas of efficiency, revenue, bidder losses, complexity, biases and the time to complete an auction.

In most discussions of the design of multi-object auctions, the primary goals have either explicitly or implicitly been high efficiency and/or high revenue. There is no fundamental conflict between efficiency and revenue. In fact, in single-item auctions, maximal revenue usually occurs by maximizing efficiency and then extracting as much of the surplus as possible (Myerson 1981).[4] It remains to be proven however that this approach works in multi-object auctions.[5] In environments without income effects, such a quasi-linear preferences, what trade-off there is can be most easily seen in the following identity:

$$\text{Efficiency} \times \text{Maximal Possible Surplus} \equiv \text{Seller's Revenue} + \text{Bidders' Profits}.$$

High efficiency and low revenue can occur if and only if bidder profits are high. And high revenue and low efficiency can occur if and only if bidders have losses.

Both the SMR and AUSM auction processes have a difficult time consistently generating 100% efficiency across a variety of environments (Ledyard, Porter, and Rangel 1997, Kwasnica, Ledyard, Porter, and Scott 1998). The SMR mechanism, because it only allows single-item bids, faces the exposure problem. The exposure problem occurs in situations where bidders' values are super-additive. In order to win a package, which the bidder values more than the sum of the individual items in the package, the bidder might need to bid above her value on the individual items. If the bidder does not end up winning the package, this will cause losses for the bidder. Or, bidders who are aware of this problem might stop bidding in order avoid the risk of losses causing low efficiencies and seller revenue. To combat the exposure problem, the FCC allowed provisional winners to withdraw with a penalty. Porter (1997) analyzes the effect of this rule and finds that although efficiencies are higher, so too are bidder losses. The AUSM mechanism, because it allows package bids, does not suffer from the exposure problem but faces the threshold problem. The threshold, or free rider, problem occurs when a number of bidders for small packages must coordinate their efforts to unseat a bidder for a big package. In this situation each bidder has the incentive to allow the other bidders to be the ones who increase their bid

[4]The optimal auction does not require that the auctioneer always sell to the highest valuing bidder. In some cases, the auctioneer may want to comit to a non-negligible reserve price. However, contingent upon completing a sale, the auctioneer wants to maximize allocative efficiency.

[5]In spite of Williams (1994) who identifies the optimal, efficient auction to be a Vickrey-Groves mechanism, we do not know that the optimal auction is always efficient. In fact Armstrong (1998) suggests it may not be so.

in order to beat the big bidder. In principle, all bidders may fail to raise their bids allowing the big package to win even if it should not have. The threshold problem may cause low efficiencies as collections of small bidders may not be able to coordinate their bids to dislodge a large, inefficient bidder. To combat the threshold problem, AUSM is often used with a standby queue – a public bulletin board on which potentially combinable bids can be displayed.[6] The use of a queue, however, shifts the computational burden to the bidders; they must now consider the bids in the queue - an NP-complete problem in itself - when making a new bid.

While it is easy to measure efficiency, seller revenue and bidder losses, it is harder to measure the complexity of a mechanism or the costs of the length of time to complete the auction.[7] Nevertheless, we can make a few observations about the performance of the SMR and AUSM designs. Because the SMR auction proceeds in measured steps and because bidders seem to have a relatively simple information processing problem at each step, most consider it a simple mechanism.[8] But, because of this slow but steady approach, SMR auctions can take a very long time to complete. The FCC's Broadband PCS D,E, and F Block auctions lasted 276 rounds spanning 85 days. AUSM proceeds in an apparent disorganized manner with bids allowed in any order, stopping when no new bids are forthcoming or the auctioneer deems the auction to be at an end. Because of this, AUSM finishes quickly. But many feel that this places a difficult information processing burden on bidders that, together with the standby queue, makes AUSM a very complex mechanism.

So each mechanism has both desirable and undesirable performance characteristics. The obvious question then is: can we do better than both? In particular can we take the successful design aspects of each, perhaps augment them a bit, and create a hybrid that dominates both? Based on the research reported in this paper, we suggest that the answer is yes.

The challenge is to take what we have learned from our experience with both the AUSM and SMR designs and to create something better. To do so requires a series of choices of design features. Here we explain the choices we made to end up with the better performing RAD mechanism.

3 The Auctions

Rather than providing a fully general framework, in this paper we will focus on the designs we evaluate. Let $I = \{1, \ldots, N\}$ represent the set of bidders, $K = \{1, \ldots, K\}$ represent the set of objects to be sold and $t = 1, 2, 3, \ldots$ represent the iterations or *rounds*. In general, a bid can be a very abstract entity involving complex contingent logic.[9] In this paper, we restrict our attention to very simple bids. A bid is a pair $b = (p, x)$ where p is positive real number representing the bid price and $x \in \{0, 1\}^K$ represents the items desired.[10] A bid signifies, "I am willing to pay up to p for the collection of objects for

[6] It is shown in Banks, Ledyard, and Porter (1989) that the queue increases both efficiency and revenue in continuous auctions.

[7] In the field, it is difficult, if not impossible, to measure any of these variables. In the lab, since we know the induced valuations of bidders, we can directly measure efficiency, revenue and losses.

[8] Formulating an optimal strategy to win packages of items when bidding is restricted to single-item bids is quite difficult.

[9] See, for example, Ishikida, Ledyard, Olson, and Porter (1998), Rassenti, Smith, and Bulfin (1982), and Grether, Isaac, and Plott (1981).

which $x_k = 1$ if and only if I get all of them."

Begin by assuming we are in round t and all N bidders have submitted their bids. The set of bids placed by bidder i in round t is B_t^i, and $B_t = \cup_{i\epsilon I} B_t^i$ is the set of all submitted bids indexed by j. An arbitrary element of B_t is expressed as $b_j = (p^j, x^j)$.

All the auction designs use a straight-forward allocation rule: provisionally award the items to the collection of bids that would yield the highest revenue. We solve the following allocation problem:

$$\max \sum_{j \in B_t} p^j \delta^j \tag{1}$$

subject to

$$\delta^j \in \{0,1\} \text{ for all } j \in B_t$$

and

$$\sum_{j \in B_t} x_k^j \delta^j \leq 1 \text{ for all } k = 1, \ldots, K.$$

If there is only single-item bidding, this simply selects the highest bidder for each item. With package bidding, the combinatorial optimization problem is equivalent to a set-packing problem on a hypergraph, which is known to be NP-complete. If the number of objects or bids are large, an optimal solution cannot be guaranteed in finite time. Computation was never an issue in any of the results reported in this paper due to the relatively small scale of the test cases examined (10 objects and 5 bidders).

Let δ_t^* be a solution to this problem. If $\delta_t^{*j} = 1$ we say that bid j is *provisionally winning* in round t. Let $W_t = \{(p^j, x^j) \in B_t \mid \delta_t^{*j} = 1\}$ be the collection of provisionally winning bids. Then i's winning bids are the set $W_t^i \equiv B_t^i \cap W_t$. An obvious initial condition is to have $W_0 = \emptyset$.

If the auction stops at this round, for each $j \in W_t^i$, bidder i will receive the items for which $x_k = 1$ and will pay p^j to the auctioneer. If the auction does not stop, then all provisional winners are automatically resubmitted in round $t+1$, or $W_t^i \subseteq B_{t+1}^i$ for all i.

3.1 The SMR Auction

The basic SMR auction design requires only a few new rules in addition to those from above. First, only single-item bids are allowed. That means for all i, and t

$$\sum_{k \in K} x_k = 1 \text{ for all } (p, x) \in N_t^i. \tag{2}$$

Second, SMR uses eligibility to force active and meaningful bidding. Introduced by Paul Milgrom as the truly unique part of the SMR design, eligibility rules are designed to encourage active bidding while not allowing the auction to stop too fast. A soft close makes an efficient allocation more likely.

Eligibility limits the number of items a bidder can bid on in a round as a function of the bidder's past bidding behavior. A bidder's eligibility is the number of distinct objects contained in all his bids in round $t-1$.[11] Let A_{t-1}^i be bidder i's eligibility in round $t-1$,

---

[10]This structure can easily be generalized to cases where there are multiple-copies of items available. We treat each $k = 1, \ldots, K$ as a single indivisible object.

[11]In the FCC spectrum auctions, a *weighted* measure of eligibility was used. Objects were weighted by their MHz Pops. Let w_k be the weight assigned to k. Let $\alpha_{t-1}^i = \{k | i \text{ has an active bid on } k \text{ in round} t-1\}$. A bidder's eligibility in t is then $A_{t-1}^i = \sum_{k \in \alpha_{t-1}^i} w_k$. A collection of bids B_t^i for i at t satisfies weighted eligibility if and only if $\sum_k \sum_{j \in B_t^i} x_k^{ij} w_k \leq A_{t-1}^i$ where $A_0^i = \sum_k w_k$.

which can be easily calculated as follows:

$$A_{t-1}^i = \sum_{k \in K} \prod_{b_j \in B_{t-1}^i} x_k^j. \tag{3}$$

Initially bidders are not restricted by the elgibility rule, $A_0^i = K$. A collection of bids B_t^i for i at t satisfies eligibility if and only if

$$\sum_{k \in K} \prod_{b_j \in B_t^i} x_k^j \leq A_{t-1}^i. \tag{4}$$

That is, a collection of bids is eligible if and only if the new bids plus last round's winning bids are placed on no more than A_{t-1}^i objects. Eligibility can easily be checked incrementally as each new bid is offered. Since eligibility limits the items a bidder bids on by the number of items they bid on in the previous round, eligibility encourages early bidding

The stopping rule is obvious once eligibility is imposed.

$$\text{Stop at the end of } t \text{ if } \sum_{i \in I} A_t^i \leq K. \tag{5}$$

If eligibility satisfies this constaint, bidders will be unable to bid on anything other than the items they are currently provisionally winning. Therefore, ownership will not change in any subsequent periods.

While eligibility encourages bidders to place bids, a bidder could repeatedly submit a small bid for the package of all items in order to maintain her eligibility. Then, $A_t^i = K$ and the auction never ends. To drive the auction to finish we also need to force new bids to be serious. So a *minimum increment* rule is imposed based on a vector of single-item prices Π^t, which are known at the start of round t. Let $N_t^i = B_t^i \setminus W_{t-1}^i$ be the set of new bids. Then we require that

$$p^{ij} \geq \sum_{k \in K} x_k^{ij}(\Pi_k^t + M) \text{ for all } (p, x) \in N_t^i. \tag{6}$$

where M is a minimum bid increment chosen by the auctioneer.[12]

So at the start of each round, each bidder $i \in I$ knows the objects for sale K, the prices on each object Π^t, her winning bids from the previous round W_{t-1}^i, and her eligibilty A_{t-1}^i. Each bidder then chooses new bids N_t^i satisfying Equations (4) and (6). By the resubmittal rule, $B_t^i = W_{t-1}^i \cup N_t^i$. Using the revenue maximizing allocation rule described by Equation (1) the auctioneer computes W_t. The auctioneer computes A_t^i for all i. Using (5) the auction is then stopped or continues to round $t + 1$.

The only remaining question is the computation of Π^{t+1}. In the SMR design, the price vector Π^{t+1} is simply the high price from t. That is,

$$\Pi_k^{t+1} = p_k \text{if} (p, x) \in W_t \text{ and } x_k = 1. \tag{7}$$

We let $\Pi_k^1 = 0$ for all k but one could allow Π^1 to be any reserve prices.

The rules given by Equations (1)-(7) fully describe what we have called the SMR design.

[12] The minimum increment could be altered over time, but we forego that degree of freedom in this paper.

3.2 The AUSM Design

Since the AUSM design is a continuous auction, it is somewhat more difficult to formally define the AUSM rules using the notation developed earlier. Banks, Ledyard, and Porter (1989) provide a more complete definition. However, attempting to be consistent with the previous notation, one could think of a continuous auction as an iterative auction where a single bid is placed in each round:

$$\#(\bigcup_{i \in I} N_t^i) = 1. \tag{8}$$

As in the SMR design, in the basic AUSM design, this new bid is considered along with the previous provisionally winning bids $B_t^i = W_{t-1}^i \cup N_t^i$ in solving the allocation problem given by Equation 1. This rule makes it very difficult for a new bid to win since it must independently raise the surplus of the allocation problem. The AUSM with a standby queue avoids this problem by requiring that all bids placed in the previous round are consider in each iteration $B_t^i = B_{t-1}^i \cup N_t^i$.[13]

The AUSM design does not place any restrictions on the types of bids placed. AUSM does not use an eligibility calculation, prices Π^t, or a minimum increment requirement to drive bidding. While the actual stopping rule may be at the final discretion of the auctioneer, a typical rule will take the form of a decision to stop the auction if no new bids have been submitted in a certain time period.

3.3 The RAD Design

This design represents a serious attempt to make package bidding work in a multi-object, iterative auction. It shares a number of similarities with the auction designs discussed previously. Like the SMR design, the RAD design is iterative, has an eligibility based stopping rule, forces a minimum bid increment, and computes prices for each item for sale. Like the AUSM design, the RAD design allows package bidding.

The key difference in design from the SMR approach is that package bids are allowed and a new pricing rule is introduced. Allowing package bids is accomplished by simply eliminating Equation (2) as a restriction on new bids. Pricing is a bit more subtle. Let $L_t = B_t \setminus W_t$ be the losing bids at t. We would ideally like a set of prices, Π^t, such that $p^j = \sum_{k \in K} \Pi_k^t x_k^j$ for all $j \in W_t$ and $p^j \leq \sum_{k \in K} \Pi_k^t x_k^j$ for all $j \in L^t$. If prices satisfies these equations, then the winning bidders would be paying their bid and losing bids would see that the prices were greater than their bid. While this is certainly true of the SMR prices (Equation 7), once package bidding is allowed and Equation (1) is used to decide winners, it can no longer be guaranteed that such prices exist. So we must turn to an approximation of the ideal. To compute prices Π^{t+1}, we begin by solving the following problem:

$$\min_{\Pi^t, Z, g} Z \tag{9}$$

Subject to

$$\sum_{k \in K} \Pi_k^t x_k^j = p^j \quad \text{for all} \quad b^j = (p^j, x^j) \in W_t$$

[13] In practice, the auctioneer might only consider a subset of those bids previously placed or he might require the new bidder to declare the combination of previously placed, but not provisionally winning, bids that beat the current winning bids when combined with her new bid. In either case, the computational requirements could be enormous for the auctioneer or the bidders.

$$\sum_{k\in K} \Pi^t_k x^j_k + g^j \geq p^j \quad \text{for all} \quad b^j = (p^j_i, x^j) \in L_t$$

$$0 \leq g^j \leq Z \quad \text{for all} \quad b^j \in L_t$$

$$\Pi^t \geq 0.$$

At the prices Π^t there may be some losing bids for which $\sum_{k\in K} \Pi^t_k x^j_k \leq p^j$, falsely signaling a possible winner. Such is the nature of package bidding. On the positive side, such bids can be resubmitted if $p^j - (\sum_{k\in K} \Pi^t_k x^j_t)$ is large enough. Further, Equation (9) is designed to minimize the number of such bids. In fact, if ideal prices exist, they will be the solution and $g^* = 0$ for all $b^j \in L_t$.

The prices from (9) may, however, not be unique. Also, it may be possible to further lower some of the g^j that, in the first solution, satisfy $0 \leq g^j \leq Z^*$. So, to complete the computation of Π^t, a sequence of iterations of Equation (9) is performed. We lexicographically lower as many g^j as possible. The formalities are provided in Appendix A. But, even this may not produce unique prices. So next, we try to maximize the minimum prices that satisfy the constraints of Equation (9).[14]

The purpose of prices is to convey information to bidders about opportunities in the next bidding round. Otherwise, any prices such that $\sum_{k\in K} \Pi^t_k x^j_k = p^j$ for all winners would suffice. The following examples help explain the ability of the RAD pricing rule to convey such information.[15]

3.1 Example Let there be two objects labeled A, B and three bidders labeled 1,2, and 3. Suppose that the following is true:

- Bidder 1 is the high bidder on the package A, B with a bid of 10.
- Bidder 2 bid 8 for A.
- Bidder 3 has not bid but is willing to pay 4 for B.

In this situation, bidder 1 holds the provisionally winning bid, but bidder 2 and 3 could combine to outbid the current standing bid. Any prices such that $\Pi_A + \Pi_B = 10$ and $\Pi_A \geq 8$ will satisfy Equation (9). However, if we choose $\Pi_A = 10$ and $\Pi_B = 0$, then bidder 3 may bid 1 for B in the next round only to find out that they lose. If $\Pi_A = 8$ and $\Pi_B = 2$, bidder 3 will know that they have to bid at least 2 in order to become provisionally winning. If bidder 3 bids 3 on B and bidder 2 resubmits his bid, then the new provisionally winning bids will be bidder 2 with object A and bidder 3 with object B. △

3.2 Example Let there be two objects labeled A, B and three bidders labeled 1,2, and 3. Suppose that the following is true:

- Bidder 1 is the high bidder on the package A, B with a bid of 10.
- Bidder 2 bid 4 for A.
- Bidder 3 has not bid but is willing to pay 6 for B.

[14] This is equivalent to minimizing the maximum prices.

[15] The following examples assume the minimum increment is zero. It is possible that a large minimum increment might upset some the usefulness of these prices, but having at least a small minimum increment is important to force realistic bidding.

If we select $\Pi_A = 4$, then it must be that $\Pi_B = 6$. Given this information, bidder 3 will assume that is not profitable for them to bid. In a sense, it puts all the burden of ousting the current standing bid on bidder 3. This could exacerbate the threshold problem. The more nature and fair decision is to 'split the difference' by setting $\Pi_A = 5$ and $\Pi_B = 5$. △

The appropriate prices identified in Examples 3.1 and 3.2 are obtained, when ideal prices exist, by minimizing the maximum of Π_A, Π_B subject to the prices satisfying Equation 9.

3.3 Example Let there be three objects labeled A, B, C and four bidders labeled 1, 2, 3, and 4. Suppose that the following is true:

- Bidder 1 is the high bidder on the package A, B, C with a bid of 30.
- Bidder 2 bid 25 for A, B.
- Bidder 3 bid 25 for B, C.
- Bidder 4 bid 22 fir A, C.
- Bidder 3 is willing to pay 15 for C but has not bid.

Bidder 1 is the provisional winner. The price we want, if they exist, satisfy $\Pi_A+\Pi_B+\Pi_C = 30$, $\Pi_A + \Pi_B \geq 25$, $\Pi_B + \Pi_C \geq 25$, and $\Pi_A + \Pi_C \geq 22$. No such prices exist. We try to get a close as possible. We choose Π_A, Π_B, Π_C and g^1, g^2, g^3 such that $\Pi_A + \Pi_B + g^1 \geq 25$, $\Pi_B + \Pi_C + g^2 \geq 25$, and $\Pi_A + \Pi_C + g^3 \geq 22$, and we want g^1, g^2, and g^3 to be small. We could minimize $g^1 + g^2 + g^3$, or we could minimize the maximum of g^1, g^2, g^3. Suppose we pick $\Pi_A = \Pi_B = \Pi_C = 10$ yielding $g^1 = g^2 = 5$ and $g^3 = 3$. We could also pick $\Pi_A = \Pi_B = 11$ and $\Pi_C = 8$ yielding $g^1 = g^2 = 6$ and $g^3 = 0$. In the second case, relative to the first, $g^1 + g^2 + g^3$ is less but the maximum is more. In the second case, a bidder for B knows exactly how much they must bid to become provisionally winning. However, the prices overvalue what someone must bid on either A or C to become provisionally winning. The prices in the first case overstate the value of the bid required for all single-items to become winning, but the difference is less for A or C as compared to the second case. The RAD pricing rule picks the first case. In either case, bidder 3 can bid for C and, assuming bidder 2 resubmits her bid, become a provisionally winning bid.

If, on the other hand, bidder 3's value for C was only 10.5, then bidder 3 would only find it profitable to bid when the price is $\Pi_C = 10$. If bidder 3's value was 7, then he would not be willing to bid in either case despite the fact that a bid of 7 could unseat the bidder 1's current high bid. △

As the reader will see in the data below, this combination of pricing and stopping rule work very well together to eliminate strategic problems caused by the threshold problem. Changing the SMR design to allow package bidding with the pricing rule generates a significant increase in performance in environments with multiple objects with complementarities. There is also no degradation of performance in environments with no complementarities.

4 The Experimental Design

The environment used as a test bed for all auctions in this paper was created by combining features of the *spatial fitting* environment originally utilized by Ledyard, Porter, and

Table 1: 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
	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	bef	abcdef
	Values:	29	25	117	125	142

Rangel (1997) and an additive environment. The five participants were allowed to bid on ten heterogeneous items labeled $a, b, c, d, e, f, g, h, i,$ and j. Bidder values for the first six items were highly super-additive. Five separate draws of valuations were determined in the following manner:

- The single-item packages, (a, b, c, d, e, f), had integer values drawn independently form a uniform distribution with support $[0, 10]$.
- The two-item packages,$(\{a, b\}, \{a, c\}, \ldots \{e, f\})$, took integer values drawn independently from a uniform distribution with support $[20, 40]$.
- The three-item packages, $(\{a, b, c\}, \ldots \{d, e, f\})$, had integer values determined independently by draws from a uniform distribution with support $[140, 180]$.
- The value for the six-item package, $\{a, b, c, d, e, f\}$, was drawn uniformly from $[140, 180]$.

For each period, a total of 25 unique packages and valuations were generated by the previous steps. Each bidder was randomly given five packages. In general, a combination of two three-item packages formed the largest total value. However, the optimal package configuration is typically overlapped by other competing packages. Therefore, these valuations were meant to be a difficult test of any allocation mechanism. An indicator of that difficulty is that, in period 3 and 5 *competitive* equilibrium prices did not exist. Table 1 provides a sample set of spatial fitting valuations (Period 2). In this example, the efficient package combination is $\{a, b, d\}$ for bidder 2 and $\{c, e, f\}$ for bidder 5.

The valuations for the remaining four objects (g, h, i, j) were determined in an additive manner. Each bidder had a valuation for each individual object between 40 and 180. If a bidder obtained more than one of these items, they received the sum of their valuations. Therefore, competitive equilibrium prices lie between the highest and the second highest valuation for each of the objects. These items were added to the spatial fitting environment for two reasons. First, since we suspected that under some institutions bidders would be making net losses on the first six objects, these objects would serve as a convenient tool to ensure that bidders' overall payoffs for the auction was not negative.[16] Second, performance in these markets could provide a quick check of any auction's proficiency in the easiest of environments.

All sessions were conducted using members of the Caltech community, primarily undergraduates. Five subjects participated in each experimental session. In each session, the

[16] In reality, in many of the SMR auction sessions even these four additive objects were not enough.

Table 2: Experimental Sessions Completed

Period	Institution	
	SMR	RAD
1	3	5
2	3	5
3	5	5
4	3	5
5	3	5

number of auctions (or periods) actually completed varied. No session lasted longer than three hours. Subjects received new redemption value sheets at the beginning of each new auction. Bidder values were kept private. At the end of each auction, subjects calculated their profits and converted the token values into dollars. Subjects were paid privately at the end of the experimental session. In addition to participating in a practice auction, all subjects had prior experience with the general auction format; they had all participated in training sessions that utilized simplified auction rules and environments. Instructions for the RAD treatment can be found in Appendix **??**.

A total of 42 auctions were completed in 15 experimental sessions. Table 2 reports the distribution of experiments across the two mechanisms and the five parameter sets. The AUSM data comes from previous experiments reported by Ledyard, Porter, and Rangel (1997).

5 Performance Measures

When choosing an auction design, a variety of criteria and measures may be used. In general there will be trade-offs between these measures and different auctions will perform better depending on which measure one focuses on. For example, high efficiency can sometimes come at the cost of seller revenue and the time to complete the auction.

5.1 Efficiency

Efficiency is the most obvious choice of a performance measure. It was, in fact, to be the original policy goal of the FCC PCS auction design. In any environment, each bidder has a set of valuations that can be indicated as a (payoff) function $V^i : \{0,1\}^K \to \mathbb{R}$ where $V^i(y)$ is bidder $i's$ redemption value, the amount the experimenter will pay that bidder, if they hold the combination of objects indicated by y at the end of the auction. The maximal possible total valuation is:

$$V^* = \max \sum_{i=1}^{I} V^i(y^i)$$

subject to

$$\sum_{i=1}^{N} y^i_k \leq 1 \;\; k = 1,.,K y^i \in \{0,1\}^K.$$

If $\{\hat{y}^i\}_{i=1}^I$ is the final allocation chosen by an iterative auction, the *efficiency* of that auction is

$$\frac{\sum_i V^i(\hat{y}^i)}{V^*}$$

It is true (see Ledyard, Porter, and Rangel (1997)) that the absolute level of efficiency can be deceptive since one can increase the percentage by simply doubling the agents' values and leaving the allocation unchanged. However, we will only use efficiency to compare performance across institutions in the same environment. So this is not a problem for us.

5.2 Seller's Revenue

If the auction designer happens to also be the seller, he may be interested in maximizing revenue: the sum of the final bids. Since that maximum can vary significantly across environments, we used the percentage of the maximum possible revenue actually captured by the seller as our measure of seller revenue. This statistic also measures the amount of the maximum surplus captured by the seller.

(add notation)

5.3 Bidder Profit

Bidder profit is another possible measure. With the presence of significant complementarities, some auction mechanisms can cause some bidders to lose money (Bykowsky, Cull, and Ledyard 1995). A high probability of losses can lead to a variety of performance failures. Bidders may be unwilling to participate in auctions they know they are likely to lose money in. They may not bid aggressively and thereby cause efficiency losses. Losses may also lead a bidder to default on payment contracts, which in turn undermines the credibility of the auction. Increasing the surplus to the bidders can, however, conflict with a goal of high revenue for the seller. All other things being equal (including efficiency of the auction), any increase in bidder profits must come at the expense of seller revenue. Therefore, while it may not be clear why a designer would want to maximize bidder profitability, there does seem to be a compelling reason to avoid bidder losses. In all of the experimental sessions we report on in this paper, a bidder's profit will be

$$V^i(\hat{y}^i_t) - \Pi^* \cdot \hat{y}^i$$

where Π^* is the vector of final prices.

5.4 Length

When analyzing iterative auctions the length of the auction becomes a relevant concern.[17] In this paper we measure auction length by the number of iterations (rounds) before the auction is completed. Increased iterations can reduce seller profitability because each iteration typically has some fixed administrative cost as well as the possible opportunity costs of foregone rental revenue on the objects. Obviously, one could hold an auction in one iteration as a sealed bid auction. But there is a possible trade-off between auction length and efficiency. Increased iterations may allow high value bidders to *find* the right package thus increasing efficiency. An auction that ends very quickly may lead to inefficient allocations.

[17] The FCC PCS auctions have been criticized for the time it took to complete one auction. The D,E, and F Block PCS Auction lasted 276 rounds or 141 days.

Since the spatial fitting and additive environments were run simultaneously, the number of iterations until the entire auction closed is not necessarily an accurate performance measure of auction length for either environment. In order to determine the auction length for the additive markets, we identified the round that these four markets would have closed in if there was no spatial fitting markets. For example, an auction may have lasted 12 iterations. However, the last new bid on any of the additive valued items occurred in the sixth iteration. Then, the auction for the additive environment would be said to have ended in iteration seven since, assuming bidding would have been identical, the auction for just the additively valued objects would have ended after no new bids were placed in the seventh round. While it is possible that the addition of the spatial fitting environment may have altered bidding behavior on the additive items, and vice versa, this measure seems to be a reasonable proxy for the speed of the auction in the additive environment. The same procedure was used for the spatial fitting environment. By necessity, the auction length for one environment is identical to the length of the actual auction.

6 Results

In this section, we compare the performance of the RAD mechanism with two other auction designs: the SMR auction and AUSM.

6.1 Results From the Spatial Fitting Test Bed

1 Conclusion *RAD yields efficiency at least as high as other designs tested.*

The average auction efficiency across periods under the SMR design was 66.95%, significantly lower than the average efficiencies of 90.42% for the RAD design. The continuous AUSM obtained an average efficiency of 94%, which is not significantly different from the results for RAD. Table 3 gives the results of Wilcoxon-Mann-Whitney Rank-Sum pair-wise comparisons of these three institutions. The fact both auctions that allow package bidding dramatically out perform the SMR design suggests an obvious conclusion.

2 Conclusion *Package bidding significantly increases efficiency.*

These results appear to provide compelling evidence that package bidding is an essential part of an auction if complementarities exist and one desires allocative efficiency. As further evidence of this, in only 4 out of 17 (24%) auctions does the Milgrom design lead to full efficiency as opposed to 20 out of 25 (80%) for RAD.

	AUSM	RAD
SMR	$z = 3.29$	$z = 3.55$
	$\alpha = .000$	$\alpha = .000$
AUSM		$z = .332$
		$\alpha = .371$

Table 3: Spatial Fitting – Auction Efficiency Rank-Sum Test

3 Conclusion *Package bidding significantly increases average bidder profits and reduces individual losses.*

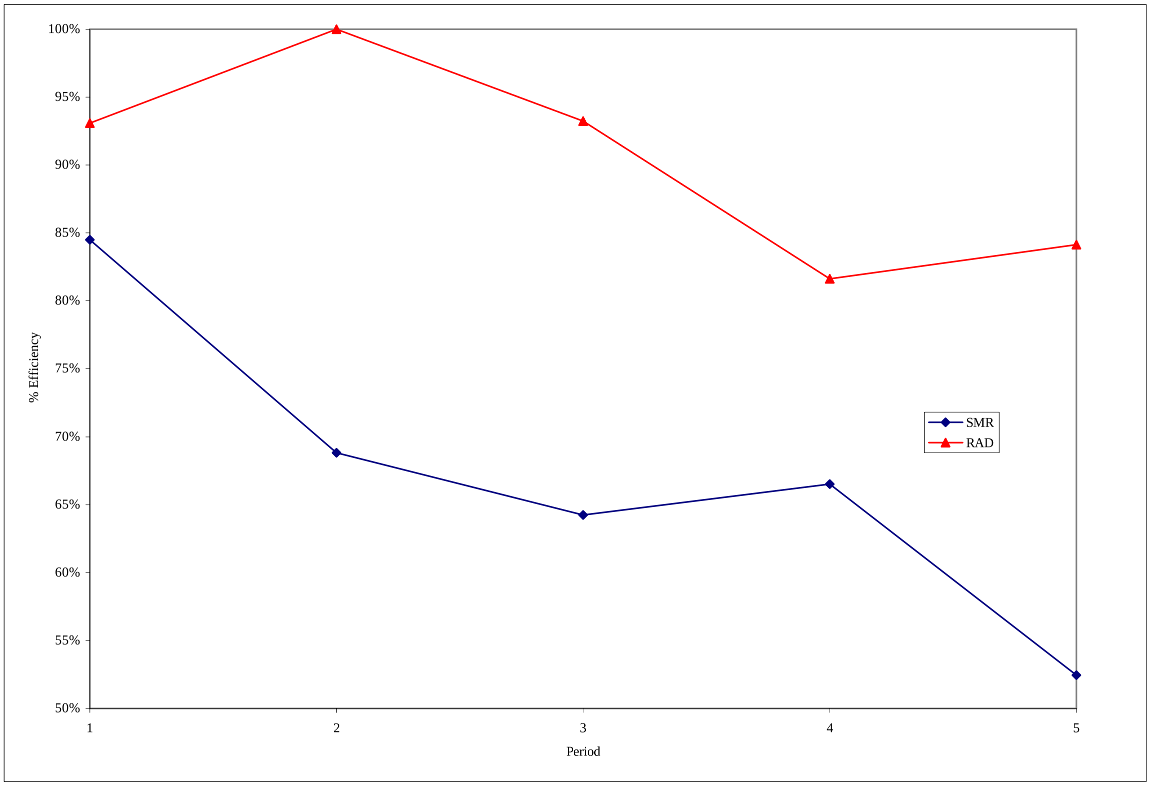

Figure 1: Spatial Fitting – Average Efficiency Per Period

When package bidding was not allowed (SMR) bidders, as a whole, averaged losses of $7.73 in each period for the markets with complementarities. However, in both institutions where package bidding was permitted, bidders earned positive profits on average. A Wilcoxon-Mann-Whitney Rank-Sum Test indicates that the level of profits is significantly higher in the RAD design than in the no packaging auction ($z = 2.83$, $\alpha = .006$). Total bidder profit averaged $4.23 in RAD. On an individual level, 30 out of 85 (35%) bidders lost money under the SMR auction. Under RAD, only 4 out of 125 (3.2%) bidders ended an auction with losses.[18]

While the number of bidders with losses decreased when package bidding was allowed, it is surprising that any bidders made losses. Without package bidding losses are to be expected. In order to win a package, bidders must put themselves at risk of only obtaining part of the package. However, when package bidding is allowed, bidders have no incentive to bid for packages above their values. After closely examining the data from experiments where bidders were allowed to bid on packages, we have some conjectures as to why losses occurred. First, eligibility encouraged bidders to bid on as many items as possible. It is possible that bidders thought that an easy and relatively risk free method to keep their eligibility high was to place small bids on single-item packages. They may have thought that it was very likely that someone would value the object above their small bid and therefore they would not lose money from this bid. However, at times, these small bids were sufficiently large to be winners. In a few experiments, we observed behavior consistent with this rationale. The strategic implications of "eligibility" remains to be seriously studied. However, it is clear that it leads to bids that are inconsistent with

[18]Similar data is not currently available for AUSM.

short-run value maximization. Second, if a bidder makes a mistake in bidding in early iterations, it may be difficult to escape from it. For whatever reasons, bidders occasionally placed bids that are inconsistent with their valuations. A simple example of this occurs if a bidder had a value of 100 for the package A,B,C, and a value of no more than 25 for any 2 item subset. If that bidder intended to place a bid of 50 on A,B,C but, through negligence, missed indicating C, they would have a bid of 50 on A,B yielding a loss of 25 if no one ever bids higher.[19] If those bids were sufficiently high, no other bidder would be able to *rescue* them by out bidding them.[20] One interesting but little studied aspect of practical auction design is the prevention of "typos": unintentional errors in data entry. The hard part is separating "typos" from strategic moves later claimed to be mistakes. We do not pursue this here.

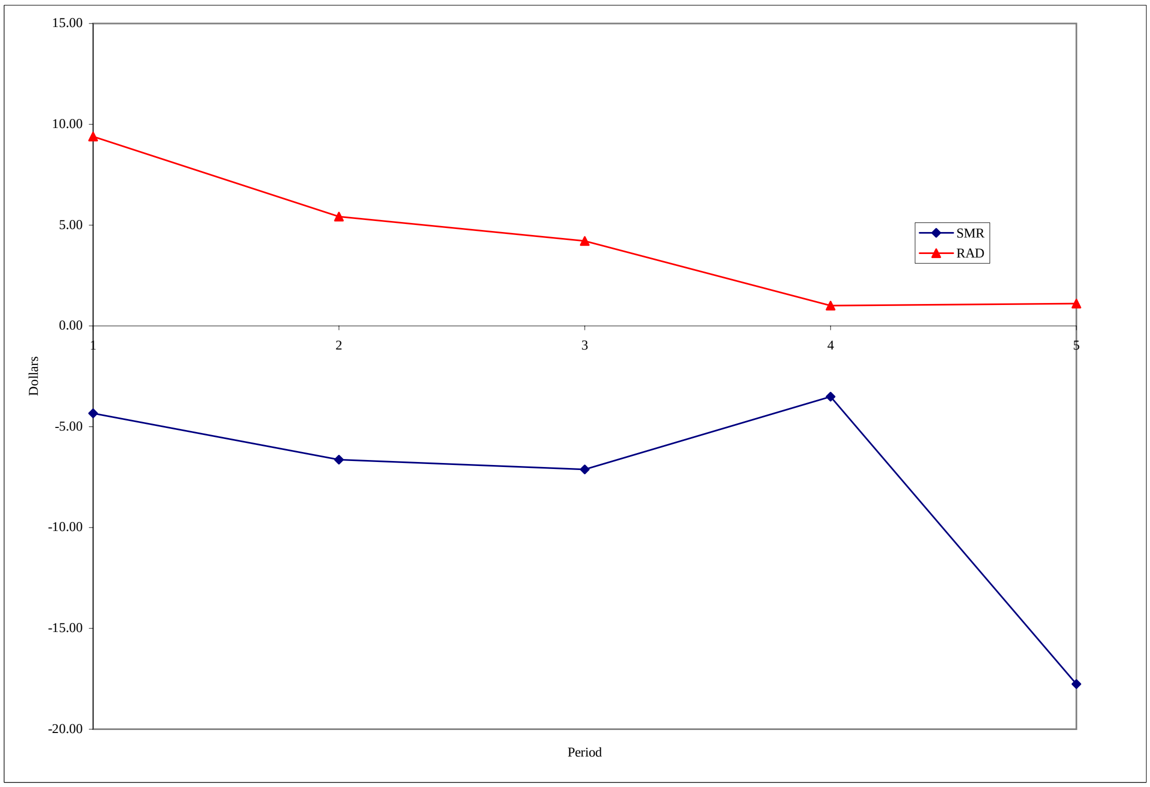

Figure 2: Spatial Fitting – Average Total Bidder Profits Per Period

4 Conclusion *Auction length is shortest under RAD.*

Auction length, measured as the number of iterations before completion of the auction, was significantly shorter in the institution where package bidding was allowed. The SMR auctions averaged 16.2 iterations in length as compared to 3.32 for RAD. A Wilcoxon-Mann-Whitney Rank-Sum Test indicates that the auction length is significantly shorter under the RAD design than in the SMR auction ($z = 4.98$, $\alpha = .000$). In fact, it often took longer to complete the additive markets than the spatial fitting items(see Conclusion 9).

[19] This actually happened to one of the authors during early software tests.

[20] However, under the iterative design, if a bidder realized his mistake before the completion of that round, he could delete the bid. It is easy to imagine that a similar errors could be made in a continuous auction without any hope for correction.

Since the AUSM mechanism was a continuous auction, it is obviously not possible to directly compare the speed of these tow formats.

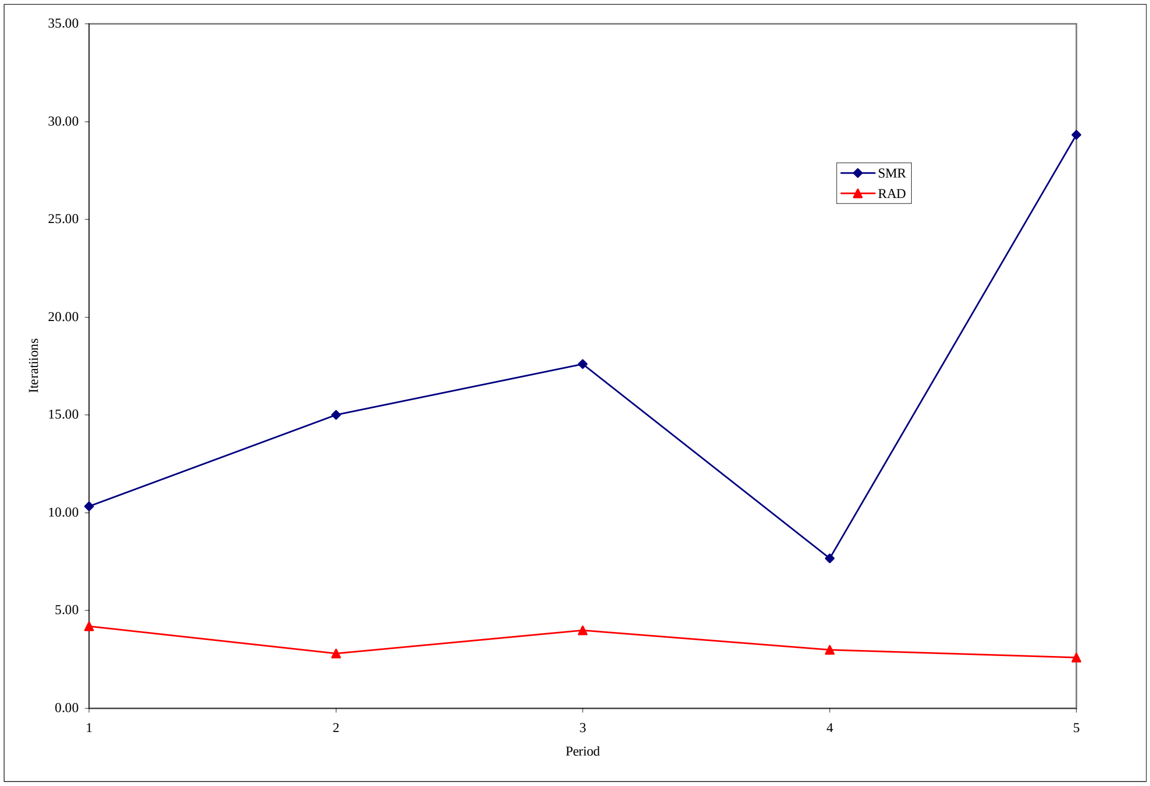

Figure 3: Spatial Fitting – Average Auction Length Per Period

5 Conclusion *Package bidding significantly reduces seller revenues. However, among package bidding designs, RAD does best.*

A simple comparison of average seller revenue does not bear out significant differences between auction institutions. But a comparison of seller's revenue as a percentage of maximum possible revenue demonstrates single-item only bidding institutions yield a significantly higher revenue percentage for the seller than any of the package bidding mechanisms with an average across periods of 96.25% as compared to 79.4%, and 71% for RAD and AUSM respectively. RAD yields significantly higher revenue than AUSM, which implies that the RAD design minimizes the potential trade-off between bidder profits and seller revenue. Table 4 gives the results of Wilcoxon-Mann-Whitney Rank-Sum pair-wise comparisons of these three institutions.

	AUSM	RAD
SMR	$z = 3.06$	$z = 2.88$
	$\alpha = .001$	$\alpha = .004$
AUSM		$z = 1.62$
		$\alpha = .100$

Table 4: Spatial Fitting – Seller Revenue Rank-Sum Test

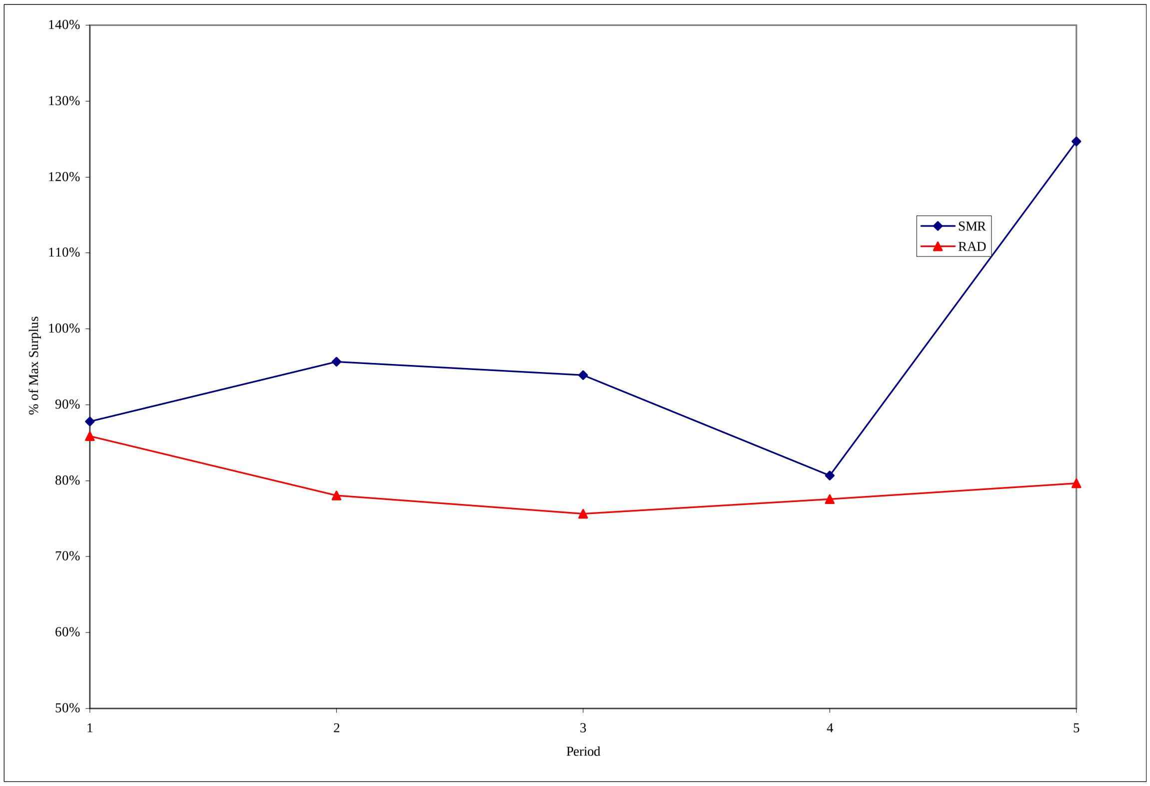

Figure 4: Spatial Fitting – Average Seller Revenue Per Period

A strong caveat is necessary before one leaps to the conclusion that the SMR auction is optimal for a seller. The result that the revenue percentage is high in SMR auctions is driven by the high level of bidder losses (see Conclusion 3). A measure that identifies this is the difference between realized revenue as a percentage of maximum possible revenue and the amount of losses as a percentage of the maximum possible surplus. By subtracting the former by the later, we have a measure of surplus captured by the auctioneer if a "default" option for the bidders was allowed. Table 5 reports these calculations. While the package bid auctions are unaffected, the Milgrom auction is reduced to 67%, which is now below that captured by the three other designs. Any revenue above 66.95% (the average efficiency), must be coming from bidder losses. In any practical situation, if a bidder faces losses by accepting the items won, that bidder will most likely default and not pay.[21] This was not possible in our experiments, although future tests could certainly include a default option. It is reasonable to conjecture that if we were to test the auctions with a zero losses option, the revenue percentage would be closer to the level of efficiency. The SMR auction revenue percentage would then be lower than the observed revenue percentages for auctions that allow package bidding.[22]

6.2 Results from the Additive Test Bed

In this section, we report on the results for the four objects that had additive valuations for all bidders. In general, the efficient allocation would require only single-item bids

[21] An extreme example can be found in the FCC C Block auctions. See "Airwave Auctions Falter as a Source of Funds for U.S." in *The New York Times* April 3, 1997.

[22] Imposing a zero loss rule could of course also significantly affect bidder strategies.

	Revenue %	Profit %	Rev.-Losses
SMR	96%	-29%	67%
AUSM	71%	23%	71%
RAD	79%	11%	79%

Table 5: Surplus Captured by the Seller as a result of profitable purchases

among the additive objects. So package bids would occur only if bidders were attempting a sophisticated strategy[23] to capture a larger share of the objects. But this rarely happened.

6 Conclusion *Package bids only rarely occur among the winning bids in the additive environment.*

In only 3 of 25 RAD auctions do the final winning bids contain packages of additive objects. Further, in these 3 auctions the final allocations involve package bids across additive and other objects.[24] Although bidders are clearly willing to bid on packages in the additive environment, they are only rarely able to use this ability to their advantage as is evidenced by the extremely high levels of efficiency achieved in the additive environment.

7 Conclusion *In the additive environment, under both designs efficiency is very near 100%. There are no discernible differences between the two auctions.*

A 100% efficient auction indicates that all the possible gains from trade (surplus) have been captured by either the bidders or the seller. As expected, all auctions did quite well in terms of efficiency in this environment. In most of the auctions, the four objects were allocated to the highest valuing bidders: 20 of 25 (80%) for RAD and 15 of 17 (88%) auctions for the SMR auction. The AUSM design also leads to high efficiencies in the additive environment. There were no significant differences in the level of efficiency achieved by any of the mechanisms. Therefore, package bidding auctions, specifically RAD, do not seem to degrade auction performance in simple settings.

8 Conclusion *In the additive environment, the auction institutions yield similar seller revenue.*

In the additive environment, there was very little difference in the revenue collected by the seller under the two auction institutions. The SMR and RAD mechanisms averaged 69.96% and 71.96% of the maximum possible revenue respectively. A rank-sum test also showed no significant difference between the observed revenues.

9 Conclusion *In the additive environment, auction length is shortest under the RAD design.*

RAD yielded lengths that were significantly shorter than the SMR. The average auction length under the SMR auction was 11.7 iterations but only 6.1 for RAD.[25] A Wilcoxon-Mann-Whitney Rank-Sum Test indicates that the auction length is significantly shorter

[23] Such a strategy might be to create an artificial *threshold* that would yield a possible *free rider* problem for others allowing the bidder to get the items even if it were not an efficient allocation.

[24] We used the final prices to estimate the portion of the bid occurring in the additive environment.

[25] These results are, of course, confounded by the fact that the length the additive part of the auction is the round after which no new bid is made on an additive object. This is not necessarily independent of the existence of the spatial part of the auction.

under the RAD design than in the SMR auction ($z = 4.67$, $\alpha = .000$). As before, there is no direct comparison with the speed of AUSM.

Since package bidding has no advantage in the additive environment, and assuming all bids in the additive markets are on the individual items the prices for the objects should be the same, we were surprised by this result. A potential explanation for this difference is that RAD allows bidders to quickly learn about the outcomes in the spatial fitting portion of the market. They can then turn their attention to the additive markets.

7 Conclusions and Open Issues

The experimental test results point to two clear conclusions:

i. The option to bid for packages clearly improves performance in difficult environments, and does not degrade performance in simple environments.

ii. The RAD redesign of the SMR rules outperforms both the iterative SMR as well as the continuous AUSM with queue.

The general principle that package bidding is an important option for multi-object auctions in environments with significant complementarities is reaffirmed by the evidence. Auctions that only allow bidding on single-items almost always exhibit lower levels of allocative efficiency and higher bidder losses. When auctions are run in an iterative mode, single-item only bidding can also lead to much longer auctions. The only redeeming feature of these auctions seems to be their revenue generating capabilities. Unfortunately, much of that revenue comes from losses to bidders as opposed to increased surplus extraction. This may be acceptable in the short run. However, if the design is used repeatedly, bidders will learn to avoid these losses, and efficiency and revenue will ultimately suffer.

But we have gone further here than simply asserting that package bidding is sensible. We have provided a new auction design, RAD, which clearly outperforms others. Relative to the SMR design, RAD produces higher efficiency, higher revenue, higher bidder profits and a much quicker time to completion. It even produces similar efficiencies to and higher revenues than the continuous AUSM with a standby queue. Since RAD uses a pricing rule instead of a queue to mitigate the threshold problem, it appears to be no more complex from a bidder's point of view than the SMR auction and significantly simpler than the continuous AUSM. Finally, there is no evidence of degradation in performance when RAD is used in simple, additive environments.

Why do we think RAD worked so well? We believe it is the decentralizing influence of the prices. Under the SMR mechanism, prices were only calculated from single-item bids. Therefore, if bidders were not bidding above their valuations, in this environment, it is guaranteed that the single-item prices would be much lower than the actual winning bids for the packages. If we consider the sample parameters given in Table 1, the maximum prices for bidders unwilling to expose themselves to potential losses are: 3, 8, 2, 8, 16, and 9 for the first six items.[26] The 'competitive' prices, however, are 38, 49, 30, 43, 38, and 49. If we examine the data for this parameter set (period 2), we find that this difference in prices is prevalent experimentally. Figure 5 lists the average prices under RAD. The RAD prices are closer to the competitive prices. In general we would expect final prices to be somewhat

[26]These are simply the maximal single-item values. The only way prices under the SMR design could be higher is if a someone bid on a single-item above their value.

lower than the competitive prices due to the bid increment requirement, which made the true price higher than that reported here. Once the mandatory bid increment of three francs is considered, the RAD prices are not significantly different than the competitive prices for five of the six objects. In the RAD mechanism, prices are calculated using all bids. Therefore, in general, they will more closely represent the level of competition for an item. Since the prices are typically calculated in order to indicate the level of competition below the winning packages, they can indicate to bidders markets where bidding is thin. Thus, prices should aid in finding an appropriate fit.

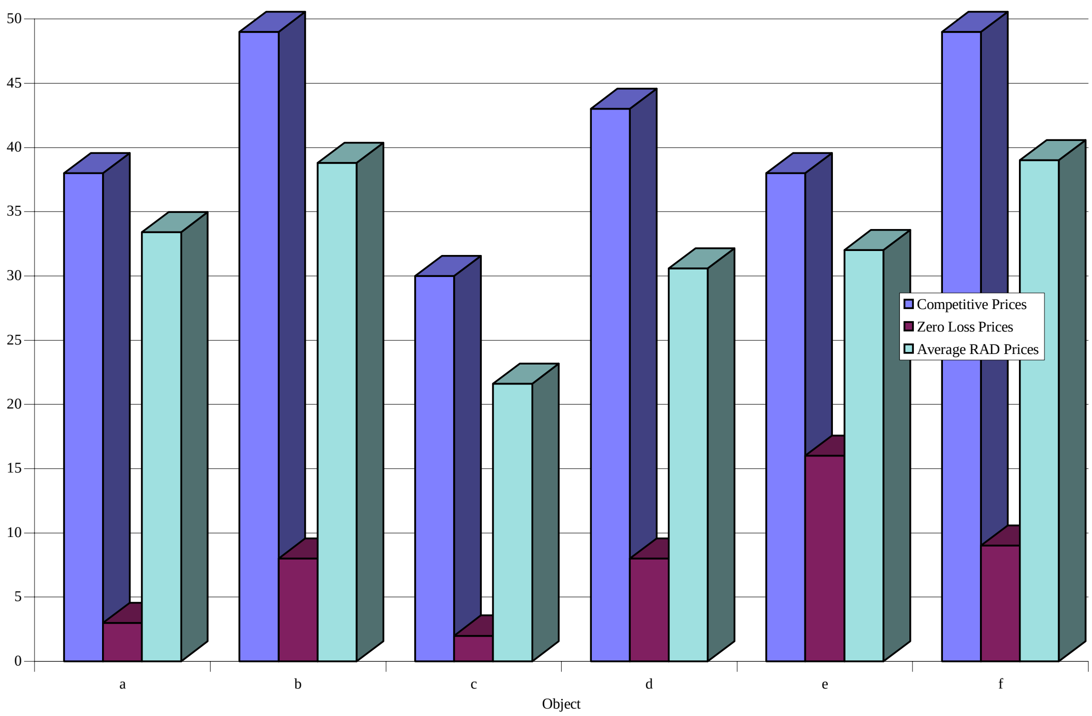

Figure 5: Average Final Prices – Period 2

It would seem that the RAD design would be a natural candidate for use as a multi-object iterative auction in its current form. However, in spite of the excellent performance in our tests, there are at least three problem areas that might be considered for redesign. The first, and simplest to fix, is a result of the eligibility rule. If bidders have budgets for items, they may find themselves bidding for and even winning items which have little value to them simply to preserve eligibility. While this problem has generally been recognized even when there is no package bidding, we know of no papers that purport to provide a solution. There is, nevertheless, a straight-forward solution: the use of "OR" bids in an iterative auction. Bidders would be allowed to place a bid saying "I bid $1000 for A,B and C OR I bid $800 for D and E". The appropriate constraints would be added to the allocation problem (1) and the rest of the mechanism would be left as is. This could be done to the SMR rules as well as to RAD and others. OR bids may appear to increase a bidder's problem complexity a bit, but such bids do eliminate the anxiety and confusion raised by the need to find "safe places" to preserve eligibility. We have not tested the effect of adding OR bids to the RAD design, something that should clearly be done before our recommendation is adopted in practice.

A second problem with the RAD design is that, although the pricing rule seems to guide and coordinate small bidders to solve the threshold problems, it can also orphan some bidders at early stages even through they belong in the efficient allocation. An example will illustrate. Suppose in round 5 there are 4 bids submitted

Table 6:

bid #	$	Items
1	99	ABC
2	75	AB
3	75	AC
4	75	BC

Under the RAD design bid #1 wins and the prices[27] are $(\Pi_c \Pi_b Pi_c) = (33, 33, 33)$. Now suppose there is a bidder who is willing to pay 30 for A. Had they bid 28 for A in round 5 they would have been a winning bid along with bid #4. But now they can't bid since $\Pi_a > 30$. This may lower efficiency. There are several features of RAD that work against such *orphaning*. First, if this bidder had bid in round 5, they would not have been orphaned. So aggressive participation helps. Second, suppose in round 6, bid #3 is not resubmitted, but 1,2, and 4 are. Then 1 still wins and $\Pi = (24, 51, 24)$.[28] If it gets to this stage our bidder for A can reenter the fray if they still have the eligibility to do so. Of course, if the auction stops in round 5, which it will if there are no additional new bids, it will end at an inefficient allocation.

A Resolving Price Ambiguity

In Section 3.3, we indicated that the RAD auction pricing algorithm (9) might not yield a unique price vector.[29] We use the following routines to eliminate that ambiguity. Let g^*, Π^*, and Z^* solve (9). If $Z^* = 0$ then go to problem (11) below. If $Z^* > 0$, let $J^* = \{j \in L_t \mid Z^* = g{*}^j\}$. If $J^* = L_t$ then go to (11) below. Otherwise,

$$\min_{\Pi^t, Z, g} Z \tag{10}$$

Subject to

$$\begin{aligned}
\sum_{k \in K} \Pi_k^t X_k^j &= p^j \text{ for all } b^j = (p^j, x^j) \in W_t \\
\sum_{k \in K} \Pi_k^t X_k^j + g^{*j} &= p^j \text{ for all } b^j = (p_i^j, x^j) \in J^* \\
\sum_{k \in K} \Pi_k^t X_k^j + g^j &= p^j \text{ for all } b^j = (p_i^j, x^j) \in L_t \setminus J^* \\
0 \le g^j \le Z & \text{ for all } b^j \in L_t \setminus J^* \\
\Pi^t &\ge 0.
\end{aligned}$$

[27]Notice that these are not separating prices, which is what causes a problem.

[28]These are separating prices. This also shows that prices are not necessarily monotonically increasing (since $24 < 33$). The sum of prices is increasing.

[29]An alternative approximation would minimize $\sum_j (g^j)^2$ which would avoid iteration. We chose to stick with linear programs for computational simplicity and a desire to minimize the number of bids missed rather than the total size of the miss.

Let $\hat{Z}, \hat{g}, \hat{\Pi}$ be the solution to (10). If $\hat{Z} = 0$ go to 11 below. Otherwise, let $\hat{J} = \{j \mid \hat{Z} = \hat{g}^j\}$. If $J^* \cup \hat{J} = L_t$ then go to problem (11) below. Otherwise, let $J^* = J^* \cup \hat{J}$ and go to (10) again.

When the iteration on (10) is complete we will have prices which approximate our "ideal" but not always obtainable prices. They may still not be unique. So, we go through a sequence of iterations which eliminate non-uniqueness and which create prices to guide bidders to solve the threshold problem. Let $\hat{Z}, \hat{g}, \hat{\Pi}$ be the solution from the last iteration of (10). Let $\hat{K} = K$.

$$\max Y \tag{11}$$

subject to

$$\sum_{k \in K} \Pi_k^t x_k^j = p^j \text{ for all } (p^j, x^j) \in W_t$$
$$\sum_{k \in K} \Pi_k^t x_k^j + g^j = p^j \text{ for all } (p^j, x^j) \in L_t$$
$$\Pi_k^t \geq Y \text{ for all } k \in \hat{K}. \tag{12}$$

Let Y^*, Π^* solve (11). Let $K^* = \{k \in K \mid \Pi_k^t = Y^*\}$. Let $\tilde{K} = \hat{K} \setminus K^*$. If $\tilde{K} \neq \emptyset$, return to (11) and solve it replacing (12) with

$$\begin{aligned} \Pi_k^t &\geq Y \text{ for all } k \in \tilde{K} \\ \Pi_k^t &= \Pi_k^{*t} \text{ for all } k \in K \setminus \tilde{K}. \end{aligned}$$

When $\tilde{K} = \emptyset$, we are done and the prices $\Pi^* = \Pi^{t+1}$. These are unique, approximate the ideal prices, and provide signals about thresholds.

References

Armstrong, M. (1998): "Optimal Multi-Object Auctions," Unpublished.

Banks, J. S., J. O. Ledyard, and D. P. Porter (1989): "Allocating Uncertain and Unresponsive Resources: an Experimental Approach," *Rand Journal of Economics*, 20(1), 1–25.

Bykowsky, M. M., R. J. Cull, and J. O. Ledyard (1995): "Mutually Destructive Bidding: The FCC Auction Problem," Social Science Working Paper 916, California Institute of Technology, Pasadena, CA.

Coppinger, V., V. L. Smith, and J. Titus (1980): "Incentives and Behavior in English, Dutch and sealed-bid Auctions," *Economic Inquiry*, 43, 1–22.

Dasgupta, P., and E. Maskin (2000): "Efficient Auctions," *Quarterly Journal of Economics*, 115(2), 341–388.

Ferejohn, J., R. Forsythe, and R. Noll (1979): "An Experimental Analysis of Decision Making Procedures for Discrete Public Goods: A Case Study of a Problem in Institutional Design," in *Research in Experimental Economics*, ed. by V. L. Smith. JAI Press, vol. 1, pp. 1–58.

Grether, D., M. Isaac, and C. Plott (1981): "The allocation of Landing Rights by Unanimity among Competitors," *American Economic Review*, 71, 166–171.

Hurwicz, L. (1960): "Optimality and Informational Efficiency in Resource Allocation Processes," in *Proceedings of a Symposium on Mathematical Methods in the Social Sciences, 1959*, pp. 27–46, Stanford. Stanford University Press.

Ishikida, T., J. Ledyard, M. Olson, and D. Porter (1998): "Implementation of a Combinatorial Market: The Experiments behind the Automated (environmental) Credit Exchange (ACE)," Unpublished.

Kagel, J. H., R. M. Harstad, and D. Levin (1987): "Information Impact and allocation rules in auctions with affiliated private values: A laboratory study," *Econometrica*, 55, 1275–1304.

Kwasnica, A. M., J. O. Ledyard, D. Porter, and J. Scott (1998): "The Design of Multi-Object Multi-Round Auctions," Social Science Working Paper 1045, California Institute of Technology, Pasadena, CA, Forthcoming.

Ledyard, J. O., M. Olson, D. Porter, J. A. Swanson, and D. P. Torma (forthcoming): "The First Use of a Combined Value Auction for Transportation Services," *Interfaces*.

Ledyard, J. O., D. Porter, and A. Rangel (1997): "Experiments Testing Multi-object Allocation Mechanisms," *Journal of Economics and Management Strategy*, 6, 639–675.

Milgrom, P. R., and R. J. Weber (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50(5), 1089–1122.

Milgrom, P. R. (1995): "Auctioning the Radio Spectrum," in *Auction Theory for Privitization*. Cambridge University Press, chap. 1.

Myerson, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

Plott, C. R. (1997): "Laboratory Experimental Testbeds: Applications to the PCS Auction," *Journal of Economics & Management Strategy*, 6, 605 – 638.

Porter, D. P. (1997): "The Effect of Bid Withdrawl in a Multi-Object Auction," Social Science Working Paper 982, California Institute of Technology, Pasadena, CA.

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 (1998): "Computationally Manageable Combinatorial Auctions," *Management Science*, 44(8), 1131–1147.

Sandholm, T., S. Suri, A. Gilpin, and D. Levine (2001): "CABOB: A Fast Optimal Algorithm for Combinatorial Auctions," *International Joint Conference on Artificial Intelligence*, pp. XXXX–XXXX.

SMITH, V. L. (1979): "Incentive Compatible Experimental Processes for the Provision of Public Goods," in *Research in Experimental Economics*, ed. by V. L. Smith. JAI Press, vol. 1, pp. 59–168.

——— (1982): "Microeconomic Systems as an Experimental Science," *American Economic Review*, 72, 923–955.

WILLIAMS, S. (1994): "A Characterization of Efficient Bayesian Incentive Compatible Mechanisms," Math Center Working Paper 1100, Northwestern University, Evanston, IL, Forthcoming in Economic Theory.