

A New and Improved Design for Multiobject Iterative Auctions

Anthony M. Kwasnica

Smeal College of Business, Pennsylvania State University,
University Park, Pennsylvania 16802, kwasnica@psu.edu

John O. Ledyard

Division of the Humanities and Social Sciences, California Institute of Technology,
Pasadena, California 91125, jledyard@hss.caltech.edu

Dave Porter

Interdisciplinary Center for Economic Science, George Mason University, Fairfax, Virginia 22030,
dporter4@gmu.edu

Christine DeMartini

RAND, Santa Monica, California 90407, martini@rand.org

In this paper we present a new improved design for multiobject 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 used by the FCC to auction the electromagnetic spectrum and the Adaptive User Selection Mechanism (AUSM) of Banks et al. (1989, "Allocating uncertain and unresponsive resources: An experimental approach," *RAND Journal of Economics*, Vol. 20, No. 1, pp. 1–25). 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. Our experiments demonstrate that the RAD auction achieves higher efficiencies, lower bidder losses, higher net revenues, and faster times to completion without increasing the complexity of a bidder's problem.

Key words: auctions; experimental economics; combinatorial auctions

History: Accepted by G. Anandalingam and S. Raghavan, special issue editors; received May 22, 2002. This paper was with the authors 4 months for 2 revisions.

1. Introduction

Theory, experiment, and practice suggest that, when bidder valuations for multiple objects are superadditive, combinatorial auctions are needed to increase efficiency, seller revenue, and bidder willingness to participate (Bykowsky et al. 2000, Rassenti et al. 1982, Ledyard et al. 2002). A combinatorial auction is an auction in which bidders are allowed to express bids in terms of packages of objects. The now-famous FCC spectrum auctions are a good example of the relevance of these issues. In 41 auction events from 1994 to 2003, the FCC used what is known as a Simultaneous Multiple Round (SMR) auction to allocate spectrum and raise over \$40 billion in revenue. This auction format does not allow package bidding. The FCC auctions also divide the spectrum by geographic location. It is reasonable to expect that some bidders might receive extra benefits by obtaining larger, more contiguous portions of the spectrum. 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 and revenue of the auction. If the bid-

der 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.¹ In response to these difficulties, the FCC plans to allow package bidding in future auctions (Federal Communications Commission 2002, Dunford et al. 2001). In particular, the FCC in its Auction #31 for the upper 700 MHz band, affords bidders the ability to submit bids for packages of licenses. The particular design presented in this paper was developed prior to the FCC package auction design. Indeed, one of the major features of the FCC design was clearly influenced by the pricing rules we developed herein. Specifically, the FCC will use a "current price estimate" in Auction #31 that will provide a price for each license and these prices will be used to determine the minimal acceptable bids in the next round of bidding. We discuss this in more detail below in §3.3.

¹ An extreme example can be found in the PCS C Block auctions, where there were \$874 million in defaults (in Landler 1997).

While the potential utility of combinatorial auctions is considerable, combinatorial auctions have not yet reached their full potential in practice.² The successful implementation of a combinatorial auction requires one to overcome a number of hurdles. Two widely recognized and discussed issues are:

- (1) the computational complexity of the winner determination problem, and
- (2) the complexity of the bidding environment for the bidder.

Computational 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. Rothkopf et al. (1998) have shown that computational issues can be reduced via limitations on acceptable bids and other strategies. Others have found promising algorithms (Sandholm et al. 2001), and Andersson et al. (2000) have shown that CPLEX software is fairly effective in solving the winner determination problem in combinatorial auction simulations. For problems of reasonable size, the computational complexity of the winner determination problem is simply not the limiting factor.³

The computational complexity of the bidders' problems is more of an issue in practice. Bidders must determine their valuations for all subsets of items they are interested in (up to a maximum of 2^K values if they are interested in K items). Then they must formulate an optimal bidding strategy given those valuations. If the bidders make incomplete or incorrect calculations, the efficiency or revenue an auction will generate can be significantly reduced.

Several approaches have been taken to reduce these difficulties. Some have proposed using the Vickrey sealed-bid auction. Under a Vickrey auction, bidders have a dominant strategy to truthfully report their

values to the auctioneer. While this would eliminate strategic complexity, it would not reduce the complexity associated with valuation determination.⁴ Some have suggested a pay-what-you-bid, sealed-bid, one-shot auction (Rassenti et al. 1982), but this brings back the strategic complexity without reducing the valuation computation complexity. Others have suggested using progressive auctions, similar to an English auction, to reduce the cognitive burden on the bidders in both valuation and strategic computation (Banks et al. 1989, Parkes 1999). Two candidate auction designs in this area are the continuous package bid auction called the Adaptive User Selection Mechanism (AUSM) first proposed by Banks et al. (1989) and the SMR auction used by the FCC (Milgrom 2000). The idea in each of these is that bidders need only compute valuations when necessary, and that bidders have time to focus on and compute strategies.

In this paper, we take the best features of these auctions, add a new element, and create a new design we call Resource Allocation Design (RAD) that produces in our experiments higher efficiencies, higher revenues, and a shorter duration than the original designs. We take the issues of computational and cognitive complexity seriously and formulate a combinatorial auction mechanism that attempts to ease those burdens. The features we borrow are: (1) package bidding from AUSM, and (2) an iterative format, eligibility, and minimum bid increments from SMR. The feature we add is (3) a method to provide *prices* that will guide bidders to desirable outcomes. We use an iterative auction that gives bidders, at each iteration, a vector of prices—one for each object—that new bids must beat in order to be accepted. Bidders need not consider separate prices for each subset of objects—a subset price will simply be the sum of the prices for each item in the subset—so the information the bidders need to process is, in some sense, small. (See Nisan and Segal 2003 for a precise analysis of the size of the messages.) Bidders need not bother computing valuations for items whose prices are obviously higher than the valuations would be, thus reducing that dimension of computational complexity.

We use the test bed approach of experimental economics to establish the performance improvements of RAD over SMR and AUSM. There is no theory to use to compare auctions, especially if one wants to take account of the limitations of bidders' cognitive skills. We have found that computer simulations also fail to

² There are an increasing exceptions, including Sears Logistics Services (Ledyard et al. 2002), the Automated Credit Exchange (Ishikida et al. 2001), the course registration auction at the Chicago Business School (Graves et al. 1993), and the Mars IBM procurement auction (Davenport et al. 2003).

³ We can also report some additional data from Net Exchange (nex.com), based on simulations with a 200 MHz PC. They created test runs where the number of bids were four times the number of items, where 1/3 are multi-item bids, where 5% of the bids involve more than three items, and where OR groups were allowed. For problems based on hundreds of items (including runs in the 500–800 range), using DASH Express, the optimum was always found very quickly. For 1,000 items, 50% of the problems computed to the optimum in less than 30 seconds. 90% computed to the optimum in less than 30 minutes. 95% computed to within 2% of the best upper bound (the relaxed linear programming solution) in less than 30 minutes. And, in one other observation, Ledyard et al. (2002), the winner-determination problems for the Sears logistics auctions for 850 items always solved in less than 30 minutes, using now totally outdated 1992 technology. With modern technology and algorithms, fairly large problems are almost always easily solved in a timely fashion.

⁴ There are other problems with Vickrey auctions, many of which were originally noted, discussed, and explained by Groves and Ledyard (1977a, b) in the context of public goods. Banks et al. (1989) also discuss the problems and investigate an iterative version of Vickrey. It was the unsatisfactory performance of that mechanism that led to the development of AUSM.

capture many of the important details of human cognition. One cannot use the data from auctions that occur in practice because it is not possible to know the fundamentals—the true values of the items for each bidder.⁵ Therefore, if progress is to be made we must adopt the approach of an engineer’s wind tunnel and turn to the laboratory for data. The use of the laboratory as a test bed for complex auctions in complex environments began with Ferejohn et al. (1979), Smith (1979), Grether et al. (1981), and Rassenti et al. (1982). This methodology has proven to be fairly successful in providing guidance for the design of a variety of implemented auctions (Plott 1997, Ishikida et al. 2001, Ledyard et al. 1997). Building on knowledge from theoretical and practical experience, one can create test bed environments in the laboratory that 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 a simple test bed. We are able to compare the performance of RAD to a version of the SMR auction used by the FCC. Because we used the same test bed as in previous experiments, we are also able to compare the performance of RAD, when possible, with that of the AUSM combinatorial auction proposed by Banks et al. (1989), which is widely regarded as one of the first combinatorial auction mechanisms.

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

2. The Context

As most theorists realize, it is relatively simple to describe a demand-revealing, efficient auction. 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. If we accept that the winner-determination problem is not an issue, then

⁵ One cannot econometrically estimate them from the data unless one knows what the strategic behavior of the agents was, and that behavior is invertible.

the Vickrey auction appears to eliminate the strategic complexity facing the bidder. However, the bidder still faces the complexity of calculating and expressing these valuations.⁶ If K items are being auctioned, each agent’s bid would need to be 2^K numbers—potentially creating a very large, very complex communication 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).⁷ 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 et al. 1987).

Progressive auctions, such as the English auction, usually perform quite well in the laboratory (Coppinger et al. 1980). 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 iterative 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.

To understand the possibilities and choices facing the designer of multiobject auctions, we begin by recalling the key features of two vastly different designs: the SMR design (Milgrom 2000) and the AUSM (Banks et al. 1989). 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

⁶ In fact, Sandholm (2000) has shown that, if valuation computation is costly, the positive strategic implications of the Vickrey auction may not hold. Also, see Larson and Sandholm (2001b, a) for a further discussion of the strategic complexity of auctions.

⁷ 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 if the other bidders’ information had been revealed.

Logistics Services to procure trucking services (Ledyard et al. 2002). Three aspects of the design are the same for each: Winning bidders pay 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.

Both the SMR and AUSM auctions represent a compromise, the result of a sequence of design choices. 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, and the time to complete an auction.

In most discussions of the design of multiobject auctions, the primary goals have either explicitly or implicitly been high efficiency and/or high revenue. The goals of maximizing efficiency and maximizing revenue are not antithetical. In fact, the amount of revenue collected is generally limited by the efficiency of the auction. In single-item auctions, contingent upon sale, maximal revenue usually occurs by maximizing efficiency and then extracting as much of the surplus as possible (Myerson 1981). This approach does not always work in multiobject auctions.⁸ In environments without income effects, such as quasi-linear preferences, what trade-off there is can be most easily seen in the following identity:

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

High efficiency and low revenue can occur if and only if bidder profits are high, which might occur under collusion, and high revenue and low efficiency can occur if and only if bidders incur losses.

Both the SMR and AUSM auction processes have a difficult time consistently generating 100% efficiency across a variety of environments (Ledyard et al. 1997, Kwasnica et al. 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 superadditive. In order to win a package that 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 can expose the bidder to losses. Bidders who are aware of this problem might stop bidding in order to 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 (1999) 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 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 in order to displace 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.⁹ The use of a queue, however, shifts the computational burden to the bidders; they must now consider the bids in the queue 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.¹⁰ 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.¹¹ However, 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 a seemingly 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. However, 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,

⁸ In spite of Williams (1999), who identifies the optimal, efficient auction to be a Vickrey-Groves mechanism, we do not know that the optimal revenue-maximizing auction is always efficient. In fact Armstrong (2000) suggests it may not be so.

⁹ It is shown in Banks et al. (1989) that the queue increases both efficiency and revenue in continuous auctions.

¹⁰ In the field, it is difficult, if not impossible, to measure any of these variables. In the lab, because we know the induced valuations of bidders, we can directly measure efficiency, revenue, and losses.

¹¹ Formulating an optimal strategy to win packages of items when bidding is restricted to single-item bids is actually quite difficult.

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.

3. The Auctions

Rather than providing a fully general framework, in this paper we will focus on the particular designs we evaluate. Let $I = \{1, \dots, N\}$ represent the set of bidders, $K = \{1, \dots, K\}$ represent the set of objects to be sold and $t = 1, 2, 3, \dots$ represent the iterations or rounds. In general, a bid can be a very abstract entity involving complex contingent logic.¹² In this paper, we restrict our attention to very simple bids. A bid is a pair $b = (p, x)$, where p is a positive real number representing the bid price and $x \in \{0, 1\}^K$ represents the items desired.¹³ A bid here signifies, “I am willing to pay up to p for the collection of objects for which $x_k = 1$ if and only if I get all of them.” In the auctions we analyze, winners will actually pay what they bid.

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 $B_t = \bigcup_{i \in I} B_t^i$ is the set of all submitted bids. An arbitrary element of B_t is expressed as $b_j = (p^j, x^j)$.

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

$$\max \sum_{j \in B_t} p^j \delta^j \quad (1)$$

subject to

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

and

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

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. We acknowledge that it is well known that if the number of objects and bids is large, then one cannot necessarily guarantee that an optimal solution will be found in a reasonable amount of time. However, it should also be noted though that computation is increasingly less of an issue in combinatorial auctions (see Footnote 3). Of course, computation was

¹² See, for example, Ishikida et al. (2001), Rassenti et al. (1982), and Grether et al. (1981). Recent work by computer scientists emphasizes expressiveness (Nisan 2000). We discuss some of this in §3.3.

¹³ This structure can easily be generalized to cases where there are multiple copies of items available. We treat each $k = 1, \dots, K$ as a single indivisible object.

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 $W_t^i \equiv B_t^i \cap W_t$. There is no restriction placed on the number of winning bids for each particular bidder; each bid placed by an individual is considered independently from the other bids placed. 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 $t + 1$, so $W_t^i \subseteq B_{t+1}^i$ 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

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

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 quickly. 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. Specifically, a bidder's eligibility is the number of distinct objects he is allowed to bid on in a round.¹⁴ Let A_{t-1}^i be the number of distinct items i bid on in round $t - 1$,¹⁵

$$A_{t-1}^i = |\{k \mid x_k^j = 1 \text{ for some } (p, x) \text{ in } B_{t-1}^i\}|. \quad (3)$$

Initially bidders are allowed to bid on all items, $A_0^i = K$. In round t , a collection of bids B_t^i for i satisfies eligibility if and only if

$$A_t^i \leq A_{t-1}^i. \quad (4)$$

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.¹⁶ Eligibility can easily

¹⁴ 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 $\alpha_{t-1}^i = \{k \mid i \text{ has an active bid on } k \text{ in round } t - 1\}$. A bidder's eligibility in t is then $A_{t-1}^i = \sum_{k \in \alpha_{t-1}^i} w_k$.

¹⁵ We use $|A|$ to indicate the cardinality of the set A .

¹⁶ In the early rounds of the FCC spectrum auctions bidders were allowed to bid on more than A_{t-1}^i items by multiplying A_{t-1}^i by $r > 1$.

be checked incrementally as each new bid is offered. Because 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.

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

If eligibility satisfies this constraint, no bidder will be able to bid on anything other than the items they are currently provisionally winning. Therefore, ownership will not change in any subsequent period.¹⁷

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. It is based on a vector of single-item prices λ^t , which are known at the start of round t . In the SMR design, the price vector λ^{t+1} is simply the high price from t . That is,

$$\lambda_k^{t+1} = p_k \quad \text{if } (p, x) \in W_t \quad \text{and} \quad x_k = 1. \quad (6)$$

We let $\lambda_k^1 = 0$ for all k , but one could allow λ^1 to be any reserve prices. Let $N_t^i = B_t^i \setminus W_{t-1}^i$ be the set of new bids. Then we require that

$$p \geq \sum_{k \in K} x_k (\lambda_k^t + M) \quad \text{for all } (p, x) \in N_t^i, \quad (7)$$

where M is a minimum bid increment chosen by the auctioneer.¹⁸

Therefore, at the start of each round, each bidder $i \in I$ knows the objects for sale K , the prices on each object λ^t , her winning bids from the previous round W_{t-1}^i , and her eligibility A_{t-1}^i . Each bidder then chooses new bids N_t^i satisfying Equations (4) and (7). 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 rules given by Equations (1)–(7) describe what we have called the SMR design.

3.2. The AUSM Design

Because the AUSM design is a continuous auction, it is somewhat more difficult to formally define the AUSM rules using the notation developed earlier. Banks et al. (1989) provide a more complete definition. One can think of a continuous auction as an

iterative auction where only a single bid is placed in each round, or $|\bigcup_{i \in I} N_t^i| = 1$. 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 because 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 previous rounds are considered in each iteration: $B_t^i = B_{t-1}^i \cup N_t^i$.¹⁹

The basic AUSM design does not place any restrictions on the types of bids placed. AUSM does not use an eligibility calculation, prices λ^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 the context of a multiobject, 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.

Some more recent auction designs allow bidders to submit “exclusive or” (XOR) bids that allow a bidder to identify a subset of her bids and require that at most one of the bids in that subset be accepted.²⁰ Parkes (1999) and Ausubel and Milgrom (2002) are recent examples that allow this more expressive bidding language. In our structure, we allow any number of a bidder’s bids to be accepted. Although we were aware of the utility and power of XOR bids at the time we were running the experiments reported in this paper, we believed that the results would be more informative and persuasive if we modified the SMR in as few ways as possible.

¹⁹ In practice, the auctioneer posted B_{t-1}^i and required 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. Thus, the computational burden was shifted to each bidder.

²⁰ The value of XOR bids was actually recognized even in the most early combinatorial auction design (Rassenti et al. 1982). See Sandholm (2002a, b) for a formal introduction to XOR bidding languages.

¹⁷ As a referee correctly pointed out, this stopping rule can allow the auction to stop early if all bidders do not bid on some object(s).

¹⁸ The minimum increment could be altered over time, but we forego that degree of freedom in this paper.

Pricing is a bit more subtle. Game theory provides one suggestion—construct Vickrey prices for the auction. Vickrey prices are personalized prices 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. Therefore, in theory, if bidders are faced with Vickrey prices they should realize that their only strategy is to simply bid their value for each particular combination. In fact, if Vickrey prices are used, then the auction can be run in one round (effectively a sealed-bid auction). However, this approach has a number of drawbacks. First, it is not clear that bidders will interpret these bids correctly. In laboratory experiments with Vickrey auctions for only one object, bidders systematically deviate from the ideal strategy of simply stating one’s value. Second, in an auction of any size, bidders may not be able to submit all of their potentially desired packages in one round. It would require messages of 2^K numbers. 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. One could, of course, try to create an iterative Vickrey-type auction. This was done in Banks et al. (1989) with little success. The interested reader should consult the discussion there.

A second potential approach is closely related to the economic theory of competitive equilibrium. A set of prices—one for each object—is said to constitute a competitive equilibrium if, given these prices, the supply of objects equals demand (i.e., excess demand is zero). This approach was explored a bit by Bykowsky et al. (2000), where it was shown that if one simply prices the items for auction, competitive equilibrium prices may not exist because of the nonconvexities caused by complementarities. Bikhchandani and Ostroy (2002) took this further and looked at personalized (individual specific) prices on packages. They were able to provide several possibility results. However, from the point of view of auction design, both of these approaches leave something to be desired. First, from Bickchandani and Ostroy, we learn that we may need 2^K prices and so would be back in the communicatively difficult world of Vickrey. Second, competitive equilibrium is just that—an equilibrium theory. It is silent on the dynamics of price discovery—unless one wants to adopt the Walrasian tâtonnement, a process that does not work well in the laboratory because it requires recontracting, which opens up all sorts of possibilities for nonconstructive manipulation. It is our belief that the advantage of well-designed iterative auctions over one-shot auctions is that they allow orderly discovery

of alternatives and prices because important feedback information is provided to bidders between rounds.

So, we turn to a third and, ultimately, more productive approach. We restrict ourselves to only pricing items so as to keep communication complexity to a minimum. We then look for a pricing rule that will convey information to bidders about opportunities in the next bidding round. Three properties seem important for this: (a) all accepted bids should, if they were to pay these prices, pay something less than or equal to what they bid; (b) all losing bids should, if they were to pay these prices, pay something more than what they bid—indicating they needed to bid higher in order to win; and (c) new bids that are willing to pay more than the price of their bundle at those prices should have a good opportunity to win—that is, the prices ought to “guide” new bids to collections and values that can increase revenues.²¹

To insure (a), to keep computation simple, and to retain the “pay what you bid” nature of SMR, we chose to require that prices λ_k^t satisfy $\sum_{k \in K} \lambda_k^t x_k^j = p^j$ for all winners. Insuring (b) is a bit more difficult. Let $L_t = B_t \setminus W_t$ be the losing bids at t . To have (b) we would need a set of prices, λ^t , such that $p^j = \sum_{k \in K} \lambda_k^t x_k^j$ for all $j \in W_t$ and $p^j \leq \sum_{k \in K} \lambda_k^t x_k^j$ for all $j \in L_t$. If prices satisfy these equations, then the winning bidders would be paying their bid, and losing bidders would see that the prices were greater than their bid. Unfortunately, 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. There are many ways to do this. We choose one we believe, and the experimental evidence supports these beliefs, also provides good interround signals about opportunities, and lack of opportunities, for successful bids. To compute RAD prices λ^{t+1} , we begin by solving the following problem:²²

$$\min_{\lambda^t, Z, g} Z \tag{8}$$

²¹ A related approach is to design an agent that efficiently queries bidders about their valuations (Conen and Sandholm 2001). However, in our experimental setting, the issue of value elicitation was not an issue. One still needs to design an auction that helps bidders know what to do with their valuations.

²² One must recognize the pioneering work of Rassenti et al. (1982). They proposed a sealed-bid combinatorial auction (RSB auction) to solve an airport slot allocation problem and introduced the use of prices computed from a relaxed problem. They also provided experimental evidence as to the capability of their mechanism. The computation we use is a bit different from theirs, primarily for one reason. Because theirs is not an iterative process, the RSB prices are simply a way to collect dollars from the bidders and are not used as an information device. We need our prices to generate good signals during iteration.

subject to

$$\begin{aligned} \sum_{k \in K} \lambda_k^t x_k^j &= p^j \quad \text{for all } (p^j, x^j) \in W_t \\ \sum_{k \in K} \lambda_k^t x_k^j + g^j &\geq p^j \quad \text{for all } (p^j, x^j) \in L_t \\ 0 \leq g^j &\leq Z \quad \text{for all } (p^j, x^j) \in L_t \\ \lambda^t &\geq 0. \end{aligned}$$

Problem (8) selects a set of prices that ensures that revenue collected from the prices exactly equals the dollar amount for each winning bid; for losing bids it attempts to set prices that keep the package out with as little distortion as possible. The variable g^j is the amount for each losing package that ensures that the bid is not affordable.²³ We then want to find the smallest such value across all losing packages so that the distortion of the information is minimized. Let g^* and Z^* be a solution to (8). At the prices λ^t there may be some losing bids for which $\sum_{k \in K} \lambda_k^t x_k^j \leq p^j$, falsely signaling a possible winner. Such is the nature of package bidding. On the positive side, such bids can be resubmitted if $p^j - (\sum_{k \in K} \lambda_k^t x_k^j)$ is large enough. Further, Equation (8) is designed to minimize the maximum violation of the inequalities for losing bundles. In fact, if ideal prices exist, they will be the solution and $g^{j*} = 0$ for all $b^j \in L_t$.

At this point it may still be possible to further lower some of the g^j that, in the first solution, satisfy $0 \leq g^j \leq Z^*$. Therefore, to further the computation of λ^t , a sequence of iterations of Equation (8) is performed. We lexicographically lower as many g^j as possible. Therefore, at this point we have satisfied the desired property (a) and have done about as well as we can on property (b). What about (c), which asks that prices provide good signals about new bids with good opportunities to win? If the solution to Equation (8) after lexicographic minimization is unique, there is no more we can do. However, in many cases, the solution will not be unique, and we have an opportunity to improve. We know that the prices indicate, for all of the packages that were submitted in the previous round, what one would have to bid to have any chance of inclusion in the next round, assuming all other bids are resubmitted. As a result, the only way to improve on this is to signal where a new package might be successful. New packages will be successful if they can be combined with losers from the last round to bump a winner from the last round out of the solution, so we will finish the price computation in a way that provides relevant information.

²³ This is a little bit like two-part pricing, a well-known solution to pricing with nonconvexities, but because the losers never pay, the second part—the g —are never really collected.

For each winning bundle we lexicographically maximize the minimum price in the bundle subject to the constraints of Equation (8) at the g^* we solved for earlier. The formalities are provided in the appendix.²⁴

Why this works may seem mysterious, so we turn to three examples that illustrate what is happening here. The following examples help explain the ability of the RAD pricing rule to convey such information.²⁵

EXAMPLE 1. 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 Bidders 2 and 3 could combine to outbid the current standing bid. Any prices such that $\lambda_A + \lambda_B = 10$ and $\lambda_A \geq 8$ will satisfy Equation (8). However, if we choose $\lambda_A = 10$ and $\lambda_B = 0$, then Bidder 3 may bid 1 for $\{B\}$ in the next round only to find out that they lose. If $\lambda_A = 8$ and $\lambda_B = 2$, 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 provisional winners will be Bidder 2 with object $\{A\}$ and Bidder 3 with object $\{B\}$. The prices that would be generated by RAD would be 8 for A and 2 for B .

EXAMPLE 2. 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 $\{B\}$.

If we select $\lambda_A = 4$, then it must be that $\lambda_B = 6$. Given this information, Bidder 3 will assume that it 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 natural and fair decision is to “split the difference” by setting $\lambda_A = 5$ and $\lambda_B = 5$.

The appropriate prices identified in Examples 1 and 2 are obtained, when ideal prices exist, by minimizing the maximum of λ_A, λ_B subject to the prices satisfying Equation (8).

²⁴ A side benefit of this procedure is that we end up with a unique set of prices. This is important for “respectability.” It is important for bidder confidence that we get the same answer if we rerun the algorithms.

²⁵ The following examples assume the minimum increment is zero. It is possible that a large minimum increment might upset some of the usefulness of these prices.

EXAMPLE 3. 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 for $\{A, C\}$.
- Bidder 3 is willing to pay 15 for $\{C\}$, but has not bid.

Bidder 1 is the provisional winner. The prices we want, if they exist, satisfy $\lambda_A + \lambda_B + \lambda_C = 30$, $\lambda_A + \lambda_B \geq 25$, $\lambda_B + \lambda_C \geq 25$, and $\lambda_A + \lambda_C \geq 22$. Because the last three inequalities imply that $\lambda_A + \lambda_B + \lambda_C \geq 36$, no such prices can exist. We try to get as close as possible. We choose $\lambda_A, \lambda_B, \lambda_C$ and g^1, g^2, g^3 such that $\lambda_A + \lambda_B + g^1 \geq 25$, $\lambda_B + \lambda_C + g^2 \geq 25$, and $\lambda_A + \lambda_C + g^3 \geq 22$, 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 . We could pick $\lambda_A = \lambda_C = 9$, and $\lambda_B = 12$, yielding $g^1 = g^2 = g^3 = 4$. We could also pick $\lambda_A = \lambda_C = 11$ and $\lambda_B = 8$ 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 the same, 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 8.5, then Bidder 3 would only find it profitable to bid when the price is $\lambda_C = 8$. 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 Bidder 1's current high bid.

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 presented with a price for each object. There is one possible complaint one might register about our pricing rule; individual prices will not necessarily be increasing over time. This is because, over time, the opportunities for new packages to combine with old rejected packages to displace pro-

visional winners will change. This is an unavoidable feature of environments with complementarities when only prices on items are used. It is important to remember, however, that the sum over all prices is always increasing, and our experimental tests of RAD reported below indicate that subjects had no problem with this feature.

There is no reason, in theory, to expect these prices to work well or badly. However, we demonstrate, through the use of human subject experiments, that RAD can perform quite well across a number of reasonable performance measures. As the reader will see in the data below, this combination of pricing and stopping rule works very well together to eliminate strategic problems caused by the threshold problem. Changing the SMR design to allow package bidding with the particular pricing rule we designed generates a significant increase in performance in environments with multiple objects with complementarities.²⁶ There is also no degradation of performance in the goods 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 et al. (1997) and an additive environment. Because the two value environments are combined into one environmental test bed, we can see if there are spillover problems from items with complementarities to those without complementarities. Specifically, the five participants were allowed to bid on 10 heterogeneous items labeled $A, B, C, D, E, F, G, H, I$, and J . Bidder values for the first six items were highly superadditive. 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 from a uniform distribution with support $[0, 10]$.
- The two-item packages,

$$(\{A, B\}, \{A, C\}, \dots, \{E, F\}),$$

took integer values drawn independently from a uniform distribution with support $[20, 40]$.

- The three-item packages,

$$(\{A, B, C\}, \dots, \{D, E, F\}),$$

had integer values determined independently by draws from a uniform distribution with support $[140, 180]$.

²⁶ Cybernomics (2000) reports on experiments in which a package bidding extension for the SMR was tested that did not use prices, but provided eligibility benefits for certain bids. In a different class of environments, they find efficiency results similar to ours, but the amount of time required to complete an auction was generally longer. They were aware of our results prior to their work, but chose not to adopt the pricing feature of our design.

- 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 of the packages. In order to obtain the value for a particular package, the bidder had to obtain all objects in the package, and the bidder could not include the same object in multiple packages. The bidder's value for a disjoint combination of packages was given by the summation of the package values. All other packages had zero value to the bidder. 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 Periods 3 and 5, *competitive* equilibrium prices did not exist. In Table 1, we provide 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, as we suspected that under some auction designs 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 were not negative.²⁷ 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 number of auctions (or periods) 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

²⁷ In reality, in many of the SMR auction sessions, even these four additive objects were not enough.

Table 1 Values in a Spatial Fitting Example

Bidder 1					
Packages:	{F}	{C, D}	{B, C, F}	{B, D, E}	{A, B, E}
Values:	9	22	128	130	120
Bidder 2					
Packages:	{B}	{D, F}	{A, E}	{A, F}	{A, B, D}
Values:	8	28	24	27	130
Bidder 3					
Packages:	{C}	{A}	{D}	{B, D}	{A, B, F}
Values:	2	3	8	20	119
Bidder 4					
Packages:	{E}	{A, B, C}	{A, D, F}	{B, D, F}	{A, E, F}
Values:	10	117	112	128	125
Bidder 5					
Packages:	{C, F}	{D, E}	{C, E, F}	{B, E, F}	{A, B, C, D, E, F}
Values:	29	25	117	125	142

auction format; they had all participated in training sessions that utilized simplified auction rules and environments.

A total of 25 RAD and 17 SMR auctions were completed in 15 experimental sessions. The AUSM data come from 12 auctions completed in previous experiments reported by Ledyard et al. (1997).²⁸

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; different auctions will perform better depending on which measure one focuses. For example, high efficiency may 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, 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 $V^i: \{0, 1\}^K \rightarrow \mathbb{R}$, 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

$$\sum_{i=1}^N y_k^i \leq 1 \quad \text{for all } k = 1, \dots, K$$

$$y^i \in \{0, 1\}^K.$$

²⁸ Due to the secondhand nature of the AUSM data, we could not compare AUSM to RAD and SMR in all cases.

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

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

It is true (see Ledyard et al. 1997) that the absolute level of efficiency can be deceptive because one can increase the percentage by simply adding a constant amount to each V^i function. This leaves the efficient allocation unchanged, but increases E when $E < 1$. However, we will only use efficiency to compare performance across auctions in the same environment. Therefore, 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. Because revenue can vary significantly across environments, we used the percentage of the maximum possible surplus (V^*) that is actually captured by the seller as our measure of seller revenue. Revenue as a percentage of maximum possible surplus is given by

$$R = \frac{\sum_i \sum_k \lambda_k^* \hat{y}_k^i}{V^*},$$

where λ^* is the vector of final prices. As with efficiency, it is not the absolute value we care about, but relative performance across auctions.

5.3. Bidder Profit

Bidder profit is another possible performance measure. With the presence of significant complementarities, some auction mechanisms can cause some bidders to lose money (Bykowsky et al. 2000). A high probability of losses can lead to a variety of performance failures. Bidders may be unwilling to participate in auctions in which they know they are likely to lose money. 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 on the bid i is given by

$$P^i = V^i(\hat{y}^i) - \lambda^* \cdot \hat{y}^i,$$

where λ^* is the vector of final prices.

5.4. Net Revenue

Because of the possibility of bidder losses, we also measure what the auctioneer might expect to actually

collect at the end of an auction. It is likely that bidders who made losses would default on their payments after the auction is over. Assume that any bidder that would experience losses by completing the deal does default on at least the portion of their bid that is not profitable.²⁹ What the auctioneer would actually collect under these circumstances is given by net revenue as a percentage of maximum possible revenue,

$$NR = \frac{(\sum_i \sum_k \lambda_k^* \hat{y}_k^i) + \sum_i L^i}{V^*},$$

where

$$L^i = \begin{cases} P^i & \text{if } P^i < 0 \\ 0 & \text{otherwise.} \end{cases}$$

In other words, revenue is only generated from the portion of sales that are profitable for the bidder as well.

5.5. Auction Duration

When analyzing iterative auctions, the duration of the auction becomes a relevant concern. In this paper we measure auction duration 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 costs 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 that generally leads to lower efficiency and revenue. There is a possible trade-off between auction duration and efficiency. Increased iterations may allow high-value bidders to find the right package, thus increasing efficiency.

Because 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 duration for either environment. In order to determine the auction duration for the additive markets, we identified the round in which these four markets would have closed if there were no spatial fitting markets. For example, an auction may have lasted 12 iterations, but 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 7 because, 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 envi-

²⁹ Our measure of net revenue is designed to be conservative in favor of revenue generation in auctions with losses. Therefore, we only subtract the actual losses from the revenue amount. It might be reasonable to assume that a bidder would default on their entire payment if they ended the auction at a loss. This would obviously greatly reduce the net revenue calculation.

ronment 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 symmetric measure was used for the spatial fitting environment.

6. Results

In this section, we compare the performance of the RAD, SMR, and AUSM mechanisms. The bottom line is that in complex environments, RAD yields higher efficiencies, higher net revenues, and lower bidder losses than does SMR, and RAD does it in many fewer iterations. Also, we find that AUSM lies somewhere in between RAD and SMR in performance in complex environments. There is no performance difference in simple additive environments.

6.1. Results from the Spatial Fitting Test Bed

We begin by considering efficiency. The average auction efficiency across periods under the SMR design was 67%. The average efficiency for the RAD design was 90%. The continuous AUSM obtained an average efficiency of 94%, which is not significantly different from the results for RAD. Table 2 gives the results of Wilcoxon-Mann-Whitney rank-sum pairwise comparisons of these three institutions.³⁰ RAD significantly improves the efficiency of the allocation over SMR in all periods.

CONCLUSION 1. RAD yields efficiency at least as high as AUSM and significantly higher than SMR.

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, 20 out of 25 (80%) auctions under RAD and 10 out of 12 (83%) auctions under AUSM led to full efficiency. This is true in only 4 out of 17 (24%) auctions using the SMR design. The fact that both auctions that allow package bidding yield dramatically higher efficiencies than the SMR design suggests an obvious conclusion.

CONCLUSION 2. Package bidding significantly increases efficiency.

When package bidding was not allowed, SMR bidders, as a whole, averaged losses of \$7.73 in each period for the markets with complementarities. In RAD, where package bidding was permitted, bidders earned positive profits on average ($z = 2.83$, $p = 0.006$).

³⁰ The Wilcoxon-Mann-Whitney rank-sum test is a powerful non-parametric substitute to the standard t -test when data has at least ordinal measurement (Siegel and Castellan 1998). When examining data generated from human subjects, it is typical to assume that the data do not meet the assumptions required for a t -test. A high test statistic, z , indicates that the second institution is stochastically larger (in terms of the performance measure) than the first. The reported p s are the p -values associated with the null hypothesis that the first institution is greater than or equal to the second institution.

Table 2 Spatial Fitting Wilcoxon-Mann-Whitney Rank-Sum Test Results

Performance measure	Institutions compared		
	SMR vs. AUSM	SMR vs. RAD	AUSM vs. RAD
Efficiency	$z = 3.29$ $p = 0.000$	$z = 3.55$ $p = 0.000$	$z = 0.332$ $p = 0.371$
Bidder profits	$z = 3.05$ $p = 0.002$	$z = 2.83$ $p = 0.006$	$z = 0.584$ $p = 0.280$
Net revenue	$z = 1.28$ $p = 0.100$	$z = 2.23$ $p = 0.013$	$z = 1.23$ $p = 0.109$

Total bidder profit averaged \$4.23 in RAD and \$5.68 in AUSM. Table 2 gives the results of Wilcoxon-Mann-Whitney rank-sum pairwise comparisons of these three institutions. 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. Under AUSM, only 1 out of 60 (1.7%) bidders ended an auction with losses.

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

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. To win a package, bidders must put themselves at risk of obtaining only 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 management encourages bidders to bid on as many items as possible in order to keep an option open. 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 even if that bid were higher than its true value. 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 management remains to be seriously studied. However, it is clear that it leads to bids that are inconsistent with short-run value maximization.³¹

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 were inconsistent with their valuations. A simple example of this occurs if a bidder had a

³¹ The use of XOR bids is one obvious solution to this problem.

value of 100 for the package $\{A, B, C\}$, and a value of no more than 25 for any two-item subsets. 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.³² If those bids were sufficiently high, no other bidder would be able to *rescue* them by out bidding them.³³ 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.

Seller net revenue as a percentage of maximum possible revenue was 74% and 69% under RAD and AUSM, respectively. The SMR auction, on the other hand, averaged a net revenue of only 61%. Both RAD and AUSM yield significantly higher net revenue than does SMR. In Table 2, we provide Wilcoxon-Mann-Whitney rank-sum pairwise comparisons of the net revenue of the three auction types. The package-bid auctions generally are expected to collect all of the revenue generated by the auction. However, the revenue for SMR is reduced from the apparent amount of 96% to the realistic expectation of 61% due to the substantial bidder losses. Because there are few instances of bidder losses under package bidding, actual revenue under RAD and AUSM is close to net revenue, at 79% and 71%, respectively.

CONCLUSION 4. Package bidding significantly increases seller net revenues—those revenues the seller can expect to collect.

RAD yields somewhat higher net and absolute revenue than AUSM, which implies that the RAD design is able to extract more of the surplus from the bidders. Using a Wilcoxon-Mann-Whitney rank-sum test, we find that revenue as a percentage of maximum surplus is significantly greater under RAD ($z = 1.62, p = 0.055$). We conjecture that this is because the second-highest bidders, the ones whose bids drive the winners to increase their bids, are more easily able to find and express their willingness to pay in the iterative mode than in a continuous mode.

Auction duration, measured as the number of iterations before completion of the auction, was significantly shorter under RAD. Using the same price increment rule in both SMR and RAD, the SMR auctions averaged 16.2 iterations as compared to 3.32 for RAD. A Wilcoxon-Mann-Whitney rank-sum test indicates that the auction duration is significantly shorter under

the RAD design than in the SMR auction ($z = 4.98, p = 0.000$). In fact, it often took longer to complete the additive markets than the spatial fitting items (see Conclusion 9). Because the AUSM mechanism was a continuous auction, it is obviously not possible to directly compare the speed of these two formats.

CONCLUSION 5. Auction duration is shortest under RAD.

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 among the additive objects, so package bids would occur only if bidders were attempting a sophisticated strategy³⁴ to capture a larger share of the objects. However, this rarely happened. In only 3 of 25 RAD auctions do the final winning bids contain packages of additive objects. Further, in these three auctions the final allocations involve package bids across additive and other objects.³⁵

CONCLUSION 6. Package bids rarely occur among the winning bids in the additive environment.

Although bidders are clearly willing to bid on packages in the additive environment, they are rarely able to use this ability to their advantage, as is evidenced by the extremely high levels of efficiency achieved in the additive environment. 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 lead to full efficiency in the additive environment in 10 out of 12 (83%) auctions. 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.

CONCLUSION 7. In the additive environment, under RAD, SMR, and AUSM, efficiency is very near 100%. There are no discernible differences between the auctions.

In the additive environment, there was very little difference in the revenue collected by the seller under SMR and RAD. The SMR and RAD mechanisms averaged 69.96% and 71.96% of the maximum possible

³² This actually happened to one of the authors during early software tests.

³³ 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.

³⁴ Such a strategy might be to create an artificial *threshold* that would yield a possible problem for others, allowing the bidder to get the items even if it were not an efficient allocation.

³⁵ We used the final prices to estimate the portion of the bid occurring in the additive environment.

revenue, respectively. A rank-sum test also shows no significant difference between the observed revenues.

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

RAD yielded lengths that were significantly shorter than the SMR. The average auction duration under the SMR auction was 11.7 iterations, but was only 6.1 for RAD.³⁶ A Wilcoxon-Mann-Whitney rank-sum test indicates that the auction duration is significantly shorter under the RAD design than in the SMR auction ($z = 4.67$, $p = 0.000$). As before, there is no direct comparison with the speed of AUSM.

CONCLUSION 9. In the additive environment, auction duration is shortest under the RAD design.

Because package bidding has no advantage in the additive environment, and assuming all bids in the additive markets are on the individual items the prices, 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

7.1. Conclusions

The experimental test results point to two clear conclusions.

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

2. RAD dominates SMR in efficiency, net revenue auction duration, and protecting bidders from losses.

The general principle that package bidding is an important option for multiobject 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 if it can be collected. However, if the design is used repeatedly, bidders will learn to avoid these losses, perhaps by avoiding the auction altogether, and efficiency and revenue will ultimately suffer.

However, we have gone further here than simply establishing that package bidding is sensible. We have

³⁶ These results are, of course, confounded by the fact that the length of 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.

provided a new auction design, RAD, which clearly outperforms others. Relative to the SMR design, RAD produces higher efficiency, greater net revenue, greater 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. Because RAD uses a pricing rule instead of a queue to mitigate the threshold problem, it is 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.³⁷ The competitive prices, which do exist in this case, are 38, 49, 30, 43, 38, and 49. If we examine the data for this parameter set (Period 2), we find that this difference between stand-alone and competitive prices is prevalent experimentally. The RAD prices, however, are close to the competitive prices. In general, we would expect final prices to be somewhat 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 is considered, the RAD prices are not significantly different from 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. As 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 aid in finding an appropriate fit.

7.2. Open Issues

It would seem that the RAD design would be a natural candidate for use as a multiobject iterative auction in its current form. However, in spite of the excellent performance in our tests, there are at least two 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

³⁷ These are simply the maximal single-item values. The only way prices under the SMR design could be higher is if someone bid on a single item above their value.

may find themselves bidding for, and even winning, items that 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 straightforward solution: the use of exclusive “XOR” bids in an iterative auction. Bidders would be allowed to place a bid saying “I bid \$1,000 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. XOR 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.

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 though they belong in the efficient allocation. An example will illustrate.

EXAMPLE 4. Suppose in Round 5 there are four bids submitted as follows:

- Bid #1 for {A, B, C} at 99.
- Bid #2 for {A, B} at 75.
- Bid #3 for {A, C} at 75.
- Bid #4 for {B, C} at 75.

Under the RAD design, Bid #1 wins and the prices³⁸ are $(\lambda_A, \lambda_B, \lambda_C) = (33, 33, 33)$. 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 cannot bid because $\lambda_A > 30$. 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 that Bid #3 is not resubmitted, but Bids 1, 2, and 4 are. Then, 1 still wins and $\lambda = (24, 51, 24)$.³⁹ If it gets to this stage, our bidder for A can re-enter 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.

The crucial point to remember is that, in spite of these problems, RAD attains high efficiency and outperforms SMR. Clearly, the problems facing bidders in SMR, like the exposure problem, are more severe than those problems that face them in RAD.

³⁸ Notice that these are not separating prices, which is what causes a problem.

³⁹ These are separating prices. This also shows that prices are not necessarily monotonically increasing (since $24 < 33$). The sum of prices is, however, always increasing.

Acknowledgments

The authors thank the associate editor and two anonymous referees for careful, constructive reviews. The authors thank Elena Katok, Evan Kwerel, Paul Milgrom, Charles Plott, and Rakesh Vohra for their helpful comments. This is a significantly revised version of DeMartini et al. (1999).

Appendix

In §3.3, we indicated that the RAD auction-pricing algorithm (8) might not yield a unique price vector. We use the following lexicographic routines to eliminate that ambiguity.⁴⁰

Let g^* , λ^* , and Z^* solve (8). If $Z^* = 0$, then go to Problem (10) below. If $Z^* > 0$, let $J^* = \{j \in L_t \mid Z^* = g^{*j}\}$. If $J^* = L_t$, then go to (10) below. Otherwise,

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

subject to

$$\begin{aligned} \sum_{k \in K} \lambda_k^t X_k^j &= p^j \quad \text{for all } (p^j, x^j) \in W_t \\ \sum_{k \in K} \lambda_k^t X_k^j + g^{*j} &= p^j \quad \text{for all } (p^j, x^j) \in J^* \\ \sum_{k \in K} \lambda_k^t X_k^j + g^j &= p^j \quad \text{for all } (p^j, x^j) \in L_t \setminus J^* \\ 0 \leq g^j &\leq Z \quad \text{for all } (p^j, x^j) \in L_t \setminus J^* \\ \lambda^t &\geq 0. \end{aligned}$$

Let \widehat{Z} , \widehat{g} , $\widehat{\lambda}$, be the solution to (9). If $\widehat{Z} = 0$, go to 10 below. Otherwise, let $\widehat{J} = \{j \mid \widehat{Z} = \widehat{g}^j\}$. If $J^* \cup \widehat{J} = L_t$, then go to Problem (10) below. Otherwise, let $J^* = J^* \cup \widehat{J}$ and go to (9) again.

When the iteration on (9) is complete we will have prices that approximate our “ideal” but not always obtainable prices. They may still not be unique. Therefore, we go through a sequence of iterations that eliminate nonuniqueness and that create prices to guide bidders to solve the threshold problem. Let \widehat{g} be the solution from the last iteration of (9). Let $\widehat{K} = K$.

$$\max_{Y, \lambda} Y \tag{10}$$

subject to

$$\begin{aligned} \sum_{k \in K} \lambda_k^t x_k^j &= p^j \quad \text{for all } (p^j, x^j) \in W_t \\ \sum_{k \in K} \lambda_k^t x_k^j + \widehat{g}^j &= p^j \quad \text{for all } (p^j, x^j) \in L_t \\ \lambda_k^t &\geq Y \quad \text{for all } k \in \widehat{K}. \end{aligned} \tag{11}$$

Let Y^* , λ^* solve (10). Let $K^* = \{k \in K \mid \lambda_k^t = Y^*\}$. Let $\widetilde{K} = \widehat{K} \setminus K^*$. If $\widetilde{K} \neq \emptyset$, return to (10) and solve it by replacing (11) with

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

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

⁴⁰ An alternative approach would minimize $\sum_i (g^i)^2$ in (8), 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.

References

- Andersson, A., M. Tenhunen, F. Ygge. 2000. Integer programming for combinatorial auction winner determination. *Proc. 4th Internat. Conf. Multiagent Systems, Boston, MA*.
- Armstrong, M. 2000. Optimal multi-object auctions. *Rev. Econom. Stud.* 67(3) 455–481.
- Ausubel, L., P. Milgrom. 2002. Ascending auctions with package bidding. *Frontiers Theoret. Econom.* 1(1) www.bepress.com.
- Banks, J. S., J. O. Ledyard, D. P. Porter. 1989. Allocating uncertain and unresponsive resources: An experimental approach. *RAND J. Econom.* 20(1) 1–25.
- Bikhchandani, S., J. M. Ostroy. 2002. The package assignment model. *J. Econom. Theory* 107(2) 377–406.
- Bykowsky, M. M., R. J. Cull, J. O. Ledyard. 2000. Mutually destructive bidding: The FCC auction problem. *J. Regulatory Econom.* 17(3) 205–228.
- Conen, W., T. Sandholm. 2001. Preference elicitation in combinatorial auctions: Extended abstract. *Proc. ACM Conf. Electronic Commerce, Tampa, FL, ACM, New York*.
- Coppinger, V. M., V. L. Smith, J. A. Titus. 1980. Incentives and behavior in English, Dutch and sealed-bid auctions. *Econom. Inquiry* 43 1–22.
- Cybernomics. 2000. An experimental comparison of the simultaneous multiple round auctions and the CRA combinatorial auction. Technical report, Report to the Federal Communications Commission, FCC, Washington, D.C.
- Dasgupta, P., E. Maskin. 2000. Efficient auctions. *Quart. J. Econom.* 115(2) 341–388.
- Davenport, A. J., C. An, E. Ng, G. Hohner, G. Reid, H. Soo, J. R. Kalagnanam, J. Rich. 2003. Combinatorial and quantity-discount procurement auctions benefit Mars, Incorporated and its suppliers. *Interfaces* 33(1) 23–35.
- DeMartini, C., A. Kwasnica, J. Ledyard, D. Porter. 1999. A new and improved design for multiobject iterative auctions. Social science working paper 1054, California Institute of Technology, Pasadena, CA.
- Dunford, M., M. Durbin, K. Hoffman, D. Menon, R. Sultana, T. Wilson. 2001. Issues in scaling up the 700 MHz auction design. Presentation, FCC, Washington, D.C., <http://wireless.fcc.gov/auctions/conferences/combin2001/presentations.html>.
- Federal Communications Commission. 2002. Auction of licenses in the 747-762 and 777-792 MHz bands scheduled for June 19, 2002: Further modification of package bidding procedures and other procedures for auction no. 31. Technical report AUC-02-31-B, Federal Communications Commission, Washington, D.C.
- Ferejohn, J., R. Forsythe, R. Noll. 1979. An experimental analysis of decision making procedures for discrete public goods: A case study of a problem in institutional design. V. L. Smith, ed. *Research in Experimental Economics*, Vol. 1. JAI Press, Greenwich, CT, 1–58.
- Graves, R., L. Schrage, J. Sankaran. 1993. An auction method for course registration. *Interfaces* 23(5) 81–92.
- Grether, D., M. Isaac, C. Plott. 1981. The allocation of landing rights by unanimity among competitors. *Amer. Econom. Rev.* 71 166–171.
- Groves, T., J. Ledyard. 1977a. Some limitations of demand-revealing processes. *Public Choice* 29(2) 107–124.
- Groves, T., J. Ledyard. 1977b. Some limitations of demand-revealing processes—Reply. *Public Choice* 29(2) 139–143.
- Ishikida, T., J. Ledyard, M. Olson, D. Porter. 2001. The design of a pollution trading system for Southern California's RECLAIM emissions trading program. R. M. Isaac, ed. *Research in Experimental Economics*, Vol. 8. JAI Press, Greenwich, CT.
- Kagel, J. H., R. M. Harstad, 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, J. Scott. 1998. The design of multi-object multi-round auctions. Social science working paper 1045, California Institute of Technology, Pasadena, CA.
- Landler, M. 1997. Airwave auctions falter as source of funds for U.S. *New York Times* (April 3).
- Larson, K., T. Sandholm. 2001a. Computationally limited agents in auctions. *Proc. Internat. Conf. Autonomous Agents, Workshop on Agent-Based Approaches to B2B, Montreal, Canada*. 27–34.
- Larson, K., T. Sandholm. 2001b. Costly valuation computation in auctions. *TARK: Theoret. Aspects Reasoning Knowledge* 8 169–182.
- Ledyard, J. O., D. Porter, A. Rangel. 1997. Experiments testing multiobject allocation mechanisms. *J. Econom. Management Strategy* 6 639–675.
- Ledyard, J. O., M. Olson, D. Porter, J. A. Swanson, D. P. Torma. 2002. The first use of a combined value auction for transportation services. *Interfaces* 32(5) 4–12.
- Milgrom, P. R. 2000. Putting auction theory to work: The simultaneous ascending auction. *J. Political Econom.* 108(2) 245–272.
- Milgrom, P. R., R. J. Weber. 1982. A theory of auctions and competitive bidding. *Econometrica* 50(5) 1089–1122.
- Myerson, R. B. 1981. Optimal auction design. *Math. Oper. Res.* 6 58–73.
- Nisan, N. 2000. Bidding and allocation in combinatorial auctions. *ACM Conf. Electronic Commerce*. 1–12.
- Nisan, N., I. Segal. 2003. The communication requirements of efficient allocations and supporting lindahl prices. Forthcoming.
- Parkes, D. C. 1999. iBundle: An efficient ascending price bundle auction. *Proc. ACM Conf. Electronic Comm.* ACM, New York, 148–157.
- Plott, C. R. 1997. Laboratory experimental testbeds: Applications to the PCS auction. *J. Econom. Management Strategy* 6 605–638.
- Porter, D. P. 1999. The effect of bid withdrawal in a multi-object auction. *Rev. Econom. Design* 4(1) 73–97.
- Rassenti, S. J., V. L. Smith, R. L. Bulfin. 1982. A combinatorial auction mechanism for airport time slot allocation. *Bell J. Econom.* 13 402–417.
- Rothkopf, M. H., A. Pekec, R. M. Harstad. 1998. Computationally manageable combinatorial auctions. *Management Sci.* 44(8) 1131–1147.
- Sandholm, T. 2000. Issues in computational Vickrey auctions. *Internat. J. Electronic Commerce* 4(3) 107–129.
- Sandholm, T. 2002a. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence* 135(1–2) 1–54.
- Sandholm, T. 2002b. eMediator: A next generation electronic commerce server. *Comput. Intelligence* 18(4) 656–676.
- Sandholm, T., S. Suri, A. Gilpin, D. Levine. 2001. Cabob: A fast optimal algorithm for combinatorial auctions. *Internat. Joint Conf. Artificial Intelligence, Seattle, WA*. AAAI Press, Menlo Park, CA, 1102–1108.
- Siegel, S., N. J. Castellan. 1988. *Nonparametric Statistic for the Behavioral Sciences*. McGraw-Hill, New York
- Smith, V. L. 1979. Incentive compatible experimental processes for the provision of public goods. V. L. Smith, ed. *Research in Experimental Economics*, Vol. 1. JAI Press, Greenwich, CT, 59–168.
- Williams, S. 1999. A characterization of efficient, Bayesian incentive compatible mechanisms. *Econom. Theory* 14(1) 155–180.