

Combinatorial Auction Design

**David Porter, Stephen Rassenti, and Vernon Smith
Interdisciplinary Center for Economic Science
George Mason University**

**Anil Roopnarine
Cybernomics Inc.**

March 1, 2003

Introduction

Combinatorial auctions enhance our ability to efficiently allocate multiple resources in complex economic environments. They explicitly allow buyers and sellers of goods and services to bid on packages of items with related values or costs. For example, “I bid \$10 to buy 1 unit of item A and 2 units of item B, but I won’t pay anything unless I get everything.” They also allow buyers, sellers and the auctioneer to impose logical constraints that limit the feasible set of auction allocations. For example, “I bid \$12 to buy 2 units of item C OR \$15 to buy 3 units of item D, but I don’t want both.” Finally, they can handle functional relationships amongst bids or allocations, such as budget constraints or aggregation limits that allow many bids to be connected together. For example, “I won’t spend more than a total of \$35 on all my bids” or “This auction will allocate no more than a total of 7 units of items F, G and H.”

There are several reasons to prefer to have the bidding message space expanded beyond the simple message space used in traditional single commodity auctions. As Bykowsky et al. (2000) point out, when values have strong complementarities, there is a danger of ‘financial exposure’ that results in losses to bidders if combinatorial bidding is not allowed. For example, in the case of complementary items such as airport take-off and landing times, the ability to reduce uncertainty to the bidder by allowing him to precisely declare his object of value, a cycle of slots for an entire daily flight pattern, is obvious: failing to acquire one component slot ruins the value of the flight cycle.¹ In the same situation substitution possibilities would also be important to consider: if flight cycle A is not won, cycle B may be an appropriate though less valuable substitute for the crew and equipment available. Allocation inefficiencies due to financial exposure in non-combinatorial auctions have been frequently demonstrated in experiments beginning with Rassenti et al. (1982) (see also Porter (1999), Banks et al. (1989), Ledyard et al. (2000) and Kwasnika et al (2003)).

This paper is concerned with designing a transparent, efficient and practical combinatorial buyers’ auction for multiple distinct items with perhaps multiple units of each available. We focused this endeavor around the past auctions, debate, proposals, and tests precipitated by the Federal Communications Commission’s (FCC) quest for the design of the ‘perfect’ spectrum auction. Many of the design principles ultimately implemented by us are easily extensible to more complex two sided² negotiations with environmental³ constraints. We begin by mimicking the original FCC debate: describing several combinatorial auction designs and their potential problems. We next present some pertinent laboratory results of auction design already tried. We then describe a new

¹ In some instances, such risk may reduce the aggressiveness with which bidders bid and, in so doing, may lead to a misassignment of items. In other instances, bidders may bid too aggressively and, in so doing, obtain items at prices that exceed the value the bidder places on those items.

² A buyers’ auction becomes two sided when sellers may participate as buyers of their own offerings.

³ For example, conservation of flow in network distribution systems or storage and channel capacity limits can be applied as allocation constraints in the auction proposed.

combinatorial auction design that we invented to eliminate many of the problems found in previous designs. We finish with the description and testing of three auctions previously untested in the laboratory, the one currently used by the FCC in the field, a combinatorial auction design solicited by the FCC but not yet implemented, and our own new design.

FCC Auction Design Process

The FCC conducted several high priced auctions for communications' spectra during the 1990s. The Simultaneous Multi-Round Auction (SMR) design employed by the FCC has several notable features:

- Simultaneous auction of the various licenses offered
- Iterative procedure allowing round by round update of bids
- Activity rules demanding continuous participation at minimum bid increments

After much discussion and debate, but no controlled testing, the FCC chose to implement the SMR. The FCC did not originally opt for a combinatorial form of auction for three stated reasons:

1. Computational uncertainty. At any point during a combinatorial auction, the selection of the winning bids and what it would cost for competition to displace them typically requires the solution of integer programming (IP) problems. These problems are notorious for being computationally burdensome: technically described as NP complete or hard problems. At issue here is that there is no guarantee that the solution for such a problem can be found in a "reasonable" amount of time when the number of bidders and items becomes larger.
2. Bidding complexity. Combinatorial auctions would be burdensome and difficult for participants and the auctioneer for at least two reasons. First, there are inconceivably many packages on which a bidder might want to place bids, and selecting any subset may be strategically awkward and provide the auctioneer with incomplete information.⁴ Second, there is a computational problem on the bidder to determine how much to bid to be successful.
3. Threshold problem. Combinatorial auctions present the following strategic impediment to efficient outcomes. Suppose each of two small bidders is bidding on a separate item, but a third bidder is bidding on a package that contains both items. Then the two small bidders must implicitly coordinate through their bidding to ascertain what price each will pay in order that the sum of both bids exceeds the package bid.

⁴ This is the supermarket problem also referred to at various FCC discussions as the 2^N boogie-man since there are that many combinations of items a bidder must consider. Suppose you have \$100 in your pocket and you are standing at the entrance to the supermarket. You despondently realize that you may never come out as you must first consider every possible way you might fill your shopping cart.

The FCC chose to implement the SMR design in all of its spectrum auctions rather than confront the above issues. Bidders in an SMR auction who had superadditive values for a particular package of licenses were subject to financial exposure if they miss acquiring at least one of those licenses. The FCC auctions implemented a withdrawal rule to “reduce” the potential financial exposure of the bidders: that is, bidders had the right to withdraw bids on particular licenses subject to the fact that they were obligated to pay the difference between their own and the final winning bid if it was less. This and other SMR rules led to results in various FCC auctions that revealed some interesting perverse strategies. In particular, individuals would withdraw and then bid just below their withdrawn bid to signal a willingness to not compete. To manage eligibility, bidders would bid on items they did not have value for in order to maintain activity without showing their hand on what they were interested in bidding for (this was called “parking”). That prompted several auction rule adjustments. Two notable rule changes implemented were restrictions on jump bidding (the allowable increment when raising your bid) and restrictions on the number of withdrawals allowed (this rule was precipitated by concerns that bidders were using such withdrawals to “game” the auction, as opposed to using them to eliminate financial risk).

Based on the previous problems associated with the SMR design and spurred by the FCC sponsored debate, several new designs have emerged.

Auction Designs for Complex Environments

It is quite natural for an auction participant to wish to reveal as little as possible concerning his interest and values for particular items or packages in an auction that simultaneously sells multiple items. In general, auction systems that provide feedback and allow bidders to revise their bids seem to produce more efficient outcomes. This feedback feature is the cornerstone of the recent FCC SMR spectrum auctions, and many others worldwide. Experience in both field and laboratory suggests that in complex economic environments iterative auctions, which enhance the ability of the participant to detect keen competition and learn when and how high to bid, produce better results than sealed bid auctions.

An iterative combinatorial auction could allow bidders to explore the bid space without having to place bids for all possible items. Amongst the iterative designs so far tried, two specific timing rules have been examined:

1. Continuous Auction (Banks et al. (1989)). A timer is started and bids can be submitted in real time. The best bid combination that fits within the logistic constraints of the auction is posted. New bids can be tendered at any time and can either be placed on a standby list to be “combined” with others or used immediately to directly replace tentatively winning bids. The standby list is a place where participants can put non-winning bids to signal to others a

willingness to combine to outbid a large package bid. The auction ends if no allocation changing bids are submitted during a fixed time period.

2. Multi-round (Kwasnika et al. (2003), CRA (1998)). A round begins and sealed bids are submitted. An integer program (IP) is solved to find the highest valued combination of winning bids. The winners are posted and a new round is started. New sealed bids can be submitted and the IP is run again to see if a new solution, with some new winners, is found. When there are no new winners or no new bids the auction is over.

Another issue in the design of a combinatorial auction involves providing bidders the opportunity to solve the threshold problem. The most obvious approach is to use a sealed bid auction with Vickrey type pricing, in which it is theoretically in the interest of the participant to reveal his true values. That is, the incentives (your bid does not determine what you pay, that is determined by the bids of others) provided by a Vickrey auction overcomes the threshold problem. In particular, Isaac and James (2000) use the standard Vickrey auction which requires the running of an IP for each winner to determine their price. Rassenti et al. (1982) eliminate the need to run individual IPs for each winner and instead use pseudo competitive prices by running a specialized Linear Program to get prices for each item.

Since the iterative and Vickrey sealed bid approaches are contrary, natural evolution led to an auction design based on combining an iterative process with some form of revelation-inducing pricing. This design change was spurred by the fact that in the above-mentioned iterative processes, the value of the information as a basis for determining where and how much to bid in the next round was unknown since participants still had to solve an IP problem based upon all previous bids to determine what it would take to combine with others to replace one or more current winning bids. Without prices to guide bidders, the iterative process can become extremely burdensome and, moreover, may make it difficult for bidders to overcome the threshold problem.

Banks et al. (1989) implemented a computationally intensive iterative process where in each round Vickrey prices are computed for each item and are delivered to each bidder along with a set of bids that, if submitted in the next round of the auction, would displace the bids submitted by one or more current high bidders assuming that no other bidders change their bids. Following the pseudo-competitive pricing of Rassenti et al. (1982), Kwasnika et al. (2003) use a procedure to calculate prices as feedback information to bidders. They find that provision of this price information allows bidders to more easily identify competitive bids and, in so doing, generate greater auction revenue and a more efficient assignment of items.

While more efficient, these new designs still require solving multiple IP problems each round and, thus, are still subject to computational complexity and scaling up concerns. To solve this problem, Rothkof et al. (1998) would simply limit the type and number of bids that can be submitted. Banks et al. (1989) would eliminate the computational

burden of the auctioneer through decentralization: requiring the bidders to execute any computations to determine what to bid on and how much to bid.

More recently, in response to an FCC RFP for a new combinatorial auction design, CRA (1998) proposed and we tested a hybrid auction system that combines multi-round and continuous bidding periods. It also implements various activity rules that mimic those in the original SMR auction along with some new variable and more complicated activity rules.

The systematic results so far derived from the body of laboratory-controlled studies of the above mechanisms are:

1. There is little evidence that the threshold problem is of great concern.
2. Iterative processes seem to be of assistance to bidders.
3. Price feedback information provides for bidder transparency.⁵
4. Computational burden must be taken seriously.

Relying on these results, its own experience with previous real auctions, and consulting advice, the FCC, for its upcoming 700 MHz auction, has designed a new multi-round process that limits the number of package bids that each bidder can submit⁶ and the prices at which he can submit them, provides computationally intensive feedback prices similar to the pricing approach in Kwasnika et al. (2003), and includes a typical cluster of activity rules. These proposed FCC auction rules for the 700 MHz auction have not been tested in the laboratory.

The Combinatorial Clock Auction

During February through October of 1999, while we were testing versions of the FCC's SMR and the CRA proposal, we had the chance to carefully observe hundreds of auction participants' behavior and the outcomes their strategies generated. Taking seriously the problem of computational burden associated with generating prices in moderately sized combinatorial auctions, yet not willing to impose activity constraints that interfere with bidding transparency by limiting types and numbers of bid revisions, we developed a simple new auction design: the Combinatorial Clock (CC) auction. Our objective was efficiency in achieving all gains from exchange, task simplicity *for the bidders*, efficacy in handling complexity in the allocation problem, and computational feasibility.

In experiments conducted in 1988 McCabe et al. demonstrated that occasionally simple ascending English auctions can become inefficient because of overstated ("jump") bids intended to forestall or signal competition. In the laboratory, we have found that the threat of financial exposure aggravates this behavior during an SMR auction for multiple items. The McCabe et al. (1988) solution is a simple upward ticking clock controlled by

⁵ By transparency, we mean that bidders can easily determine what bid amount would be the current winner of a set of items.

⁶ 12 items and 12 package bids.

the auctioneer to remove active bidder price control. The CC, following this lesson, created a unique price for each item that is controlled by that item's clock.⁷ The price ticks upward only when there is excess demand for an item.

The process we devised is simpler and quite transparent. The clocks, one for each item, are started at low prices. Each round bidders are given a fixed amount of time to submit which packages or individual items they would like to purchase at the current clock prices. A simple algorithm then counts up the demand for each item by each bidder, making sure not to double count licenses that appear in multiple bids. The item demands are then aggregated across participants. For items that have more than one bidder demanding more units than are available, the clock price is raised. A new round is started and new bids are requested.

The CC is transparent because no bidder has to submit a price, and she is free to indicate the contents of any packages that she will buy at the given clock prices. Her previous bids for packages at previous prices are remembered and eligible for consideration by the auctioneer unless they are purposefully withdrawn.

The auction continues as long as there is excess demand for one of the items being offered. If the auction reaches the point where there is exactly one bid left for each unit of each item available, then it ends and the standing bidders are awarded the items at the current clock prices: no computation is required!

But it can happen that after a particular clock price increases, the demand for that item becomes less than is available. In the case where there is excess supply for at least one item, and demand exactly equals supply for every other item available, then the auctioneer must compute the solution to an IP to find the allocation of items that would maximize his revenue. He includes all bids at current and previous clock prices that have not been withdrawn. The solution will tend to use old bids to allocate units for which there is excess supply at the current clock prices. If the overall solution does not seek to displace those holding the standing winning bids on items where supply equals demand, the auction is over and bidders pay the prices bid. If the solution does seek to displace at least one bidder holding the standing winning bid on an item where supply equals demand, then that item is considered to have excess demand, the clock price is ticked upward, and the auction continues.

Therefore, the final allocation is one in which all standing bids win items at the final clock prices, and items with excess supply at the final clock prices are awarded at previous clock prices. The auction process that generates this allocation is simply a greedy algorithm to discover pseudo-dual upper bound prices⁸: the lowest prices (final clock prices) at which everyone who submitted a bid is definitely declared a winner. In

⁷ If there is more than one unit of an item available then the clock price for that item applies identically to all the units available.

⁸ See Rassenti, Smith and Bulfin (1982) for a thorough description of pseudo-dual price computation.

complex environments these prices are often not unique, and this allows the auction mechanism some flexibility in achieving an efficient outcome⁹.

The beauty of the CC auction is in its simplicity to the participant, and its minimal computational requirements of the auctioneer. It trivially accommodates the sale of multiple units of multiple items. Bidders have complete freedom to move in and out of the auction bidding, on any packages at will. It allows the bidder to impose logical constraints without increasing the computational burden during the auction. For example, a bidder may submit mutually exclusive bids and ‘if and only if’ bids: the auction simply computes his demand for an item as the maximum number of units he could possibly win. The bidder is also free to blend current and previous clock prices in a current compound bid as long as part of her bid is at current clock prices.¹⁰ The CC auction is the most flexible combinatorial auction we know of. Strategic behavior is further controlled by feeding back only that information bidders need to know (item prices) in order to avoid bidding more than their maximum willingness to pay. For this purpose bidders do not need to know who is bidding, how many are bidding, and on which items or packages.

Auction Tests

After having conducted a series of experiments that tested versions of the CRA combinatorial auction and the FCC’s SMR auction, in the fall of 1999 we continued by independently testing our own invention: the CC auction. All of the test auctions offered only one unit of each of ten items for sale. Ten bidders each had different values for various packages of items. The value distributions that we created used variations on the following conditions to test the efficacy of the auction designs:

Condition 1: The Joint Value Factor. This measures the relative difference between the value of the optimal allocation (V^*),¹¹ which was constructed to include several smaller bidder packages, and the next highest value allocation (\underline{V}),¹² which was constructed to include a single bidder’s package covering the optimal set of smaller packages. We define the Joint Value Factor as the ratio $J = (\underline{V}/V^*)$. The larger the value of J , the greater is the need for bidders to coordinate their bids in order to “defeat” the bidder with the highest package bid. This coordination is made difficult by the incentive bidders have to “free-ride” off the bid increases of other bidders. But note that as J approaches 1, the relative loss in efficiency from having the big bidder win approaches 0. J directly

⁹ For example, if a winning bid is for a package containing one unit of item C and one of item D, then the bidder isn’t concerned with whether the winning prices are 40 and 60 or 60 and 40 respectively, as long as they eliminate his competition and have a total cost less than he is willing to pay.

¹⁰ For example a bid on a package containing items C and D at current clock prices in round t can be resubmitted as a package in round $t+1$ by bidding on one item at the previous (*pre*) price as long as at least one item is bid at the current new price: for example, $\{C, D\}_t$ can become $\{preC, D\}_{t+1}$.

¹¹ V^* is the value of the maximization problem $\max \sum v_j \cdot x_j$ subject to $x_j \in \{0,1\}$ and $\sum x_j \cdot y_j \leq Y$, where v_j is the value for package j , y_j is the vector of elements of package j and Y is the vector of total supply.

¹² \underline{V} is the value of the $\max \{v_{1Y}, \dots, v_{nY}\}$ where v_{iY} is participant i ’s value for the entire supply Y .

measures the best possible market efficiency that can result when the auction fails to solve the threshold problem.

Condition 2: The *Own Effect*. This effect is one that is coupled with the join factor. It occurs when b is the large package bidder who demands $\frac{V}{2}$, but b is also one of the small package bidders included in V^* . To achieve the optimal allocation b must forego his large package to be included in the optimal allocation of smaller winning packages. Efficiency may be hurt because b may not collaborate in his role as a small package bidder for two reasons: first, he may feel he is in a stronger negotiating position (since he owns the large package) and may demand more of the surplus than other small package bidders; or second, he may think that displacing himself, even if it is apparently profitable, may create unpredictable dynamics in the subsequent bidding.

Tables 1 and 2 show optimal and second best allocations for the various boundary experiments we conducted to sell ten licenses (A through J) with different values to different bidders. Complete value parameters for these experiments can be found at http://ices3.gmu.edu/FCC_Parameters .

Table 1: Case1 Values generating Join Factors and Own Effects

Optimal Allocation of 10 Licenses											
A	B	C	D	E	F	G	H	I	J	\$ Value	Bidder ID
◆	◆		◆							100	1
		◆						◆		80	2
				◆			◆			80	3
					◆	◆				120	4 or 6
									◆	50	5
2 nd Best Allocation of 10 Licenses											
◆	◆	◆	◆	◆	◆	◆	◆	◆	◆	350 or 301	6

In Case 1 (Table 1) bidder 6 has value for the super-package and is either included in the optimal allocation or not depending on whether he **owns** the package (F,G). The Join is either **high** .81 (350/430) or **low** .70 (301/430), depending on bidder 6’s assigned value for the super-package.

Table 2: Cases 2a and 2b, Values generating Join Factors and Own Effects

Optimal Allocation of 10 Licenses											
A	B	C	D	E	F	G	H	I	J	\$ Value	Bidder ID
◆										17.37	4
	◆									36.27	5
		◆	◆	◆						88.59	1
					◆					24.00	2
						◆				30.00	4
							◆			36.00	3
								◆		48.00	1
									◆	54.00	5 or 6
2 nd Best Allocations of 10 Licenses											
					◆	◆	◆	◆	◆	180 or 153	5
◆	◆	◆	◆	◆						114	1

In Case 2 (Table 2) the licenses were divided into two separable groups of five, 2a and 2b, with separate join factors and own effects within each group and with no bidders valuing packages containing licenses in both groups. For licenses A-E several bidders had values for the entire license group from A to E, while some of them had values for packages that were part of the optimal allocation. In the second group of licenses F-J, each bidder had values for single licenses except bidder 5 who had value for the entire group of licenses F-J. We divided the separable outcomes of the two groups into Cases 2a and 2b respectively.

In total 55 test auctions were executed. Student participants, bidding anonymously through a local area network of computers, earned significant cash payments for their bidding successes. Table 3 shows the allocation efficiencies for all of the CRA, SMR and CC auctions.

Table 3: CCC vs CRA vs SMR

Case	Join	Own	Auction	% Allocation Efficiency
1	.81	Yes	CC	100,100,100
			CRA	78, 79, 78
			SMR	59
1	.81	No	CC	100,100,100
			CRA	97, 79
			SMR	63
1	.70	Yes	CC	100,100,100,100
			CRA	100, 100
			SMR	70
2a	.80	Yes	CC	100,100,99,100,99,100
			CRA	99, 99, 99, 95, 94, 95, 95
			SMR	100, 99, 95, 95
2b	.94	Yes	CC	100,100,100
			CRA	91, 94, 94
			SMR	100
2b	.94	No	CC	100,100,100
			CRA	95, 95
			SMR	100
2b	.80	Yes	CC	100,100,100
			CRA	100, 91
			SMR	100

When comparing the outcomes of the non-combinatorial SMR (used until now by the FCC) against the outcomes of the combinatorial auction originally solicited by the FCC from CRA we find:

Result 1: As the join factor increases, efficiency falls. The own effect also reduces efficiency. Even though efficiencies are low for the CRA auction in Case 1, they are higher than for the SMR auction. In Case 2b in which bidders only have values for single license packages, except for one bidder who has a value for all the licenses, the SMR auction outperforms the CRA auction.

But if we compare the more transparent and simpler Combinatorial Clock auction use we find:

Result 2: The combinatorial clock auction was uniformly more efficient than both the other auction mechanisms across all the test environments. In addition, it yielded 100% efficiency in all but two of the test auctions.

Conclusion

The consistency of the CC test auction outcomes is remarkable given the typically somewhat noisy results of laboratory experiments with human subjects. As simple as it is, the CC auction is extremely promising, and merits continuing study in a variety of other auction environments.

References

- Banks, J., Ledyard, J., Porter, D., 1989. Allocating uncertain and unresponsive resources: an experimental approach. *Rand Journal of Economics* 20, 1–25.
- Bykowsky M., Cull, R., Ledyard, J., 2000. Mutually destructive bidding: the Federal Communications Commission auction design problem. *Journal of Regulatory Economics* 17(3), 205-228.
- Charles River and Associates Inc. and Market Design Inc., 1998. Report 1A: Auction design enhancements for non-combinatorial auctions. Charles River and Associates No. 1351-00.
- Charles River and Associates Inc. and Market Design Inc., 1998. Report 1B: Package bidding for spectrum licenses. Charles River and Associates No. 1351-00.
- Charles River and Associates Inc. and Market Design Inc., 1998. Report 2: Simultaneous ascending auctions with package bidding. Charles River and Associates No. 1351-00, March 1998.
- Isaac, R., James, D. 2000. Robustness of the incentive compatible combinatorial auction. *Experimental Economics* 3(1): 31-53.
- Kwasinka, T., Ledyard, J., Porter, D., DeMartini, C., 2003 A new and improved design for multi-objective iterative auctions. *Management Science*.
- Ledyard, J., Olson, M., Porter, D., Swanson, J., Torma, D., 2000. The Sears logistics auction. *Inform*s.
- McCabe, K., Rassenti, S., Smith, V., 1988. Testing Vickrey's and other simultaneous multiple unit versions of the English auction, revised 1989. In: R. M. Isaac (Ed). *Research in Experimental Economics*, Vol. 4, 1991. Greenwich, CT: JAI Press.
- Porter, D. 1999. An experimental examination of bid withdrawal in a multi-object

auction. *Review of Economic Design*.

Rassenti, S., Smith, V., Bulfin, R., 1982. A combinatorial auction mechanism for airport time slot allocation. *Bell Journal of Economics*, 13, 402–417.

Rothkopf, M., Peke, A., Harstad, R. 1998. Computationally manageable combinatorial auctions. *Management Science*, (44) no. 8.