

The Continuous Combinatorial Auction Architecture[†]

By CHARLES. R. PLOTT, HSING-YANG LEE, AND TRAVIS MARON*

The paper reviews the prominent features of a continuous, combinatorial auction. A brief discussion of background material helps draw distinctions between auction architectures. Background definitions, rules, and operational features define the auction in Section II. Testing and performance are outlined in Section III.

I. Background

The use of auctions to solve allocation problems among items exhibiting complementarities was first addressed in 1978 by Grether, Isaac, and Plott (1989) who studied the allocation of airport landing rights to airlines that operated at multiple airports and placed value on routes as opposed to operations at a single airport. Rassenti, Smith, and Bulfin (1982), motivated by the Grether, Isaac, and Plott (1989) research, connected bids with “either/or” and “and” constraints thus combining bids into a package. The ideas were generalized by Banks, Ledyard, and Porter (1989) to include the concept of a “standby queue.” These auctions operated as sealed bid auctions or as a series of sealed bid auctions that operated in discrete rounds.

The first example of a continuous combinatorial auction is found in Brewer and Plott (1996), who represented the allocation problem in terms of binary conflicts of packages that created both the flexibility for widespread application and the computational speed required

to support a continuous auction. They applied the mechanism to the allocation of the rights to operate trains on a fixed set of tracks in Sweden. A package was a function in time and space describing the scheduled location of a train. The intersection of two packages represented a scheduled crash between trains that must be avoided, hence the term “binary” in the paper title. The computer could quickly compute non-intersecting packages that maximized the value of the sale and permitted the auction to proceed as a type of continuous, simultaneous, ascending price auction. The non-intersecting packages that produced the most revenue from the auction were declared the leading bids at each instant of time. That first mechanism was followed by slight generalization to a procurement problem in which the buyer organized sellers to minimize procurement cost and sellers could offer endogenous packages of backhaul services (Brewer and Plott 2002). The organization was a simultaneous, decreasing price auction.

In the 1990s, the FCC was considering the adoption of a combinatorial auction as a replacement for the simultaneous, rounds-based, ascending price auction that the FCC had used to auction parts of the electromagnetic spectrum (Plott and Salmon 2004). One option studied was a hybrid process that consisted of rounds followed by a continuous phase (Plott 2000). Experiments with the hybrid revealed that most of the adjustment and efficiency came from the continuous phase, a discovery that led to the renewed study of combinatorial auctions that operated only in continuous time.

II. Definitions and Rules

Items.—The set of n items is denoted by $Y = \{1, 2, \dots, n\}$; Let $S \in \{0, 1\}^n$ be a combination, set, or package of items.

Individuals.—The set of m individuals is $M = \{1, 2, \dots, m\}$.

*Plott: Division of Humanities and Social Sciences, California Institute of Technology, 1200 E. California Blvd., Pasadena, CA 91125 (e-mail: cplott@hss.caltech.edu); Lee: Division of Humanities and Social Sciences, California Institute of Technology, 1200 E California Blvd., Pasadena, CA 91125 (e-mail: hyl@hss.caltech.edu); Maron: Division of Humanities and Social Sciences, California Institute of Technology, 1200 E. California Blvd., Pasadena, CA 91125 (e-mail: maron@caltech.edu). We thank Eiichiro Kazumori, Nicola Lansdell, Bill Stevenson, and Arch Roberts, Jr.

[†] Go to <http://dx.doi.org/10.1257/aer.104.5.452> to visit the article page for additional materials and author disclosure statement(s).

Bids.—A bid is a price and a package of items of the form $b_j^i(S)$ where i is the index of the individual submitting the bid, j is the bid number as recorded in the system, and $S \subseteq Y$ is a package of items. So, $b_q^i(S_q)$ is bid number q , where i is the bidder and $S_q \in \{0, 1\}^n$. That is, the q th bid was placed by i , for a dollar amount b_q^i for a package of items S_q .

Bid Properties.—Bids are submitted under “all or none” conditions. Either the entire package is accepted as a provisional winner or none of it is accepted. Multiple bids can be tendered within “either/or” constraints. All bids remain in the system and can be selected as provisional winners unless cancelled. Provisional winning bids cannot be cancelled.

Provisional Winners.—When the auction ends the provisional winners become winners and pay the amount of the bid. After each bid is submitted, the system publishes the set of provisionally winning bids.

Provisional Winner Determination.—The provisional winners are bids in the set of bids that would maximize the value of the sale if the auction concluded at that moment subject to the condition that no item is contained in more than one provisionally winning bid.

B = all bids submitted and not cancelled. $x_q \in \{0, 1\}$ indicates whether or not the q th bid was accepted as a provisional winner of the auction.

W = Provisionally winning bids. Provisional winners are the subset of all bids, B , that maximize the value of the sale subject to the fact that no item is contained in more than one provisional winning bid. That is, provisional winning bids are $q: x_q = \arg \max R$.

$$R = \max_{x_q: q \in B} \sum_{x_q b_q^i} \text{ Subject to } \sum_{x_q b_q^i} \leq (1, 1, 1, \dots, 1)$$

Non-Provisional Winning Bids.—The bids that are not provisional winners remain in the system and play an important role. The computation of the provisional winners includes an examination of all bids in the system. Non-winning bids exist in the system as potential partners or as the pieces of a complex coalition that can be assembled to replace large package bids as provisional

winners. Depending on the timing of the placement of a non-winning bid, the bidder is revealing a maximum willingness to pay.

Increment Requirements.—The function of the increment rule is to encourage bids to move the system to an equilibrium and to do so at a fast pace. Let $v(S)$ be the maximum value for which the set S could be sold given the bids in the system. It is determined by computing the winning bids from all bids submitted given that the sale of items was restricted to the set S . Let k be the (constant) increment required for bidding on a single item. For a bid $b(S)$ to meet the increment, it must meet the condition $b(S) \geq v(S) + k |S|$.

Stopping Rules.—Two clocks are used: a new bid clock and a new provisional winner clock. The new bid clock resets with each new bid and starts a countdown. The new winner clock resets each time a bid is placed that determines a new pattern of provisional winners. The auction ends if either clock reaches zero. Basically, the new bid clock forces a flow of bids, similar to offers in a negotiation and does so under the threat of the auction ending. The new winner clock forces concessions of sufficient magnitude to advance the value of the sale. In essence, the new bid clock says “You must make an offer” and the new winner clock says “You must make an offer sufficiently to get a deal done or the auction ends anyway.” Thus, the ultimatum feature of game theory is operational in both clocks. The system is constantly pressing for revenue gains by using the threat of ending the auction. While bidders who are not provisional winners do face a dominant strategy of bidding as the clock counts down, bids need not be large and there is no advantage to waiting until the last moment to bid. Last moment bids just give competitors more time.

Information, feedback and tools for fashioning bids are important.

Provisional Winners Table.—A table is published that contains the provisional winner of each item, whether the provisionally winning bid is a package bid or a single and the highest bid placed on the item as a single. The table is updated with each new provisionally winning bid accompanied by a small red dot that disappears in a few seconds. New non-winning bids are shown as a small black dot on the

provisionally winning table. The dots also disappear after a few seconds. These black dots signal that a bidder that wants part of a package has not bid enough to become a provisional winner and seeks partners to bid on the rest of the package of which the bidder wants a portion. The black dots carry information that serves to coordinate bidding by coalitions of small bidders who want to break up a large package bid.

All Bids Table.—A page of all bids in the system is published.

Query.—The query system and related functionality serve to replace the role of item prices in the fashioning of bids. When a bidder selects a package for a potential bid, the bidder is immediately shown both the minimum amount that can be bid as dictated by the increment requirements and the minimum amount it would take for the bid to become a provisional winner. By adding or removing items from a package, the query can be used to determine the marginal cost of adding items to a package. By removing items from a package, the bid can “fit” better with other bids and the price required to become a provisional winner can be significantly reduced.

Show as Winning.—After a bid is fashioned but before it is submitted, the bidder can choose this option to display the pattern of winning and non-winning bids that will be the consequence of the submission. It will show all new provisional winners, all bids that were provisional winners and remain as such, all bids that were provisional winners and now are not, and all bids that were not provisional winners and would be provisional winners if the bid was submitted. This supports coordination by allowing bidders to search more efficiently for partners and avoid adding items to a package that would be too costly.

III. Performance

Mechanism testing used experimental economics testbed procedures in which the experimenter has control of the number of items, the number of bidders, and the preferences of individual bidders that dictate the nature of synergies, nonconvexities, and the coordination required for efficiency. Since individual preferences are known to the experimenter, the efficiency of an allocation resulting from the

auction can be computed as the sum of values of items to the acquiring bidders divided by the maximum possible.

Testbed methods operate under conditions in which theory is suggestive but limited. Because the theory is incomplete and is vulnerable to a variety of sources of rejection, two experimental questions are posed in a testbed. (i) Does the mechanism do what it is supposed to do? The question asks for a demonstration of proof of principle. (ii) Does the mechanism do it for understandable reasons? The question asks about a test of design consistency. The question asks if the result reflects the (limited) theory that was used in the design or are they simply random? Clearly, this is a basic question because it asks about the possibility that the design will scale. Of course, this second question is a key. It asks about the robustness of the theory when applied to possibly unknown conditions. It calls for tests under a variety of environments that could challenge the performance.

The test environments might look nothing at all like “the real world” because the real world in which one might imagine the mechanism being deployed might not have conditions that theory suggests are stressful. On the other hand, testing in environments that might closely resemble the application environment might prove valuable in uncovering interactions with institutions and aspects of the environment that might not be anticipated by theory. Institutional facts and environmental features can interact in surprising ways and have negative effects on performance. Examples of both types of environments are reported in the second section.

Four classes of parameters existed for stress tests of the mechanism. The easiest test involved no synergies. A second test involved twenty items to be allocated to five bidders. Each bidder had preferences for all items but received only four items in the efficient allocation. A third environment added three bidders to the second environment with the resulting efficient allocation assigning parts of each of the five to accommodate additional bidders. This third environment is viewed as a complex fitting problem. A fourth environment expanded scale by replicating the smaller environments.

Experiments revealed that understanding induced preferences and familiarity with the functionality of the auction process were major

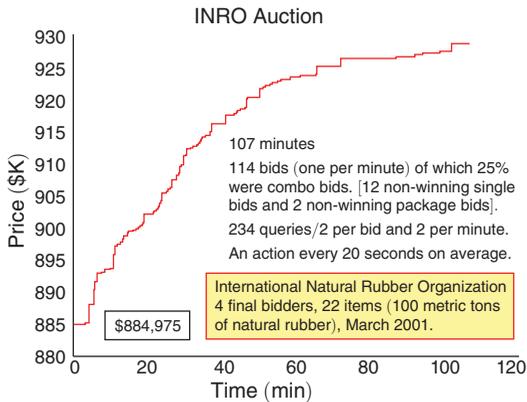


FIGURE 1

contributors to any observed lack of efficiency. Trained subjects operated at near 100 percent efficiency and scale presented no problems.

The time series of total revenues of two field applications are reported in Figures 1 and 2. The results of an auction for 100 metric ton pallets of natural rubber are contained in Figure 1. Four Internet bidders located around the world competed for 22 pallets of natural rubber located in a warehouse in Vietnam. The auction was conducted by the United Nations International Natural Rubber Organization, which had accumulated the pallets as part of a price stabilization program and was prepared to release the natural rubber back to private companies. Buyer identities were not public information. The pallets were from different plantations. Natural rubber from a given plantation is a homogeneous product, but rubber from different plantations has different and well known qualities. Starting bids were tendered by bidders as sealed bids and are similar to other sealed bids that the administrator for the INRO auction had observed were approximately market prices that exist in public markets. Package bidding followed quickly after the initial bids. Some bidders expressed values for rubber from a limited set of plantations and others seemed to be interested in a mix with some sensitivity to price and quantity. The black dots appeared throughout the auction, signaling a bid on part of an existing larger package bid. If the starting revenue of \$884,975 is assumed to be the revenue that would have been produced by a sealed bid when compared with the \$927,000 auction revenue, the combinatorial auction produced about 5.5 percent more revenue.

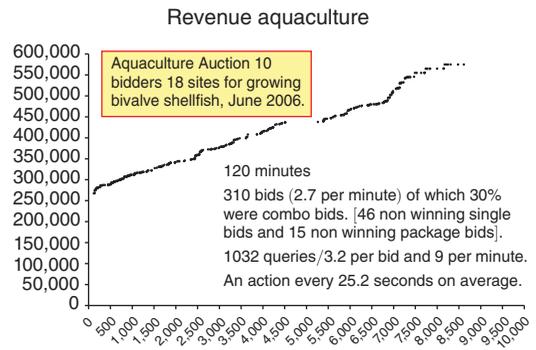


FIGURE 2

The second field application is an auction for aquaculture sites located in Port Phillip Bay near Melbourne, Australia. The sites are appropriate for the growing of bivalve shellfish. The state of Victoria decided to auction 18 sites. A total of ten bidders participated and bid for 18 sites. Seven bidders were winners producing \$575,000 in revenue. The sites were scattered across six locations. Bidders were interested in scale since they must meet regular demand for deliveries. They are also interested in a portfolio of sites reflecting a diversity of location due to currents, winds, possible diseases, and location relative to home base and delivery points. Thus, multiple synergies existed and package bids were used frequently.

The two field applications have similar properties. Both lasted about two hours. Bidding in the aquaculture auction was about three times the speed of bidding in the natural rubber auctions. There was also a larger percentage of non-winning bids (21 percent versus 13 percent) suggesting more aggressive attempts to coordinate. The larger number of bidders and the facts that the bidders were better trained and all at one location could account for the difference. The proportion of combination bids is about 25 percent–30 percent in both auctions so the package bidding was a valuable tool. The numbers of queries per bid are 2 per bid for the natural rubber auction and 3.2 per bid in the aquaculture auction suggesting a harder “fitting” or “coordination” problem existed in the aquaculture auction. The pace of bidding starts fast in both auctions and slows dramatically at the end. This property is evident in both figures.

IV. Observations

Several features of the mechanism are worth emphasis. The reader should notice that concepts typical of rounds bids such as eligibility, activity, and package prices based on sums of items prices are all absent. The absence of a concept of a price per item is a departure from tradition. Replacing the measures contained in prices are queries and displays that can respond to human pattern recognition and crafted information needs. The use of clocks carries key public information and creates the proper level of incentives for coordination. A bidder needs only to meet an increment requirement to keep a negotiation alive before facing an all or none choice of implementing a “contribution” to the public good of breaking up a large bid or collections of bids and becoming a provisional winner. The dots provide feedback by calling attention to actions of others and the possible intentions that underlay the actions of others play. This type of information that contributes to coalition formation plays a key role and renewal of the “free rider” problem. Computational problems can clearly pose problems as the size of the auction grows. The testbed methods have some departures from what an untutored theorist might expect. The methods are designed to address problems where the theory is not complete and might be no more than suggestive. Classical statistical tests of such theories do not make much sense if research is confronted by a scale of limited budget, limited time, and an unbounded infinity of variables. Yet, the role of theory plays a fundamental role. Theory, regardless of how incomplete it might be, is the tool that takes the analysis from the limited observations under controlled conditions to the substantially unknown conditions of the field.

The theory must be robust and the testbeds help with the necessary judgments.

REFERENCES

- Banks, J. S., J. O. Ledyard, and D. P. Porter.** 1989. “Allocating Uncertain and Unresponsive Resources: An Experimental Approach.” *RAND Journal of Economics* 20 (1): 1–25.
- Brewer, Paul J., and Charles R. Plott.** 1996. “A Binary Conflict Ascending Price (BICAP) Mechanism for the Decentralized Allocation of the Right to Use Railroad Tracks.” *International Journal of Industrial Organization* 14: 857–86.
- Brewer, Paul J., and Charles R. Plott.** 2002. “A Decentralized, Smart Market Solution to a Class of Backhaul Transport Problems: Concept and Experimental Testbeds.” *Interfaces* 32 (5): 13–36.
- Grether, David M., R. Mark Isaac, and Charles R. Plott.** 1989. *The Allocation of Scarce Resources: Experimental Economics and the Problem of Allocating Airport Slots*. Boulder, CO: Westview Press, 1989.
- Plott, Charles R.** 2000. “A Combinatorial Auction Designed for the Federal Communications Commission.” Presented at the Federal Communications Commission Bidding Conference Why River, May 2000.
- Plott, Charles R., and Timothy C. Salmon.** 2004. “The Simultaneous, Ascending Auction: Dynamics of Price Adjustment in Experiments and in the U.K. 3G Spectrum Auction.” *Journal of Economic Behavior and Organization* 53 (3): 353–83.
- Rassenti, S. J., V. L. Smith, and R. L. Bulfin.** 1982. “A Combinatorial Auction Mechanism for Airport Time Slot Allocation.” *Bell Journal of Economics* 13 (2): 402–17.