

Combinatorial Auctions
Edited by Peter Cramton, Yoav Shoham, and Richard Steinberg
Cambridge and London: MIT Press, 2006

Review for the *Journal of Economic Literature*
by Michael Ostrovsky, Stanford University

This book is a collection of essays on combinatorial auctions—allocation problems in which one agent has a number of objects for sale and several other agents are interested in buying these objects. What makes such allocation problems particularly interesting and challenging is the possibility that the buyers have complex preferences over packages of objects: the value of a package to a potential buyer is not necessarily equal to the sum of the values of individual items; the buyer may view some items as substitutes for one another and some items as complements. A famous example of such problem is the allocation of licenses for electromagnetic spectrum. A license typically covers a specific geographic region and a specific range of frequencies, giving the buyer a right to use that range of frequencies in that region. Some buyers may only be interested in licenses for a particular region and at the same time may not care what specific frequency range they get; others may want to assemble a package of licenses covering several regions, placing an extra premium on packages of licenses with the same frequency range, and at the same time may have very low values for individual licenses in the package. Another classic example is the problem of allocating takeoff and landing rights in congested airports: an airline will not place a high value on a takeoff right at the airport of origin if it does not have a corresponding landing right at the airport of destination.

Designing a good combinatorial auction is hard; sometimes, it is not even exactly clear what makes such an auction “good”, since the seller may care about several objectives: expected revenue; allocative efficiency; simplicity, transparency, and robustness of the rules; the speed of execution; and so on. At the same time, this problem has an immediate and obvious practical importance. This combination has attracted the attention of researchers in several academic disciplines, most notably economics and game theory, computer science, and operations research. Each group has looked at combinatorial auctions in a particular light, each has focused on a particular set of issues, and each has made progress in recent years. At the same time, they often neglected the issues raised by the researches in other fields: e.g., computer scientists and operations researchers frequently assumed “straightforward” bidding by the buyers, ignoring the possibility of strategic behavior, while economists mostly assumed that as long as the valuations of bidders were reported truthfully, or all package prices were known, optimal allocations could be found quickly, ignoring the potential complexity of such computations.

This book successfully brings together the different traditions in the combinatorial auctions literature and makes both the basic results and the latest advances in these fields very accessible. It opens with Vernon Smith's brief account of early experiments on combinatorial auctions and the editors' introduction to basic auction theory, and then proceeds to the 23 contributed essays, divided into five parts. Part I describes various auction mechanisms proposed and used to sell packages of objects, such as the Vickrey auction, simultaneous ascending auctions, clock and proxy auctions, and several other dynamic procedures. Part II studies the efficiency properties of some of these mechanisms and discusses one of the potential bottlenecks for achieving efficiency in combinatorial auctions: the complexity of expressing and communicating the buyers' bids. In an auction with m distinct items, a fully general bid would have to consist of $2^m - 1$ numbers; one for each subset of the full set of items. Communicating such general bids would be infeasible even with a couple dozen items. Part III discusses another potential concern: the complexity of computing an optimal allocation given a set of bids, i.e., the winner determination problem. Even if the number of bids is small, selecting the subset of bids that maximizes the total value of an allocation or the revenue of the seller may be computationally hard. The authors of this part's essays present various ways of mitigating this complexity in practice. Part IV discusses implementation details for one of the proposed mechanisms and also describes a software suite that allows researchers to generate sets of bidder valuations resembling those that are likely to occur in specific real-world applications. In Part V, the focus of the book shifts to applications. Chapter 20 by Michael Ball et al. introduces a particular allocation problem in need of a well-designed mechanism—the problem of assigning takeoff and landing rights at congested airports. Chapters 21–23 by Chris Caplice and Yossi Sheffi, Estelle Cantillon and Martin Pesendorfer, and Martin Bichler et al. describe three types of “reverse” combinatorial auctions that are already used in practice: industrial procurement auctions; trucking services auctions; and London bus routes auctions, where private operators compete for the right to carry bus services. In these auctions, the roles of agents are flipped: a buyer needs to procure several goods or services and solicits bids from multiple sellers. (Chapter 4 by Peter Cramton could also be placed in Part V of the book—it discusses simultaneous ascending auctions and their use by the FCC to sell electromagnetic spectrum.)

Notably, the same main message emerges from the papers in all three traditions: There is no single, “magic” combinatorial auction that would work well for all markets and in all situations. For economists, a natural candidate for such universal auction is the Vickrey-Clarke-Groves mechanism, in which bidders report values for all packages and the auctioneer then chooses the allocation that maximizes the total surplus and charges each bidder the amount by which that bidder's presence reduces the surplus of others. In this mechanism, it is a dominant strategy for each bidder to report his value truthfully (and VCG is essentially a unique general mechanism with this property). However, as Lawrence Ausubel and Paul Milgrom point out in Chapter 1, in the presence of

complementarities VCG suffers from a variety of serious problems: e.g., the seller's revenues may be extremely low and may decrease as the number of bidders or the amounts of their bids increase, losing bidders may collude and change the outcome of the auction in a way that is profitable to them, and, in fact, a single losing bidder may increase his payoff by shill bidding through multiple identities (this point is also made by Makoto Yokoo in Chapter 7). And even if bidders were non-strategic, there would be no hope of finding a universal mechanism: Chapter 11 by Ilya Segal shows that for general valuations of bidders, exponential amount of communication is required to find an efficient (or even approximately efficient) allocation and Chapter 12 by Daniel Lehmann, Rudolf Müller, and Tuomas Sandholm explains that even if each bidder is allowed to submit a small number of bids, with a large number of bidders the problem of finding the surplus-maximizing allocation is computationally infeasible.

At the same time, by taking into account the special features of a particular setting, it is often possible to design a well-functioning mechanism, a compact and reasonably accurate representation of preferences, and a fast winner determination algorithm. In the simultaneous ascending auctions used by the FCC to allocate spectrum (Chapter 4), the bidders are not allowed to submit package bids, despite the fact that some of them may view some of the items as complements: they have to assemble these packages by participating in several auctions for individual items that run in parallel. Nevertheless, the fact that the FCC continues to use such auctions suggests that this is a reasonably good (though, of course, not necessarily ideal) mechanism. In procurement auctions (Chapters 21—23), a very different mechanism—a combinatorial analogue of the first-price sealed bid auction—is often used. The buyer solicits quotes from the sellers, who tell him how much they will charge for various packages of goods or services, and then simply chooses the combination of bids that he finds optimal, paying the suppliers he chose the amounts they requested. Again, the continued use of such auctions suggests that the auctioneers find them acceptable, are able to overcome the communicational and computational complexity, and do not suffer substantially from the fact that the mechanism is not incentive-compatible and hence the suppliers are forced to strategize. For some special classes of valuations, it is also possible to construct mechanisms that are approximately efficient, incentive compatible, and computationally fast (see, e.g., Section 15.7 by Amir Ronen).

Given the non-existence of a universal mechanism, the success of some auctions in specific practical applications, and the large number of mechanisms proposed in Part I, it is somewhat disappointing that the essays in the book provide little guidance as to which of the mechanisms are likely to work best in which settings. Much of our intuition about efficient and optimal auction design comes from the literature on single-object auctions, and is thus of limited help when we need to choose from many combinatorial auctions that all become very similar when only one object is for sale. As several chapters show, some intuition also comes from the literature on two-sided matching markets, but there

most of the positive results come from settings with substitutable goods, whereas in combinatorial auctions most interesting and challenging issues arise precisely because of complementarities. Can we rely on the intuition developed in those models for the analysis of combinatorial auctions with complementarities? Which of the conclusions are robust? Which are not? The essays in the book do not provide many answers to these questions.

The absence of answers is most likely due to the complexity of characterizing equilibrium strategies and outcomes in fully general models of combinatorial auctions—very little is known about them. In fact, the book does not present any theorems describing the equilibria of any of the proposed general combinatorial auctions (with the exception of VCG-based mechanisms) in environments with complementary goods, incomplete information, and fully strategic bidders and contains only two explicitly worked out examples of equilibrium behavior in such environments (an example for the proxy auction and an example for the pay-as-bid auction). At the same time, the authors of the essays emphasize the potential importance of a variety of issues that are hard to analyze formally without characterizing the equilibria: interdependencies in valuations, costly valuation of packages and acquisition of information, optimal bundling of objects for sale, etc.

There are several potential approaches to dealing with the complexity of combinatorial auctions. Some special classes of valuations, corresponding to particular real-world settings, may be tractable analytically. In other cases, experimental studies may be helpful. Finally, equilibrium strategies and outcomes can be obtained computationally; this approach, while lacking the elegance of an analytic solution, may be just as useful in studying the issues mentioned above and in helping auctioneers pick the most appropriate auction format. Future progress in these areas will likely depend on successful collaboration between economists, computer scientists, operations researchers, and practitioners. By bringing their work together and by establishing the common language, the book makes this collaboration easier. Anyone with an interest in auction theory, market design, and more broadly, practical applications of game theory, will find this book extremely valuable both as a reference to the existing work on combinatorial auctions and as a source of topics and ideas for new research.