More Expressive Market Models and the Future of Combinatorial Auctions

DANIEL GOLOVIN
Carnegie Mellon University

Research on auctions and other engineered markets is undergoing a renaissance, spurred by the growth of computing power and communication networks, and the new opportunities for social and economic interaction that they enable. Previous research has been applied to good effect, perhaps most famously by search engines to allocate advertising space in keyword auctions. Other case studies of combinatorial auctions being used in practice can be found in the book of Cramton *et al.* [Cramton et al. 2006]. Despite these success stories, there is much more to do. In this letter, I will give my opinion as to the ultimate determinants of how successful research on engineered markets can be, with implications for where the community can focus attention to maximize its impact. I will use [Golovin 2007] to illustrate initial progress in the direction I have in mind.

Categories and Subject Descriptors: K.4.4 [**Electronic Commerce**]: Payment schemes

General Terms: Economics, Algorithms

Additional Key Words and Phrases: Combinatorial Auctions, Mechanism Design, Engineered Markets

Recall that in a *combinatorial auction* buyers bid on subsets of goods offered up for sale by a single seller, instead of bidding on individual items that are sold one by one. Why are combinatorial auctions useful? One common answer is that combinatorial auctions allow the bidders to express complimentarities and/or substitution effects among items. Often, these may be exploited by a centralized market mechanism to increase economic efficiency over what decentralized market mechanisms and single item auctions would give.

The key quality here is *expressiveness*. By more accurately modeling the desires of the market participants and giving participants the means to express more nuanced preferences, a clever mechanism may do a better job in optimizing the outcome. Unfortunately, there is a price to be paid for greater expressiveness; the mechanisms themselves must be more intricate, and running them can become much harder computationally (e.g., as expressiveness grows, the winner determination problem may go from being polynomial-time solvable to being **NP**-hard, or it may transition to being easy to approximate well to being **NP**-hard to approxi-

Email: dgolovin@cs.cmu.edu. Author's address: D. Golovin, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Supported in part by NSF ITR grants CCR-0122581 (The Aladdin Center) and IIS-0121678.

mate well). Ultimately, the success of engineered markets depends on how well our mechanisms address the following questions.

(1) How faithfully does the mechanism's model of the world conform to reality?
(2) How effective is the mechanism at obtaining accurate (or at least useful) information from the various market participants?
(3) How effective is the mechanism at finding "good" market outcomes in its own model of the world, for various notions of "good"?

There are various tradeoffs here. As stated, as the model becomes more realistic finding good outcomes in it may become computationally intractable. Also, a more realistic model may put heavier communication burdens on market participants, and may even rule out the use of some models if such burdens become too onerous. Finally, obtaining truthful preference information from the market participants may require a loss of efficiency in some cases.

Given these tradeoffs, where should we focus our research efforts? I would argue in favor of increasing the fidelity of the market models, and doing our best to address the resulting increased difficulty in dealing with questions (2) and (3). Limiting the accuracy of the model places fundamental limitations on how well a market mechanism can do in reality. Increasing the accuracy of the model may lead to limitations on how well we can collect information and optimize market outcomes. In contrast to the first case, however, these limitations are due to our incomplete understanding of mechanism design as well as results from complexity theory. The former difficulty might very well be overcome with more research. The latter may be surmountable in practice, with clever heuristics and ever increasing computational power. Thus, while it remains to be seen where the "sweet-spot" is with respect to these tradeoffs, it is worth focusing on more accurate models.

The following are concrete examples of how we might make our model more accurate and increase the participants' expressiveness.

(1) Treat repeated auctions as a repeated game. For example, if a firm runs a procurement auction every year, and many of the same bidders are likely to participate year after year, it is likely advantageous to consider the sequence of auctions as a whole, rather than as independent auctions.
(2) Incorporate notions of reputation and trust, as well as the costs borne by the market for agents' defaulting on their obligations. This is particularly important for combinatorial exchanges.
(3) Incorporate probabilistic events. Demand and supply might be significantly affected by the weather, stock market fluctuations, the tides of war, election outcomes, and so on. In some cases a centralized mechanism with some probabilistic information about the future and how this affects the market participants could achieve greater efficiency than current approaches.

As an example of how one might incorporate probabilistic events, in [Golovin 2007] a protocol for auctions with probabilistic demand and supply is presented. For simplicity, assume that there is some set of items to be auctioned off, and the supply is deterministic and known in advance to the auctioneer. Participants are allowed to express probabilistic demands, and submit bids of the form (S_i, b_i, p_i), where S_i

is a set of items, b_i is the value of the bid, and p_i is the probability that this bid will correspond to an actual demand. Once these bids are collected, the auctioneer decides on a subset of bids to *notify*. Notified bids are eligible to participate in the auction. All other bids are discarded. Each notified bid (S_i, b_i, p_i) then *appears* with probability p_i, by which we mean their corresponding demand actually materializes. The auctioneer must then solve a winner determination problem on the bids that appear, and determine the payments each participant must make. Participants whose bids were not notified, or whose bids were notified but did not appear, receive nothing and pay nothing. Additionally, we imagine that participants with bids that were notified and appeared but who won nothing suffer some disappointment, and we pay such participants a *compensation cost* to offset this disappointment.

The advantage of this protocol is that bidders can express probabilistic demands to the auctioneer without committing to a bid, and the auctioneer does not commit to supplying all appearing bids, but can "buy back" excess demand at auction time. By cleverly selecting which bids to notify, the auctioneer can achieve greater social welfare than a traditional combinatorial auction, especially if there are many bids with low p_i values. The challenge is to decide which bids to notify, while being mindful of how decisions at this stage affect the standard mechanism design problem. In [Golovin 2007], we first consider the problem of designing an approximation algorithm for social welfare in such a system, and obtain an $O(k)$ approximation, where k is the maximum cardinality of any set of items bid upon. This matches up to constant factors the current best approximation guarantee for Maximum k-Set Packing, which our problem generalizes. We then give a truthful in expectation mechanism for the problem that achieves an $O(k)$ approximation to the best expected social welfare, assuming the compensation costs exactly offset the disappointment of bidders that appeared but did not win anything at auction. (If the probabilities are off by as much as a multiplicative factor of α, the approximation guarantee becomes $O(\alpha k)$.) This represents some modest initial progress in the quest to incorporate probabilistic supply and demand into market mechanisms. By itself this is a large research agenda, and the far broader problem of designing mechanisms for richer market models is sure to generate many interesting problems and applications for many years to come.

REFERENCES

Cramton, P., Shoham, Y., and Steinberg, R., Eds. 2006. *Combinatorial Auctions.* MIT Press.

Golovin, D. 2007. Stochastic packing-market planning. In *EC '07: Proceedings of the 8th ACM conference on Electronic commerce.* ACM Press, New York, NY, USA, 172–181.