Reducing Costly Information Acquisition in Auctions

Extended Abstract

Kate Larson
Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada

Abstract

Most auction research assumes that potential bidders have private information about their willingness to pay for an item. In reality, bidders often have to go through a costly information-gathering process in order to learn their valuations. Recent attempts at modelling this phenomena has brought to light complex strategic behaviour arising from information-gathering, and has shown that traditional approaches to auction and mechanism design are not able to overcome it.

1 Introduction

Most research on auctions assumes that the problem facing bidders is how to bid given their private preferences for the item(s) being auctioned. In reality, however, bidders often have to go through a costly information-gathering process in order to determine what their actual preferences are. This process may involve such things as asking questions about a product in order to learn whether it is of high or low quality, solving optimization problems in order to ensure that the bidder bids on the minimum amount of materials needed to complete some job, or conducting research on a firm in order to learn whether its assets complement the assets currently held by the potential bidder. One model that has been studied allows a bidder the choice between participating in the auction without knowing its true valuations or paying a fee to learn them. Questions asked using this model include what sort of incentives are required for bidders to acquire information about their valuations [1], and how does information acquisition depend on the rules of the auction [3, 12, 13].

It has been noted that a bidder's decision as to whether to compute or gather information about it's valuations can depend on the preferences of others [15]. We propose explicitly modeling the information-gathering actions of bidders along with the decisions they make when deciding how to use their information-gathering resources. We place this *deliberative-bidder* model into a game-theoretic framework, and analyze auctions in order to gain an understanding of the impact that computational and information-gathering constraints have on bidders' strategic behaviour [8]. This brings to light new forms of strategizing on the part of bidders and provides a possible explanation for behaviour seen in practice [14].

We also study the problem of designing auctions specifically for bidders who must gather information in order to determine their actual valuations and formulate their bids [10, 7].

Dagstuhl Seminar Proceedings 07271
Computational Social Systems and the Internet
http://drops.dagstuhl.de/opus/volltexte/2007/1159

We argue that auctions designed for deliberative bidders should exhibit, along with standard economic properties, desirable deliberative properties.

2 Deliberative Bidders

In this section we describe our model of a *deliberative bidder* and explain how we analyze the strategic behaviour of such bidders. We assume that the reader has a basic understanding of elementary game theory and auction theory. While we do explain key concepts as they are needed, a more thorough treatment on this topic can be found in microeconomic texts or books specializing on auctions [11, 6]. We also focus our attention to single-item, private-value auctions. However, the models we present can be extended to multi-item settings in the obvious way.

A deliberative bidder is one who must compute or gather information in order to determine how much it values the item(s) being auctioned. We are interested in settings where the bidders have restrictions on their computing or information-gathering capabilities, and who must carefully consider how to use their available resources given these limitations.

We assume that a deliberative bidder has a set of deliberation resources, which we will simply refer to as the bidder's resources. We denote the resources of bidder i by R_i. A bidder is able to use its resources to determine how much value it places on winning a particular item in an auction. We also allow a bidder to use its resources to learn how much value *another* bidder places on winning an item. We let $r = (r_1, \ldots, r_m) \in R_i^m$ denote the situation where bidder i has devoted r_j resources to scenario j where a scenario is a situation where bidder j is allocated the item in the auction.

We model deliberation-resource limitations through cost functions. The cost function of bidder i is $\mathrm{cost}_i : R_i^m \rightarrow \mathbb{R}^+$. The only restrictions placed on the cost functions are that they must be additive and non-decreasing.

One of the challenges faced by a deliberative bidder is how to allocate its resources given its cost function. To this end, we assume that a deliberative bidder is endowed with a set of *models*, Ω_j, one for each scenario j. These models describe how resource investment (*i.e.* computing or gathering information) changes the information available to the bidder. They allow the bidder to make decisions as to where to invest its resources optimally, given its cost constraints, and provide a way of predicting how future investment will change the bidder's information. While in this extended abstract we do not provide additional details on the models, one particularly useful model is the *performance profile tree* from the resource-bounded reasoning literature [9].

To summarize, a deliberative agent is defined by

$$\langle R_i, \mathrm{cost}_i(\cdot), \{\Omega_j\}\rangle$$

where R_i is the set of deliberation resources of agent i, $\mathrm{cost}_i : R_i^m \rightarrow \mathbb{R}$ is a cost function which limits the amount of resources the agent can use, and $\{\Omega_j\}$ is the set of models.

A deliberative bidder needs to decide how and whether to gather information and then how to use the information to bid. Figure 1 illustrates this process. Motivating the bidders is the desire to maximize their utility. In particular, the utility of a deliberative bidder depends on its valuation of winning the auction, whether or not it has won the auction,

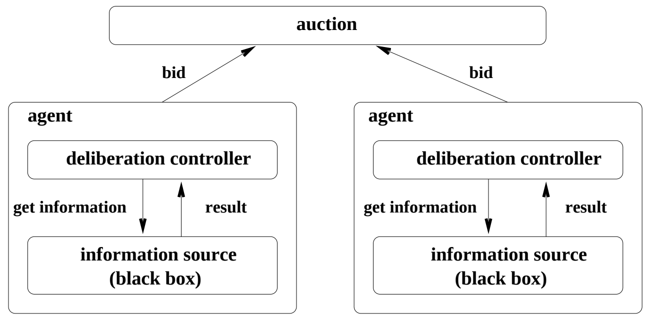

Figure 1: An auction with two deliberative bidders. Each bidder must gather information in order to determine how much it values the item that is being auctioned. A bidder may also have incentive to gather information on other participants in the auction.

and the price that it pays if it was the winner. That is, if a bidder i has used resources $r = (r_1, \ldots, r_m)$, has determined that its valuation is $v_i(r_i)$, has won the item, and has to pay price p, then its utility would be

$$u_i = v_i(r_i) - p - \mathrm{cost}_i(r).$$

If the agent did not win the auction then its utility would be

$$u_i = -\mathrm{cost}_i(r).$$

Note that the bidder may use its deliberation resources on problems which do not directly affect its own valuation. However, this still influences its utility since the bidder incurs a cost for the total resource usage.

A strategy for a deliberative bidder is a policy which specifies what actions (deliberative and bidding) to execute at every stage in the auction. We define a *history* at stage t, $H(t) \in \mathcal{H}(t)$, as the set which includes all actions (both deliberative and bidding) that the bidder itself has taken up to stage t, the results of the deliberation actions, as well as all actions other bidders may have taken. A (deliberation) *strategy* is a mapping from the set of histories to the set of actions (deliberative or bidding) for each stage in the game. That is, $S_i = (\sigma_i^t)_{t=0}^{\infty}$ and

$$\sigma_i^t : \mathcal{H}(t) \mapsto A_i$$

where A_i is the set of all actions available to bidder i.

To clarify this definition we present a simple example. In a direct auction, bidders submit their bid directly to the auctioneer at some deadline T. The strategies of a deliberative bidder take the form; $S_i = (\sigma_i^t)_{t=0}^{\infty}$ where

$$\sigma_i^t(H_i(t)) = \begin{cases} d_i^j & \text{if } t < T \text{ or } t > T \\ \hat{v} & \text{if } t = T \text{ and } \hat{v} \in \mathbb{R} \end{cases}$$

where $\hat{v}$ is the bid submitted to the auctioneer, and d_i^j is a deliberative action.

In this new, enlarged, strategy space we look for equilibria, which we call *deliberation equilibria*. We are particularly interested in understanding how deliberation equilibria differ from equilibria when bidders are no deliberative.

3 Deliberation Equilibria in Auctions

We studied three standard auction mechanisms in order to gain an understanding of the difference in strategic behaviour between deliberative and non-deliberative bidders. We investigated a first-price sealed bid, a second-price sealed bid and an ascending auction, all in private-value settings.

In a first-price sealed bid auction, all bidders submit their bid to the auctioneer at some point in time T. The bidder with the highest bid wins the auction and pays an amount equal to what it bid. As is well known, there is no dominant strategy equilibrium for non-deliberative bidders; instead in the Bayes-Nash equilibrium the bid of one bidder depends on the other bids submitted. When we studied this auction with deliberative bidders, we observed that in equilibrium sometimes bidders are best off actively using some of their resources in order to discover the valuations of other bidders in the auction, instead of merely speculating as to what they might be. We called this behaviour *strategic deliberation*. The knowledge gained by strategic deliberation helped the bidders to form low bids so that they could win the auction.

In the second-price sealed bid auction, all bidders submit their bid to the auctioneer at some point in time T. The bidder with the highest bid wins the auction, but pays an amount equal to the second high bid. As is well known, in this auction non-deliberative bidders have a dominant strategy which is to simply submit a bid equal to their actual valuation of winning the item being auctioned. Interestingly, when we studied deliberative bidders we noticed that in equilibrium strategic deliberation can occur. That is, bidders are sometimes best off using some of their resources to gather information about other bidders, even though when it comes to bidding they can ignore the other bids and simply reveal their valuation. We observed that in the second-price sealed bid auction with deliberative bidders, it was important for a bidder to learn its likelihood of winning the auction. For example, if the bidder learned that it was unlikely to actually win the auction, then there was little incentive to invest in learning its own valuation. The result is that in the second-price sealed bid auction, deliberative bidders do not have dominant strategies, since their deliberation actions depend on other bidders' deliberation results.

Finally, we studied an ascending auction in which non-deliberative bidders have dominant strategies.[1] Again, if the bidders were deliberative, strategic deliberation occurred in equilibrium, even though the auction mechanism itself provided a lot of information. This was not enough to remove the incentive for bidders to discover, by strategically deliberating, whether they were likely to win the auction or not. Table 1 summarizes our findings.

4 Reducing Strategic Deliberation

We believe that strategic deliberation is a undesirable phenomenon. When bidders strategically deliberate they use their own costly information-gathering resources on valuation

[1] The version of ascending auction we analyzed was one where the auctioneer sets the price at zero and gradually raises it. All bidders observe the current price and indicate whether to buy at the price. Once a bidder drops out of the auction (by indicating that it is no longer willing to buy at the current price) it is not allowed to re-enter. An important feature of this auction is that all bidders know what the current price is, and all bidders know who the remaining active bidders are.

Auction Type	Equilibrium (Non-deliberative bidders)	Equilibrium (Deliberative bidders)
First-price sealed bid	Bayes-Nash	Strategic deliberation
Second-price sealed bid	Dominant	Strategic deliberation
Ascending	Dominant	Strategic deliberation

Table 1: Strategic deliberation is prevalent in many types of standard auctions.

problems that do not directly relate to their own preferences, thus wasting these resources. Additionally, the strategic burden placed on the bidders is potentially overwhelming. They must take into consideration the bidding actions of competitors, the deliberation actions of competitors, and the results of their competitors' deliberation actions, in order to determine how to best deliberate and bid for themselves.

The issue with the three auctions that we discussed in the previous section was that the auction mechanisms were all blind to the fact that the bidders were deliberative. We hypothesize that if the auctioneer has some additional information about the bidders' deliberative processes then it might be possible to design an auction that reduces or eliminates strategic deliberation. To this end we propose allowing the auctioneer to know the models and cost functions of the bidders participating in the auction. The key insight is that if the auction designer has information about the bidders' cost functions and models, then it can use this information to sort and classify the bidders. The auction can then process the bidders in a particular order, getting each one in turn to gather information and bid until some criterion has been met. Such an approach has been proposed by Burguet [2] and Cremer *et al* [4] for the problem of auctioning of an item when the *auctioneer* incurs a cost from contacting prospective bidders. The auction of Cremer *et al* can be adapted to our setting where the bidders are incurring the cost, instead of the auctioneer. By using an optimal search procedure from the operations research literature [16], coupled with carefully designed reservation prices, it is possible to create an auction where the most promising bidders (*i.e.* the ones most likely to have a high valuation and thus win the auction) are asked to participate early in the process, and, secondly, where bidders are provided with enough information that they no longer have incentive to gather information about others [7].

While the use of an optimal-search auction avoids strategic deliberation, it is not clear that it is a reasonable solution for auctions for deliberative bidders. While it removes the strategic burden from the bidders, it places a high computational demand on the auctioneer. The auctioneer must interpret the models of the bidders appropriately, and is required to compute an index, based on the model and cost function, for each bidder. Since these indices are closely related to Gittins indices [5], the practical feasibility of such an auction is questionable in general.

5 Summary

A common assumption in auction research is that bidders have private information about their willingness to pay for the item being auctioned and that they use this information strategically when formulating their bids. In reality, bidders often have to undergo a costly information-gathering process in order to determine their valuations for the item being auctioned. Recent attempts at modelling this phenomenon has brought to light complex strategic behaviour that had previously been overlooked.

In recent work we showed that by providing the auction designer with some information about the information-gathering processes of the bidders, it is possible to design an auction such that the bidders only gather information on their own valuation problems, and, when asked, report the results truthfully to the auctioneer. The insight is that by carefully ordering the agents using an optimal search procedure, and then asking them sequentially to gather information and reveal their results, it is possible to remove the uncertainty that drives the complex strategic behaviour of deliberative agents in standard auctions.

Acknowledgements

Some of this work was joint with Tuomas Sandholm.

References

[1] D. Bergemann and J. Välimäki. Information acquisition and efficient mechanism design. *Econometrica*, 70:1007–1034, 2002.

[2] R. Burguet. Optimal repeated purchases when sellers are learning about costs. *Journal of Economic Theory*, 68:440–455, 1996.

[3] O. Compte and P. Jehiel. Auctions and information acquisition: Sealed-bid or dynamic formats? Working paper, May 2004.

[4] J. Cremer, Y. Spiegel, and C. Zheng. Optimal search auctions. *Journal of Economic Theory*, 134(1):226–248, 2007.

[5] J. C. Gittins. *Multi-Armed Bandit Allocation Indices,*. Wiley, 1989.

[6] V. Krishna. *Auction Theory.* Academic Press, 2002.

[7] K. Larson. Reducing costly information acquisition in auctions. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1167–1174, Hakodate, Japan, May 2006.

[8] K. Larson and T. Sandholm. Costly valuation computation in auctions. In *Theoretical Aspects of Rationality and Knowledge (TARK VIII)*, pages 169–182, Siena, Italy, July 2001.

[9] K. Larson and T. Sandholm. Using performance profile trees to improve deliberation control. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 73–79, San Jose, CA, USA, July 2004.

[10] K. Larson and T. Sandholm. Mechanism design for deliberative agents. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Utrecht, The Netherlands, 2005.

[11] A. Mas-Colell, M. Whinston, and J. R. Green. *Microeconomic Theory.* Oxford University Press, 1995.

[12] D. Parkes. Auction design with costly preference elicitation. *Annals of Mathematics and AI*, 44:269–302, 2004.

[13] N. Perisco. Information acquisition in auctions. *Econometrica*, 68(1):135–148, January 2000.

[14] M. Rothkopf. On misusing auctions to value stranded assets. *The Electricity Journal*, 10(10):10–17, December 1997.

[15] T. Sandholm. Issues in computational Vickrey auctions. *International Journal of Electronic Commerce*, 4(3):107–129, 2000. Special Issue on Applying Intelligent Agents for Electronic Commerce.

[16] M. Weitzman. Optimal search for the best alternative. *Econometrica*, 47(3):641–654, 1979.