Strategic Deliberation and Truthful Revelation: An Impossibility Result

Kate Larson
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
klarson@cs.cmu.edu

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
sandholm@cs.cmu.edu

ABSTRACT

In many market settings, agents do not know their preferences *a priori*. Instead, they may have to solve computationally complex optimization problems, query databases, or perform expensive searches in order to determine their values for different outcomes. For such settings, we have introduced the *deliberation equilibrium* as the game-theoretic solution concept where the agents' deliberation actions are modeled as part of their strategies.

In this paper we lay out auction design principles for deliberative agents. We propose a set of intuitive properties which are desirable in auctions for deliberative agents. First, we propose that auctions should be *non-deliberative*: the auction should not actively do the deliberation for the agents. Second, auctions should be *deliberation-proof*: in equilibrium agents should not have an incentive to deliberate on each others' valuation problems. Third, the auction should be *non-deceiving*: agents should not have incentive to strategically misrepresent. We show that it is impossible to design interesting auctions which have these three properties.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Economics Algorithms Theory

Keywords

Mechanism design, Bounded rationality

1. INTRODUCTION

Game theory, and mechanism design in particular, have long been successfully used in economics and have recently drawn a lot of research interest from computer scientists (e.g., [3] [4]). In most of this work it is assumed that the participants, or agents, know their preferences and the goal of the mechanism (for example, an auction) is to extract this information to a sufficient extent, and select an outcome such that desirable properties are achieved. However, there are many settings where agents do not know their preferences *a priori*. Instead they may, for example, have to solve computationally complex optimization problems, query databases, or perform complicated searches in order to determine the worth of an outcome. We call the actions taken to determine preferences *deliberation*.

While the commonly used auction mechanisms from the economics literature have been designed to have desirable economic properties, they ignore computational and deliberation issues. In this paper we ask the question "Is it possible to design auction mechanisms which have both desirable economic and deliberative properties?"

2. DELIBERATIVE AGENTS

In order to participate in an auction, agents need to know their valuations for the item(s) up for auction. In this paper we do not assume that agents know their valuations *a priori*. Instead, we assume that each agent has anytime algorithms that it can use to determine valuations, has cost functions which limit deliberation, and can do performance-profile based deliberation control. Anytime algorithms are algorithms which have the desirable property that they can be stopped at any point in time and return a result, and if allocated additional resources, will return a better solution. These algorithms allow agents to make tradeoffs between the quality of the solution (for example, the value an agent has for an item) and the cost of acquiring the solution. Many algorithms have this anytime property, including iterative refinement algorithms, which always return a solution and improve upon it if allowed to run longer. Similarly, many search and information gathering applications can be viewed as anytime algorithms. As more information about an item is gathered, an agent's value for the item is refined.

While anytime algorithms allow agents to make tradeoffs between solution (valuation) quality and cost of obtaining the solution, they do not provide a complete answer since they do not specify how the tradeoff should be made. Instead, agents use *performance profiles* to help make this decision. Performance profiles describe how deliberation changes the solution returned by the algorithm.[1] Agents use the performance profiles to determine deliberation policies. The actual deliberation actions that agents take depend on the policies, the partial results of deliberation and also the actions taken by other agents.

[1]Performance profile based deliberation control has been well studied by AI researchers. See for example, [1] for an overview of this research area.

Copyright is held by the author/owner.
EC'04, May 17–20, 2004, New York, New York, USA.
ACM 1-58113-711-0/04/0005.

We have proposed explicitly including the deliberation actions of agents into their strategies, and then analyzing games for *deliberation equilibria* which are fixed points in the space of strategies in this enlarged strategy space [1]. Using this approach we have studied common auction mechanisms and found, in all of them, the existence of interesting strategic behavior. In equilibrium, bidders have an incentive to use some of their deliberation resources in order to partially determine *other agents' valuations* for the item(s) being auctioned. We coined this phenomenon *strategic deliberation*.

3. MECHANISM DESIGN FOR DELIBERATIVE AGENTS

Mechanism design has been of particular use when it has come to designing auctions for ecommerce settings. The mechanism design problem is to implement a set of "rules" so that an optimal system-wide solution is chosen, despite agents acting in their own self-interest.[2] A mechanism describes the sets of permitted strategies for the agents (for example, the bidding rules in an auction), and an outcome rule which is a function of the strategies played by the agents (*e.g.*, the allocation and payments in an auction). A mechanism implements a social choice function (where a social choice function maps agents' values to outcomes) if there is an equilibrium of the game induced by the mechanism which results in the same outcome as the social choice function.

While the tools available from the mechanism design literature are very powerful, they do not always give adequate solutions for deliberative agents. For example, the Revelation Principle states that if a social choice function can be implemented in dominant strategies, then it can be implemented by a mechanism where agents truthfully reveal their values to the mechanism [2]. While it is possible to construct a Revelation Principle for deliberative agents, by having agents submit deliberative tools to the mechanism and having the mechanism do all deliberation for all agents, this is not a reasonable solution. First, it is unreasonable to assume that agents are able to supply the mechanism with the algorithms, performance profiles, cost functions and problem instances. Second, there is no reason to assume that the mechanism would have enough deliberative resources of its own to solve all agents' deliberation problems even if the agents could reveal all relevant information.

We believe that mechanisms for deliberative agents should have the following, desirable, deliberative properties.

PROPERTY 1 (NON-DELIBERATIVE). *A mechanism should not solve the agents' individual deliberation problems.*

If a mechanism is non-deliberative, agents are responsible for solving their own deliberation problems.

PROPERTY 2 (DELIBERATION-PROOF). *Strategic deliberation should not occur in equilibrium.*

If agents strategically deliberate then they use some of their own deliberation resources to determine the values of their competitors, leading to complex, non-dominant equilibrium strategies. Finally, we believe that mechanisms should have incentives for agents to truthfully reveal their (partial) deliberated values.

PROPERTY 3 (NON-DECEIVING). *In equilibrium no agent should have incentive to lie about the results of their deliberation.*

Our first result states that it is possible to design auctions such that these three properties are satisfied.

THEOREM 1. *There exist mechanisms which are non-deliberative, deliberation-proof and non-deceiving.*

At first glance, this seems like a positive result. However, the mechanisms which are non-deliberative, deliberation-proof, and non-deceiving are often *not sensitive* to the agents' strategies in that the outcome of the mechanism does not depend on the strategies that the agents play. For example, random auctions, where the item is allocated to an agent at random, independent of the bids, are not sensitive. We claim that these are not particularly desirable auctions.

Our next theorem looks at what is possible in the space of *sensitive* mechanisms.

THEOREM 2. *There exists no sensitive non-deliberative mechanism (direct or multi-stage) that is both deliberation-proof and non-deceiving across all instances. (Where an instance is defined by agents' algorithms, performance profiles, cost functions).*

This result is negative in that it states that we must have auctions which either do some deliberation for the agents, or where agents strategically deliberate on competitors problems. However, there is some hope. It may be possible to weaken one of the properties slightly, while still achieving the others. For example, it may be possible to design markets where agents who can deliberate cheaply and efficiently do all deliberation for agents, and then truthfully share their results in order to get social welfare maximizing outcomes.

[2]We often want to maximize social welfare, i.e. in an auction we wish to allocate the item to the bidder who values the item the most.

4. ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under CAREER Award IRI-9703122, Grant IIS-9800994, ITR IIS-0081246, and ITR IIS-0121678.

5. REFERENCES

[1] Kate Larson and Tuomas Sandholm. Costly valuation computation in auctions. In *TARK VIII*, pages 169–182,July 2001.

[2] Andreu Mas-Colell, Michael Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[3] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.

[4] Christos Papadimitriou. Algorithms, games and the Internet. In *STOC*, pages 749–753, 2001.