Dominant-Strategy Auction Design for Agents with Uncertain, Private Values

David R.M. Thompson Kevin Leyton-Brown

University of British Columbia

Outline

Valuation Uncertainty

- A strong assumption from classical auction theory: agents know their own valuations
- Imagine going to a foreclosure auction to buy a house
  - Large purchase: you'll think carefully about your strategy
  - Can you identify a real value x, such that you'd be happy to buy the house for $x - \$0.01$, and that you'd prefer to keep the money if offered the price $x + \$0.01$?

Valuation Uncertainty

- A strong assumption from classical auction theory: agents know their own valuations
- Imagine going to a foreclosure auction to buy a house
  - Large purchase: you'll think carefully about your strategy
  - Can you identify a real value x, such that you'd be happy to buy the house for $x - \$0.01$, and that you'd prefer to keep the money if offered the price $x + \$0.01$? (I can't.)

How can we model such settings?

Deliberative agents: must pay a cost to learn about their own values. Will only pay if the expected benefits outweigh the cost.

- Thinking hard is costly


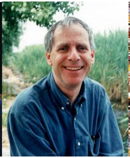
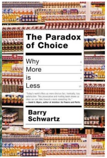

- Must solve a computational problem to determine my value

Modeling deliberative agents

Definition (general model; informal)

We require that agents have independent, private values.

Deliberative agents can nevertheless be quite complex:

- may be able to choose among a wide range of deliberations
- available deliberations may depend on the agent's current belief state
- deliberations may be noisy
- agents may be unable ever to discover their values perfectly
- agents may be able to learn about each other's valuations as well as their own
- ...

A formal model appears in our paper.

Modeling deliberative agents

Useful also to consider the simplest possible deliberative model:

Definition (Simple deliberative agent)

A simple deliberative agent i has two equally likely possible valuations (v_i^L, v_i^H). Values are independent and private. At any time, the agent can pay cost $c_i > 0$ to discover his true valuation.

Background

- Second-price auctions don't have DS [Sandholm, 2000].
- Second-price auctions give rise to a (mis-)coordination problem; don't have symmetric PSNE [Thompson & L-B, 2007].

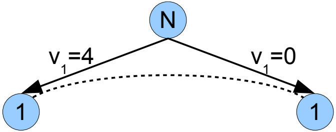

Background

- Second-price auctions don't have DS [Sandholm, 2000].
- Second-price auctions give rise to a (mis-)coordination problem; don't have symmetric PSNE [Thompson & L-B, 2007].

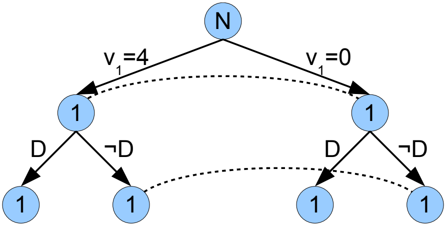

Background

- Second-price auctions don't have DS [Sandholm, 2000].
- Second-price auctions give rise to a (mis-)coordination problem; don't have symmetric PSNE [Thompson & L-B, 2007].

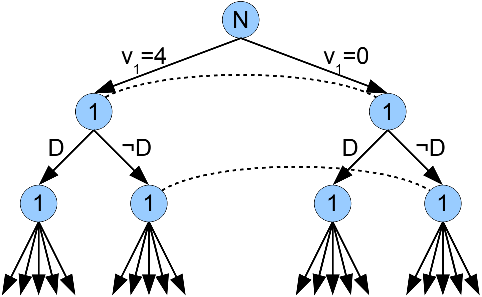

Background

- Second-price auctions don't have DS [Sandholm, 2000].
- Second-price auctions give rise to a (mis-)coordination problem; don't have symmetric PSNE [Thompson & L-B, 2007].

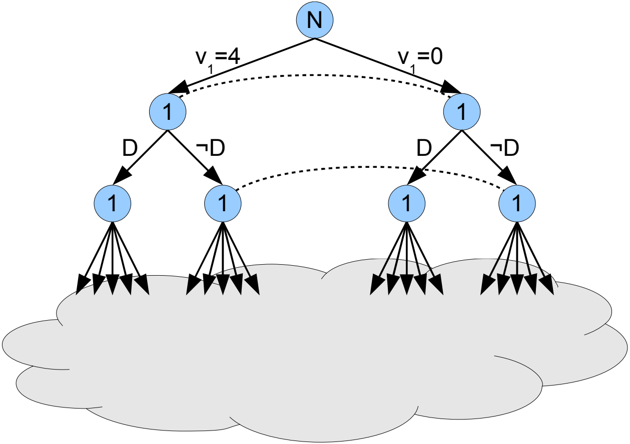

Background

- Second-price auctions don't have DS [Sandholm, 2000].
- Second-price auctions give rise to a (mis-)coordination problem; don't have symmetric PSNE [Thompson & L-B, 2007].

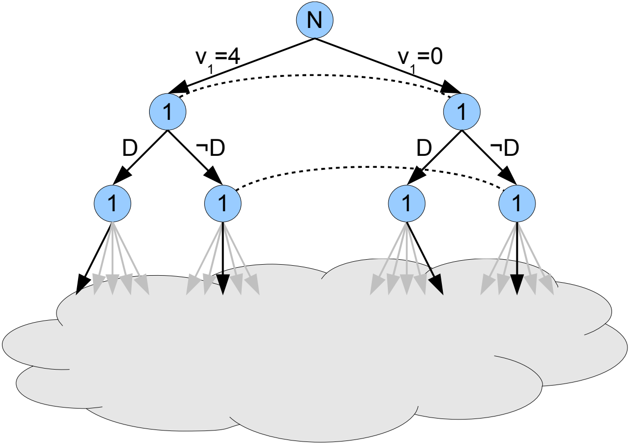

Background

- Second-price auctions don't have DS [Sandholm, 2000].
- Second-price auctions give rise to a (mis-)coordination problem; don't have symmetric PSNE [Thompson & L-B, 2007].

	D	$\neg D$
D	$1-c, 1-c$	$1-c, 1$
$\neg D$	$1, 1-c$	$0, 0$

$$D, D\text{: } \tfrac{1}{4}(4-4) + \tfrac{1}{4}(4-0) + \tfrac{1}{2}(0) = 1$$
$$D, \neg D\text{: } \tfrac{1}{2}(4-2) + \tfrac{1}{2}(0) = 1$$
$$\neg D, D\text{: } \tfrac{1}{2}(2-0) + \tfrac{1}{2}(0) = 1$$

Background

- Second-price auctions don't have DS [Sandholm, 2000].
- Second-price auctions give rise to a (mis-)coordination problem; don't have symmetric PSNE [Thompson & L-B, 2007].

	D	$\neg D$
D	$1-c, 1-c$	$1-c, 1$
$\neg D$	$1, 1-c$	$0, 0$

- Similarly, Japanese and eBay (ascending proxy) auctions don't have dominant strategies, and neither is equivalent to Second-price. [Compte & Jehiel, 2001; Rasmusen, 2006]

Mechanism design in deliberative settings

- Bayes-Nash:
  - The second-price auction is the most efficient sealed-bid auction [Bergemann & Valimaki, 2006], but is strictly worse than the Japanese auction [Compte & Jehiel, 2001].
  - The social-welfare maximizing single-good auction is known [Cavallo & Parkes, 2008].
- Dominant strategies:
  - Impossibility result for general (non-IPV) valuations [Larson & Sandholm, 2004].

Question

For deliberative agents with IPV valuations,
what (if any) single-item auctions offer dominant strategies?

Outline

Revelation Principle for Deliberative Agents?

- Revelation principle is very useful for characterizations.
- However, in our setting direct mechanisms can't simulate indirect mechanisms [Larson & Sandholm, 2004]
  - the mechanism can't deliberate for agents
  - agents' decisions about whether to deliberate may be conditional, so they can't be asked to deliberate up front.
- Larson & Sandholm's negative result proven without appeal to a revelation principle.

Revelation Principle

Definition (Dynamically direct mechanism)

Dynamically direct mechanisms ask one agent to deliberate and report the result, repeat this process an arbitrary number of times, and then choose an outcome.

Revelation Principle

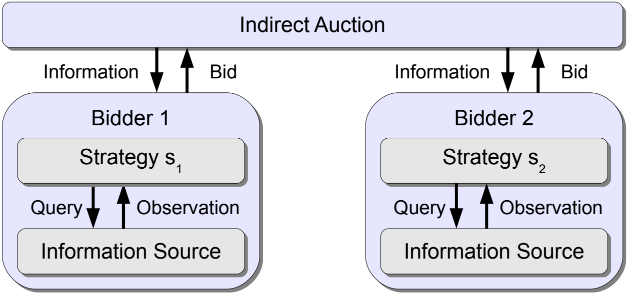

Revelation Principle

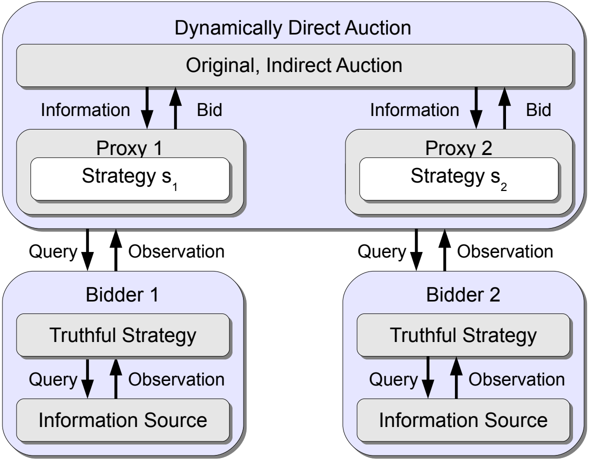

Revelation Principle

Definition (Truthful)

In a truthful strategy, the agent deliberates when asked and reports his true value.

Definition (Social choice function)

A social choice function is a (possibly randomized) mapping from (true) valuation profiles to allocations.

Theorem (Revelation principle for deliberative agents[1])

If social choice function χ is implementable in dominant strategies, then χ is implementable in truthful dominant strategies by a dynamically direct mechanism.

[1] Holds even for the general, IPV model.

Outline

Main Result

Definition (Sequential Posted-Price Auction (SPPA))

A sequential posted-price auction is a multistage auction in which at every stage, the auctioneer makes a posted-price, take-it-or-leave-it offer to a single agent. Each agent gets at most one offer.

Theorem (Characterization)

In our model, a deterministic social choice function χ is implementable in dominant strategies if and only if it is implementable by a SPPA.

Proof

- Proof for the if direction (SPPA $\implies$ dominant strategies) is straightforward:
  - Because values are IPV, the agent currently being offered a price is indifferent to everything that could happen after, and learns nothing from what happened before.
  - We obtain dominant strategies even in our general IPV model.
- The proof for the only-if direction is more complicated; I'll sketch it here. This (negative result) holds even under the simple IPV model.

Only-If Proof: Dominant Strategies $\implies$ SPPA

- Assume that we have some truthful, deterministic, and dynamically direct mechanism M that implements χ.
- Assume the simple IPV model.
- We show that χ is implementable by an SPPA.

Lemma (Information Availability)

The outcome chosen by the mechanism is completely determined by the types of the agents who deliberate.

Only-If Proof: Influence

Lemma (Influence)

An agent only deliberates when doing so always makes the difference between winning and losing (i.e., when he would always win if he reported the high type and lose otherwise).

- Otherwise, for some strategies of the other agents:
  - i always wins or always loses: would strictly prefer not to pay the deliberation cost
  - i loses with the high type and wins with the low type: violates DS truthfulness.
- Note: relies on our assumption of determinism.

Only-If Proof: Dominant Strategies $\implies$ SPPA

Consider social choice function χ's value on different inputs.

v	$\chi(v)$
H H H H H H H H	?

Only-If Proof: Dominant Strategies $\Longrightarrow$ SPPA

Consider social choice function χ's value on different inputs.

v	$\chi(v)$
H H H H H H H H	$?$

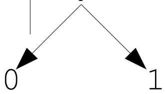

Either nobody wins, or an arbitrary agent (say, 1) wins.

Only-If Proof: Dominant Strategies $\implies$ SPPA

Consider social choice function χ's value on different inputs.

v	$\chi(v)$
H H H H H H H H	1

If nobody wins in this case, nobody ever deliberates, and so by information availability nobody ever wins. This is a trivial SPPA. Now consider the case where 1 wins.

Only-If Proof: Dominant Strategies $\implies$ SPPA

Consider social choice function χ's value on different inputs.

v								$\chi(v)$
H	H	H	H	H	H	H	H	1
H	$*$	$*$	$*$	$*$	$*$	$*$	$*$	1

By influence, χ can't depend on the valuations of any agent other than 1, because the others can't be asked to deliberate.

Only-If Proof: Dominant Strategies $\implies$ SPPA

Consider social choice function χ's value on different inputs.

v								$\chi(v)$
H	$*$	$*$	$*$	$*$	$*$	$*$	$*$	1

By information availability, the mechanism must set 1's price independently of the other agents' valuations, and by DS it must set the price independently of 1's valuation. Thus, it must be equivalent to a posted price.

Only-If Proof: Dominant Strategies $\Longrightarrow$ SPPA

Consider social choice function χ's value on different inputs.

v								$\chi(v)$
H	$*$	$*$	$*$	$*$	$*$	$*$	$*$	1
L	H	H	H	H	H	H	H	?

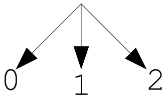

Here nobody could win, 1 could win again, or some other arbitrary agent (say, 2) could win. The first two cases are trivial.

Only-If Proof: Dominant Strategies $\Longrightarrow$ SPPA

Consider social choice function χ's value on different inputs.

v								$\chi(v)$
H	*	*	*	*	*	*	*	1
L	H	*	*	*	*	*	*	2

When 2 wins, by influence χ can't depend on the valuations of any agent > 2. By information availability, 2's payment can't depend on these values; by DS it cannot depend on 2's declaration. Thus 2 is asked to pay a posted price.

Only-If Proof: Dominant Strategies $\Longrightarrow$ SPPA

Consider social choice function χ's value on different inputs.

v								$\chi(v)$
H	$*$	$*$	$*$	$*$	$*$	$*$	$*$	1
L	H	$*$	$*$	$*$	$*$	$*$	$*$	2
L	L	H	$*$	$*$	$*$	$*$	$*$	3
L	L	L	H	$*$	$*$	$*$	$*$	4
L	L	L	L	H	$*$	$*$	$*$	5
L	L	L	L	L	H	$*$	$*$	6
L	L	L	L	L	L	H	$*$	7
L	L	L	L	L	L	L	H	8

We proceed by induction, completing the proof.

Discussion

- Our result circumvents Larson & Sandholm's impossibility result in the IPV setting
  - we satisfy all of their desiderata that make sense under IPV (dominant strategies, strategy dependence, non-misleadingness, preference formation independence; strategic-deliberation-proofness does not apply)
- Our result adds to recent arguments in favor of SPPAs
  - similar posted-price mechanisms appear often in practice
  - they have been shown to have good revenue and efficiency properties (e.g., [Blumrosen & Holenstein 2008; Chawla, Hartline, Malec & Sivan 2010; Kleinberg & Leighton 2003; Shakkottai, Srikant, Ozdaglar, & Acemogluet 2008])
  - they limit info revelation (e.g., [Sandholm & Gilpin 2006]).

Summary

- Our model: single-item auction with “deliberative” agents who must pay to learn about their own IPV valuations.
- Our contributions:
  - Revelation principle: still works, but must generalize to multi-stage (“dynamically direct”) mechanisms
  - Characterization: dominant strategies $\iff$ sequential posted-price auction
- Interesting open questions:
  - Maximizing revenue or welfare
  - More general domains (e.g., multi-unit auctions)
  - Randomized mechanisms