

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

David R.M. Thompson and Kevin Leyton-Brown
University of British Columbia
[daveth,kevinlb]@cs.ubc.ca

Abstract

We study the problem of designing auctions for agents who incur a cost if they choose to learn about their own preferences. We reformulate the revelation principle for use with such *deliberative agents*. Then we characterize the set of single-good auctions giving rise to dominant strategies for deliberative agents whose values are independent and private. Interestingly, this set of dominant-strategy mechanisms is exactly the set of sequential posted-price auctions, a class of mechanisms that has received much recent attention.

Introduction

We consider the problem of designing auctions for settings in which bidders have to pay a cost to learn about their preferences, and hence can face tradeoffs between the cost and accuracy of their preference information. Such bidders are called *deliberative agents*, and have featured in a wide variety of auction models. For example, costly deliberation can model solving a hard computational problem to discover the optimal use for a good (Cavallo and Parkes 2008), making an R & D investment to decrease the cost of fulfilling a contract in a procurement auction (Tan 1992), or even just “thinking hard” (Rasmusen 2006). There is an extensive body of work showing that when deliberative agents bid in common auctions, surprising and often undesirable outcomes can result (see e.g. (Bergemann and Valimaki 2002; Larson and Sandholm 2001; Persico 2000; Thompson and Leyton-Brown 2007)). Other work has identified new mechanisms—based on optimal search algorithms (Cremer, Spiegel, and Zheng 2003; Larson 2006) or derived from VCG (Bergemann and Valimaki 2002; Cavallo and Parkes 2008)—that do exhibit desirable behavior with deliberative agents, but only in Bayes-Nash equilibrium implementations.

It is most desirable to design auctions that give rise to dominant strategies. For example, such auctions are robust to irrational behavior by a subset of the agents, and do not require coordination to equilibrium. Such auctions also do not require agents to have common knowledge of the setting. One line of work has investigated the extension to deliberative settings of auction mechanisms that have dominant strategies in standard settings, where indeed such mechanisms are often equivalent to each other. It has been shown that second-price

auctions (Sandholm 2000), Japanese auctions (Compte and Jehiel 2001), and English auctions with proxies (Rasmusen 2006) are all inequivalent in deliberative settings, and that none of them gives rise to dominant strategies.

A more comprehensive approach to the investigation of dominant-strategy auctions in deliberative settings is to characterize the space of dominant-strategy mechanisms, rather than checking individual mechanisms to see whether they maintain dominant strategies. In this vein, Larson and Sandholm (Larson and Sandholm 2004a; 2005) examined the space of dominant-strategy implementable mechanisms under four assumptions (defined below). They obtained the following negative result.

Theorem 1 (Dominant strategy impossibility (Larson and Sandholm 2004a)). *There does not exist any mechanism that is strategic deliberation-proof, strategy-dependent, non-misleading, and preference-formation independent in dominant-strategy equilibrium across all possible quasi-linear deliberative-agent settings.*

We now formally define the four properties upon which Theorem 1 depends, since they play a role in what follows. We begin with an issue that has not yet featured in our discussion, nor indeed in most of the related work that we have discussed. In some models, agents can deliberate about others’ valuations as well as about their own; doing so is termed *strategic deliberation*. When such deliberations are possible, we might want to guarantee that in equilibrium, strategic deliberations do not occur.

Definition 2 (Strategic deliberation-proofness). *In equilibrium, no agent deliberates to learn about another agent’s preferences.*

The second property is a weakening of so-called *consumer sovereignty* (Feigenbaum et al. 2003): each agent must be able to affect the mechanism at least sometimes.

Definition 3 (Strategy dependence). *For each agent i there exists some pair of strategy profiles s', s'' that differ only in i ’s strategy and that produce different outcomes.*

Third, the following definition can be understood as a weakening of truthfulness.

Definition 4 (Non-misleadingness). *In equilibrium, no agent reports an (expected) valuation that the other agents believe is impossible. In particular, no agent reports a valuation that is greater than his highest possible valuation.*

It is worth discussing why Larson and Sandholm required non-misleadingness rather than truthfulness, which can usually be achieved without restriction via the revelation principle. In deliberative-agent settings, an agent’s knowledge about his own preferences depends on which deliberations he has performed, which can depend in turn on what information the mechanism reveals to the agent. In such settings, Larson and Sandholm argued, direct mechanisms are not without loss of generality—the agent cannot decide whether to deliberate without knowing what information will be revealed by the mechanism, which can depend in turn on other agents’ declarations—and so the standard revelation principle cannot be applied (Larson and Sandholm 2005).

Our final property, preference-formation independence, is conceptually related to “detail-freeness” (Wilson 1987).

Definition 5 (Preference-formation independence). *The mechanism is not involved in the process by which agents form their preferences and requires no knowledge of the details of that process.*

An auction that only solicits bids from agents with easy-to-compute valuations is not preference-formation independent. An agent reducing uncertainty about his own preferences is understood as a kind of “preference formation.”

In the literature, Theorem 1 has been understood as very discouraging news about the applicability of dominant-strategy mechanisms to deliberative settings. It is indeed very natural to want an auction mechanism to be strategy dependent, non-misleading and preference-formation independent. On the other hand, it is not inevitable that agents will be able to deliberate about each other’s valuations. We thus consider it important to understand the extent to which Larson and Sandholm’s negative result continues to hold when strategic deliberation is impossible. We believe this to be an open question: none of the positive results of which we are aware identifies a dominant-strategy auction, while Theorem 1 is the most relevant negative result, and its proof depends critically on the possibility of strategic deliberation.

Our paper makes two key contributions. First, we show that it is possible to recover the revelation principle in a useful sense, although it is indeed true in deliberative settings that not every multistage mechanism is strategically equivalent to some direct mechanism. Second, leveraging our first result, we characterize dominant-strategy mechanisms for the specific case of independent-private-value (IPV) auction settings in agents can deliberate only about their own valuations. Specifically, we show every dominant strategy truthful auction must implement the same social choice function as some sequential posted-price auction.

We note that our second result may be interesting even beyond the literature on deliberation, because it adds to a recent accumulation of arguments in favor of sequential posted-price auctions. Not only do similar auctions appear often in practice, but posted-price mechanisms have been shown to have good properties in terms of revenue and efficiency (e.g., (Blumrosen and Holenstein 2008; Chawla et al. 2010; Kleinberg and Leighton 2003; Shakkottai et al. 2008)) while limiting information revelation (e.g., (Sandholm and Gilpin 2006)). Our work is the first to study sequential posted-price auctions for deliberative agents, and to show that they also

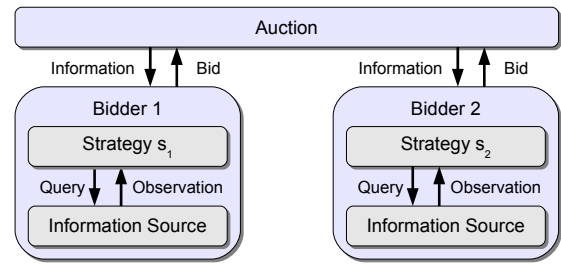


Figure 1: (Adapted from (Larson 2006)) In a deliberative setting, instead of an epistemic type, each agent has his own black-box information source. Information asymmetry arises through agents’ interactions with their information sources.

have (uniquely) good properties in such settings.

Model

We now present a formal model of deliberative-agent settings. Instead of knowing all his private information before he acts, each agent starts with no private information, but has a black-box information source that he can query at non-negative cost to receive new information.¹ Such a model is illustrated in Figure 1. Observe that such a model differs from a standard epistemic-type Bayesian game in which agents are (perfectly) aware of their own types; instead, it has more in common with the standard common-value auction model in which agents are imperfectly informed about their own valuations.

The key to a deliberative-agent setting is the availability of *deliberation* actions.

Definition 6 (Deliberation). *A deliberation is an action by which an agent queries his black-box information source and observes the response, allowing him to update his beliefs about all agents’ realized valuations.*

Throughout we assume that bidders have independent private values (IPV). Independence is the standard assumption that the joint value distribution is a product of individual value distributions. In the context of deliberative agents, *privacy* means that strategic deliberation is impossible.

Definition 7 (Privacy). *In a deliberative setting where values are private, no agent’s value depends on another agent’s signal; when an agent deliberates, he can only update his beliefs about his own value.*

Our black-box information-source model is extremely general and expressive, but has the drawback that its full description is complex and notation heavy. We present such a general model in an appendix; all of our results in this paper hold under that model with the same proofs.

In this section we present a much simpler, special case of that model. All of our results also hold (albeit with more limited implications) under the simpler model, which we call the single-step deliberative setting.

The single-step deliberative setting is like a standard IPV auction setting (e.g. from (Myerson 1981)), except that agents have to decide whether or not to deliberate. Each

¹We can handle the case where agents *do* start with private information with a zero-cost deliberation that reveals this information.

bidder starts out knowing only the common prior over his valuation. There is only a single deliberation action available, and after it is performed an agent knows his exact valuation.

Definition 8 (Single-step deliberative setting). *The single-step deliberative setting is a 4-tuple $\langle N, f, V, c \rangle$. N is a set of agents. f is a vector of valuation distributions: f_i is the commonly-known prior over i 's valuation, with domain V_i ; $V = \prod_i V_i$. c is a cost vector, where c_i is a strictly positive amount that i must pay to discover his true valuation.*

In a single-step deliberative setting, an agent's utility depends on his valuation v_i , whether or not he wins ($w_i = 1$ if he wins and 0 otherwise), his total transfers t_i , and whether or not he performed a deliberation ($d_i = 1$ if he did and 0 otherwise). We assume that utility is quasilinear and that there are no externalities: $u_i(v_i, w_i, t_i, d_i) = w_i v_i - t_i - d_i c_i$. We further assume that the entire setting is common knowledge. Finally, we make two technical assumptions: (1) the auction is deterministic, and (2) losers make no payments.

Any combination of a deliberative setting and a mechanism induces an imperfect-information extensive-form game with chance nodes. The game begins with a move by nature: each agent's valuation v_i is drawn according to distribution f_i . The agents then interact with the mechanism in a standard extensive-form game. However, any time the mechanism asks an agent to bid, he has the option of deliberating first.

Revelation principle

The revelation principle is a classical tool for understanding the space of possible mechanisms that meet given criteria. Essentially, it says that without loss of generality we can restrict our attention to direct, truthful mechanisms (Myerson 1981). This spares us the complexity of reasoning about indirect mechanisms such as multistage auctions. As described above, Larson and Sandholm showed that this argument does not work in deliberative settings: deliberation is an inherently multistage process, and considering only direct mechanisms is restrictive (Larson and Sandholm 2005).

Here we show that it is still possible to obtain a kind of revelation principle for deliberative agents. The key idea is to restrict the space of mechanisms not to direct mechanisms, but to a larger space, *dynamically-direct mechanisms*. These are multistage mechanisms that only interact with the agents by requesting that they perform specific deliberations, and asking them to report the results.

Definition 9 (Dynamically direct mechanism). *In a dynamically direct mechanism, every message from the auctioneer specifies a deliberation for the bidder to perform, and every reply must be an observation that could result from that deliberation.*

Definition 10 (Truthfulness). *A strategy is truthful if it specifies performing all requested deliberations and truthfully reporting the resulting observations. A dynamically direct mechanism is truthful if truthfulness is a dominant strategy.*

Definition 11 (Social choice function). *A social choice function is denoted by (χ, p) where χ maps from (true, potentially unknown) preferences to distributions over allocations ($\chi : V \mapsto \Delta(N \cup \emptyset)$) and p is an n -element vector, with each $p_i : V \mapsto \Delta\mathbb{R}$ serving as i 's payment function.*

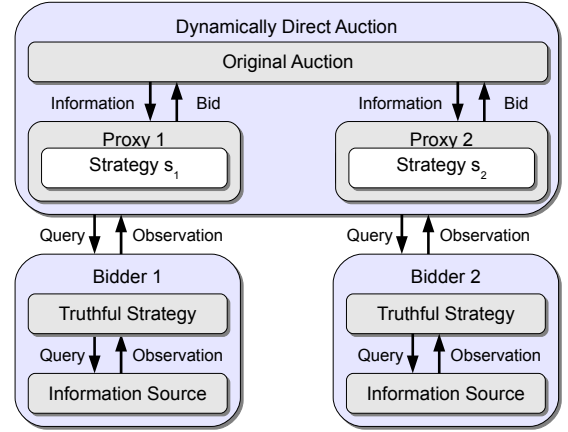


Figure 2: In a dynamically direct auction, agents act as intermediaries between their information sources and proxies.

Note that the domain of each function is V , the set of possible valuation profiles. This means that we might expect that pooling social choice functions will often arise. For example, if in equilibrium some agent i never deliberates, then the social choice function cannot depend on i 's true preferences. When deliberation can give noisy signals (as is possible in the general setting, though not in the model proposed in the previous section), then even a deterministic mechanism can implement randomized social choice functions. This explains why Definition 11 maps to distributions over outcomes, even though we only consider deterministic mechanisms.

Theorem 12 (Revelation principle for deliberative agents). *In any deliberative setting, for any mechanism M with dominant strategy equilibrium s , there exists a dynamically direct, truthful mechanism M^D that implements the same social choice function as M .*

*Proof Sketch.*² Construct a mechanism M^D where each bidder i reports to a proxy that plays s_i , but uses bidder i as the information source. (See Figure 2.) Any way of misreporting to a proxy corresponds to some deviating strategy s'_i in M . Because s_i dominates all s'_i , truthfulness must be a dominant strategy in M^D . \square

Dynamically direct mechanisms are impractical: in order to tell each agent when and how to deliberate, they must encode complete knowledge of how the agents deliberate and the associated costs. In other words, they violate Larson and Sandholm's desideratum of preference-formation independence. Nevertheless, they constitute a useful analytic tool

²We note that existing revelation principles for dynamic mechanisms use similar proof techniques; however, none can be directly applied to yield our result. Myerson (1986) gave a dynamic revelation principle for communication equilibria (a correlated-equilibrium-like solution concept for extensive-form games). However, in his model, the mechanism does not control the outcome and cannot impose payments. More recently, a similar technique has been used in dynamic mechanism design (e.g., Pavan, Segal, and Toikka 2008). However, these results assume that the mechanism chooses an allocation at each time step, and that agents' private information can change only as a result of that allocation.

for characterizing the set of social choice functions that can be implemented in deliberative settings. Observe that every preference-formation-independent mechanism also has some dynamically direct analogue; thus, using our revelation principle to characterize a space of mechanisms does not mean giving up on preference-formation independence.

Auction Design

We now turn to the question of which mechanisms give rise to dominant strategies for IPV deliberative agents. We begin by defining sequential posted-price auctions.

Definition 13 (Sequential posted-price (SPP) auction). *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. The agent can either accept the offer, in which case he immediately wins the good and pays the posted price, or reject the offer, in which case he does not interact with the auctioneer again and does not win the good.*

We can now present our main result, that dominant strategy auctions in IPV deliberative settings are characterized by sequential posted-price auctions.

Theorem 14 (Characterization). *In deliberative auction settings with independent, private values, the set of auctions with dominant strategies is characterized by the set of sequential posted-price auctions, as follows:*

1. *Every SPP auction has a dominant strategy equilibrium in every IPV deliberative-agent setting;*
2. *There exist IPV deliberative-agent settings in which the only dominant-strategy implementable social choice functions are those implemented by SPP auctions.*

Remark. Intuitively, the fact that all SPP auctions offer dominant strategies follows from the IPV auction model: an agent is indifferent to what happened before he received an offer (because values are independent and private, he cannot learn anything from the timing of his selection) and is indifferent to what happens after he rejects (because he has no externalities). The argument that SPP auctions are necessary for dominant strategies is more involved, but essentially revolves around the concept of value of information. If deliberation is part of an agent's dominant strategy, then the value of the information gained must always outweigh its cost. We show that there exist settings in which agents deliberate in dominant strategies only when they receive posted-price offers.

Proof. Part 1: Every sequential posted-price auction has a dominant strategy equilibrium in every setting.

In any given SPP auction, each bidder i can only receive an offer with one specific price p_i , though whether or not i receives an offer at all can depend on other agents' strategies. Consider the one-player SPP auction in which a deliberative agent i is offered posted price p_i . As there are no other agents, trivially i has a dominant strategy in this auction.³ Let us call

³Note that this dominant strategy can be extremely complex. For example, under the preference model given in the appendix, it can involve conditional, multistage deliberation policies. See Larson and Sandholm (2004b) for how to construct such a deliberation policy using an MDP representation of the one-player game.

that dominant strategy s_i .

Now we argue that s_i is also a dominant strategy for i in any multi-player SPP auction in which the offer i can receive has price p_i . Let s_{-i} be an arbitrary strategy profile for the other agents $N \setminus \{i\}$. If agent i does not receive an offer, he cannot influence the outcome; in this case, s_i is trivially a best response. Because the agents' valuations are independent and private, other bidders cannot learn anything about i 's value through their own deliberations. Therefore, i cannot infer anything about his own value from the fact that he receives an offer. Because i has no externalities, he does not care which agent (if any) will win the auction in the event that he loses. Thus, i faces the same strategic situation every time he receives a given offer p_i , regardless of which agents have received offers before him, the amounts of those offers, and the deliberations performed by those agents. Agent i also thus faces the same strategic situation as in the one-player SPP auction, and so s_i is still a dominant strategy for i .

Part 2: There exist settings in which the only dominant-strategy implementable social choice functions are those implemented by sequential posted-price auctions.

If a mechanism implements some social choice function (χ, p) in dominant strategy equilibrium, by Theorem 12 (the revelation principle) some other dynamically direct mechanism M implements (χ, p) in truthful dominant strategy equilibrium. Thus, without loss of generality we consider only truthful, dynamically direct mechanisms. We consider a single-step deliberative setting in which each agent i has only two possible valuations (v_i^L and v_i^H).

The rest of the proof proceeds in five steps.

Step 1 (Information Availability): The outcome chosen by the mechanism is completely determined by the types of the agents who deliberate.

Because M is deterministic, because the agents' strategies are deterministic (by dominance), and because the observation of every deliberation is deterministic (by the single-step setting), the social choice function must map to deterministic outcomes ($\chi : V \mapsto N \cup \emptyset$ and each $p_i : V \mapsto \mathbb{R}$). Let $\delta_i(v) = 1$ if i deliberates in equilibrium and $\delta_i(v) = 0$ otherwise. Observe that if an agent i does not deliberate ($\delta_i(v) = 0$) then χ, p and δ cannot depend on his type. This is because the dynamically direct mechanism will not ask i to report v_i (the mechanism is truthful) and so the mechanism cannot condition on v_i when choosing the outcome or when instructing another agent $j \neq i$ to deliberate. If only a subset $S \subset N$ of agents deliberates given type profile v ($i \in S$ iff $\delta_i(v) = 1$), then the outcome does not depend on the types of the other agents $N \setminus S$. ($\chi(v) = \chi(v_S, *)$, $p_j(v) = p_j(v_S, *)$ and $\delta_j(v) = \delta_j(v_S, *)$, where $*$ refers to any possible types of the agents who were not explicitly listed. Thus, e.g., the proposition " $\delta_j(v) = \delta_j(v_i, *)$ " is shorthand for " $\forall v'_{-i}, \delta_j(v) = \delta_j(v_i, v'_{-i})$."

Step 2 (Influence): An agent i only deliberates when doing so makes the difference between winning and losing ($\forall i, \forall v_{-i}, \delta_i(v_i, v_{-i}) = 1$ implies $\chi(v_i^L, v_{-i}) \neq i$ and $\chi(v_i^H, v_{-i}) = i$).

Assume for contradiction that for some v_{-i} , i deliberates ($\delta_i(v_i, v_{-i}) = 1$) and either (1) always wins ($\chi(v_i^L, v_{-i}) =$

$\chi(v_i^H, v_{-i}) = i$); (2) always loses ($\chi(v_i^L, v_{-i}) \neq i$ and $\chi(v_i^H, v_{-i}) \neq i$); or (3) wins with the low type and loses with the high type ($\chi(v_i^L, v_{-i}) = i$ and $\chi(v_i^H, v_{-i}) \neq i$). In case (3), truthfulness is not a dominant strategy: if $p_i(v_i^L, v_{-i}) \leq v_i^L$ then whenever i has the high type, he prefers to misreport that he has the low type (winning and paying $p_i(v_i^L, v_{-i}) \leq v_i^L < v_i^H$) instead of reporting honestly (losing and paying nothing); if $p_i(v_i^L, v_{-i}) > v_i^L$ then whenever i has the low type, he prefers to misreport that he has the high type (losing and paying nothing) rather than reporting honestly (winning and paying more than his value). Thus, assume that we are in case (1) or (2). If i 's price depends on his reported type ($p_i(v_{-i}, v_i^L) \neq p_i(v_{-i}, v_i^H)$), then i always strictly prefers not to deliberate and to report the type that minimizes p_i . Otherwise, i strictly prefers not to deliberate (reporting arbitrarily), obtaining the same allocation and payment but avoiding the cost of deliberation.

Step 3 (Base case): If the outcome is not constant, then there exists some designated agent (without loss of generality, agent 1) who always wins when he has the high valuation, regardless of the valuations of the other agents ($\exists v, v', \chi(v) \neq \chi(v')$ implies that $\exists i, \chi(v_i^H, *) = i$).

Let v^H denote the type profile (v_1^H, \dots, v_n^H) in which every agent has the high type. If given v^H , agent 1 deliberates ($\delta_1(v^H) = 1$) then, by influence, 1 must win ($\chi(v^H) = 1$). Because 1 wins even when all the other agents also have the high type, by influence he must be the only agent to deliberate ($\chi(v^H) = 1$ implies $\forall_{j \neq 1} \delta_j(v^H) = 0$). By information availability, the allocation must not depend on the non-deliberating agents' types ($\chi(v_1^H, *) = 1$).

Step 4 (Induction step): When every agent in $\{1, \dots, k\}$ has the low valuation, either the auction terminates ($\chi(v_1^L, \dots, v_k^L, *) = \emptyset$), or some designated agent (without loss of generality, agent $k+1$) wins whenever he has the high valuation, regardless of the valuations of agents $\{k+2, \dots, n\}$ ($\chi(v_1^L, \dots, v_k^L, v_{k+1}^H, *) = k+1$).

Let v^{LH} denote the type profile $(v_1^L, \dots, v_k^L, v_{k+1}^H, \dots, v_n^H)$ in which all agents from 1 up to k have the low type, and all agents from $k+1$ up to n have the high type. As in Step 3, under v^{LH} if some agent (assume $k+1$) in $\{k+1, \dots, n\}$ deliberates then, by influence, he must win ($\chi(v^{LH}) = k+1$) and must be the only agent in that group to deliberate. By influence, the losing agents with high types ($\{k+2, \dots, n\}$) do not deliberate. By information availability, the allocation must be the same regardless of their types ($\chi(v_1^L, \dots, v_k^L, v_{k+1}^H, *) = k+1$).

If none of the agents in $\{k+1, \dots, n\}$ deliberates given v^{LH} , then the auction can choose any outcome in which no bidder from $\{1, \dots, k\}$ wins. Allocating the good to an agent $i \leq k$ who already received an offer would violate influence: when the other agents report according to v^{LH} , i would be asked to deliberate even though he always wins.

Step 5 (Necessity): Dynamically direct mechanism M implements the same social choice function as some SPP auction.

If M asks no agent to deliberate, then the auctioneer cannot condition on any agent's preferences. In this case, because M is deterministic, it must yield a constant outcome

($\forall v, v', \chi(v) = \chi(v')$). M thus implements the same social choice function as a degenerate SPP auction that makes no offers and either always awards the good to some agent j or never awards the good.

Otherwise, by the argument in Steps 3 and 4, at each stage, a single designated agent is asked to deliberate, and that agent wins iff he reports the high valuation. By the influence property, we know that no further agents deliberate after some agent has reported the high valuation. Because of this, the price that agent pays must be a constant (it cannot vary depending on the deliberations of any other agents, because the earlier-ordered agents all must reveal low types for i to be queried, and because none of the later-ordered agents deliberate; it cannot vary depending on i 's own deliberation, because he only wins when he reveals the high type). M therefore implements the same social choice function as an SPP that makes offers to the agents in ascending order, with each agent i being offered price $p_i(v_i^H, v_{-i}^L)$. \square

We now relate our result to Larson and Sandholm's impossibility result (Theorem 1). We find that, in IPV settings, SPP auctions have dominant strategies and satisfy all of Larson and Sandholm's desiderata except for strategic-deliberation proofness (which doesn't apply to IPV settings).

Corollary 15. *SPP auctions are strategy dependent, non-misleading, and preference-formation independent in dominant-strategy equilibrium in IPV auction settings.*

Proof. Any SPP auction, provided that it does not halt until an offer is accepted or all the agents have rejected offers, is strategy-dependent: each bidder can affect the outcome by accepting or rejecting his offer. All SPP auctions are preference-formation independent: the auctioneer does not specify how the agents should deliberate, and does not need to condition the rules of the auction on details of the setting. All SPP auctions are non-misleading: in equilibrium, no agent would never accept an offer that exceeds his valuation. \square

Conclusions

This paper studied the problem of designing auctions for deliberative agents without strategic deliberation, a setting in which the classical revelation principle does not apply and no (non-trivial) dominant strategy auctions had been identified. We (1) gave a novel reformulation of the revelation principle which applies even in the presence of dynamic information gathering, and (2) leveraged this result to characterize the set of dominant strategy auctions in IPV settings, showing that this set is precisely the sequential posted price auctions.

We foresee three strands of future work. First, we aim to investigate the application of SPP auctions to (non-IPV) settings with strategic deliberation. Second, we are interested in investigating whether a randomized auction could guarantee a sufficiently high expected value of information to yield dominant strategies *in expectation* without being an SPP auction. Third, we aim to design SPP auctions that maximize revenue or social welfare in deliberative settings, leveraging recent positive results that use SPP auctions to maximize these quantities in non-deliberative settings (e.g. (Blumrosen and Holenstein 2008; Chawla et al. 2010; Sandholm, Conitzer, and Boutilier 2005)).

Appendix: Extended model

The single-step deliberative model used in the body of the paper is expressive enough to describe interesting settings, e.g., in which indirect mechanisms can implement social choice functions that cannot be implemented by direct mechanisms, and in which SPP auctions are the unique dominant strategy auction design. However, our positive results (Theorem 12 and Theorem 14 Part 1) apply to a much broader class of settings. In this expanded model, deliberation (1) can yield noisy signals; (2) can be a multistage process in which a bidder gets progressively more information about his valuation; and (3) can change an bidder's valuation (e.g., if he is running an anytime optimization algorithm to identify a policy for using the good, finding a better feasible solution increases the value that he can derive from the good).

Definition 16 (General deliberative Setting). *A general deliberative setting is a 8-tuple $\langle N, X, O, V, f, D, c, P \rangle$.*

- N is a set of agents numbered $1, \dots, n$.
- O is a set of possible observations an agent can make. \bar{O} is the set of all finite-length vectors of observations, including the zero-length vector, denoted $\langle \rangle$.
- V is a set of possible valuation-function profiles. Every element $v \in V$ is a n -length vector of valuation functions (one per agent). Each valuation function v_i has the form $v_i : \bar{O} \mapsto \mathbb{R}$, where $v_i(\bar{o}_i)$ is i 's value for winning the good given that he has made observations \bar{o}_i .
- $f : V \mapsto \mathbb{R}$ is the prior probability distribution over valuation-function profiles. Agents get their asymmetric information only by deliberating, so an agent who does not deliberate knows nothing about his own preferences beyond this prior.
- D is a set of deliberations. Every $d \in D$ is a function $d : V_i \mapsto \Delta O$, where ΔS denotes the set of probability distributions over set S . \bar{D} is the set of all finite-length vectors of deliberations, including the zero-length vector, denoted $\langle \rangle$.
- c is a vector of n cost functions. $c_i : \bar{D} \mapsto \mathbb{R}$ is the cost i must pay to perform the set of deliberations \bar{D} . Every cost function $c_i(\cdot)$ is non-decreasing; thus, agents cannot reduce their total deliberation costs by increasing the number of deliberations they perform. We assume that $c_i(\langle \rangle) = 0$: the cost of not deliberating is zero.
- P is a vector of n functions specifying the agents' possible deliberation actions. $P_i : \bar{D} \times \bar{O} \mapsto 2^D$ is set of deliberations i is capable of performing, given the deliberations he has previously performed and his observations.

In a general deliberative setting, an agent's utility depends on his valuation function v_i , whether or not he wins ($w_i = 1$ if he wins and 0 otherwise), his total transfers t_i , the deliberations he has performed \bar{d}_i and the observations he has made \bar{o}_i . We assume that utility is quasilinear: $u_i(v_i, w_i, t_i, \bar{d}_i, \bar{o}_i) = w_i v_i(\bar{o}_i) - t_i - c_i(\bar{d}_i)$. (Note that this expression is slightly different from the one given earlier in the paper; in particular, an agent's utility can depend on what observations he has made.) As before, we assume that the entire setting is common knowledge, that the auction is deterministic, and that losers make no payments.

References

- Bergemann, D., and Valimaki, J. 2002. Information acquisition and efficient mechanism design. *Econometrica* 70(3).
- Blumrosen, L., and Holenstein, T. 2008. Posted prices vs. negotiations: an asymptotic analysis. In *EC*.
- Cavallo, R., and Parkes, D. 2008. Efficient metadeliberation auctions. In *AAAI*.
- Chawla, S.; Hartline, J. D.; Malec, D. L.; and Sivan, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *STOC*.
- Compte, O., and Jehiel, P. 2001. Auctions and information acquisition: Sealed-bid or dynamic formats. Working paper.
- Cremer, J.; Spiegel, Y.; and Zheng, C. 2003. Optimal selling mechanisms with costly information acquisition. Technical report, Working paper, August.
- Feigenbaum, J.; Krishnamurthy, A.; Sami, R.; and Shenker, S. 2003. Hardness results for multicast cost sharing. *Theoretical Computer Science* 304(1–3).
- Kleinberg, R., and Leighton, T. 2003. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *FOCS*.
- Larson, K., and Sandholm, T. 2001. Costly valuation computation in auctions. In *TARK*.
- Larson, K., and Sandholm, T. 2004a. Strategic deliberation and truthful revelation: An impossibility result. In *ACM-EC*.
- Larson, K., and Sandholm, T. 2004b. Using performance profile trees to improve deliberation control. In *AAAI*.
- Larson, K., and Sandholm, T. 2005. Mechanism design and deliberative agents. In *AAMAS*.
- Larson, K. 2006. Reducing costly information acquisition in auctions. In *AAMAS*.
- Myerson, R. 1981. Optimal auction design. *Mathematics of Operations Research* 6(1).
- Myerson, R. 1986. Optimal auction design. *Econometrica* 52(2).
- Pavan, A.; Segal, I.; and Toikka, J. 2008. Dynamic mechanism design: Revenue equivalence, profit maximization and information disclosure.
- Persico, N. 2000. Information acquisition in auctions. *Econometrica* 68(1).
- Rasmusen, E. 2006. Strategic implications of uncertainty over ones own private value in auctions. *Advances in Theoretical Economics* 6(1).
- Sandholm, T., and Gilpin, A. 2006. Sequences of take-it-or-leave-it offers: near-optimal auctions without full valuation revelation. In *AAMAS*.
- Sandholm, T.; Conitzer, V.; and Boutilier, C. 2005. Automated design of multistage mechanisms. In *Workshop on Incentive-Based Computing*.
- Sandholm, T. 2000. Limitations of the Vickrey auction in computational multiagent systems. *International Journal of Electronic Commerce* 4(3).
- Shakkottai, S.; Srikant, R.; Ozdaglar, A. E.; and Acemoglu, D. 2008. The price of simplicity. *IEEE Journal on Selected Areas in Communications* 26(7).
- Tan, G. 1992. Entry and R & D in procurement contracting. *Journal of Economic Theory* 58(1).
- Thompson, D., and Leyton-Brown, K. 2007. Valuation uncertainty and imperfect introspection in second-price auctions. In *AAAI*.
- Wilson, R. 1987. Game-theoretic approaches to trading processes. *Advances in Economic Theory: Fifth World Congress*.