

Approximately Revenue-Maximizing Auctions for Deliberative Agents

L. Elisa Celis
 ecelis@cs.washington.edu
 University of Washington

Anna R. Karlin
 karlin@cs.washington.edu
 University of Washington

Kevin Leyton-Brown
 kevinlb@cs.ubc.edu
 University of British Columbia

C. Thach Nguyen
 camnguyen@fb.com
 Facebook *

David R. M. Thompson
 daveth@cs.ubc.edu
 University of British Columbia

Abstract

In many real-world auctions, a bidder does not know her exact value for an item, but can perform a costly deliberation to reduce her uncertainty. Relatively little is known about such deliberative environments, which are fundamentally different from classical auction environments. In this paper, we propose a new approach that allows us to leverage classical revenue-maximization results in deliberative environments. In particular, we use Myerson (1981) to construct the first non-trivial (i.e., dependent on deliberation costs) upper bound on revenue in deliberative auctions. This bound allows us to apply existing results in the classical environment to a deliberative environment. In addition, we show that in many deliberative environments the only optimal dominant-strategy mechanisms take the form of sequential posted-price auctions.

Introduction

Consider the following example:

An agent is considering buying a used car for \$8,000. The value of the car to her depends on her needs and preferences. She initially believes the value is uniformly between \$5,000 and \$10,000. However, she can “deliberate”: that is, she can act to reduce her uncertainty about this value. For example, she can hire a mechanic to examine the car, or take it for a test drive. Each deliberation has a different cost (in money or time), and reveals different information. As a rational agent, she evaluates the cost and value of information for each deliberation, and chooses the best one. She then decides whether or not to purchase the car, based on what she learned.

This example introduces a deliberative agent, who is uncertain about her preferences, but can take actions to reduce the uncertainty. Judging the value of a good is difficult since it depends on many parameters. Additionally, there may be computational constraints that prevent an agent from achieving certainty about her valuation.

Previous work shows that auction design for deliberative agents is fundamentally different from classical auction design due to the greater flexibility in the agents’ strategies. In

classical mechanism design, an agent only has to decide how much information to reveal. In deliberative-agent mechanism design, an agent first has to decide how much information to acquire and then how much to reveal. This affects equilibrium behavior. For example, in second-price auctions, deliberative agents do not have dominant strategies (Larson and Sandholm 2004) and must coordinate their information gathering (Thompson and Leyton-Brown 2007). Furthermore, the standard revelation principle, which asserts that every multi-stage auction is equivalent to some sealed-bid auction, no longer holds. For example, in a Japanese auction, bidders can condition their information gathering on information revealed at earlier stages, coordinating in ways that are not possible in sealed-bid auctions (Compte and Jehiel 2005).

There has been considerable interest in designing novel auctions for deliberative agents. This research has mostly focused on maximizing social welfare subject to various constraints (Bergemann and Valimaki 2002; Cavallo and Parkes 2008; Larson 2006), with some research on revenue maximization in Bayes-Nash equilibrium (Cramer, Spiegel, and Zheng 2003; Bikhchandani 2009). More recently, Thompson and Leyton-Brown (2011) demonstrated that dominant strategy auctions are also possible. However, they proved that, for single-item auctions with binary-valued agents, the space of dominant strategy mechanisms is limited, in that it is equivalent to a sequence of posted-price offers.

In this paper, we study the design of revenue-maximizing auctions with dominant strategies in deliberative settings, and make two contributions. First, we show that posted-price-based auctions characterize the space of dominant strategy auctions in significantly more general deliberative settings. Second, we show how to design auctions that obtain revenue that is within a small constant factor of the maximum possible revenue in these settings.

Background

In this paper, we consider a simple model of deliberative agents, for whom deliberation is an all-or-nothing decision.

Definition 1 (Simple Deliberative Agent). *A simple deliberative agent i is represented by a tuple $(v_i; F_i; c_i)$, where*

- v_i is the true value of an agent (this value is unknown, even to the agent herself).

*Work completed while studying at University of Washington
 Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

- F_i is the publicly known distribution from which v_i is drawn. (Values are independent and private, and F_i has bounded expectation.)
- c_i is the cost of deliberating, i.e., the price agent i must pay to learn the value v_i .

We consider “single-parameter” settings where an auctioneer is offering goods or a service, and based on the outcome of a mechanism, chooses a subset of the agents to be served (we call them “winners”) and a price p_i for each winning agent i . Depending on the particular setting, only certain subsets of agents can be feasibly served.¹ The utility of agent i for the outcome of a mechanism is $v_i - p_i - c_i$ if the agent deliberates and wins, $-c_i$ if the agent deliberates and loses, $E(v_i) - p_i$ if the agent doesn’t deliberate and wins, and 0 otherwise.

Definition 2 (Deliberative environment). A deliberative environment is a tuple $(N; \mathcal{S}; \mathbf{F})$ where

- N is a set of deliberative agents, (each with her own value distribution and deliberation cost),
- \mathcal{S} is a collection of subsets of N . Each set $S \in \mathcal{S}$ represents a set of agents that can be served at the same time.
- \mathbf{F} is the joint distribution of agents’ values.

Throughout, we will assume that the environment is common knowledge to the agents and to the auctioneer. We consider only deterministic mechanisms. Also, we stipulate that mechanisms never allocate to any agent who is indifferent between winning and losing given the price. (That is, we assume that if there exists a threshold t for agent i (given the reported values of the other agents) where if $v_i > t$ then i is served and if $v_i < t$ then i is not served, then i is not served when $v_i = t$.) Without this technical assumption, the characterization of dominant strategy mechanisms is more complicated, in uninteresting ways. We discuss this point in more detail in an online appendix available at <http://www.cs.ubc.ca/research/deliberation/>.

As in the classical analysis of auctions, mechanisms can be complicated multi-stage processes (e.g., Japanese auctions). Thus a pure strategy can be a complex policy conditioning on information revealed over time, and a dominant strategy is a policy that maximizes the agent’s expected utility regardless of the other agents’ policies. However, deliberative environments differ from classical environments in that the revelation principle cannot reduce every mechanism down to a strategically equivalent single-stage mechanism. This is because a deliberative agent might want to defer her deliberation until she learns something about her competitors (e.g., how many bidders are still standing in a Japanese auction). Nevertheless, Thompson and Leyton-Brown (2011) provide a revelation principle for deliberative agents, showing that any dominant-strategy mechanism is equivalent to a

¹For example, in a single-item auction only one agent can be served. In a k -unit auction, any subset of k agents can be served. In a single-minded combinatorial auction, any subset of agents whose desired item sets do not overlap can be served. See Hartline and Karlin (2007) for a discussion of single-parameter settings.

dominant-strategy truthful “dynamically direct mechanism” (DDM).²

Definition 3 (Dynamically Direct Mechanism (DDM) for simple deliberative agents). A dynamically direct mechanism is a multi-stage mechanism where at each stage a single agent is asked to deliberate and report her true value.

Although DDMs “request” that bidders report true values, the bidders can respond by bidding in any way they like. This is analogous to direct mechanisms in the classical setting; e.g., note that first-price and second-price auctions are both direct mechanisms, but only the latter is truthful. Truthful DDMs, used in later parts of the paper, are DDMs for which truthful reporting is a dominant strategy.

In this paper, we consider DDMs under two communication models, the *private-communication model*, where no agent observes the interaction between the mechanism and any other agent, and the *public-communication model*, where all agents observe all such interactions. Of course, in either case, the mechanism itself is common knowledge.

Sequential posted-price mechanisms (SPPs) play an important role in what follows.

Definition 4 (Sequential Posted Price mechanism). A sequential posted-price mechanism offers posted prices to agents, one at a time. If an agent accepts an offer, then the mechanism must serve her and charge her exactly the offered price. An agent who rejects an offer will not be served. (No agent gets a second offer.) The mechanism is forbidden from making offers to agents it can no longer feasibly serve given the offers that have already been accepted.

Characterization

In this section, we prove that SPPs characterize dominant strategy mechanisms for two classes of deliberative environments. The if direction (i.e., that SPPs have dominant strategies in any deliberative setting) follows trivially from Thompson and Leyton-Brown (2011). Thus, we focus our attention on proving that only mechanisms that are equivalent to SPPs have dominant strategies.

Theorem 5 (Characterization: only-if direction).

1. In any single-item deliberative environment with private or public communication, every dominant strategy mechanism \mathcal{M} is equivalent to an SPP \mathcal{N} .
2. In any single-parameter deliberative environment with public communication, every revenue-maximizing dominant strategy mechanism \mathcal{M} is equivalent to an SPP \mathcal{N} .

In the online appendix, we provide complete proofs, as well as discussion and examples showing why various assumptions we have made are necessary. Here, we sketch the proof of Theorem 5 and the lemmas that lead up to it.

From the revelation principle for deliberative agents, which characterizes the set of dominant strategy mechanisms, it suffices to characterize the class of truthful DDMs. As in classical settings, any truthful mechanism must satisfy a standard monotonicity condition.

²We refer the reader to Thompson and Leyton-Brown (2011) for details on dominant strategy mechanisms and other aspects of deliberative agents in the single-item setting.

Lemma 6 (Monotonicity). *In any truthful DDM, for any agent i who is asked to deliberate (given whatever values are reported by other agents, \hat{v}_{-i}) there exists a “critical value” $t_i(\hat{v}_{-i})$ where if $v_i > t_i(\hat{v}_{-i})$ then i is served and pays exactly $t_i(\hat{v}_{-i})$, and where otherwise she is not served and pays nothing.*

The proof is straightforward and follows the proof of the equivalent lemma in the classical setting (Myerson 1981).

The monotonicity lemma implies that if an agent deliberates and wins, the price she is charged depends only on values reported by other agents. For characterization via SPPs, we need a slightly stronger result: that the price of an agent depends only on the values of the agents which deliberated before her, i.e., the information the mechanism has when she is asked to deliberate. Thompson and Leyton-Brown (2011) proved such a result for binary-valued agents using an “Influence Lemma,” which states that if an agent with two possible values deliberates, then she must be served when she reports the higher value and not served when she reports the lower. We generalize that lemma to show that there is a range of values (strictly inside the support of F_i) that are potential critical values if an agent is asked to deliberate in a truthful DDM. For simplicity, we present the lemma here for the case of atomless distributions.

Lemma 7 (Generalized Influence Lemma). *Every deliberative agent i has low and high deliberation thresholds ℓ_i and h_i , where $\ell_i \leq h_i$, both in the interior of the support of F_i , such that: In any truthful DDM, if i is asked to deliberate given v_{-i} , then her critical value is between those thresholds: $t_i(v_{-i}) \in [\ell_i, h_i]$.*

Proof. Given the reported values of other agents \hat{v}_{-i} and by monotonicity, agent i is effectively faced with a posted price offer of $p = t_i(\hat{v}_{-i})$. Faced with this offer, she has three possible strategies: she can accept the offer without deliberation, reject it without deliberation, or deliberate and only accept when her value is greater than p . Denote the expected utility of these three strategies $u_i^a(p)$, $u_i^r(p)$ and $u_i^d(p)$ respectively. (See Figure 1.) It is easy to see that:

$$\begin{aligned} u_i^a(p) &= \mathbb{E}[v_i] - p = \int_0^\infty v f_i(v) dv - p; \\ u_i^d(p) &= \mathbb{E}[v_i - p | v_i \geq p] \Pr[v_i \geq p] - c_i \\ &= \int_p^\infty v f_i(v) dv - p(1 - F(p)) - c_i; \\ u_i^r(p) &= 0. \end{aligned}$$

Intuitively, given the cost c_i of deliberating, ℓ_i denotes the price where i is indifferent between accepting and deliberating, and h_i denotes the price where she is indifferent between rejecting and deliberating. Thus, we have only to show that $\ell_i \leq h_i$. To see this, observe that if $\mathbb{E}[v_i] \leq c_i$ then $u_i^d(p) \leq 0$ for any p , hence i never deliberates, and $\ell_i = h_i = \mathbb{E}[v_i]$. On the other hand, if $\mathbb{E}[v_i] > c_i$, we can easily check that $u_i^a(p) - u_i^d(p)$ is a strictly decreasing function, and $u_i^r(p) - u_i^d(p)$ is a strictly increasing function. Hence, there are unique points t_1 and t_2




Figure 1: The relationships between u_i^a, u_i^d and u_i^r in Lemma 7. (a) $t_1 > t_2$; (b) $t_1 \leq t_2$.

such that $u_i^a(p) - u_i^d(p) \geq 0$ if and only if $p \leq t_1$, and $u_i^r(p) - u_i^d(p) \geq 0$ if and only if $p \geq t_2$.

If $t_1 \leq t_2$ then it is easy to check that $\ell_i = t_1$ and $h_i = t_2$ satisfy the lemma’s conditions. Otherwise, we must have $t_2 \leq \mathbb{E}[v_i]$ since $u_i^a(t_2) \geq u_i^d(t_2) = u_i^r(t_2)$, and $t_1 \geq \mathbb{E}[v_i]$ since $u_i^a(t_1) = u_i^d(t_1) \leq u_i^r(t_1)$. In this case, $\ell_i = h_i = \mathbb{E}[v_i]$ satisfy the lemma’s conditions. \square

We now use the generalized influence lemma to prove that for our settings, the price offered to an agent depends only on the values of the agents that deliberated before her.

Lemma 8. *Let \mathcal{M}*

1. *be a truthful DDM for a single-item auction, or*
2. *a revenue maximizing truthful DDM for a public-information single-parameter setting*

where all agents have bounded expected values. If \mathcal{M} asks i to deliberate and decides to sell to i , the price it charges i depends only on the values reported by agents asked to deliberate before i .

Proof. Let \mathbf{v}_{-i}^b be the values reported by the agents asked to deliberate before i , and \mathbf{v}_{-i}^a be the values of the remaining agents, except for i .

Let us first prove (1). By Lemma 7, if $v_i > h_i$ then \mathcal{M} has to sell to i . In this case, \mathcal{M} cannot ask anyone else to deliberate after i does. Otherwise, another agent j may report that $v_j > h_j$, forcing \mathcal{M} to sell to both i and j , which is impossible because \mathcal{M} has only one item to sell. Since \mathcal{M} cannot ask anyone else to deliberate, the price p_i at which it sells to i if $v_i > h_i$ does not depend on \mathbf{v}_{-i}^a . For i to report truthfully, it cannot depend on v_i either. Hence, p_i only depends on \mathbf{v}_{-i}^b .

Next, suppose $v_i \in [\ell_i, h_i]$. If there is no \mathbf{v}_{-i}^a such that \mathcal{M} sells to i then the lemma holds. Otherwise, fix some \mathbf{v}_{-i}^a such that \mathcal{M} sells to i at price p'_i . We argue that $p'_i = p_i$. Indeed, if $p'_i > p_i$ then i has incentive to lie that her value is more than h_i to win at a lower price. On the other hand, suppose $p'_i < p_i$. Then in the case where i ’s valuation is $v'_i > h_i$ (and other agents report according to \mathbf{v}_{-i} if asked to), i has incentive to lie that her valuation is v_i to still win at a better price. We conclude that to incentivize i to report truthfully, the mechanism must set $p'_i = p_i$.

Let us now prove (2). Since the agents have full information, the agents in \mathbf{v}_{-i}^a can condition their strategies on v_i ; this is the primary difference between the two models. First, assume there is some \mathbf{v}_{-i}^a such that when $v_i \in (\ell_i, h_i]$, \mathcal{M} sells to i at price p'_i . By Lemma 7, recall that if a mechanism asks agent i to deliberate and $v_i > h_i$ then \mathcal{M} must sell to agent i . Now, consider \mathbf{v}_{-i}^a such that conditioned on $v_i > h_i$, i is charged price p_i .³

Thus, a strategy exists where the agents in \mathbf{v}_{-i}^a reach an outcome that allocates to i at price p'_i if i reports $v_i \in (\ell_i, h_i]$, and reaches the outcome that allocates to i at price p_i if i reports $v_i > h_i$. Fix this strategy for \mathbf{v}_{-i}^a . We will now show that $p'_i = p_i$. First, if $p'_i > p_i$, then i has incentive to lie that her value is greater than h_i to win at a lower price. On the other hand, suppose $p'_i < p_i$. Then in the case where i 's valuation is $v'_i > h_i$ i has incentive to lie that her valuation is v_i to again win at a better price. Thus, to incentivize i to report truthfully, the mechanism must set $p'_i = p_i$. Note that this is true of any price p_i which can be reached when $v_i > h_i$. Therefore, if $v_i > h_i$ she is always charged p'_i . Similarly, for any situation in which i wins after reporting $v_i \in (\ell_i, h_i]$ we have $p'_i = p_i$. Hence, the price i is charged does not depend on \mathbf{v}_{-i}^a .

Now assume there is no \mathbf{v}_{-i}^a such that \mathcal{M} sells to i when $v_i \in (\ell_i, h_i]$, thus we cannot use the argument above. Instead, we assume for sake of contradiction that there are two different outcomes for $v_i \geq h_i$ that charge p_i and p'_i respectively. Without loss of generality assume $p_i < p'_i$. Note that if either price was not acceptable to v_i , then she would have incentive to lie and report $v_i < \ell_i$ to avoid being overcharged. Thus, we can define a mechanism \mathcal{N} that is identical to \mathcal{M} except that it charges p'_i instead of p_i . Since \mathcal{N} attains strictly more revenue, this contradicts the optimality of \mathcal{M} . This completes the proof for both settings. \square

With these lemmas, we can now prove Theorem 5.

Proof. Given \mathcal{M} , we will construct \mathcal{N} .

Consider the valuation profile where $v_i > h_i$ for all agents i . If the item is not sold in \mathcal{M} for this case, then no agent is asked to deliberate; hence \mathcal{M} is equivalent to the vacuous SPP \mathcal{N} that does not make any offer.

Now, assume \mathcal{M} is not vacuous. For any given set of agents and priors, since \mathcal{M} is deterministic, there is some agent whom \mathcal{M} will address first. Without loss of generality, let us refer to her as agent 1. Recall from Lemma 7 that if $v_1 > h_1$ then \mathcal{M} must sell to agent 1, and let the price of this sale be p_1 .

From Lemma 8 we know that whenever the item is sold to agent 1, it must be sold this same price p_1 . Let \mathcal{N} offer a fixed price p_1 to agent 1. If \mathcal{M} sells to 1 at price p_1 without asking her to deliberate, then, p_1 is clearly a price that agent 1 prefers to accept without deliberating. Hence, she will also accept \mathcal{N} 's offer.

³Note that, contrary to the single-item case, \mathcal{M} may continue asking agents to deliberate as long as the addition of the agent would not invalidate the feasibility of the winning set. Thus, we have not yet ruled out the possibility of multiple outcomes with different p_i . We will address this shortly.

Instead, assume \mathcal{M} asks 1 to deliberate. Then, since \mathcal{M} is truthful, it is in the agent's interest to deliberate and learn her value v_1 . Similarly, when \mathcal{N} offers her price p_1 it will be in her interest to deliberate. Then, from Lemma 7, \mathcal{M} awards her the item if and only if p_1 is less than v_1 . This exactly matches the scenario under which she accepts \mathcal{N} 's offer.

Thus, the behavior of \mathcal{M} and \mathcal{N} is the same after agent 1. The theorem then follows by induction on the next agent \mathcal{M} addresses. \square

Approximate Revenue Maximization

We now turn to the question of designing mechanisms to maximize expected profit in deliberative environments. Mechanisms that do this are called *optimal* mechanisms. In some cases, optimal SPPs can be computed directly and efficiently. However, since our characterization via SPPs does not hold for private communication single-parameter auctions, we pursue an alternative approach that is more broadly applicable: we present a transformation from a deliberative environment to a related classical environment and then show that the expected revenue of the optimal mechanism in this classical environment is an upper bound on the optimal revenue in the deliberative setting. This allows us to use near-optimal SPPs known for classical environments to obtain near-optimal mechanisms in the deliberative setting.

Upper Bound in Classical Environment

For the ease of presentation, we assume the agents' distributions are continuous, and that for each distribution F_i , the probability density function f_i exists.

In truthful DDMs, by Lemma 7, no agent has a critical value (or pays a price) outside of $[\ell_i, h_i]$. This motivates the following definition of the effective value of a deliberative agent, the value she would have if she were a classical agent.

Definition 9. Consider an agent i with value v_i and low and high deliberation thresholds ℓ_i and h_i respectively. The effective value of i is v'_i , defined by

$$v'_i = \begin{cases} \ell_i & \text{if } v_i \leq \ell_i \\ v_i & \text{if } v_i \in (\ell_i, h_i) \\ h_i & \text{if } v_i \geq h_i. \end{cases}$$

The classical agent with value v'_i is the representative of agent i . We will use i' to denote the representative of i , and use v'_i to denote the effective value of i , i.e., the value of i' .

It is clear that v'_i is drawn from i 's effective distribution.

Definition 10. Let i be a deliberative agent with values drawn from a distribution F_i with low and high deliberation thresholds ℓ_i and h_i . Then the effective distribution of i is defined over the interval $[\ell_i, h_i]$ by $G_i(x) = F_i(x)$ for all $\ell_i \leq x < h_i$ and $G_i(h_i) = 1$. Hence, G_i has a point mass of $F_i(\ell_i)$ at (ℓ_i) and a point mass of $1 - F_i(h_i)$ at h_i . Furthermore, if $\ell_i = h_i$ then the support of G_i contains a single value.

We say that $E' = (N', S', \mathbf{G})$ is the *representative environment* of $E = (N, S, \mathbf{F})$ if E' is obtained from E by replacing each agent i with value v_i drawn from F_i by an agent i' with value v'_i drawn from G_i (which induces the

corresponding set of agents N' and feasible subsets \mathcal{S}'). The following theorem relates the expected revenue of truthful mechanisms in the two environments.

Theorem 11. *For any truthful DDM \mathcal{M} in E , there is mechanism \mathcal{N} where:*

1. \mathcal{N} is a truthful (in expectation) mechanism in E' , and
2. the expected revenue of \mathcal{M} is at most the expected revenue of \mathcal{N} , i.e., $\mathbb{E}_{\mathbf{v} \sim \mathbf{F}} [\text{rev} \mathcal{M} E, \mathbf{v}] \leq \mathbb{E}_{\mathbf{u} \sim \mathbf{G}} [\text{rev} \mathcal{N} E', \mathbf{u}]$.

Before proving Theorem 11, we introduce some notation and definitions. For a deliberative agent i whose value comes from a distribution F_i and whose low and high deliberation thresholds are ℓ_i and h_i respectively, we denote by F_i^ℓ the distribution over $[0, \ell]$ defined by $F_i^\ell(0) = 0$ and $F_i^\ell(x) = F_i(x)/F_i(\ell_i)$ for all $x \in (0, \ell]$. Similarly, F_i^h denote the distribution over $[h, \infty)$ defined by $F_i^h(h) = 0$ and $F_i^h(x) = (F_i(x) - F_i(h))/(1 - F_i(h))$.

Definition 12. *Let $\mathbf{v}' \in \text{Support}(\mathbf{G})$ be a valuation profile of the representative agents in E' . We say that \mathbf{v} is an originator of \mathbf{v}' if*

- $v'_i = \ell_i$ implies $v_i \leq \ell_i$,
- $v'_i = h_i$ implies $v_i \geq h_i$, and
- $v'_i \in (\ell_i, h_i)$ implies $v_i = v'_i$.

Given a valuation profile \mathbf{v}' of the representative agents, we can construct a random originator \mathbf{v} of \mathbf{v}' by setting (i) $v_i = v'_i$ if $v'_i \in (\ell_i, h_i)$, (ii) v_i equal to a random number drawn from F_i^ℓ if $v'_i = \ell_i$, and (iii) v_i equal to a random number drawn from F_i^h if $v'_i = h_i$. A random originator constructed this way is called a *sampled originator* of \mathbf{v}' .

Proof of Theorem 11. Given a truthful DDM \mathcal{M} on E , we construct the following mechanism \mathcal{N} on E' :

1. Solicit a bid vector \mathbf{b} from the agents.
2. Construct a sampled originator \mathbf{u} of \mathbf{b} .
3. Run \mathcal{M} on (E, \mathbf{u}) .

First, we show that \mathcal{N} is truthful in expectation. Consider a representative $i' \in N'$ with value $v_{i'}$. We show that submitting $b_{i'} = v_{i'}$ yields the best expected utility for i' . To this end, we show that this is the case for any fixed \mathbf{u}_{-i} . Once \mathbf{u}_{-i} is fixed, there are three cases regarding whether i is served.

1. \mathcal{M} does not ask i to deliberate and does not serve her. In this case, i' is not served by \mathcal{N} and her bid does not matter, therefore she is truthful.
2. \mathcal{M} asks i to deliberate. Then by Lemma 6, there is a threshold $t(\mathbf{u}_{-i})$ that does not depend on b_i such that i is served if and only if $u_i \geq t(\mathbf{u}_{-i})$. Moreover, $t(\mathbf{u}_{-i})$ is the price of i , hence i' , if she is served. By Lemma 7, we have $t(\mathbf{u}_{-i}) \in (\ell_i, h_i)$. If $v_{i'} \geq t(\mathbf{u}_{-i})$ then i' prefers to buy at this price, and bidding truthfully makes sure that this happens. On the other hand, if $v_{i'} < t(\mathbf{u}_{-i})$ then i' prefers not to buy, and bidding truthfully also ensures that this outcome is chosen.

3. \mathcal{M} does not ask i to deliberate but serves i and charges her some price p_i , which is independent of i 's value. Since \mathcal{M} is a truthful DDM, by Lemma 7, we must have $p_i \leq \ell_i \leq v_{i'}$, therefore i' prefers buying to withdrawing from the mechanism, and hence bids truthfully.

The second part of the theorem follows from the fact that if \mathbf{b} is randomly drawn from \mathbf{G} then \mathbf{u} is a random draw from \mathbf{F} . \square

As an immediate corollary of Theorem 11, we get an upper bound on the revenue of all truthful DDMs on E .

Corollary 13. *For any truthful DDMs \mathcal{M} , the expected revenue of \mathcal{M} in E is at most the expected revenue of the optimal revenue-maximizing auction in E' .*

Approximation

Corollary 13 suggests that in order to approximate the expected revenue of the optimal mechanism in a deliberative environment, we can design a truthful DDM that approximates the optimal auction in their representative environment.

For this we apply the following theorem of Chawla et al. (2010), which shows how, in classical single parameter environments, to obtain a constant factor approximation of the optimal auction (Myerson 1981) with an SPP.

Theorem 14 (Chawla et al. 2010). *There are SPPs that approximate the expected revenue of the optimal mechanism in various classical environments. In particular,*

- *For any general matroid environment, there is an SPP whose expected revenue is at least 1/2 of the optimal expected revenue;*
- *For any uniform matroid or partition matroid environment, there is a SPP whose expected revenue is at least $(e-1)/e$ times the optimal expected revenue;*
- *For any environment whose feasible set system is the intersection of two matroids, there is a SPP whose expected revenue is at least 1/3 times the optimal expected revenue.*

Moreover, it is immediate that any SPP outputs the same outcome on a instance (E, \mathbf{v}) of a deliberative environment E and its representative instance (E', \mathbf{v}') . This gives us the following approximation result.

Corollary 15. *There are SPPs that approximate the expected revenue of the optimal truthful DDMs in the single-parameter settings of Theorem 14. These settings include multi-unit auctions, single-minded combinatorial auctions, and many other natural settings.*

Tightness of the upper bound and approximation ratio

Unfortunately, the upper bound given by Corollary 13 is sometimes unachievable. Example 16 shows a case where the optimal mechanism, an SPP, gets strictly less revenue than the bound from the representative environment. Note that the gap example in Blumrosen and Holenstein (2008), demonstrating that SPPs are not optimal, does not apply. They stipulate that the value distribution cannot have an atom at its lower bound, which is not true of representative environments. However for auctions with two or more units, not only do truthful DDMs other than SPPs exist, but they can also get strictly

greater revenue (as demonstrated by an example in the online appendix).

Example 16. Consider a single-item auction with two bidders whose values are drawn from the uniform distribution over $[0, 1]$ and whose costs of deliberation are both 0.01. The optimal mechanism for this environment is an SPP (offering the first bidder a price of 0.625, and the second a price of 0.5) which gets a revenue of ~ 0.391 . In the representative environment, the optimal auction gets ~ 0.416 .

Future Work

There is much more work to be done to understand deliberative environments. While the characterization of DDMs by SPPs shows that deliberative settings are fundamentally different than their classical counterparts, our approximation results show that they are related, at least in this simple deliberative setting. The future is wide open for mechanism design in deliberative environments with many different directions to follow.

There are still numerous open problems concerning revenue maximization in dominant strategies. As the examples in our online appendix show, there are other dominant-strategy-truthful mechanisms (for private communication, single parameter environments) that get strictly more revenue than the optimal SPP. Characterizing these mechanisms remains open. Also, there is still a gap between public and private communication settings. It may be that even revealing a small amount of information (e.g., how many goods have already been allocated) is enough to violate dominant-strategy truthfulness in non-SPP mechanisms.

Another direction is to consider alternate solution concepts, such as Bayes-Nash equilibria, implementation by iterative removal of dominated strategies, or randomized truthful (in expectation) mechanisms. Either of these relaxations allows for mechanisms that get strictly more revenue than the optimal SPP. In fact, such mechanisms can get more revenue than is possible in any dominant-strategy-truthful mechanism. Note that most of our results, including the revenue upper bound, only apply to deterministic dominant-strategy mechanisms. Hence, to study approximation for these looser solution concepts, we would need to derive a different upper bound.

Additionally, although there exist social-welfare optimizing mechanisms (Bergemann and Valimaki 2002; Cavallo and Parkes 2008) they rely on restricted environments (e.g., single-good auctions or models in which all deliberations must happen simultaneously) and are only Bayes-Nash incentive compatible, not dominant-strategy truthful. It may be possible to use existing SPP results in classical settings (Blumrosen and Holenstein 2008) to get approximate SPPs for deliberative environments, but not using the proof techniques we used here.

Lastly, more general (i.e., not single-step) deliberative models are still untouched. Here, the question of defining a meaningful class of deliberations is an important one. One potential model allows for agents to choose from a variety of noisy deliberations, trading off accuracy against cost. Another model allows for agents to do multiple stages of

deliberation, for example, getting tighter and tighter bounds on their true value. Still another allows for the possibility of one agent deliberating about another agent's value (so called "strategic-deliberation"). Almost nothing is known about dominant-strategy mechanism design in these settings.

Acknowledgments

The authors would like to thank the Institute for Advanced Studies at the Hebrew University of Jerusalem, which hosted the authors as part of the Institute's algorithmic game theory special semester, and at which much of this work was done.

References

- Bergemann, D., and Valimaki, J. 2002. Information acquisition and efficient mechanism design. *Econometrica* 70(3):1007–1033.
- Bikhchandani, S. 2009. Information acquisition and full surplus extraction. *Journal of Economic Theory*.
- Blumrosen, L., and Holenstein, T. 2008. Posted prices vs. negotiations: an asymptotic analysis. In *ACM Conference on Electronic Commerce (ACM-EC)*, 49.
- Cavallo, R., and Parkes, D. C. 2008. Efficient metadeliberation auctions. In *Conference on Artificial Intelligence (AAAI)*, 50–56.
- Chawla, S.; Hartline, J.; Malec, D.; and Sivan, B. 2010. Sequential posted pricing and multi-parameter mechanism design. In *STOC*, 311–320.
- Compte, O., and Jehiel, P. 2005. Auctions and information acquisition: Sealed-bid or dynamic formats? Levine's Bibliography 78482800000000495, UCLA Department of Economics.
- Cramer, J.; Spiegel, Y.; and Zheng, C. 2003. Optimal selling mechanisms with costly information acquisition. Technical report.
- Hartline, J., and Karlin, A. 2007. Profit maximization in mechanism design. In Nisan, N.; Roughgarden, T.; Tardos, É.; and Vazirani, V., eds., *Algorithmic Game Theory*. Cambridge University Press. chapter 13, 331–362.
- Larson, K., and Sandholm, T. 2004. Strategic deliberation and truthful revelation: an impossibility result. In *ACM Conference on Electronic Commerce (ACM-EC)*, 264–265.
- Larson, K. 2006. Reducing costly information acquisition in auctions. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1167–1174.
- Myerson, R. 1981. Optimal auction design. *Mathematics of Operations Research* 6:58–73.
- Thompson, D. R. M., and Leyton-Brown, K. 2007. Valuation uncertainty and imperfect introspection in second-price auctions. In *Conference on Artificial Intelligence (AAAI)*, 148–153.
- Thompson, D. R. M., and Leyton-Brown, K. 2011. Dominant-strategy auction design for agents with uncertain, private values. In *Conference on Artificial Intelligence (AAAI)*, 745–750.