

Valuation Uncertainty and Imperfect Introspection in Second-Price Auctions

David R. M. Thompson and Kevin Leyton-Brown
Department of Computer Science, University of British Columbia
2366 Main Mall, Vancouver, B.C., Canada, V6T 1Z4
{daveth, kevinlb}@cs.ubc.ca

Abstract

In auction theory, agents are typically presumed to have perfect knowledge of their valuations. In practice, though, they may face barriers to this knowledge due to transaction costs or bounded rationality. Modeling and analyzing such settings has been the focus of much recent work, though a canonical model of such domains has not yet emerged. We begin by proposing a taxonomy of auction models with valuation uncertainty and showing how it categorizes previous work. We then restrict ourselves to single-good sealed-bid auctions, in which agents have (uncertain) independent private values and can introspect about their own (but not others') valuations through possibly costly and imperfect queries. We investigate second-price auctions, performing equilibrium analysis for cases with both discrete and continuous valuation distributions. We identify cases where every equilibrium involves either randomized or asymmetric introspection. We contrast the revenue properties of different equilibria, discuss steps the seller can take to improve revenue, and identify a form of revenue equivalence across mechanisms.

Introduction

Imagine deciding to bid for a particular car at an auction, and trying to determine what price you would be prepared to pay. You would probably start out very uncertain—is the new car worth more than \$10,000? \$15,000? \$18,526? You would probably have actions available to you that could increase your level of certainty, such as test driving the car, checking online reviews, and so on. But these actions would consume quite a lot of effort: pretty quickly you might feel compelled to make up your mind one way or the other, even if you were not yet entirely certain of your valuation.

Similarly, imagine the problem of submitting a bid on behalf of your new startup company in an online advertisement auction. The value of a customer's click on your ad would not be obvious: at the moment of entering the ad you might not yet have accurate information about demographics, conversion rates, and so on. As before, it would be possible to collect additional information that would allow you to make a better decision; also as before, the cost of gathering such information would probably lead you to submit your bid even despite some residual uncertainty about your valuation.

Both of these scenarios illustrate a fact that most canonical models of auctions fail to capture: bidders are often un-

certain about their own valuations. While it is often possible for bidders to introspect (or otherwise gather information) in order to reduce this uncertainty, doing so is costly. Furthermore, such introspection is usually imperfect, in the sense that it does not entirely eliminate uncertainty. Bidders must be prepared to place bids despite residual uncertainty. In this paper, we explore the problem of building formal game-theoretic models to describe such settings.

In what follows, we define terminology and then offer a taxonomy of auction models in which bidders are uncertain about their valuations. We use this taxonomy to survey related work from the literature. We propose a novel auction model that aims to capture settings like those described above. Applying this model to the special case of second-price auctions, we offer several theoretical results, identifying equilibria and making revenue comparisons.

Terminology

It is common in auction theory to refer to an agent's *type*, meaning both the agent's private information (or *signal*) and the agent's *valuation*. Even in the classical literature there are settings in which this is not reasonable, such as common values: here agents all have the same valuation for the good, but they are uncertain about what it is and all have (potentially) different private signals. We will follow Bergemann and Morris (2006) in dividing type into *belief type* (private information) and *payoff type* (valuation).

We must also be careful with the notions of *ex ante*, *ex interim* and *ex post*. Over the course of an auction, an agent might receive several signals, updating his beliefs after each. In such settings, it isn't immediately clear which belief type *ex interim* should refer to. We will use *ex interim* to refer to all of them (so *ex interim* individually rational becomes a stronger condition: no sequence of signals should cause an agent to strictly prefer to drop out of the auction), while *ex ante* refers to taking an expectation with respect to an agent's beliefs given no private information. By *ex post* we refer to taking an expectation conditioned on all agents' final belief types (so *ex post* efficiency means maximizing the expected social welfare). We will introduce *ex interim perfect* to refer to perfect knowledge of a single agent's valuation, and *ex post perfect* to refer to such knowledge of all agents' valuations (so *ex post* perfect efficiency means maximizing the actual social welfare).

Definition 1 (Deliberation) *An action that updates an agent's beliefs.*

Parameter	Possible values
Valuation distribution:	Independent, common, interdependent
Secrecy:	Secret, non-secret
Perfection:	Perfect certainty, residual uncertainty
Volatility:	Volatile, non-volatile
Costliness:	Costly, free
Limitations:	Limited, unlimited
Separability:	Separable, inseparable

Table 1: The features of our taxonomy

Because so-called “deliberative agents” must make decisions about how to deliberate and how to bid conditioned on the results of that deliberation, the underlying game is extensive-form. We distinguish two types of deliberation.

Definition 2 (Introspection) *A deliberation where the signal provides information about the deliberator’s valuation.*

Definition 3 (Strategic deliberation; Larson and Sandholm (2001b)) *A deliberation where the signal depends on another agent’s valuation.*

Taxonomy

Deliberating agents have been studied under a variety of assumptions, models and names. In order to survey the literature, we first introduce a unifying taxonomy with a number of free parameters (see Table 1).

Definition 4 (Valuation distribution) *If all agents’ valuations are independent, we call the setting independent value. If every agent has the same valuation, we call the setting common value. We call all other cases interdependent value.*

Definition 5 (Secret) *In secret settings, the only deliberations agents can perform are introspections.*

Notably, researchers have studied independent non-secret values (for example, Larson (2006)).

Definition 6 (Perfection) *A setting with perfect certainty requires agents to discover their exact valuation. Otherwise, the setting allows residual uncertainty.*

One example of a perfect setting is when agents must solve a constrained optimization problem (e.g., vehicle scheduling and routing), to discover how to use the good and thus its value. Without a feasible solution, the agent can’t consume the good. The seller can also impose perfection by requiring bidders to commit to such a solution while bidding.

Definition 7 (Volatile) *In volatile settings, deliberation can change an agent’s underlying valuation.*

Like perfection, this can be motivated by the example of agents solving a constrained optimization problem. Feasible solutions tell the agent how to use the good and thus its value. By finding another solution, the agent might be able to update this value. Perfection is a special case of volatility where agents’ valuations are 0 until they perform a perfect introspection.

Definition 8 (Costly) *Performing costly deliberations costs an agent some utility. Costly settings involve costly deliberations.*

Definition 9 (Limited) *Limited settings have hard constraints about when any particular deliberation is possible.*

For example, agents might be limited as to how many deliberations are possible during a particular stage.

Definition 10 (Separability) *Where no deliberations are possible once the mechanism begins. (This is a special case of limited settings.)*

All sealed-bid mechanisms are trivially separable. Some staged auctions are also separable. (For example, in some auctions for antiques, bidders are only free to examine the goods until the English auction begins.)

Because the problem of deliberation is equivalent whether the source of the information is internal or external, there is an equivalence between costly deliberation (with secrecy and non-volatility) and costly communication to a proxy.

Previous work

The earliest work on modeling agents with valuation uncertainty is probably Wilson’s (1967) work on common values. This problem has been extensively studied since then, but primarily in settings in which agents are not deliberative. Another related literature considers tractably expressing bidder preferences in multi-good auctions (Ausubel & Milgrom, 2002; Blumrosen & Nisan, 2005; Nisan & Segal, 2005; Sandholm & Boutilier, 2006).

We are specifically interested in deliberative agents. We will classify the existing work on this topic according to our taxonomy, though the classification isn’t always precise: some of the existing work spans and contrasts classes, and some results are more general than their authors state.

A number of researchers have considered the problem of deliberative agents in settings requiring perfect certainty. Parkes (2005) stated that handling residual uncertainty “appears beyond the scope of current methods (either analytic or computational), even for simple auctions such as an ascending-price auction for a single item.” His paper focused on the dual problem to our own: auction and proxy design with costly communication. Parkes demonstrated that incremental revelation mechanisms can achieve the same allocations as direct mechanisms with lower communication costs. Cremer *et al.* (2003) described how to extract full surplus from agents who commit to participating in the auction prior to deliberating, using a sequential auction based on optimal search. Larson and Sandholm (2005) proved that no (interesting) mechanism exists in which agents have no incentive to strategically deliberate or to mislead the seller and the mechanism does not depend on knowledge of agents’ deliberation costs and limits. The proof assumes costly deliberations, volatility (or non-volatile, perfect certainty), and non-secrecy. Larson (2006) described an optimal search auction in which agents never strategically deliberate or mislead the seller. Larson and Sandholm (2001a; 2001b) provided a very general model for costly and/or limited deliberations in auctions and showed that under costly deliberation models and in a number of common auction types (Vickrey, English, Dutch, first price and (for multiple goods) VCG) bidders perform strategic deliberation in equilibrium.

There has been a small number of papers analyzing deliberative agents in settings that allow residual uncertainty. (Blumrosen & Nisan, 2002) showed that in auctions with severely limited communication (equivalent to severely limited deliberation), social welfare can be improved by using asymmetric proxies (equivalent to asymmetric deliberation strategies). Bergemann and Valimaki (2002) showed that in independent secret settings, VCG is *ex post* efficient and

Paper	Values	Secret	Perfect	Volatile	Costly	Limits	Separable
Cremer et al (2003)	ind	yes	yes	no	yes	no	no
Parkes (2005)	all	yes	both	no	yes	no	no
Larson & Sandholm (2005)	ind	no	yes	yes	yes	no	yes
Larson (2006)	ind	no	yes	no	yes	no	no
Larson & Sandholm (2001a)	all	no	no	yes	no	yes	both
Larson & Sandholm (2001b)	all	no	no	yes	yes	no	both
Blumrosen & Nisan (2002)	ind	yes	no	no	no	yes	yes
Bergemann & Valimaki (2002)	all	partial	no	no	yes	no	both
Persico (2000)	inter	partial	no	no	yes	yes	yes
Larson & Sandholm (2003)	all	no	no	yes	yes	yes	yes
Larson & Sandholm (2004)	ind	no	no	yes	yes	yes	yes
Compte & Jehiel (2001)	ind	yes	no	no	yes	no	both
Sandholm (2000)	ind	yes	no	no	yes	no	yes
Rasmusen (2006)	ind	yes	no	no	yes	yes	no
Our Paper	ind	yes	no	no	yes	yes	yes

Table 2: Classification of previous work on deliberative agents according to our taxonomy

provides *ex ante* incentives to deliberate. For common value settings, they showed that no mechanism can achieve both properties. Larson and Sandholm (2003) demonstrated that in Vickrey auctions the ratio of computing costs between the social optimum and the worst case Nash equilibrium is unbounded, but that this can be improved by mechanism design. Larson and Sandholm (2004) described an experimental evaluation of deliberations on Vickrey auctions in non-secret settings. In simulation, no strategic deliberations occurred when agents had symmetric cost functions. Compte and Jehiel (2001) showed that sellers can get more revenue from Japanese-like ascending auctions than from second price auctions. Sandholm (2000) demonstrated that in second price auctions, bidding one’s true expected valuation is a dominant strategy for risk neutral bidders but not for risk averse ones. He also demonstrated that second price auctions do not always have dominant deliberation strategies when deliberation is costly. Rasmusen (2006) presented residual uncertainty as a motivation for sniping on eBay auctions: buyers don’t have time to deliberate in the last seconds of an auction, but don’t deliberate earlier because of the cost. Persico (2000) demonstrated that, when valuations are correlated (i.e. secrecy isn’t absolute), the value of information is greater in first price auctions than in second price auctions, because bidding strategies are conditioned on beliefs about opponents’ valuations.

Model

In our work we make the strong assumption of separability. Existing work shows that staged auctions are better than sealed bid for maximizing revenue and minimizing deliberation costs (Compte & Jehiel, 2001; Parkes, 2005; Bergemann & Valimaki, 2002). However, separability can arise when deliberations are too time consuming to occur between bidding rounds: e.g., the deliberation might involve analysis of a broadcast spectrum market or drilling control wells in an offshore oil patch. Furthermore, many other real-world auctions are inherently separable (e.g., sealed-bid). We make a number of less unusual assumptions: independence, secrecy, symmetry and non-volatility.

Our setting is a six-tuple: (N, f, Q, A, p, c) . N is a set of agents, each of whom has a valuation v_i drawn from distribution f , which has support on the interval $[v, \bar{v}]$. Q is the set of possible introspections, from which each agent i chooses one, $q_i \in Q$; A is the set of possible signals the agent can receive, according to conditional probability dis-

tribution $p(a_i|q_i, v_i)$; and $c(q_i, a_i)$ is the cost of the signal $a_i \in A$. Since agents can choose not to deliberate, we assume the existence of a special deliberation q_\emptyset that costs nothing and is uninformative. No agent (including the seller) knows how any other agent deliberated. This model restricts the model of Larson and Sandholm (2001b) to the case of secrecy, non-volatility and separability. It is without loss of generality with respect to a number of important features.

Proposition 1 *This model is equivalent to one in which separability holds and agents can perform multiple introspections.*

Proof Sketch. An agent has a choice of deliberations given the results of previous introspections, yielding a tree structure (that can include restrictions on what sequences of introspections are possible) with some nodes representing deliberation actions and others random signals from Nature. When the game is separable we can construct new deliberation actions that correspond to deliberation policies in the original game, making multiple introspections unnecessary. ■

Proposition 2 *This model is equivalent to one in which agents can start off with useful information.*

This can be achieved by adding costless deterministic or stochastic introspections to the tree before any of an agent’s action nodes.

Proposition 3 *This model is equivalent to one in which the cost of introspection depends on an agent’s true valuation.*

In the latter setting an agent would condition his beliefs on his observations of these costs, which should be part of the signal. Thus, for every distinct signal, there is a known cost.

Proposition 4 *Risk neutral agents with residual uncertainty should always bid as though their expected valuations were their exact known valuations, as long as independence, secrecy and separability hold.*

Proof. Having won the good, agent i would be indifferent between the good and a fixed transfer of $\mathbb{E}[v_i|a_i, i \text{ wins auction}]$. However, when independence and secrecy hold, the event of winning the auction is uninformative about v_i (because all the seller knows about v_i was revealed by i). Thus the agent can bid as though $v_i = \mathbb{E}[v_i|a_i]$, once he decides not to perform any more deliberations. From separability, the agent cannot bid before this point. ■

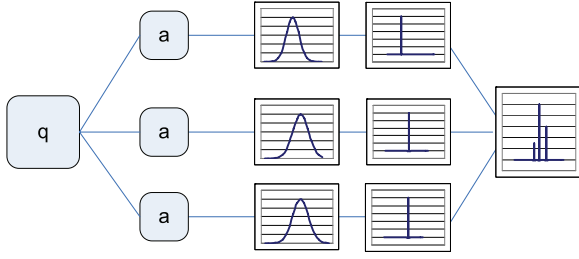


Figure 1: Induced valuation distribution of a deliberation

	q^*	q_\emptyset
q^*	.25 - c, .25 - c, .25	.25 - c, .25, .25
q_\emptyset	.25, .25 - c, .25	0, 0, .5

Figure 2: Induced normal-form game for two bidders in a second-price auction who have valuations of 0 or 1 with equal probability. The third payoff is the seller's revenue.

Because agents facing residual uncertainty bid as though their signals informed them of an exact valuation, we can define an *induced valuation distribution*: the valuation distribution an agent acts as though he had, given his deliberation. The induced valuation distribution f_q of deliberation q is defined as

$$f_q(v'_i) = \sum_{a_i \in A} p(a_i|q) \delta(\mathbb{E}[v_i|a_i, q] = v'_i),$$

where $\delta(e)$ is the Kronecker delta: $\delta(e) = 1$ iff e is true and $\delta(e) = 0$ otherwise.

Second price auctions

In this section we find and characterize the equilibria of second price auctions for some simple value distributions and sets of possible introspections. We show cases where there are two qualitatively different classes of equilibria having different revenue and efficiency characteristics. Throughout, we assume (for reasons of simplicity and tractability) that agents are risk neutral and always follow the dominant strategy of bidding their expected valuations truthfully (see Proposition 4). This means that the second price auction is trivially *ex post* efficient (for any set of deliberation strategies), though not necessarily *ex post* perfect efficient.

We begin by the case of two bidders with valuations equally likely to be either 0 or 1. (We will subsequently scale up the number of bidders.) They can choose between a perfect introspection $q_i = q^*$ having fixed cost c or not introspecting, $q_i = q_\emptyset$. This yields the induced normal form game given in Figure 2.

We will only consider the non-trivial values of $0 < c < 0.25$: in the other cases information is either free or prohibitively expensive. This game has a number of noteworthy features: The value of information for deliberation q^* depends on the other agent's deliberation action, and equilibrium involves either randomization or asymmetry. These two flavors of equilibria lead to very different revenues. Under the asymmetric pure strategy equilibria, the deliberation cost doesn't affect bidding strategies or revenue. Under a symmetric mixed strategy equilibrium (each bidder introspecting with probability $p = 1 - 4c$), the probability of introspection will vary continuously with costs, meaning that revenue will vary with costs as well. Both agents weakly prefer the pure strategy equilibria, which also max-

imize the social welfare, but face a coordination problem to reach them. If they follow symmetric strategies, they will sometimes miscoordinate and so expected social welfare decreases. By providing the agents with a correlated signal to condition their deliberations on, the seller would enable the agents to coordinate (i.e. correlated equilibrium). In this particular case, this would also cause the seller's revenue to decrease, but that is an artifact of having so few bidders.

We can show that these two sets of equilibria also exist for versions of this game involving larger numbers of agents.

Lemma 5 *In every pure strategy equilibrium of this game exactly $\min(n - 1, \lfloor -\log_2(c) \rfloor - 1)$ agents introspect while the other agents bid their (uninformed) expected valuations.*

Proof Sketch. Let k denote the number of agents that introspect. We can write expressions for the expected utility of any agent given k :

$$\mathbb{E}[u_i|k, q_i = q_\emptyset] = \begin{cases} (1/2)(1/2)^k & k = n - 1 \\ 0 & o.w. \end{cases}, \text{ and}$$

$$\mathbb{E}[u_i|k, q_i = q^*] = \begin{cases} (1/2)^k - c & k = n \\ (1/2)(1/2)^k - c & o.w. \end{cases}.$$

A set of equilibria exist with k agents introspecting iff:

$$\mathbb{E}[u_i|k, q_i = q_\emptyset] \geq \mathbb{E}[u_i|k + 1, q_i = q^*], \text{ and}$$

$$\mathbb{E}[u_i|k - 1, q_i = q_\emptyset] \leq \mathbb{E}[u_i|k, q_i = q^*].$$

Solving this system of inequalities for k gives that $\min(n - 1, \lfloor -\log_2(c) \rfloor - 1)$ agents introspect. ■

Theorem 6 *When introspection is costly, there are symmetric settings where the second price auction has no symmetric pure strategy Nash equilibrium.*

This follows from Lemma 5 since there is no equilibria for $k = n$ or $k = 0$.

Lemma 7 *This game has a symmetric mixed strategy where every agent introspects with probability p .*

Proof Sketch. Each agent's expected utility under a symmetric mixed strategy profile where they each introspect with probability p is

$$\sum_{j=0}^{n-1} B(j; n, p) (p \mathbb{E}[u_i|j + 1, q^*] + (1 - p) \mathbb{E}[u_i|j, q_\emptyset]),$$

where $B(k; n, p)$ is the binomial distribution. Since the game is symmetric, it must have a symmetric Nash equilibrium at some p^* , which must satisfy

$$\mathbb{E}[u_i|q_i = q^*, p^*] = \mathbb{E}[u_i|q_i = q_\emptyset, p^*].$$

An analytic expression for the equilibrium can be found by solving the system of equations for p given a particular n . ■

Theorem 8 *In second price auctions with costly introspection, the pure strategy equilibria can yield different revenue from the symmetric equilibria.*

Proof Sketch. The seller's expected revenue under the pure strategy equilibria is:

$$\mathbb{E}[\text{revenue}|k] = \begin{cases} 1 - 2^{-k} - k2^{-k} & k = n \\ 1 - 2^{-k} - k2^{-k-1} & k = n - 1 \\ 1 - 2^{-k-1} - k2^{-k-1} & o.w. \end{cases}.$$

Under the symmetric (mixed) strategy equilibrium strategy profile p , the seller's expected revenue is:

$$\sum_{j=0}^n p^j (1 - p)^{n-j} \binom{n}{j} \mathbb{E}[\text{revenue}|k = j].$$

Under the equilibria above, these two quantities are not equal. ■

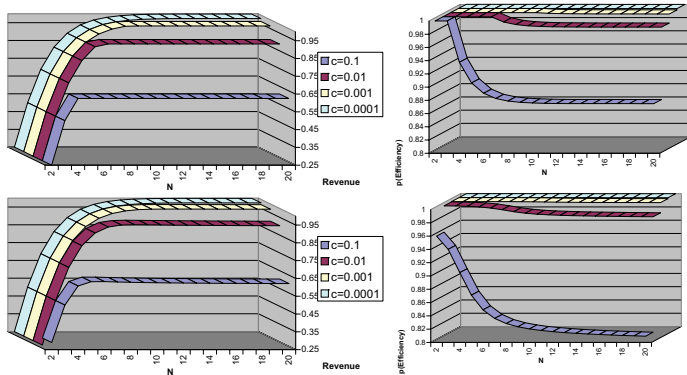


Figure 3: Exact expected revenue under pure (top) and symmetric (bottom) strategies




Figure 4: Exact probability of *ex post* perfect efficiency under pure (top) and symmetric (bottom) strategies

Using these analytic expressions for equilibrium revenue, we can compare the revenue of the different classes of Nash equilibria as n varies. The qualitative revenue properties of the equilibria are consistent with the intuition from two bidder example (see Figure). For pure strategy equilibria, all costs have the same revenue for low values of n . When n becomes too large, adding more agents will have no effect on the number of agents that introspect, causing the revenue to plateau. Notably, this threshold only increases linearly for exponential decreases in costs, so that this limit applies to relatively small numbers of bidders, even when the cost is orders of magnitude smaller than the agents' valuations. For mixed strategy equilibria, the revenue decreases continuously as costs increase. As in the two-bidder case, symmetric mixed strategy equilibria involve a probability of miscoordinating which increases with the number of agents. As n gets large the probability of introspection will tend to 0, and the expected revenue will tend to the uninformed expected valuation (though in this case it does so extremely slowly).

Although the second price auction is always *ex post* efficient, the equilibrium reached will affect the probability of *ex post* perfect efficiency (the probability of actually maximizing social welfare). Mixed strategy equilibria tend to have worse probabilities of *ex post* perfect efficiency (see Figure 4) because of the risk of miscoordination.

The analytic expression for the probability of *ex post* perfect efficiency of pure strategy equilibria is

$$e_p(k, n) = 1 - \left(2^{-k} \sum_{j=1}^{n-k-1} 2^{k-n} \binom{n-k}{j} \frac{j}{n-k} \right).$$

The analytic expression for the probability of *ex post* perfect efficiency of mixed strategy equilibria is

$$e_m(p, n) = \sum_{k=0}^n \binom{n}{k} p^k (1-p)^{n-k} e_p(k, n).$$

Now we consider a new problem: limited but free deliberations. Let agents have valuations drawn from a uniform distribution on the interval $[0, 1]$. Agents can perform one deliberation of the following type: for any value on the interval $[0, 1]$, they can discover (by a 1-bit signal) whether their valuation is above or below it. Since the action space of deliberations is continuous (and purely cooperative, though this is a coincidence due to $n = 2$), we graph the (continuous) induced normal form in Figure 5. Again,




Figure 5: Exact induced normal form of the two-bidder limited deliberation case. The peaks (pure equilibria) do not fall on the $q_1 = q_2$ line, so they are not symmetric.

the only pure strategy equilibria— $[q_1 = 1/3, q_2 = 2/3]$ and $[q_1 = 2/3, q_2 = 1/3]$ —are asymmetric, even though deliberations are no longer costly. This finding mirrors those of Blumrosen and Nisan (2002), that agents benefit from asymmetry when severe limitations apply (though in this case, the asymmetric meaning of the 1-bit signal is chosen by the bidder rather than by the mechanism or proxy). The agents could again randomize across these two equilibria, though again this would cause them to lose expected utility through miscoordination. By ordering the bidders by weakly increasing q_i , we can write an expression for the expected utility of any number of agents under any pure strategy profile (see Figure 6).

Revenue relationships and bounds

This section contains theoretical results: we will show a limited application of revenue equivalence, an upper bound on the revenue of separable auctions with costly deliberation and an impossibility result.

Theorem 9 *In all symmetric, separable IPV settings where for every deliberation q the induced valuation distribution f_q is differentiable on the interval $[\underline{v}, \bar{v}]$, all *ex post* efficient auctions are revenue equivalent under symmetric equilibria.*

Proof. Any convex combination of the induced valuations will itself be differentiable on the interval $[\underline{v}, \bar{v}]$. If all agents play the same mixed deliberation strategy s , they will all have the same induced valuation distribution, f_s . Fixing this deliberation strategy profile, the problem of how to bid is strategically equivalent to how to bid in the same auction if agents had perfectly-known valuations drawn from distribution f_s . Since the problems are equivalent, the equilibria of the bidding subgame will be the same and revenue equivalence will apply to any efficient auction for a fixed f_s . Since agents get the same expected utility from s regardless of the efficient auction used, the same deliberation strategies will be equilibria for all efficient auctions. ■

We can also show an extremely general (though loose) revenue bound, which applies to all separable auction settings and all allocation rules.

Theorem 10 *In any *ex ante* individually-rational, separable auction, the expected revenue is bounded from above by $\bar{v} - \sum_i \mathbb{E}[c_i]$ where $\mathbb{E}[c_i]$ is i 's expected cost of deliberation.*

$$\mathbb{E}[u_i|q_{1..n}] = \frac{1}{2} \left[\sum_{j=1}^{i-1} \left((q_i - q_j)(1 - q_i)(1 - q_j) \left(\prod_{k=i+1}^n q_k \right) \left(\prod_{k=j+1}^{i-1} q_k \right) \right) + \left(\prod_{k \neq i} q_k \right) \left\{ \begin{array}{ll} (1 - q_i)(1 + q_i) + q_i^2 - q_{n-1} & i = n \\ (1 - q_i)(1 + q_i - q_n) & o.w. \end{array} \right. \right]$$

Figure 6: The expected utility of agent i in the continuous case

Proof. Let $\gamma_i(s)$ be the marginal probability of i receiving the good given strategy profile s . Let $t_i(s)$ be i 's expected transfer to the seller given s .

$$\begin{aligned} \mathbb{E}[u_i|s] &= \gamma_i(s)\mathbb{E}[v_i|s, i \text{ wins}] - t_i(s) - \mathbb{E}[c_i|s_i] \\ t_i(s) &\leq (\gamma_i(s)\bar{v} - \mathbb{E}[c_i|s_i]) \\ \sum_{i \in N} t_i(s) &\leq \sum_{i \in N} (\gamma_i(s)\bar{v} - \mathbb{E}[c_i|s_i]) \\ \mathbb{E}[\text{revenue}|s] &\leq \bar{v} - \sum_i \mathbb{E}[c_i|s_i] \quad \blacksquare \end{aligned}$$

Although previous work has already shown that Vickrey auctions don't necessarily have dominant strategies in the case where deliberations are costly, this bound allows us to generalize that finding into a broader impossibility result.

Corollary 11 *No budget-balanced, individually-rational, separable mechanism can have a dominant strategy that involves an unbounded number of agents performing costly deliberations.*

Proof. There must exist some n where $nc > \bar{v}$, where c is the cost of the cheapest deliberation. If more than n agents perform a deliberation with cost $\geq c$, then the seller's expected revenue becomes negative. \blacksquare

Theorem 12 *In ex post efficient, separable auctions with independent secret values, the maximum value of information for any deliberation q falls off exponentially in the number of agents performing it (denoted by k).*

Proof. Suppose we only allowed agents who performed deliberation q to bid. For any signal a_i , i 's probability of receiving the good is proportional to $F_q^{k-1}(\mathbb{E}[v_i|a_i])$, where $F_q(v)$ is the cumulative induced valuation distribution of q . Hence, i 's expected surplus (and the value of information) must fall off exponentially in k . Allowing agents to bid without performing deliberation q weakly decreases the expected utility of q (by reducing the probability of receiving the good and increasing the expected payment) and weakly increases the value of not performing q . \blacksquare

Conclusion

We have expanded on previous work that shows that the problem of how to deliberate prior to bidding adds a significant extra layer of strategic complexity to auctions. We showed that limited and costly deliberation can impose a coordination problem on bidders, that the chance of miscoordination has an impact on revenue and the probability of efficiency, and that these issues are present even when deliberation costs are small relative to valuations. Future work could relax assumptions about distribution, secrecy and separability. For example, a model could describe the costs to the seller (and other bidders) of waiting for slow deliberations, allowing a continuous trade-off of the speed of sealed auctions against the lower deliberation costs and higher revenue of staged auctions.

References

- Ausubel, L., & Milgrom, P. (2002). Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1(1), 1–42.
- Bergemann, D., & Morris, S. (2006). Robust implementation: The case of direct mechanisms. Coweles foundation discussion paper 1561.
- Bergemann, D., & Valimaki, J. (2002). Information acquisition and efficient mechanism design. *Econometrica*, 70(3), 1007–1033.
- Blumrosen, L., & Nisan, N. (2002). Auctions with severely bounded communication. *FOCS* (pp. 406–415).
- Blumrosen, L., & Nisan, N. (2005). On the computational power of iterative auctions. *ACM-EC* (pp. 29–43).
- Compte, O., & Jehiel, P. (2001). Auctions and information acquisition: Sealed-bid or dynamic formats. working paper.
- Cremer, J., Spiegel, Y., & Zheng, C. (2003). *Optimal selling mechanisms with costly information acquisition* (Technical Report). Working paper, August.
- Krishna, V. (2002). *Auction theory*. Academic Press.
- Larson, K. (2006). Reducing costly information acquisition in auctions. *AAMAS* (pp. 1167–1174).
- Larson, K., & Sandholm, T. (2001a). Computationally limited agents in auctions. *AGENTS*.
- Larson, K., & Sandholm, T. (2001b). Costly valuation computation in auctions. *TARK* (pp. 169–182).
- Larson, K., & Sandholm, T. (2003). Miscomputing ratio: The social cost of selfish computing. *AAMAS* (pp. 273–280).
- Larson, K., & Sandholm, T. (2004). Experiments on deliberation equilibria in auctions. *AAMAS*.
- Larson, K., & Sandholm, T. (2005). Mechanism design and deliberative agents. *AAMAS* (pp. 650–656).
- Nisan, N., & Segal, I. (2005). The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*.
- Parkes, D. (2005). Auction design with costly preference elicitation. *Annals of Mathematics and Artificial Intelligence*, 44(3), 269–302.
- Persico, N. (2000). Information acquisition in auctions. *Econometrica*, 68(1), 135–148.
- Rasmusen, E. (2006). Strategic implications of uncertainty over ones own private value in auctions. *Advances in Theoretical Economics*, 6(1).
- Sandholm, T. (2000). Limitations of the Vickrey auction in computational multiagent systems. *International Journal of Electronic Commerce*, 4(3), 107–129.
- Sandholm, T., & Boutilier, C. (2006). Preference elicitation in combinatorial auctions. In Cramton, Shoham and Steinberg (Eds.), *Combinatorial auctions*, chapter 10.
- Wilson, R. (1967). Competitive bidding with asymmetric information. *Management Science*, 13(11), 816–820.
- Wilson, R. (1969). Competitive bidding with disparate information. *Management Science*, 15(7), 446–448.