Valuation Uncertainty and Imperfect Introspection in Second-Price Auctions

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.

1 Introduction

Imagine deciding to bid for a particular car at a car auction, and trying to determine what price you would be prepared to pay. At first 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 100% 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. Again, it would be possible to collect additional information that would allow you to make a better decision; again, also, the cost of gathering such information would probably lead you to submit your bid even despite some residual uncertainty about the value of a click.

Both of these scenarios illustrate a fact that most canonical models of auctions fail to capture: bidders are often uncertain 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 thus be prepared to place bids even in the face of 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 a variety of related work in the literature. We propose a novel auction model that aims to reflect the settings described above. Applying this model to the special case of second-price auctions, we offer several theoretical results, identifying equilibria and making revenue comparisons.

1.1 Terminology

It is common in auction theory to refer to an agent's *type*, by which is meant both the agent's private information (or *signal*) and the agent's *valuation*. However, even in the classical literature there are settings in which this modeling assumption is not reasonable. A key example is the common value setting: here agents all have the same valuation for the good, but they are uncertain about what this valuation is and all have (potentially) different private signals about it. In our work 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 purchase several signals, changing and refining his beliefs after each one. This agent would have several different belief types, and it isn't immediately clear of these which *ex interim* refers to. We will take the position that *ex interim* refers to all of them (so that *ex interim* individually rational becomes a much stronger condition: at no stage should any signal should cause an agent to prefer to drop out of the auction), while *ex ante* refers to an agent's beliefs given no private information. We will use *ex post* to refer to the aggregation of all agents' final belief types (so that *ex post* efficiency means maximizing the expected social welfare). Unfortunately, this doesn't leave us with any ways of talking about payoff types. We will introduce *ex interim perfect* to refer to perfect knowledge of a single agent's valuation, and *ex post perfect* to refer to the aggregation of all agents' valuations (so that *ex post* perfect efficiency means maximizing the actual social welfare).

Definition 1 (Deliberation) *When the agent selects a signal, we call this a* deliberation. *We call agents capable of this* deliberative agents.

Because agents can choose from a menu of potential signals, the resulting game is extensive form: Agents must make decisions about how to deliberate and how to bid conditioned on that the results of that deliberation. We will distinguish two special types of deliberation based on the information revealed.

Definition 2 (Introspection) *An* introspection *is a deliberation where the information is independent of any other agent's valuation given the deliberator's valuation.*

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

Table 1: The axes of our taxonomy

Definition 3 (Strategic deliberation) *A* strategic deliberation *(coined by Larson and Sandholm (2001)) is the opposite of introspection; the information is independent of the deliberator's valuation given all the other agents' valuations.*

1.2 Taxonomy

The existing literature has studied the problem of deliberating agents under a variety of assumptions and models (and names). In order to study and compare the existing 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 (Privacy) *If agents are only capable of performing introspections, we call the setting* private.

Notably, researchers have studied independent non-private values (for example, Larson (2006)) and interdependent private values (for example, Persico (2000)).

Definition 6 (Perfection) *A* perfect introspection *is one that reveals an agent's exact valuation, while an* imperfect introspection *does not. Many mechanisms or settings require* perfect certainty*: agents must perform a perfect introspection before bidding or receiving or consuming the good. Settings that do not are said to allow* imperfect certainty *(or* residual uncertainty *(Rasmusen, 2006)).*

One justification for saying that a setting requires is perfect certainty is that agents must solve a constrained optimization problem (for example, vehicle scheduling and routing), to discover how to use the good and what value it will be worth for that use. Without a feasible solution, the agent can't consume the good. Perfect certainty could also be a property of the mechanism: the seller could require that the agents reveal and commit to a particular solution before bidding or before the auction closes.

Definition 7 (Volatility) *We call any setting* volatile *if there is some possibility of a deliberation action changing the agent's valuation rather than just his belief type.*

Like perfection, this can be motivated by the example of agents solving a constrained optimization problem, to determine how to use the good. Feasible solutions can tell the agent how to use the good and the exact utility for that use. By finding another solution, the agent might be able to get a different exact utility. Notably, perfect certainty settings are a special case of volatility where agents' valuations are 0 until they perform a perfect introspection.

Definition 8 (Costliness) *If the agent has disutility for performing a deliberation, we call that deliberation* costly. *Settings are also called costly if they include even the possibility of costly deliberations.*

We will model costs as another term in quasilinear utility:

$$u_i = v_i(x) - t_i - c_i(q)$$

where $v_i(x)$ is an arbitrary function for outcome x, t_i is the transfer the agent makes to the seller and $c_i(q)$ is an arbitrary function for deliberation policy q. Costly deliberation is not equivalent to an entry fee because it is not transferred to the seller.

Definition 9 (Limited) *We say that a setting has* limited deliberation *if agents face 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 of the mechanism.

Definition 10 (Separability) Separability *is a special case of limited deliberation, where no deliberations are possible once the mechanism begins.*

All sealed-bid mechanisms are trivially separable. Some settings with staged mechanism are also separable. (For example, in auctions for antiques, bidders are often only free to closely 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 privacy and non-volatility) and costly communication to a proxy.

We will make one strong assumption: separability. While the seller can impose separability (for example, by sealed-bid auction), existing work shows that staged auctions perform better than sealed bid, in terms of maximizing revenue and minimizing deliberation costs (Compte & Jehiel, 2001; Parkes, 2005; Bergemann & Valimaki, 2002). However, separability can also occur because deliberations are too time consuming to allow between bidding rounds: the deliberation might involve analysis of a particular market potential for a broadcast spectrum, or drilling control wells to evaluate an offshore oil patch. Further, since separability is a feature of many common real-world auction types, we believe there is value in an analysis of separable auctions. We will also make a number of less unusual assumptions: independence, privacy, symmetry and non-volatility.

Paper	Valuations	Privacy	Perfection	Volatility	Costliness	Limitations	Separability
Cremer et al. (2003)		**private**	perfect certainty	**non-volatile**	**costly**	unlimited	inseparable
Parkes (2005)	**independent**	**private**	perfect certainty	**non-volatile**	**costly**	unlimited	inseparable
Larson and Sandholm (2005)	**independent**	non-private	perfect certainty	volatile	**costly**	unlimited	**separable**
Larson (2006)	**independent**	non-private	perfect certainty	**non-volatile**	**costly**	unlimited	inseparable
Larson and Sandholm (2001)		non-private			**costly**		
Blumrosen and Nisan (2002)	**independent**	**private**	**residual uncertainty**	**non-volatile**	free		
Bergemann and Valimaki (2002)		**private**	**residual uncertainty**	**non-volatile**	**costly**	**unlimited**	
Compte and Jehiel (2001)	**independent**	**private**	**residual uncertainty**	**non-volatile**	**costly**	unlimited	
Sandholm (2000)	**independent**	**private**	**residual uncertainty**	**non-volatile**	**costly**	unlimited	**separable**
Rasmusen (2006)	**independent**	**private**	**residual uncertainty**	**non-volatile**	**costly**	**limited**	inseparable
Persico (2000)	interdependent	**private**	**residual uncertainty**	**non-volatile**	**costly**	**limited**	**separable**
Our Paper	**independent**	**private**	**residual uncertainty**	**non-volatile**	**costly**	**limited**	**separable**

Table 2: Classification of previous work on deliberative agents

1.3 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).

However, our specific interest is in the problem of deliberative agents in auctions. We will classify the existing work on this topic according to our taxonomy, though the classification isn't always precise. In particular, some of the existing work spans and contrasts classes. Also, some theorems and results are likely more general than their stated assumptions.

A number of researchers have considered the problem of deliberative agents in settings requiring perfect certainty. Indeed, 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." The paper from which this quote was taken focuses on the dual problem to our own: auction and proxy design with costly communication. Parkes demonstrated that incremental revelation mechanisms can achieve the same allocation as direct ones 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, agents have no incentive to mislead the seller and the mechanism not depend on knowledge of agents' possible deliberation costs and limits. The proof assumes costly deliberations, volatility (though the proof works equally well for non-volatile settings that require perfect certainty), and non-privacy. Larson (2006) described an optimal search auction in which agents never strategically deliberate or mislead the seller (though the optimal search requires some knowledge of the possible deliberations and costs). Larson and Sandholm (2001) provided a very general model for costly, limited deliberations in auctions and shows that under costly deliberation models and in a number of common auction types (Vickrey, English, Dutch, first price and (for multiple goods) Generalized Vickrey Auction) bidders would perform strategic deliberation in equilibrium. As well, all but Vickrey and English can also have strategic deliberation in limited, uncostly deliberation settings.

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 communications (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 private settings, the VCG mechanism can be *ex post* efficient and provide the *ex ante* incentives to deliberate. For common value settings, they showed that no mechanism can achieve both. They also demonstrated that staged mechanisms require less costly deliberation than direct mechanisms, to implement the same social choice function. Compte and Jehiel (2001) showed that in inseparable settings, a seller can get more revenue from a Japanese-like ascending auction than from a second price auction. Sandholm (2000) demonstrated that in second price auctions, bidding 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 necessarily 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, the value of information is greater in first price auctions than second (because bidding strategies are conditioned on opponents' valuations).

2 Model

Our setting is a six-tuple: (N, f, Q, A, p, c). N is the set of all agents, each of which has a valuation v_i drawn from distribution f (which has support on the interval $[\underline{v}, \bar{v}]$). Q is the set of possible introspections (from which each agent chooses one, q_i) and A is the set of possible signals the agent can receive, according to conditional probability distribution $p(a_i|q_i, v_i)$. $c(q_i, a_i)$ is the cost of the signal. Since the agent is capable of not deliberating, we will assume the existence of a special deliberation $q_\varnothing$ which costs nothing and is totally uninformative. No agent (including the seller) knows how any other agent deliberated.

This model is without loss of generality regarding a number of important features.

Proposition 1 *This model is equivalent to a separable model where agents are able to perform multiple introspections.*

> **Proof Sketch.** The agent has a choice of deliberations given the results of previous introspections, forming a tree structure (including limit of which sequences of introspections are possible) with some nodes being choices of the agent (including $q_\varnothing$) and others being random. The agent could chose their policy on this tree as a single choice (with conditional probabilities and costs), as with collapsing an extensive form game into normal form. This collapsing requires separability. ■

Proposition 2 *This model is equivalent to a model where agents begin with some possibility of (partial) knowledge.*

We can trivially add a free initial introspections to the tree which gives the agent imperfect (or perfect with some probability) information.

Proposition 3 *This model is equivalent to a model where the cost of introspection depends on the true valuation.*

If so the agent would want to condition beliefs on the costs, which should be part of the signal. Thus, for every distinct signal, there is a known cost.

We will show some (trivial) results needed for our later analysis.

Proposition 4 *Risk neutral agents with residual uncertainty should bid as though their expected valuations were their exact known valuations, in any setting with independence, privacy and separability.*

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 privacy hold, the event of winning the auction is uninformative regarding v_i (because the 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 it had decided to perform no more deliberations. Since the setting is separable, the agent cannot bid before this point. ■

Because agents can bid as though their signals informed them of an exact valuation, we can define an *induced valuation distribution*: the value distribution agents act as though they had given their 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.

3 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 competing classes of equilibria with different properties as far as revenue and efficiency are concerned. 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 (as shown by Proposition 1). Because of this, the second price auction is trivially *ex post* efficient (for any set of deliberation strategies), though not necessarily *ex post* perfect efficient.

We will begin by considering an extremely simple example: 2 bidders with valuations of either 0 or 1 (with equal probability). The agents can chose between a perfect introspection, $q_i = q^*$ with fixed cost c or not introspecting, $q_i = q_\varnothing$. This results in the induced normal form game given in Figure 1.

We will only consider the non-trivial values of $0 < c < 0.25$, because otherwise 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 the only equilibria require either randomization or asymmetry. These two sets of equilibria have qualitatively different revenues. Under an asymmetric pure strategy equilibrium, the deliberation is a sunk cost and doesn't affect bidding strategies or revenue. Under a symmetric mixed strategy equilibrium, 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 maximize the social welfare), but face a coordination problem to reach them. If they

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

Figure 1: 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 indicates the seller's revenue.

must follow symmetric strategies, then they risk miscoordinating (with a loss of social welfare). By allowing the agents to condition their deliberations for a correlated equilibrium, the seller would allow them to always reach the pure strategy equilibria. (In this case, the seller's revenue would decrease, but that is an artifact of having so few bidders.)

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

Theorem 5 *In every pure strategy equilibrium of this game exactly* $min(n-1, \lfloor -log_2(c) \rfloor - 1)$ *agents introspect.*

Proof Sketch. We enumerate the set of all pure strategy profiles as follows. Let k denote the number of agents that introspected. We can write expressions for the expected utility of an agent given that pure strategy profile:

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

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:

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

and

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

Solving this system of equations, we can find the values of k which represent the set of all possible pure strategy Nash equilibria. ■

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

We enumerate the set of all symmetric strategy profiles as follows. Let p be the probability of deliberating in any symmetric mixed strategy. We can write an expression for the expected utility of each agent under a symmetric mixed strategy profile:

$$\mathbb{E}[u_i|p] = \sum_{j=0}^{n-1} p^j (1-p)^{n-j-1} \left(\begin{array}{c} n-1 \\ j \end{array} \right) (p\mathbb{E}[u_i|j+1, q_i = q^*] + (1-p)\mathbb{E}[u_i|j, q_i = q_\varnothing]).$$

Since the game is symmetric, it must have a symmetric Nash equilibrium at some p^*, which much satisfy

$$p^* \in \arg\max_p \mathbb{E}[u_i|p].$$

Theorem 7 *In second price auctions with costly introspection, the pure strategy equilibria can yield different revenue from the symmetric equilibria.*

With these equilibria, we can compare revenues. We can write an expression for the expected revenue of the seller under pure strategy profiles:

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

Under symmetric (mixed) strategy profiles, the revenue is

$$\mathbb{E}[\text{revenue}|p] = \sum_{j=0}^{n} p^j (1-p)^(n-j) \left(\begin{array}{c} n \\ j \end{array} \right) \mathbb{E}[\text{revenue}|k=j].$$

With these analytic solutions, we can compare the revenue of the different classes of Nash equilibria as n varies. The qualitative revenue properties of the are consistent with the intuition from 2 bidder examples, but the quantitative impact on revenue is striking (see Figure 2). For pure strategy equilibria, all costs have the same revenue for low values of n. When n becomes too large, adding more agent 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 is 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 2-bidder case, symmetric mixed strategy equilibria involve a risk of miscoordinating, though this risk becomes greater (and has a negative impact on seller revenue for $n > 3$) as the number of agents increases. The value of information decreases, decreasing the probability of

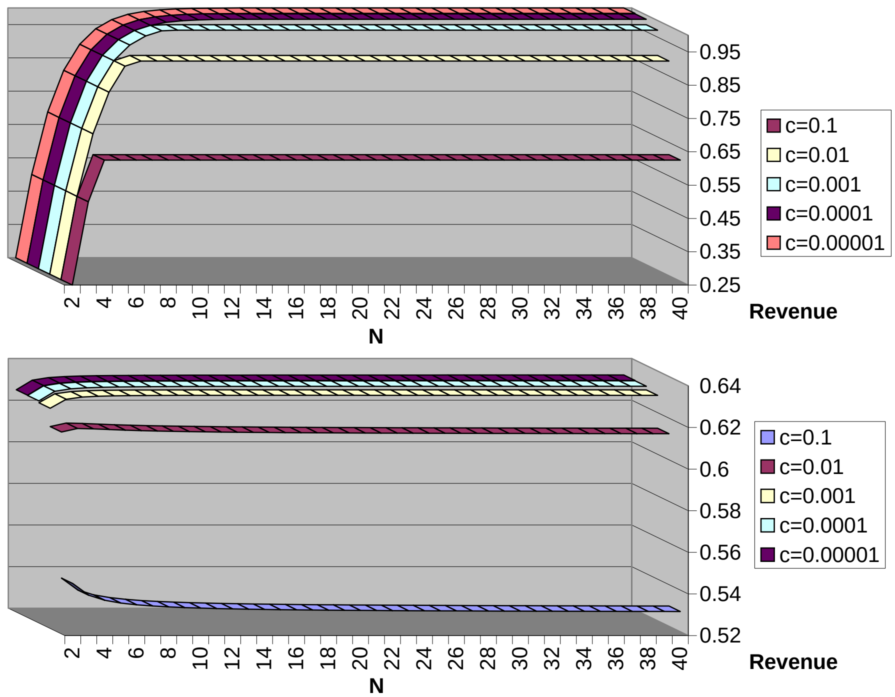

Figure 2: Expected revenue under pure (top) and mixed (bottom) strategies

deliberation. As n gets large the probability of introspection will tend to 0, and the expected revenue will tend the uninformed expected valuation (though in this case it does so extremely slowly).

Although the second price auction is always *ex post* efficient, the deliberation strategies used will affect the probability of *ex post* perfect efficiency (the probability that the allocation actually maximizes social welfare, which given our simple 0,1 valuation structure is also the expected fraction of optimal social welfare). These different equilibria have different probabilities of *ex post* perfect efficiency (see Figure 3). Again, the risk of miscoordination has a significant negative impact.

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

$$e_p(k,n) = 1 - \left((1/2)^k \sum_{j=1}^{n-k-1} (1/2)^j \left(\begin{array}{c} n-k \\ j \end{array} \right) \frac{j}{n-k} \right).$$

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

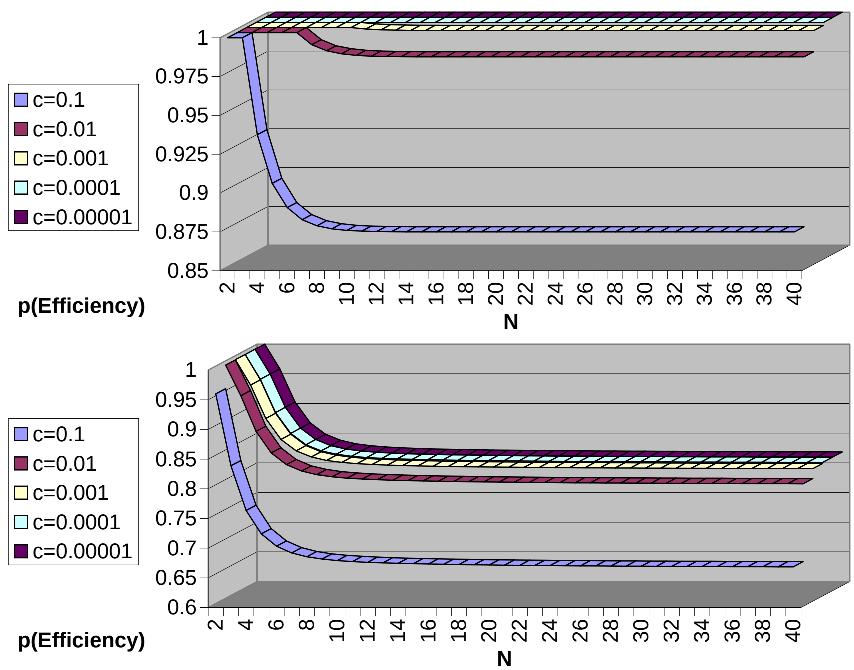

Figure 3: Probability of *ex post* perfect efficiency under pure (top) and mixed (bottom) strategies

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

Next we will consider a case with limited, free deliberations, and show that equilibria can have similar properties here. The agents have valuations drawn from a uniform $[0,1]$ distribution. 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 of $n = 2$), we show the induced normal form continuously (see Figure 4). Again, we have the problem that the only pure strategy equilibria are asymmetric (even though deliberations are no longer costly): $[q_1 = 1/3, q_2 = 2/3]$ and $[q_1 = 2/3, q_2 = 1/3]$. This reflects the results 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 the mechanism or proxy). The agents could randomize across these two equilibria, though again they will lose expected utility through miscoordination. By ordering the bidders by weakly increasing q_i, we can write a expression for the expected utility of any number of agents under any pure strategy profile:

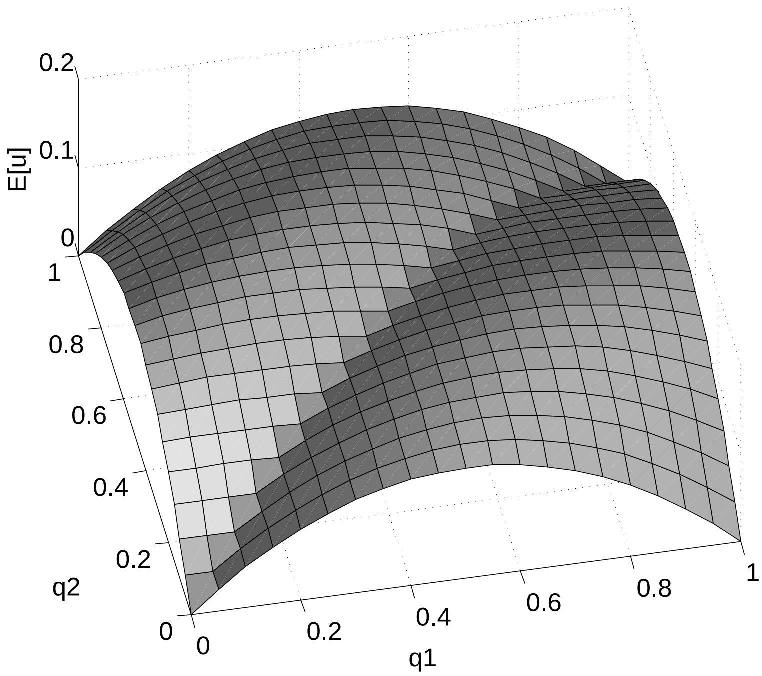

Figure 4: Graphical representation of the induced normal form of the 2-bidder limited deliberation case. The peaks (pure equilibria) do not fall on the $q_1 = q_2$ line, so they are not symmetric.

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

4 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 8 *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 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 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 a wide variety of models including independent, common and interdependent values, as well as private and non-private. It also applies volatile settings, provided the agents' valuations are always bounded above by $\bar{v}$. Finally, it makes no assumptions about the allocation rule, so it applies to all auctions, efficient or not.

Theorem 9 *Any individually-rational, separable auction the upper bound on the expected revenue is reduced by the sum of the expected costs of every agent's deliberations.*

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.

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

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

Corollary 10 *No budget-balanced, individually-rational, separable mechanism can have a dominant strategy which 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 are performing a deliberation with cost $\geq c$, then the seller's expected revenue becomes negative. ■

Theorem 11 *In efficient, separable auctions with independent private values, the 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 the deliberation to bid in the auction. 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, his expected surplus (and the value of information for deliberation q) 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. ■

5 Conclusion

We have expanded on previous work which shows that the problem of how to deliberate prior to bidding adds a significant extra layer of strategic complexity, by demonstrating that limited and costly deliberation can impose a coordination problem on those agents. We have shown that there is a significant effect on the utility of bidders and sellers. We have also shown that this impact is present even for small numbers of bidders and extremely small costs. Future work on this problem could benefit from relaxing a number of the assumptions, about distribution, privacy and separability. In particular, it might be useful to formally model the costs to the seller (and other bidders) of waiting for slow deliberations, allowing a continuous trade-off of the speed of separable auctions against the lower deliberation costs and higher revenue of inseparable auctions. Work could also be done to characterize cases deliberation acts both as a source of information and as a randomizing device.

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. (2001). Costly valuation computation in auctions. *TARK* (pp. 169–182).

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.