

Optimal Internet Auctions with Costly Communication

Yuqian Li
Department of Computer Science
Duke University
Durham, NC 27708, USA
yuqian@cs.duke.edu

Vincent Conitzer
Department of Computer Science
Duke University
Durham, NC 27708, USA
conitzer@cs.duke.edu

ABSTRACT

Iterative auctions can reach an outcome before all bidders have revealed all their preference information. This can decrease costs associated with communication, deliberation, and loss of privacy. We propose an explicit cost model that is inspired by single-item Internet auctions, such as those taking place on auction sites (eBay) or via informal communication (craigslist, mailing lists). A nonzero bid comes at a cost to both the seller and the bidder, and the seller can send broadcast queries at a cost. Under this model, we study auctions that maximize the seller's profit (revenue minus seller cost). We consider multi-round Vickrey auctions (MVAs), in which the seller runs multiple Vickrey auctions, with decreasing reserve prices. We prove that restricting attention to this class is without loss of optimality, show how to compute an optimal MVA, and compare experimentally to some other natural MVAs. Among our findings are that (1) the expected total cost is bounded by a constant for arbitrarily many bidders, and (2) the optimal MVA and profit remain the same as long as the total bid cost is fixed, regardless of which portion of it belongs to the seller and which to the buyer.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—Multi-agent systems; J.4 [Social and Behavioral Sciences]: Economics

General Terms

Algorithms, Economics, Theory

Keywords

Auctions, mechanism design, communication costs

1. INTRODUCTION

Auctions constitute a favored method for allocating scarce resources in multiagent systems. However, communication requirements can pose a bottleneck. Motivated by the revelation principle, much of the theory of mechanism design considers *direct-revelation* mechanisms, in which each agent declares its entire valuation function to the auctioneer. The corresponding overhead, not only in terms of communication per se but also in terms of the agent

having to completely *determine* its valuation function, can be prohibitive. All of this is well understood (for further discussion, see, e.g., [2]), and a significant amount of research has been devoted to the design of *iterative* auction mechanisms (e.g., [10]) and (roughly equivalently) auctions with explicit *elicitation* of agents' valuations (e.g., [11]). Such auctions aim to avoid unnecessary communication. For example, in a Vickrey auction, if it is known that bidders 1 and 2 have valuations above \$100 and bidder 3 has one below \$100, then there is no need to query bidder 3 any further.

How do we evaluate how effective a particular iterative auction mechanism is in reducing communication? Perhaps the simplest measure is the number of bits communicated (see, e.g., [9]). Within a particular query model, it may also make sense to minimize the total number of queries (see, e.g., [5]). However, in this paper, we argue that such existing models fail to capture important aspects of the cost of communication in certain types of Internet auctions. In particular, in such auctions often the most costly communication that takes place is the *first* time that a bidder gives a positive response, indicating having a nonzero value for an item. This can be the case for several reasons. One possibility is that the auction website at this point may insist that the bidder provides payment (say, credit card) information or places money in escrow. Otherwise, a malicious user may steer the auction in a particular direction and in the end refuse to pay. Alternatively, in other settings (such as an item having been posted for sale on craigslist or a similar list), at this point the bidder may wish to set up an appointment with the seller to check the item. In both cases, these actions come at costs for both the potential buyer and the seller, in terms of effort, time, loss of privacy, and so on. While participation costs have been studied before in auctions [14, 15], a distinguishing feature of our model is that a bidder can observe the proceedings of the auction at no cost, until the bidder decides to actively participate by indicating a nonzero valuation and thereby changing the course of the auction. This appears to us to be a more natural model of Internet (or other highly anonymous) auctions of the type discussed above.

The rest of this paper is organized as follows. In Section 2, we provide a formal model of communication cost motivated by the observations just discussed. In Section 3, we study the special case in which efficient allocation is a constraint and bidders experience no cost from bidding. The characterization of the optimal mechanism in this special case turns out to be closely related to earlier work on finding an optimal agent by iteratively relaxing the parameters of the search [13] (see the more detailed discussion at the end of Section 3.3). In Section 4, we drop the assumptions of efficiency and no bidder cost for communication. In both cases, the optimal auctions are proved to be multi-round Vickrey auctions (MVAs), which were previously studied in [6]. In that paper, there are no

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May 6–10, 2013, Saint Paul, Minnesota, USA. Copyright © 2013, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

explicit communication costs, but going through multiple rounds is costly for the agents because they discount the future.

2. COST MODEL AND SETTINGS

In this section, we formally define our cost model and explain its motivation.

DEFINITION 1 (SETTING). *One seller is selling one item to n buyers (bidders) whose valuations v_i ($1 \leq i \leq n$) are independently and identically distributed (i.i.d.) over $[0, 1]$ with PDF $f(x)$ and CDF $F(x)$ which are assumed to have no gaps ($F(x)$ is strictly increasing everywhere). The seller can broadcast a message to all bidders, at cost b to the seller. A bidder can reply to that broadcast, or remain silent. If a bidder replies (does not stay silent), this comes at a cost β_1 to the seller and a cost β_2 to that bidder, for a total cost of $c = \beta_1 + \beta_2$.*

Two key aspects of this model are that (1) staying silent comes at no cost and (2) replying comes at a positive cost, and this positive cost is the same no matter how complex the query and answer are. This is motivated by the settings discussed in the introduction, where a bidder can observe the process of the auction (or messages posted on a board) silently at no cost, but once the bidder acts in the auction, costs occur—e.g., the bidder has to submit credit card information, the bidder and the seller have to arrange an in-person meeting, etc.). A key aspect of such costs is that they tend to be the same regardless of the level of detail in the bidder’s answer: for example, if the bidder just reports having a valuation greater than \$10 without specifying what it is exactly (rather than reporting a valuation of exactly \$14), this is not likely to reduce any of the above costs. In particular, the seller is likely to want to verify the bidder’s authenticity at any point where the bidder’s reply affects the course of the auction from then on. This leads to the following easy proposition:

PROPOSITION 1. *In the model defined in 1, without loss of optimality, we can restrict our attention to (broadcast) queries that result in each bidder either staying silent, or immediately revealing his exact valuation to the seller.*

Another notable point is that we restrict communication from the seller to broadcast queries. This is a common restriction: any sealed-bid auction can be considered a broadcast auction with only one broadcast: the reserve price. The bisection auction [4] is an example auction with many rounds of broadcasts. In each round, it broadcasts a price and asks bidders to reply whether their valuation is above or below that price. Besides the broadcast model being simple, natural, and common in existing auction mechanisms, it is also naturally motivated in the Internet domains we consider: the bidders are entirely anonymous until their first reply, so before this point querying such bidders individually is not feasible and they can only be reached by, say, posting on a public website; and after they have replied, we will know their valuation exactly (by Proposition 1) and we no longer need to query them. Of course, there are offline cases where the set of bidders is small and explicit (e.g., the government wants to sell land or spectrum to one of three known companies); in such settings, it can indeed be helpful for the seller to communicate with bidders individually [7, 12]. Such settings do not fit our model; we explicitly focus on highly anonymous settings, and the costs that the seller incurs from the broadcast query correspond to the time, effort, and third-party charges associated with posting a public message.

DEFINITION 2 (NO BLIND ALLOCATION). *An item cannot be awarded to a bidder who has remained silent to all queries.*

DEFINITION 3. *A mechanism in our setting consists of (1) a full contingency plan for which query to broadcast at each point, depending on answers given so far, and a termination condition; and (2) an allocation and pricing rule that is defined on each terminal state, satisfying no blind allocation. A mechanism is individually rational if losing bidders never pay and winning bidders never pay more than their valuations. We say an individually rational mechanism is optimal if it has a Bayes-Nash equilibrium for the bidders that maximizes the seller’s profit (among all Bayes-Nash equilibria of all individually rational auction mechanisms). Here, seller profit is revenue minus seller’s costs. A class of mechanisms is optimal if it contains at least one optimal mechanism.*

3. OPTIMAL MECHANISMS WITH EFFICIENCY AND ONLY SELLER’S COST

In this section, we make two simplifying assumptions: (1) we restrict attention to mechanisms that allocate the item efficiently and (2) we assume that the bidder cost for replying (β_2) is zero.¹ For example, someone who is moving and selling furniture on craigslist is likely to have 0 valuation for the item and cannot commit to withhold the item or prevent re-sale between bidders. Under such circumstances, efficient mechanisms not only maximizes the social welfare but also maximizes the seller’s revenue [1] (though this does not consider communication costs). In Section 4 we drop these assumptions.

The rest of this section is organized as follows. First, we introduce a class of mechanisms called multi-round Vickrey auctions (MVA). Then, we prove that we can restrict attention to MVAs without loss of optimality. After that, we find the specific MVA that is optimal. Finally, we experimentally compare this optimal MVA to some other natural mechanisms.

3.1 Multi-round Vickrey Auctions

In an MVA, the seller runs a Vickrey auction with a reserve price; if nobody bids (above the reserve price), the seller runs another Vickrey auction with a lower reserve price, etc., until the item is sold. (In Section 4, we will also consider MVAs that can terminate without having sold the item.) For example, consider a sequence of eBay auctions (with proxy bidding) in which the seller is repeatedly lowering the reserve price.

DEFINITION 4 (MULTI-ROUND VICKREY AUCTION (MVA)). *An MVA is defined by a sequence of reserve prices r_1, r_2, \dots (which may be finite or infinite) where $r_i > r_{i+1}$. In round i , a Vickrey auction with reserve price r_i is run (if the item has not been sold).*

In an MVA, if a bidder decides to bid (above the reserve price), it is optimal to bid truthfully since doing so is dominant in a Vickrey auction. However, a bidder may choose strategically to stay silent even with a valuation above the reserve price, in the belief that nobody else will bid this round so that the price will decrease in the next round. Thus, in our (symmetric) setting, Bayes-Nash Equilibria (BNE) for MVAs can be described by a sequence of thresholds a_1, a_2, \dots where $a_i > a_{i+1}$; in round i , a bidder bids if and only if his valuation is at least a_i .

We will be interested in the following question: given desired thresholds a_i , which reserve prices r_i result in these thresholds? Then, by the revelation principle, we can convert this to a mechanism in which we query agents whether their valuations are above a_i and it is optimal for them to respond truthfully.

¹Note that if (2) does not hold, then allocating efficiently may not be possible without the seller compensating the bidders for their bid costs.

LEMMA 1. Consider a symmetric strategy profile in an MVA where in the i th round, valuation a_i is the threshold for bidding (so that bidders with lower valuations stay silent and those with higher valuations bid). Then this constitutes a Bayes-Nash equilibrium if:

$$P(a_i)(a_i - r_i) = \int_{a_{i+1}}^{a_i} (a_i - x)p(x)dx + P(a_{i+1})(a_i - r_{i+1}) \quad (1)$$

(where $P(x) := F(x)^{n-1}$ and $p(x) := P'(x) = (n-1)F(x)^{n-2}f(x)$) when i is not the last round, and either $a_i = r_i$ when i is the last round or $\lim_{i \rightarrow \infty} a_i = \lim_{i \rightarrow \infty} r_i$.

PROOF. Consider a bidder with valuation a_i ; we first show that in round i , such a bidder is indifferent between bidding and staying silent if the condition holds. If i is the last round, clearly he is indifferent between bidding and not iff $a_i = r_i$. We now consider the case where i is not the last round. If another bidder bids in round i , that bidder will bid at least a_i , and our bidder will have zero utility. Therefore, the left-hand side of (1) represents the expected utility to our bidder for bidding now. Corresponding to the right-hand side, if our bidder stays silent, and there is a next round and our bidder bids in that next round, then he will win in that round, either with another bidder bidding (first term) or not (second term).

If we now consider a bidder with valuation above a_i , a similar analysis shows that this bidder strictly prefers bidding in round i to waiting one more round; inductively, he will prefer it to waiting any number $k > 0$ more rounds; and he will also prefer this to never bidding because either $a_j = r_j$ when j is the last round or $\lim_{j \rightarrow \infty} a_j = \lim_{j \rightarrow \infty} r_j$. Finally, let us consider a bidder whose valuation lies below all the a_i . Because either $a_j = r_j$ when j is the last round or $\lim_{j \rightarrow \infty} a_j = \lim_{j \rightarrow \infty} r_j$, this bidder is best off never bidding. \square

For a mechanism that allocates efficiently, we either have $r_k = a_k = 0$ for some k , or $\lim_{i \rightarrow \infty} r_i = 0$.

THEOREM 1. Given a decreasing sequence of $a_i \in [0, 1]$ that ends at or converges to 0, let

$$r_i = \left(\int_0^{a_i} x p(x) dx \right) / P(a_i) \quad (\text{if } a_i > 0) \quad (2)$$

and $r_i = 0$ if $a_i = 0$. The corresponding MVA has a pure strategy Bayes-Nash equilibrium characterized by bidding thresholds a_1, a_2, \dots

PROOF. We show that the conditions of Lemma 1 hold. If i is the last round, then $a_i = 0 = r_i$. Also,

$$\lim_{i \rightarrow \infty} r_i \leq \lim_{i \rightarrow \infty} \left(\int_0^{a_i} a_i p(x) dx \right) / P(a_i) = \lim_{i \rightarrow \infty} a_i = 0$$

If i is not the last round, by (2) we have $r_i P(a_i) = \int_0^{a_i} x p(x) dx$ for all i (this clearly also holds if $a_i = 0$). Thus the right-hand side of Equation (1) is:

$$\begin{aligned} & \int_{a_{i+1}}^{a_i} a_i p(x) dx - \int_{a_{i+1}}^{a_i} x p(x) dx + P(a_{i+1})(a_i - r_{i+1}) \\ &= a_i P(a_i) - \underline{a_i P(a_{i+1})} - r_i P(a_i) + \underline{r_{i+1} P(a_{i+1})} \\ & \quad + \underline{P(a_{i+1}) a_i} - \underline{P(a_{i+1}) r_{i+1}} \end{aligned}$$

which equals to left-hand side. \square

This tells us that a bidder will bid in a round of an MVA if and only if the expected second-highest bid conditional on this bidder's

valuation being the highest is greater than the reserve price of that round. For example, if the distribution is uniform, $r_i = \frac{n-1}{n} a_i$.

A similar analysis to the one in this subsection appears in [6]. In that paper, the seller is unable to commit to, when no bid exceeds the reserve price, not put the item up for sale again in the next round at a lower reserve price, but because of discounting some of the value of the item is destroyed in this case. However, that paper does not consider communication costs, and so we diverge from that work in what follows.

3.2 Optimality of MVAs

Since the mechanism is required to be efficient and bidders' communications costs (β_2) are zero, any mechanism that gives utility zero to an agent with valuation zero results in the same revenue for the seller, by the revenue equivalence theorem [8].² Hence, maximizing profit is equivalent to minimizing the seller's query costs.

By the revelation principle, we can restrict our attention to mechanisms in which agents always answer truthfully in equilibrium.³ By Proposition 1, we can assume that an agent reveals his entire valuation when not staying silent. By the efficiency constraint, the mechanism must at least discover a bidder with the highest valuation. This optimization problem corresponds to Definition 5.

DEFINITION 5. A (direct-revelation) query is given by a subset $Q \subseteq [0, 1]$, such that if the agent's valuation is in Q , he replies with his exact valuation, and otherwise stays silent. A strategy for asking queries can be represented by a function

$$S(f, m, V, \mathcal{Q} = \{Q_1, Q_2, \dots, Q_{i-1}\}) = Q_i$$

which means: suppose that the set of queries asked previously is \mathcal{Q} , the set of values already reported is V , and there are m bidders left who have not responded (whose valuations were drawn i.i.d. from f); then the strategy S will next ask Q_i . We will use $Q_i = \emptyset$ as shorthand for terminating the algorithm. $C_{f,n}(S)$ is the expected cost of strategy S , and we wish to find $C_{f,n}^* = \inf_S C_{f,n}(S)$.

To find out the optimal query strategy, we first show that the minimum cost is independent of the PDF $f(x)$.

LEMMA 2. Consider the uniform PDF $f_u(x) = 1$ and let $C_n^* = C_{f_u,n}^*$. For any other PDF $f(x)$, we have $C_{f,n}^* = C_n^*$.

PROOF. For any strategy S_{f_u} for the uniform distribution, we can transform it into a strategy S_f for the distribution with PDF f by querying at the equivalent percentiles. Specifically,

$$S_f(f, m, V, \mathcal{Q}) = F^{-1}(S_{f_u}(f_u, m, F(V), F(\mathcal{Q})))$$

The performance of S_f on f is the same as that of S_{f_u} on f_u . Therefore, $C_{f,n}^* \leq C_{f_u,n}^*$. Conversely and similarly, for any strategy S_f for PDF f , we can transform it into a strategy S_{f_u} for uniform PDF f_u . Specifically,

$$S_{f_u}(f_u, m, V, \mathcal{Q}) = F(S_f(f, m, F^{-1}(V), F^{-1}(\mathcal{Q})))$$

Thus, $C_{f_u,n}^* \leq C_{f,n}^*$. Hence, we have $C_{f,n}^* = C_{f_u,n}^* = C_n^*$. Figure 1 illustrates those two constructions we used for this proof. \square

We next prove that descending query strategies, in which all bidders are asked whether their valuations are above a_i for a decreasing sequence of a_i , are optimal. (For a descending query strategy, as soon as we get a positive reply, we are done, so the contingency plan does not need to branch.)

²Moreover, by individual rationality we cannot give an agent with valuation zero less than zero utility; we could give such an agent more than zero utility (as in redistribution mechanisms), but this would only hurt revenue.

³For example, for MVAs, we can directly query the agents whether their valuations are above a_i while still charging according to r_i .




Figure 1: These two figures illustrate how to construct a strategy for an arbitrary PDF f from one for the uniform distribution f_u , and vice versa. Here, we depict a strategy as a box that takes four inputs f, m, V, Q (PDF, number of unknown values, reported values set, and set of asked queries) and returns Q (the next query).

LEMMA 3. *Every optimal strategy uses only descending queries $Q_1 = [a_1, 1)$, $Q_2 = [a_2, a_1)$, $Q_3 = [a_3, a_2) \dots$ (with the possible exception of strategies that differ only on a measure zero set and are therefore equivalent to a descending-query strategy with probability 1).*

PROOF. Let us assume, for the sake of contradiction, that there exists an optimal strategy that uses a non-descending query (meaning, one that differs from a descending query on a non-measure-zero set). In that strategy S , there must be a first non-descending query $Q_{i+1} = S(F, m, V, Q = \{Q_1, Q_2, \dots, Q_i\})$. Consider an alternative descending query $Q'_{i+1} = [a'_{i+1}, a_i)$ (presumably $a_0 = 1$) such that $|Q'_{i+1}|_f = |Q_{i+1}|_f$ where for any query Q , $|Q|_f$ denotes the probability under f of the region queried by Q .

Because Q_1 to Q_i are all descending, we have $m = n$ and $V = \emptyset$ (otherwise an optimal strategy should terminate without asking Q_{i+1}). Let C be the expected cost of using S from this point on (starting with Q_{i+1}), and let C' be the expected cost of using Q'_{i+1} and querying optimally after that. We have $C = b + \sum_{j=0}^n p_j(j \cdot c + L_j)$ and $C' = b + \sum_{j=0}^n p'_j(j \cdot c + L'_j)$ where p_j (or p'_j) is the probability that there are j reported values within Q_{i+1} (or Q'_{i+1}), and L_j (or L'_j) is the expected cost of later queries given that j values have been found in Q_{i+1} (or Q'_{i+1}).

Because $|Q'_{i+1}|_f = |Q_{i+1}|_f$, we have $p'_j = p_j$ for all j . We have $L_n = L'_n = 0$. By Lemma 2, $L_0 = L'_0 = C_n^*$ because if we know that no value lies in Q_1, Q_2, \dots, Q_{i+1} , the remaining problem is equivalent to that with the revised PDF

$$f_{i+1}(x) = \begin{cases} \lambda f(x), & x \notin Q_1 \cup Q_2 \cup \dots \cup Q_{i+1} \\ 0, & x \in Q_1 \cup Q_2 \cup \dots \cup Q_{i+1} \end{cases}$$

where λ is a constant that makes $\int_0^1 f_{i+1}(x) dx = 1$. Finally, since Q'_{i+1} is a descending query, for all $0 < j < n$, $L'_j = 0$ but $L_j > 0$ (because for these values of j , it is possible that all reported values in Q_{i+1} lie below some value that has not been queried yet). Moreover, there exists some j with $0 < j < n$ and $p_j > 0$ unless (1) the probability of each agent's reply to Q_{i+1} is 1—but in this case it differs only on a measure zero set from the descending query that asks for all the remaining values, contradicting our assumption; (2) the probability of each agent's reply to Q_{i+1} is 0—but such a measure zero query is clearly suboptimal; or (3) $n = 1$ —but in this case any query that differs by more than a measure zero set from the query that asks for all the remaining values is clearly suboptimal. Hence, $C' < C$, contradicting our initial assumption. \square

It follows that:

THEOREM 2. *Among mechanisms that are required to allocate efficiently, Multi-round Vickrey Auctions (MVAs) have minimum cost.*

PROOF. By Lemma 3, we can restrict our attention to mechanisms with descending query strategies $Q_1 = [a_1, 1)$, $Q_2 = [a_2, a_1)$, $Q_3 = [a_3, a_2) \dots$. For any such mechanism, Theorem 2 tells us how to find reserve prices r_1, r_2, \dots such that the corresponding MVA has a Bayes-Nash equilibrium that is equivalent to this descending-query mechanism. \square

By the revenue equivalence theorem, we obtain

COROLLARY 1. *If the bidders have no cost for bidding ($\beta_2 = 0$), MVAs that minimize cost for the seller are optimal among mechanisms that allocate efficiently.*

3.3 MVAs with Minimum Cost

The above analysis still leaves open what the optimal parameters (thresholds a_i , or equivalently, reserve prices r_i) of the optimal MVA are for a given setting f, n, b, c (valuation PDF, number of bidders, broadcast cost, bid cost). According to Lemma 2, the cost does not depend on f and it suffices to restrict our attention to the uniform distribution. Let $\rho = b/c$ (for cases where $c > 0$).

DEFINITION 6. *Given f , we define the α -MVA to be the MVA in which, the expected number of nonsilent bids in each round (conditional on having reached that round) is $(1 - \alpha)n$. In the case where f is uniform, the α -MVA is characterized by $a_i = \alpha^i$.*

PROPOSITION 2. *For the purpose of minimizing total cost when $\beta_2 = 0$ under the constraint of efficient allocation: If $c = 0$, the Vickrey auction (the 0-MVA) is optimal. If $\rho = b = 0$, the Dutch auction (which is approximated by the $(1 - \epsilon)$ -MVA) is optimal. Otherwise, it is optimal to use an α -MVA where α satisfies*

$$\alpha^{n-1}(\rho + (1 - \alpha)n) - (1 - \alpha^n) = 0 \quad (3)$$

PROOF. If $c = 0$, the Vickrey auction has cost b , which is optimal. If $\rho = b = 0$, the expected total cost of the Dutch auction is c (since with probability 1 only one agent will reply), which is optimal. For the remaining case, we first argue that there must be some α -MVA that is optimal. By Lemma 2, we can assume f is uniform. By Theorem 2, some MVA must be optimal. For this MVA, consider a_1 . If $a_1 = 0$, this is the 0-MVA. If $a_1 > 0$, then if at least one bidder is above a_1 , we finish after the first query; otherwise, the resulting conditional distribution for each bidder is $f|_{[0, a_1]}$. According to Lemma 2, we can rescale this conditional distribution to the uniform distribution over $[0, 1)$ to arrive back at our original problem; and an optimal mechanism for that is the same MVA that starts with query a_1 (which translates to a_1^2 without rescaling). Repeated application of this reasoning results in the α -MVA with $\alpha = a_1$.

All that remains to show is the characterization of this α . If C_α is the expected overall cost for the α -MVA, then we have $C_\alpha = b + (1 - \alpha)nc + \alpha^n C_\alpha$, or equivalently $C_\alpha = \frac{b + (1 - \alpha)nc}{1 - \alpha^n}$. If we optimize this with respect to α , we must either have

$$\frac{\partial C_\alpha}{\partial \alpha} = \frac{\alpha^{n-1} n (b + (1 - \alpha) nc)}{(1 - \alpha^n)^2} - \frac{nc}{1 - \alpha^n} = 0$$

\Downarrow

$$\alpha^{n-1}(\rho + (1 - \alpha)n) - (1 - \alpha^n) = 0$$

or have α at a boundary value (0 or 1). However, $\alpha = 1$ (or approaching 1) is clearly suboptimal unless $b = 0$, a case that we have already covered; $\alpha = 0$ can be ruled out as $\frac{\partial C_\alpha}{\partial \alpha} \Big|_{\alpha=0} < 0$. \square

By Proposition 2, the optimal α can be computed instantly using any numerical solver. This result is analogous to a result by Sarne




Figure 2: Comparisons for the optimal α and its approximations. The first approximation is $\alpha = 1 - 1/n$, the second is $\alpha = (-W(e^{-1-\rho}))^{-1/n}$. ($c = 0.1$, $b = \rho c$. Recall that $\rho = b/c$ for cases where $c > 0$.)

et al. [13]; translating⁴ their result into our setting also gives a proof that α -MVAs are optimal among MVAs and also characterizes the optimal value of α . However, they do not provide a proof that MVAs are optimal among all mechanisms; their setup implicitly restricts attention to MVAs (when translated to our setting).

3.4 Approximating the Optimal α

As we are not able to obtain a closed form for the optimal α from Equation (3), we here give some formulas to approximate it.

THEOREM 3. *Setting $\alpha = 1 - 1/n$ results in a $1/(1 - e^{-1}) \approx 1.582$ approximation of the cost of the optimal α -MVA. Also, the total expected cost of this mechanism is at most $(b + c)/(1 - e^{-1})$ (and hence constant in the number of agents).*

PROOF. We have $C_{\alpha=1-1/n} = (b + c)/(1 - (1 - 1/n)^n)$. Because $(1 - 1/n)^n \leq e^{-1}$, we have $C_{\alpha=1-1/n} \leq (b + c)/(1 - e^{-1})$. On the other hand, any mechanism requires at least one broadcast and one bid to terminate and hence $C_n^* \geq b + c$. \square

We now present a different approximation whose expected cost is guaranteed to converge to the optimal one as n grows. The proof (in the full paper online) is omitted due to the space constraint.

THEOREM 4. *Let W denote the lower branch of Lambert W function defined by $W(x)e^{W(x)} = x$. Let $\alpha = (-W(e^{-1-\rho}))^{-1/n}$. Then $\lim_{n \rightarrow \infty} C_n^* = \lim_{n \rightarrow \infty} C_{\alpha=(-W(e^{-1-\rho}))^{-1/n}}$*

3.5 Experiments

Without loss of generality, all our experiments concern the uniform distribution of valuations over $[0, 1)$. The experiments in Figure 2 compare the MVAs with (1) the optimal α (solved numerically), (2) the approximation $\alpha = 1 - 1/n$, and (3) the approximation $\alpha = (-W(e^{-1-\rho}))^{-1/n}$, all for $c = 1$ and $\rho = b = 0.2, 1, 5$. Both approximations perform well for $\rho = 1$, but the second performs much better for other values.

We now compare the α -MVA (with the optimal α calculated numerically) to some other MVAs, each of which uses at most k rounds. An advantage of such MVAs is that they can be used when there is a hard deadline (such as in a moving sale). They are:

- *uniform k -MVA*: the uniform k -MVA has its k thresholds at $0, 1/k, 2/k, \dots, (k - 1)/k$, and the *optimal uniform MVA*

⁴A detail that needs to be accounted for is that Sarne et al. motivate their model as performing increasingly broad searches for an optimal agent and thereby focus on increasing threshold search to find a minimum, rather than decreasing queries to find a maximum.




Figure 3: Comparing the expected costs of different MVAs over the number of bidders n . ($b = 0.2$, $c = 0.1$)

corresponds to the optimal k (computed by brute force search). (Compare the fixed-step search strategy in [13, 3].)

- *optimal k -MVA*: the optimal MVA using at most k rounds. (These are computed using techniques that we will discuss in Section 4.3.)
- *α -cutoff- k -MVA*: proceeds as the α -MVA, except its k th (and final) query is at 0.

The results of our experiments are shown in Figure 3.

- The optimal uniform-MVA's cost is very close to that of the optimal α -MVA, especially when n is large. This makes sense because when n grows large, the optimal α approaches 1, so the first few queries of the optimal α -MVA are at similar distances from each other.
- The optimal k -MVA's cost decreases and approaches the optimal cost quickly as k grows.
- When k is fixed to a small value, the uniform k -MVA has significant higher cost than both the optimal k -MVA and the α -cutoff- k -MVA.
- The α -cutoff-3-MVA works well (close to optimal 3-MVA).

4. THE GENERAL CASE

In this section, we drop the constraints that (1) the mechanism must allocate efficiently and (2) bidders have no cost for bidding ($\beta_2 = 0$), and prove that MVAs are optimal even without these constraints (i.e., in the fully general model). We also study methods for finding the specific optimal MVA, though we do not manage to obtain a characterization that is as elegant as in the restricted case.

An interesting effect of $\beta_2 > 0$ is that it gives bidders an incentive *not* to bid when they expect many others to bid, because if they do not win they still pay β_2 for bidding. (Also, we note that their probability of winning with a fixed value is exponentially small in the number of bidders, so that even small bidding costs, such as the effort needed to send an e-mail, can be significant.) This effect is rather opposite to the effect observed in earlier parts of this paper that bidders are less inclined to wait for later rounds when there are many bidders.

4.1 Spending Equivalence Theorem

Another effect of $\beta_2 > 0$ is that the revenue equivalence theorem (straightforwardly interpreted) is no longer applicable, because this theorem assumes that the utility is equal to valuation minus payment, which is no longer true because cost also plays a role. Intuitively, we can fold this cost into the payment to make the theorem applicable again, though we have to keep in mind that this cost does not benefit the seller, unlike (true) payments. This is how we proceed.

THEOREM 5 (SPENDING EQUIVALENCE). *The bidders' total spending (expected payment + bidding costs) is completely determined by (1) the expected utility of lowest-type bidders and (2) the allocation probability function*

$$p : (v_1, v_2, \dots, v_n) \rightarrow (p_1, p_2, \dots, p_n)$$

where p_i is the probability that bidder i will get the item.

PROOF. Any mechanism in our domain (with bidding costs) can also be used in a domain without costs, letting the seller collect the bidding costs (as well as the payments) instead. The bidders' incentives will be identical in this mechanism, and so the result follows from applying the standard revenue equivalence theorem to these transformed mechanisms. \square

This greatly simplifies our profit maximization problem:

COROLLARY 2. *To maximize profit for a fixed allocation rule $p : (v_1, v_2, \dots, v_n) \rightarrow (p_1, p_2, \dots, p_n)$, it suffices to minimize total cost (including both seller's cost and bidders' cost).*

PROOF. The seller's profit is revenue – seller's cost = (bidders' spending – bidders' cost) – seller's cost = bidders' spending – total cost, and bidders' spending is fixed by Theorem 5. \square

4.2 MVAs' Optimality in General

It is important to note that when $\beta_2 > 0$, the bidders' incentives have changed, and if we wish to achieve given thresholds a_1, a_2, \dots using an MVA, just setting the reserve prices r_1, r_2, \dots according to Lemma 1 will not suffice. For example, if we set the reserve prices in this way, it may be the case that $\beta_2 > a_i - r_i$, in which case it would be clearly suboptimal for a bidder with valuation $a_i + \epsilon$ to bid in round i . Rather than adjust the reserve prices, we can employ a simpler trick: for every bid that is placed, the seller can pay that bidder β_2 , thereby absorbing the cost of the query. By doing so, the bidders' incentives are as they were before, and by Theorem 5 this will lead to the same profit as any other mechanism with the same query costs and allocation probabilities. We will still consider a mechanism with such refunds an MVA.

By Myerson's result, to maximize expected total *bidder spending*, we need to maximize the expected virtual value of the winning bidder. Because we are considering cases with symmetric priors, this results in a Vickrey auction with a reserve price that corresponds to a virtual valuation of zero. However, we are interested in profit, not spending. Not only does this mean that we need to




Figure 4: Optimal first-query-length $x (= 1 - a_1)$ as a function of l . ($\rho = 2, n = 10$, and $[0, 1]$ is discretized into 1000 pieces.)

plan our queries carefully, but it can also affect the allocation rule. For example, increasing the final reserve price (below which we do not allocate) slightly above Myerson's reserve price may increase profit, because it will reduce communication costs. Still, we can build on Myerson's result to show that MVAs (generally with a final reserve price above 0) are optimal.

THEOREM 6. *Under a regularity condition on the PDF (implying the virtual valuation is strictly increasing [8]), there exists an optimal MVA.*

PROOF. Consider any optimal mechanism given by a query strategy $S(f, m, V, \mathcal{Q})$ (see Definition 5) and an allocation function based on V , the set of reported values. We then construct another mechanism with query strategy $S'(f, m, V', \mathcal{Q}')$ where \mathcal{Q}' consists of descending queries and

$$|S'(f, m, \emptyset, \{Q'_1, \dots, Q'_{i-1}\})|_f = |S(f, m, \emptyset, \{Q_1, \dots, Q_{i-1}\})|_f$$

If $S(f, m, \emptyset, \{Q_1, \dots, Q_{i-1}\}) = \emptyset$ (termination), then it also has $S'(f, m, \emptyset, \{Q'_1, \dots, Q'_{i-1}\}) = \emptyset$ and the item is not allocated in this case. Under S' , once at least one bidder bids, we allocate to the highest bidder (and payments are set to incentivize truthful reporting, as discussed earlier). The querying costs of S' are at most those of S because in any given round, the probability of any given number of bidders bidding is the same in both mechanisms (cf. Lemma 3), and S cannot terminate (with allocation) before at least one bidder has bid. Moreover, the expected virtual valuation of the winning bidder is at least as high under S' as it is under S . Hence, by Myerson's result, expected total bidder spending is at least as high under S' as it is under S . Therefore, $E_{S'}[\text{profit}] = E_{S'}[\text{bidder spending}] - E_{S'}[\text{total cost}] \geq E_S[\text{bidder spending}] - E_S[\text{total cost}] = E_S[\text{profit}]$, so S' must also be optimal. \square

For a given MVA, let $l = \inf_i a_i$ be the value below which we do not allocate. It is easy to see that if $l > 0$ and $b > 0$, then the sequence of a_i will be finite (and the last one will be l). This is because if the sequence were infinite, there would be a nonzero probability of asking infinitely many broadcast queries. We next discuss how to optimize, given l , the remaining a_i . We note that when l is fixed, this comes down to a query cost minimization problem similar to those studied earlier in this paper, and we can assume without loss of generality that the distribution is uniform by an appropriate transformation. To find the optimal l , we resort to brute-force search over all (discretized) values above Myerson's reserve value.

4.3 Optimal MVAs Given a Positive Low Value

We now consider how to minimize cost for a given low value $l > 0$ of the MVA, restricting our attention to uniform distributions without loss of generality. A key observation is that if all bidders

remain silent to a query, then we can renormalize the resulting conditional distribution to be over $[0, 1)$ as we did in the efficient case; however, in this case, the resulting problem is not identical to the original, but rather has a larger value of l . Using this insight, we can solve a discretized version of the problem, using dynamic programming to compute the optimal first query as a function of (a discretized set) l , starting with large l and working our way down. The result is plotted in Figure 4 and suggest that a result as elegant as for the efficient case is out of the question. Nevertheless, some sense can be made of it. The graph is *almost* piecewise linear. The rightmost piece is linear and corresponds to asking a single query at l . The second-to-rightmost piece is not exactly linear, but corresponds to asking a query that, if unsuccessful, will result in being on the rightmost piece.

We know that with $l > 0$, in an optimal query strategy, at most a finite number k of queries will be asked. (k may depend on n and l .) Let us first consider how to optimize the strategy for a given k (which will remove the discontinuities). Let $\mathbf{a} = (a_0, a_1, a_2, \dots, a_k)$ be the vector of thresholds in this optimal MVA (where $a_0 = 1$ for convenience, and $a_k = l$). Let $C(\mathbf{a}, k, b, c, n)$ be the expected total communication cost for the MVA query strategy k, \mathbf{a} .

Let C^i denote the total communication cost that we incur in round i and let $v^* = \max_i v_i$. Then:

$$\begin{aligned} C(\mathbf{a}, k, b, c, n) &= \sum_{i=1}^k E[C^i] \\ &= \sum_{i=1}^k \Pr(v^* < a_{i-1}) E[C^i | v^* < a_{i-1}] \\ &= \sum_{i=1}^k \frac{a_{i-1}^n}{a_0^n} \left(b + \frac{a_{i-1} - a_i}{a_{i-1}} cn \right) \end{aligned} \quad (4)$$

Taking the derivative, we get

$$\frac{\partial C}{\partial a_i} = \frac{n(nc + b)a_i^{n-1} - n(n-1)ca_{i+1}a_i^{n-2} - nca_{i-1}^{n-1}}{a_0^n} \quad (5)$$

Setting this to zero for all i ($1 \leq i \leq k-1$) will give the optimal solution (the boundary cases are clearly not optimal).

In special cases, we can solve this analytically: for example, for $k = 2, n = 3$ we have: $a_1 = (\sqrt{a_2^2(\rho + 3)} + a_0^2 + a_0)/(\rho + 3)$ (where $\rho = b/c$), which indeed is close to linear (when $a_0 \geq a_2$), but not quite. In general, we use an R package called BB [16] to solve these systems of equations. This requires us to first choose k and an initial guess for \mathbf{a} . We thus turn our attention to choosing k . We need the following lemma:

LEMMA 4. Consider two vectors $\mathbf{a} = (a_0, \dots, a_{k_1})$ and $\mathbf{a}' = (a'_0, \dots, a'_{k_2})$ and suppose they are optimal for $l = a_{k_1}$ and $l = a'_{k_2}$, respectively. Then it cannot be the case that $p_2 \leq q_2 < q_1 < p_1$ where $p_2 = a_i, p_1 = a_{i-1}$ are two consecutive thresholds in the first sequence and $q_2 = a'_j, p_1 = a'_{j-1}$ in the second sequence.

PROOF. Let $C(l)$ denote the optimal cost for finding the maximum value above low value l . Then by the optimality of the two sequences we have

$$\begin{aligned} C(q_1) + q_1^n [\rho + n(q_1 - q_2)] &\leq C(p_1) + p_1^n [\rho + n(p_1 - q_2)] \\ C(p_1) + p_1^n [\rho + n(p_1 - p_2)] &\leq C(q_1) + q_1^n [\rho + n(q_1 - p_2)] \end{aligned}$$

Adding these two inequalities and performing some cancellations gives $p_1^n \leq q_1^n$, which contradicts $q_1 < p_1$. \square

THEOREM 7. $k_\alpha^+ = \lceil \log_\alpha(1/l) \rceil$ and $k_\alpha^- \lfloor \log_\alpha(1/l) \rfloor$ are upper and lower bounds for the optimal k . Here, α is the optimal α from the α -MVA (when $l = 0$).




Figure 5: The profit of the optimal MVA and the α -cutoff-MVA are plotted as solid lines. (The optimal value of l is determined by brute-force search.) α and l are plotted in dashed lines, to illustrate that the the only cases where the α -cutoff-MVA performs significantly more poorly are when α (or α^2) are slightly above l , because in this case the α -cutoff-MVA is often forced to ask another query to get below l .

PROOF. For the sake of contradiction, assume $k > k_\alpha^+$ for optimal thresholds $(a_0^*, a_1^*, a_2^*, \dots, a_k^*)$ (where $a_0^* = 1, a_k^* = l$). Points $a_0 = \alpha^0, a_1 = \alpha^1, \dots, a_{k_\alpha^+} = \alpha^{k_\alpha^+}$ constitute an optimal sequence for low value $l' = \alpha^{k_\alpha^+}$ (if there were a better sequence, we could improve on the α -MVA by starting with that sequence instead). This sequence also splits interval $[l', 1)$ into k_α^+ subintervals $[a_i, a_{i-1})$, which together contain k thresholds a_1^*, \dots, a_k^* . Since $k > k_\alpha^+$, at least one subinterval $[a_i, a_{i-1})$ contains two thresholds a_j^*, a_{j-1}^* , so we have $a_i \leq a_j^* < a_{j-1}^* < a_{i-1}$. This contradicts with Lemma 4. The proof of k_α^- is similar. \square

4.4 Experiments

We continue to evaluate on a uniform distribution over valuations. By Theorem 7, we only need to try two values of k . For the initial thresholds \mathbf{a} , Lemma 4 suggests initializing $a_i = \alpha^i$. We refer to this initialization as the α -cutoff-MVA. Indeed, this initialization allows the solver to converge to the optimal solution, unlike simple heuristics such as spreading the a_i uniformly. In fact, even the α -cutoff-MVA by itself (without running the solver) already performs close to optimality, as shown in Figure 5.

Finally, we compare the profit of our mechanisms to that of mechanisms that use only a single round, such as Myerson's mechanism.⁵ These experiments give insight into whether setting l correctly and not ignoring the communication costs are important. If we ignore communication costs, the Myerson auction (a single-shot Vickrey auction with reserve price 0.5) is optimal. But when communication costs exist, this is not necessarily true. If $\beta_1 > 0$ and $\beta_2 = 0$, one may benefit from setting a somewhat higher reserve price to reduce the number of bids. On the other hand, if $\beta_1 = 0$ and $\beta_2 > 0$, bidders just barely above the reserve price will not bid, and one may want to set the reserve price lower. Hence, we also compute the optimal reserve price for a single-shot auction (by brute force) and compare to it.

The results⁶ are shown in Figure 6. On the MVA side, the profits for using the optimal $l, l = 0$, and $l = 0.5$ become close very

⁵We do not consider other, heuristic querying strategies here; intuition about how these perform can be obtained from the experiments in Section 3.5.

⁶We ran experiments for n up to 100; they are not substantially different from $n = 10$




Figure 6: The expected profit, as well as the value of l (the lowest valuation that may bid), for various mechanisms as a function of n . We have broadcast cost $b = 0.1$ and seller and bidder bid costs $\beta_1 = \beta_2 = 0.05$ (thus $c = \beta_1 + \beta_2 = 0.1$).

quickly as n grows. Hence, the choice of l is relatively unimportant for large n . The single-shot mechanisms, on the other hand, perform significantly worse, and the choice of reserve price is quite important. If we set the reserve price low so that even bidders with low valuations bid, the communication costs will get large. If we set the reserve price high, only bidders with very high valuations bid, and total spending is low. The results show that attempting strike a balance and setting the reserve price in between is optimal, but not very effective.

5. CONCLUSION

In this paper, we modeled Internet auctions in which the seller communicates to the bidders with costly broadcast queries, and if a bidder actively participates there is potentially a communication cost to both that bidder and the seller. We focused primarily on multi-round Vickrey auctions (MVAs). We first assumed that (1) efficient allocation is required and (2) the bidders' cost for bidding is zero ($\beta_2 = 0$). We proved that there is an MVA, which we call an α -MVA, that has optimal profit for the seller among all mechanisms. We gave two ways to compute an approximately optimal α . The first obtains a 1.582-approximation and its expected communication cost is bounded by a value that is constant in the number of bidders. The expected cost of the second is guaranteed to converge to the optimal expected cost as the number of bidders grows.

We then dropped the two constraints and considered the general case. We showed that bidders' expected total spending is fixed by the allocation rule (and the expected utility of bidders with value 0). Thus, for a fixed allocation rule, the profit is maximized when expected total cost is minimized. We then proved that MVAs are also optimal in the general case and focused on computing an optimal MVA. The optimal MVA in the general case turns out to be not as elegant and simple as an α -MVA. Nevertheless, we show that by using a modification that we call an α -cutoff-MVA as an initial guess, we can compute the optimal MVA easily.

6. ACKNOWLEDGMENTS

We thank NSF and ARO for support under grants CCF-1101659, IIS-0953756, W911NF-12-1-0550, and W911NF-11-1-0332. We also thank Degao Peng and Yuan Xue for useful discussions.

7. REFERENCES

- [1] L. M. Ausubel and P. Cramton. The optimality of being efficient. 1999.
- [2] V. Conitzer and T. Sandholm. Computational criticisms of the revelation principle. In *ACM EC*, pages 262–263, 2004.
- [3] J. Hassan and S. Jha. On the optimization trade-offs of expanding ring search. In *IWDC*, pages 489–494, 2004.
- [4] P. J.-J. Herings, R. Müller, and D. Vermeulen. Bisection auctions. *SIGecom Exch.*, 8(1):6:1–6:5, July 2009.
- [5] S. Lahaie and D. Parkes. Applying learning algorithms to preference elicitation. In *ACM EC*, pages 180–188, 2004.
- [6] R. McAfee and D. Vincent. Sequentially optimal auctions. *Games and Economic Behavior*, 18(2):246–276, 1997.
- [7] R. P. McAfee and J. McMillan. Search mechanisms. *Journal of Economic Theory*, 44(1):99–123, 1988.
- [8] R. B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [9] N. Nisan and I. Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192–224, 2006.
- [10] D. Parkes. Iterative combinatorial auctions. In *Combinatorial Auctions*, chapter 2, pages 41–77. MIT Press, 2006.
- [11] T. Sandholm and C. Boutilier. Preference elicitation in combinatorial auctions. In *Combinatorial Auctions*, chapter 10, pages 233–263. MIT Press, 2006.
- [12] T. Sandholm and A. Gilpin. Sequences of take-it-or-leave-it offers: Near-optimal auctions without full valuation revelation. In *AAMAS*, pages 1127–1134, 2006.
- [13] D. Sarne, S. Shamoun, and E. Rata. Increasing threshold search for best-valued agents. In *AAAI*, pages 848–853, 2010.
- [14] M. Stegeman. Participation costs and efficient auctions. *Journal of Economic Theory*, 71(1):228–259, October 1996.
- [15] G. Tan and O. Yilankaya. Equilibria in second price auctions with participation costs. *Journal of Economic Theory*, 130(1):205–219, 2006.
- [16] R. Varadhan and P. Gilbert. Bb: An R package for solving a large system of nonlinear equations and for optimizing a high-dimensional nonlinear objective function. *Journal of Statistical Software*, 32(4):1–26, 10 2009.