Incentive-Compatible Online Auctions for Digital Goods

SIMS296A-5/CS294-6 - Class Project Report

Ziv Bar-Yossef
zivi@cs.berkeley.edu

Kris Hildrum
hildrum@cs.berkeley.edu

Felix Wu
felix@cs.berkeley.edu

May 15, 2001

1 Introduction

Traditionally, auctions are useful mechanisms for allocating resources because they achieve a certain kind of economic efficiency: we are interested in transferring resources to the parties that can make the most use of these resources. Auctions attempt to achieve this by giving bidders the opportunity to declare how much they are willing to pay. In a well-designed auction, these bids should be an indication of the value the bidders place on the resources, and these valuations are in turn taken to be a measure of the usefulness to the bidders. Of course, from the seller's perspective, what matters is not the relative usefulness of the item for sale, but rather the revenue generated by the sale. Luckily, in the usual auction scenario, these goals are not in conflict, since one expects to be able to extract the most revenue from whoever values the item most.

Increasingly, however, we are interested in distributing not physical goods, but digital, or information, goods. This includes, for example, electronic books, software, and digital copies of music or movies. The key feature of such goods is that they can be duplicated with negligible marginal cost, and hence, we can think of them as being available in unlimited supply. With scarcity no longer an issue, there is no question of allocation: the economically efficient thing to do is to give a copy to anyone who asks. This course of action, though, would be unpalatable to the seller, who presumably incurred fixed costs in producing the initial copy of the digital good. Hence, we return to auctions, but this time, we use valuation information not to determine how to allocate goods per se, but to determine how to charge for them, with our only goal being to maximize revenue.

The study of revenue-maximizing auctions for goods available in unlimited supply was begun by Goldberg et al. [GHW01]. As with most work on designing auctions, they focus on auctions which are *incentive-compatible*, that is, auctions in which bidders always do best by placing bids which equal their valuations. Restricting to such auctions seems reasonable on a number of fronts: it relieves bidders of the need to make assumptions about the behavior of other bidders, and it gives the seller the most information about the bidders. In what follows, we too will be examining only incentive-compatible auctions.

In the model of Goldberg et al., they assume that bids are presented to the auction mechanism as a set, and that the mechanism makes all pricing decisions at once, after all bids have been received. Such an assumption might be reasonable for selling a digital broadcast, for instance. In this case, the item is transmitted all at once to the buyers, and hence it might be reasonable to have a cutoff time for the bidding. However, most digital goods, for instance, books, music, or software, are sold continuously, without any point at which sales end.

Furthermore, since digital goods can be delivered immediately, buyers would presumably expect a fairly rapid, if not immediate, response to a bid for digital goods. Even in the case of a digital broadcast, where the delivery is not immediate, buyers would want to know quickly whether or not their offers had been accepted. In auctioning a pay-per-view broadcast, for instance, we might want to allow bids to come in mere minutes before the broadcast, or maybe even after the broadcast has begun. It would be clearly unacceptable to delay all decisions about who would receive the broadcast until such a point.

What this suggests is that, particularly for the case of digital goods, it is important to design auctions which work in an online model. That is, we would like incentive-compatible auction mechanisms which decide whether to satisfy a bid and at what price before the next bid arrives. As is usual in the study of online algorithms, we want our auction to perform well relative to the best offline algorithm, on any possible sequence of bids.

In this setting, some care needs to be taken in defining what we mean by the best offline algorithm. One possible view is that the offline algorithm gets the same sequence of bids, but knows the sequence in advance. Note that since we are dealing with incentive-compatible auctions, knowing the sequence of bids is the same as knowing the sequence of valuations. Hence, the offline algorithm can simply offer the item to each bidder for a price which is exactly equal to that bidder's valuation. Such an algorithm achieves a revenue equal to the sum of the valuations, which is the maximum possible.

It seems rather unfair, however, to allow the offline algorithm to know the valuations, since this defeats the purpose of restricting to incentive-compatible auctions. Instead, we can view the offline algorithm as receiving the entire sequence before deciding who to sell to and at what price. If we also require the offline algorithm to be incentive-compatible, then the best offline algorithm is exactly the best algorithm in the model of Goldberg et al. They give a randomized algorithm which achieves revenue at least a constant fraction of the optimal fixed price revenue, F_{opt} (the maximum revenue one could achieve by setting a single price for all bidders). They also show that no (randomized) algorithm can do better than F_{opt}. Hence, we are primarily interested in how well our online algorithms do relative to F_{opt}.

Note that because we also require the offline algorithm to be incentive-compatible, it is actually possible for the online algorithm to do better than the offline algorithm on some sequences of bids, if we assume that the offline algorithm receives the bids as a set, and not as a sequence. The offline result that no algorithm can achieve revenue greater than F_{opt} immediately implies that for a worst-case ordering, no online algorithm can achieve revenue greater than F_{opt} either. However, if we make additional assumptions about the ordering of the bids, for example, that the bids fall into phases of sufficient length and regularity, an online algorithm may be able to take advantage of these patterns in a way that is not possible when the bids are received all at once. Many digital goods have some natural lifecycle, during which valuations as a whole may change in some consistent way. For example, pop music can have little value while it is unknown, a much higher value during the height of its popularity, and gradually decreasing value as it ages. The usual approach is to periodically adjust prices through promotions and sales; one might hope that an online auction mechanism could make these adjustments more automatically, without needing to know ahead of time the frequency or magnitude of the changes.

In what follows, we begin in Section 2 by describing our model. Then, in Section 3, we review previous work in related models. Goldberg et al. consider offline auctions for goods in unlimited supply. On the other hand, Lavi and Nisan [LN00] consider online auctions for goods in limited supply. In general, the latter results are not directly applicable to our setting, since their auctions need to balance the revenue gained by a current sale with the revenue lost from not being able to sell the same item to a later bidder.

In Section 4, we show that any deterministic auction performs poorly in our model. In Section 5, we first give an extremely simple algorithm which performs optimally relative to the total valuation. Since this algorithm completely ignores the bids, this provides further evidence that the total valuation is probably an inappropriate benchmark. Then, we give a more complicated algorithm, which is a variation on an algorithm of Goldberg et al. Our algorithm achieves revenue which is within an $O(\sqrt{\log h})$ factor of F_{opt}, where h is the ratio between the highest and lowest valuations. Hence, we are able to achieve in the online case the same result that the corresponding offline algorithm achieves in the offline case.

Translating offline algorithms into the online setting is a nontrivial task. In particular, there seems to be no clear analogue for the offline random sampling algorithm, the algorithm which achieves optimal revenue relative to F_{opt}. This algorithm randomly divides the bids in half and uses each half to set a fixed price for the other half. The power of random sampling lies in its ability to approximately compute the optimal fixed price without having to look at too many bids. In the online setting, the barrier to computing the optimal fixed price is not that incentive-compatibility limits the number of bids we are allowed to examine, but rather the fact that all future bids are completely unknown. In particular, we could decide to sample from past bids, but since decisions have already been made on all of these bids, we can only be better off using the information from all past bids. Hence, there seems to be no natural way to sample in the online setting.

2 Preliminaries

2.1 The Model

We consider auctions for goods of unlimited supply. In particular, the auctioneer can sell the good to each one of the bidders. We assume each bidder can win at most one item of the good.

We denote by n the number of bidders. Each bidder i has a *valuation* v_i for the auctioned good, which is known only to himself and stays fixed throughout the auction. By normalizing the valuations, we can assume that each v_i is a real number in the interval $[1, h]$ (we do not permit bidders to have null valuations). Notice that h is the ratio between the highest and lowest valuations. The *utility* of bidder i, $u_i(p_i)$, is $v_i - p_i$, if the bidder wins the good at price p_i, and 0 if he does not win the good. The goal of each bidder is to maximize his utility.

An *online auction* receives n bids, $b_1, \ldots, b_n$, one by one in a sequence. The auctioneer, when presented with the bid b_i of bidder i, is supposed to answer with a price p_i *before it gets to see the rest of the bids*. An online auction A can be viewed as a collection of n functions $A_1, \ldots, A_n$, where $A_i(b_1, \ldots, b_i)$ determines the price p_i based on the bids received up to step i. If $p_i \leq b_i$, bidder i wins the good and has to pay p_i; otherwise, bidder i does not win the good, and pays nothing. The functions A_i may be deterministic or randomized. The *total revenue* of an auction, $R_A(b_1, \ldots, b_n)$, is the sum of the payments received. When A is randomized, the total revenue is a random variable, and thus we speak about the *expected total revenue*.

In this paper we consider a restricted model, in which bidders cannot bid twice, may not collude, and cannot control their location in the sequence. These restrictions simplify the analysis, and are justifiable in some real-life situations (such as pay-per-view auctions). Future work will try to address more general models, where no such restrictions are imposed.

A bidder, in general, is allowed to apply an arbitrary bidding strategy to determine his bid. This strategy is (a possibly randomized) function of the bidder's valuation, of the bids seen so far, and of the prices set for these bids. A *truthful* strategy is one that always bids the true valuation: $b_i = v_i$. We concentrate on *incentive-compatible* auctions – auctions for which the dominant strategy of every bidder is the truthful one. That is, for every choice of b_i, and for any sequence of bids of the previous bidders, $b_1, \ldots, b_{i-1}$, $u_i(A_i(b_1, \ldots, b_{i-1}, b_i)) \leq u_i(A_i(b_1, \ldots, b_{i-1}, v_i))$ (for a deterministic A) and $E(u_i(A_i(b_1, \ldots, b_{i-1}, b_i))) \leq E(u_i(A_i(b_1, \ldots, b_{i-1}, v_i)))$ (for a randomized A). Notice that for incentive-compatible auctions, we can view A_i as a function of the bidder valuations.

An auction is called *bid-independent* if the price for bidder i depends only on the previous bids, and not on b_i. That is, for any sequence of bids, $b_1, \ldots, b_{i-1}$, and for any two choices for the i-th bid, b_i and b'_i, $A_i(b_1, \ldots, b_{i-1}, b_i) = A_i(b_1, \ldots, b_{i-1}, b'_i)$ (for a deterministic A) and $E(A_i(b_1, \ldots, b_{i-1}, b_i)) = E(A_i(b_1, \ldots, b_{i-1}, b'_i))$ (for a randomized A). Goldberg et al. [GHW01] show that bid-independence is equivalent to incentive-compatibility for offline auctions. This immediately implies a similar equivalence for online auctions.

2.2 Competitive Analysis

Following [LN00, GHW01], we adopt worst-case *competitive analysis* from the study of online algorithms (see, e.g., [BEY98]) as our means of evaluating the performance of online auctions. We compare the revenue of the online auction on the "worst" sequence of valuations $v_1, \ldots, v_n$ against a "benchmark". For unlimited supplies, Goldberg et al. propose two possible benchmarks:

1. The optimal revenue: $R_{opt}(v_1, \ldots, v_n) = \sum_{i=1}^{n} v_i$; this is the maximum revenue any auction can hope to achieve from the given set of bidders.
2. The optimal fixed price: $F_{opt}(v_1, \ldots, v_n) = \max\{v_i \cdot (n - i + 1) \mid i = 1, \ldots, n\}$, where $v_1 \leq v_2 \leq \ldots \leq v_n$; this is the maximum revenue that can be achieved using a single fixed price for all bidders.

An incentive-compatible online auction A is said to be *c-competitive* relative to a benchmark B, if for any sequence of valuations $\overline{v} = v_1, \ldots, v_n$, $R_A(\overline{v}) \geq B(\overline{v})/c$ (for a deterministic A) and $E(R_A(\overline{v})) \geq B(\overline{v})/c$ (for a randomized A). We call c the *competitive ratio* of A relative to B. Note that for incentive-compatible auctions $R_A(\overline{v})$ is simply $R_A(\overline{b})$, where $\overline{b} = b_1, \ldots, b_n$, and $b_i = v_i$ for all i.

2.3 Valuation Buckets

A notion that will play an important role in the auctions we present in the sequel is the partitioning of the valuation range $[1, h]$ into $\lfloor \log h \rfloor + 1$ sub-intervals, $I_0, \ldots, I_{\lfloor \log h \rfloor}$, where $I_k = [2^k, 2^{k+1})$. Given a valuation sequence $\overline{v}$, we define the *k-th bucket* of $\overline{v}$ to be the collection of valuations that fall in the k-th interval:

$$B_k(\overline{v}) \stackrel{\text{def}}{=} \{v_i \in \overline{v} \mid v_i \in I_k\}$$

The *size* of the k-th bucket is the number of valuations it contains:

$$\ell_k(\overline{v}) \stackrel{\text{def}}{=} |B_k(\overline{v})|$$

The *weight* of the k-th bucket is the sum of valuations in the bucket:

$$W_k(\overline{v}) \stackrel{\text{def}}{=} \sum_{j=1}^{\ell_k(\overline{v})} v_{i_j}$$

where $v_{i_1}, \ldots, v_{i_{\ell_k(\overline{v})}}$ are the valuations in $B_k(\overline{v})$.

3 Previous Work

3.1 Offline Auctions for Unlimited Supply

Goldberg, Hartline and Wright [GHW01] examine auctions of digital (or negligible marginal cost) goods, and present several algorithms for incentive-compatible auctions. They use competitive analysis to compare their algorithms to the optimal fixed price, and in some cases, to the total possible revenue. The paper also presents experimental results for several distributions of bids. We state here their main results.

Theorem 1 *The optimal fixed price is $O(\log h)$-competitive relative to the optimal revenue.*

We omit the proof here, but wish to note that we can also show that the competitive ratio is tight, i.e., that the optimal fixed price is $\Omega(\log h)$-competitive relative to the optimal revenue.

Theorem 2 *Every randomized auction is $\Omega(1)$-competitive relative to the optimal fixed price.*

This paper also includes a proof that any deterministic incentive-compatible auction is $\Omega(h)$-competitive relative to the optimal fixed price. However, the proof works only for "symmetric" auctions, and online auctions are inherently non-symmetric. So we provide an analogous proof for the online case.

The first auction Goldberg et al. present is based on sampling. Pick roughly half the bids at random, reject them, and determine what would have been the ideal fixed price for the rejected bids. That fixed price is then used as the price for the unrejected bids.

Since it is somewhat unsatisfactory to summarily reject half the bids, they present two modifications to this algorithm. If it is permissible to have two prices for the same item, divide the bids into two roughly equal classes. Class 1 determines the fixed price for class 2, and class 2 determines the fixed price for class 1. If a dual-price auction is not acceptable, one could reduce the sample size to about $\sqrt{n}$, which they found balances the accuracy of the sample with the revenue lost because the sample bids are rejected.

Theorem 3 *The random sampling auction is O(1)-competitive relative to the optimal fixed price.*

The second auction algorithm is *weighted pairing.* For each bid b_i, choose another bid b_j at random with probability $\frac{b_j}{\sum_{k \neq i} b_k}$. Then use b_j as p_i, the price of the item.

Theorem 4 *The weighted pairing auction is* $\Theta(\log h)$*-competitive relative to the optimal revenue and* $\Theta(\sqrt{\log h})$*-competitive relative to the optimal fixed price.*

It is not always clear how to make the offline algorithms work in the online setting. Even supposing that we can make them work online, the performance guarantees may not hold any more.

3.2 Online Auctions for Limited Supply

In [LN00], Lavi and Nisan give incentive-compatible online auctions for the case of multiple goods in limited supply.

In their model, the seller is auctioning k identical indivisible goods. At some point, each bidder learns his or her utility and makes a bid. These bids are not single prices; they are functions of the quantity q. This bid is immediately accepted or rejected by the auctioneer. For very large k, this is the same as auctioning an infinitely divisible good (such as network bandwidth). In the offline case, the natural way to hold such an auction is as a Vickrey auction, so that is what they use as their benchmark. All the theorems below are from [LN00].

Lavi and Nisan first present a characterization of incentive-compatible online auctions for goods in limited supply using the following notion of "supply curves":

Definition 5 (Supply curves) *An online auction is called "based on supply curves" if before receiving the i'th bid it fixes a function (supply curve)* $p_i(q)$ *based on previous bids and*

1. *The quantity q_i sold to bidder i is the quantity q that maximizes the sum $\sum_{j=1}^{q}(b_i(j) - p_i(j))$.*
2. *The price paid by agent i is $\sum_{j=1}^{q_i} p_i(j)$.*

Theorem 6 *Any online auction for limited supply goods is incentive-compatible if and only if it is based on supply curves.*

The paper is mainly concerned with supply curves for which $p_i(q)$ is non-decreasing.

Definition 7 (Global supply curve) *A global supply curve is a set of supply curves such that $p_i(q)$ (the curve for bidder i) is $p(q + \sum_{j=1}^{i-1} q_j)$, where $p(q)$ is a global curve and q_j is the quantity sold to the j'th bidder.*

Theorem 8 *For a properly chosen $p(q)$, an online auction based on $p(q)$ is $k \cdot h^{\frac{1}{k+1}}$-competitive relative to the Vickrey auction. If $k \geq 2\ln h$, then it is $2e(\ln h + 1)$-competitive relative to the Vickrey auction.*

Directly applying these results to the infinite supply case does not give useful results. The logical thing would be to let the number of objects equal the number of bidders. However, it is unclear how to write the function $p(q)$ when the number of bidders (and hence the number of objects) is unknown. Second, the Vickrey auction is not a good basis for comparison in the infinite supply case. The Vickrey auction would sell the item to everyone at the lowest possible price, which would result in very little revenue compared to the optimal fixed price.

Further, this auction behaves in a counter-intuitive way. As items are sold, the price goes up. When items are in limited supply, this seems reasonable. However, in an auction with an infinite number of goods with little or no marginal cost the price should not be forced to go upward.

4 Deterministic Online Auctions for Unlimited Supply

We first show that no deterministic online auction is competitive relative to the optimal fixed price:

Theorem 9 *Any deterministic incentive-compatible online auction A is $\Omega(h)$-competitive relative to the optimal fixed price, even when restricting to valuations $\overline{v}$ with $F_{opt}(\overline{v}) \geq \alpha h$ for any constant $\alpha \geq 1$.*

Proof. Let A be any deterministic incentive-compatible online auction; by [GHW01], A is bid-independent. Therefore, for all i, the price for bidder i depends only on $b_1, \ldots, b_{i-1}$.

We will construct a valuation sequence $\overline{v}$, for which $R_A(\overline{v}) \leq F_{opt}(\overline{v})/h$. $\overline{v}$ will be a *bipolar sequence*, that is, consisting of valuations only in $\{1, h\}$. Note that for such a sequence, we can assume without loss of generality that A sets prices only in $\{1, h\}$.

Fix any constant $\alpha \geq 1$. We choose the valuation sequence $\overline{v}$ as follows: if $A_i(b_1, \ldots, b_{i-1}) = 1$, pick $v_i = h$, and if $A_i(b_1, \ldots, b_{i-1}) = h$, pick $v_i = 1$. We stop when either the number of 1-valuations in $\overline{v}$ exceeds αh, or the number of h-valuations exceeds α. Note that the revenue of A from bidder i is 0 if $v_i = 1$, and 1 if $v_i = h$. Therefore, if we denote by n_h the number of h-valuations in $\overline{v}$ and by n_ℓ the number of 1-valuations in $\overline{v}$, we have:

$$R_A(\overline{v}) \;=\; n_h \cdot 1 + n_\ell \cdot 0 \;=\; n_h$$

On the other hand, $F_{opt}(\overline{v}) \geq \max\{hn_h, n_\ell\}$, implying that $R_A(\overline{v}) \leq F_{opt}(\overline{v})/h$ with $F_{opt}(\overline{v}) \geq \alpha h$. $\square$

5 Randomized Online Auctions for Unlimited Supply

5.1 Simple Randomized Auction

We first present a simple incentive-compatible randomized online auction that achieves a competitive ratio of $O(\log h)$ relative to the optimal revenue. This is comparable to the randomized *offline* auctions presented in [GHW01].

The auction we propose, named S for "Simple", has to know a priori the bound h on the range of valuations. For each bidder i, S picks a random $k \in \{0, \ldots, \lfloor \log h \rfloor\}$, and sets $p_i = 2^k$. It is clear that this auction is incentive-compatible, since the prices are totally independent of the bids. We next prove that it is competitive relative to the optimal revenue:

Theorem 10 *The auction S is $\Theta(\log h)$-competitive relative to the optimal revenue and relative to the optimal fixed price.*

Proof. We first show that S is $O(\log h)$-competitive relative to the optimal revenue. This immediately implies that it is also $O(\log h)$-competitive relative to the optimal fixed price.

Consider any valuation sequence $\overline{v}$. For each valuation v_i, denote by k_i the maximal integer k for which $2^k \leq v_i$. Note that $v_i/2 \leq 2^{k_i} \leq v_i$, and thus if S sets $p_i = 2^{k_i}$, it is guaranteed to earn a payment of at least $v_i/2$ from bidder i. Define an indicator random variable X_i to be 1 if S picks $p_i = 2^{k_i}$, and 0 otherwise. Note that

$$E(R_S(\overline{v})) \;\geq\; E(\sum_{i=1}^{n} \frac{v_i}{2} X_i) \;=\; \frac{1}{2} \sum_{i=1}^{n} v_i E(X_i)$$

Since S picks k uniformly at random, $E(X_i) = \Pr(k_i \text{ is chosen}) \geq 1/(\log h + 1)$. It follows that

$$E(R_S(\overline{v})) \;\geq\; \frac{1}{2} \sum_{i=1}^{n} v_i \frac{1}{\log h + 1} \;=\; \frac{1}{2(\log h + 1)} R_{opt}(\overline{v})$$

We now show that S is $\Omega(\log h)$-competitive relative to the optimal fixed price. This would imply that it is also $\Omega(\log h)$-competitive relative to the optimal revenue.

Fix some constant $\alpha \geq 1$, and let $\overline{v}$ be the valuation sequence consisting of $\lceil \alpha \rceil$ h's. Clearly, $F_{opt}(\overline{v}) \geq \alpha h$. However, the expected revenue of S from every bidder i is $O(h/\log h)$, and thus the total expected revenue is $O(\alpha h/\log h) \leq F_{opt}(\overline{v})/\log h$. □

It is possible to modify S to work also when h is not known in advance. The auction will start with assuming $h = 1$, and continuously update h according to the maximum bid seen so far. Note that the modified S might miss additional revenue from at most one bid in each bucket. However, the sum of these bids cannot be more than $2h$, and thus, assuming $R_{opt}(\overline{v}) \geq 4h$, S can still extract an $O(\log h)$ fraction of $R_{opt}(\overline{v})$.

5.2 Weighted Buckets Auction

We next present a slightly more involved online auction that achieves the same competitive ratio relative to the optimal revenue, but has a better competitive ratio relative to the optimal fixed price. It is based on an adaptation of the weighted pairing auction of Goldberg et al. to the online setting.

The new auction, named W, maintains the weights of the valuation buckets based on the bids seen so far. We define the k-th bucket at step i to be the set of valuations from $v_1, \ldots, v_{i-1}$ that fall in the k-th interval:

$$B_k^i(\overline{v}) = \{v_j \in \overline{v} \ \mid \ j \in \{1, \ldots, i-1\}, v_j \in I_k\}$$

As usual, $\ell_k^i(\overline{v})$ is the size of the k-th bucket, and $W_k^i(\overline{v})$ is the weight of the bucket. As we will see later, W is incentive-compatible, and can thus compute the valuation buckets based on the truthful bids.

The auction W uses the bucket weights to pick a random bucket, and sets the price for bidder i to be the lower limit of the corresponding bucket. More formally, W picks $p_i = 2^k$ with probability $W_k^i(\overline{v})/(\sum_{r=0}^{\lfloor \log h \rfloor} W_r^i(\overline{v}))$. For the first bid, W picks a random $k \in \{0, \ldots, \lfloor \log h \rfloor\}$, and sets $p_1 = 2^k$.

Note that W is bid-independent, and therefore is incentive-compatible. We next prove that it is as good as the simple auction relative to the optimal revenue, but has better performance relative to the optimal fixed price. The competitive ratios hold only for valuation sequences in which h is not too large relative to the optimal revenue and the optimal fixed price.

Theorem 11 *(1) Restricting to valuation sequences with $R_{opt}(\overline{v}) \geq 4h$, W is $\Theta(\log h)$-competitive relative to the optimal revenue.*

(2) Restricting to valuation sequences with $F_{opt}(\overline{v}) \geq 4h$, W is $\Theta(\sqrt{\log h})$-competitive relative to the optimal fixed price.

Proof. We first prove that W is $O(\log h)$-competitive relative to the optimal revenue. Fix some valuation sequence $\overline{v}$, for which $R_{opt}(\overline{v}) \geq 4h$. Let K_1 be the set of buckets k, for which $\ell_k(\overline{v}) \leq 1$, and let K_2 be the rest of the buckets. Note that the sum of the weights of the buckets in K_1 is at most $2h$, and thus $\sum_{k \in K_2} W_k(\overline{v}) \geq R_{opt}(\overline{v})/2$.

Fix some bucket $k \in K_2$, and let $v_{i_1}, \ldots, v_{i_{\ell_k}}$ be the valuations in this bucket. Assume that the bids of the corresponding bidders arrive in this order: $b_{i_1}, \ldots, b_{i_{\ell_k}}$. We will show that the expected revenue from the last $\lceil l_k/2 \rceil$ bidders is high relative to the total weight of the bucket.

Fix some $\lfloor l_k/2 \rfloor + 1 \le j \le l_k$. When bidder i_j arrives, W has already seen bidders $i_1, \ldots, i_{j-1}$, and thus has some positive probability to pick the k-th bucket. If it does so, bidder i_j wins the good at price 2^k, since $v_{i_j} \ge 2^k$. Therefore, we can bound the expected revenue from bidder i_j as follows:

$$E(R_W(v_{i_j})) \;\ge\; 2^k \cdot \Pr(\text{bucket } k \text{ is chosen}) \;=\; 2^k \cdot \frac{W_k^{i_j}}{\sum_{r=0}^{\lfloor \log h \rfloor} W_r^{i_j}} \;\ge\; 2^k \cdot \frac{\sum_{r=1}^{\lfloor l_k/2 \rfloor} v_{i_r}}{\sum_{r=0}^{\lfloor \log h \rfloor} W_r(\overline{v})} \;=\; 2^k \cdot \frac{\sum_{r=1}^{\lfloor l_k/2 \rfloor} v_{i_r}}{R_{opt}(\overline{v})}$$

Note that since all the valuations in the k-th bucket are at most a factor of two away from each other, the sum of any $\lfloor l_k/2 \rfloor$ of them cannot be less than $1/4$ of their total sum. We thus conclude:

$$E(R_W(v_{i_j})) \;\ge\; 2^k \cdot \frac{W_k(\overline{v})}{4R_{opt}(\overline{v})}$$

Thus, the total expected revenue from the k-th bucket is:

$$E(R_W(B_k(\overline{v}))) \;\ge\; \sum_{j=\lfloor l_k/2 \rfloor + 1}^{l_k} E(R_W(v_{i_j})) \;\ge\; \lceil \frac{l_k}{2} \rceil \cdot 2^k \cdot \frac{W_k(\overline{v})}{4R_{opt}(\overline{v})} \;\ge\; \frac{W_k(\overline{v})^2}{16R_{opt}(\overline{v})}$$

where the last inequality follows from the fact that since $2^k \ge v_{i_j}/2$ for any $v_{i_j} \in B_k(\overline{v})$, then $2^k \lceil l_k/2 \rceil \ge W_k(\overline{v})/4$.

Now the total expected revenue of W is:

$$E(R_W(\overline{v})) \;\ge\; \sum_{k \in K_2} E(R_W(B_k(\overline{v}))) \;\ge\; \sum_{k \in K_2} \frac{W_k(\overline{v})^2}{16R_{opt}(\overline{v})} \;\ge\; \frac{1}{16R_{opt}(\overline{v})} \cdot \frac{(\sum_{k \in K_2} W_k(\overline{v}))^2}{|K_2|} \tag{1}$$

where the last inequality follows from the Cauchy-Schwartz inequality. We now use the fact that $\sum_{k \in K_2} W_k(\overline{v}) \ge R_{opt}(\overline{v})/2$ to obtain the desired competitive ratio relative to the optimal revenue:

$$E(R_W(\overline{v})) \;\ge\; \frac{1}{16R_{opt}(\overline{v})} \cdot \frac{R_{opt}(\overline{v})^2}{4(\log h + 1)} \;=\; \frac{R_{opt}(\overline{v})}{64(\log h + 1)}$$

This bound is tight, due to the general theorem of Goldberg et al., who prove that $O(\log h)$ is the best any randomized auction can achieve relative to the optimal revenue.

In order to prove the bound on the competitive ratio relative to the optimal fixed price, we note the following characterization of the optimal fixed price. For a valuation sequence $\overline{v}$, denote by k^* the bucket with maximum weight (in case of ties, k^* is an arbitrary bucket with maximum weight). We prove:

Claim 12 $W_{k^*}(\overline{v})/2 \le F_{opt}(\overline{v}) \le 3W_{k^*}(\overline{v})$

Proof. Note that for any $v_i \in B_{k^*}(\overline{v})$, $v_i/2 \leq 2^{k^*} \leq v_i$. Thus a fixed price of 2^{k^*} will yield the following revenue:

$$F_{opt}(\overline{v}) \ \geq \ 2^{k^*} \ell_{k^*}(\overline{v}) \ \geq \ \sum_{j=1}^{\ell_{k^*}(\overline{v})} \frac{v_{i_j}}{2} = \frac{W_{k^*}(\overline{v})}{2}$$

For the other direction, assume p is the optimal fixed price, and let k be the integer such that $2^k \leq p < 2^{k+1}$. Note that for any $r > k$, $W_r(\overline{v}) \geq 2^r \ell_r(\overline{v}) \geq 2^{r-k-1} \cdot p \cdot \ell_r(\overline{v})$. Therefore:

$$\begin{aligned} F_{opt}(\overline{v}) \ &\leq \ W_k(\overline{v}) + p \cdot \sum_{r=k+1}^{\lfloor \log h \rfloor} \ell_r(\overline{v}) \ \leq \ W_k(\overline{v}) + \sum_{r=k+1}^{\lfloor \log h \rfloor} \frac{W_r(\overline{v})}{2^{r-k-1}} \\ &\leq \ W_{k^*}(\overline{v}) \cdot (1 + \sum_{r=k+1}^{\lfloor \log h \rfloor} \frac{1}{2^{r-k-1}}) \ \leq \ 3W_{k^*}(\overline{v}) \end{aligned}$$

□

We now use Claim 12 to prove the $O(\sqrt{\log h})$ bound on the competitive ratio of W relative to the optimal fixed price. Fix some valuation sequence $\overline{v}$ with $F_{opt}(\overline{v}) \geq 4h$. If $F_{opt}(\overline{v}) \leq R_{opt}(\overline{v})/\sqrt{\log h}$, then we have $E(R_W(\overline{v})) \geq R_{opt}(\overline{v})/O(\log h) \geq F_{opt}(\overline{v})/O(\sqrt{\log h})$. So assume $F_{opt}(\overline{v}) \geq R_{opt}(\overline{v})/\sqrt{\log h}$.

Let k^* be the maximum weight bucket of $\overline{v}$. The size of this bucket, ℓ_{k^*}, has to be at least two, because otherwise $W_{k^*}(\overline{v}) \leq h$, implying by Claim 12 that $F_{opt}(\overline{v}) \leq 3h$. Now, using what we proved in (1), we have:

$$E(R_W(\overline{v})) \ \geq \ \frac{1}{16R_{opt}(\overline{v})} \cdot \sum_{k \in K_2} W_k(\overline{v})^2 \ \geq \ \frac{W_{k^*}(\overline{v})^2}{16R_{opt}(\overline{v})} \ \geq \ \frac{F_{opt}(\overline{v})^2}{144R_{opt}(\overline{v})} \ \geq \ \frac{F_{opt}(\overline{v})}{144\sqrt{\log h}}$$

Since W can be viewed as a special case of the weighted pairing auction of Goldberg et al., the $\Omega(\sqrt{\log h})$ lower bound they prove for the competitive ratio of weighted pairing relative to the optimal fixed price holds also for W. □

6 Future Work

We were able to show a number of initial results about our model; more work remains to be done both to strengthen our results, and to extend our model. One primary goal would be to design a randomized algorithm which achieves a constant competitive ratio with respect to F_{opt}. We had a number of ideas for ways to modify our existing algorithms. Straightforward modifications of the simple algorithm yield a competitive ratio of $\Omega(\log h)$ relative to the fixed price. We also considered a number of ways to modify the weights in the weighted buckets algorithm. Some ideas along these lines seem promising, and at the very least, we believe it should be possible to improve the current bound of $O(\sqrt{\log h})$.

There are also a number of extensions to the model which could be considered. In our current results, we assume that bidders bid only once, and that they bid independently of each other. In real life, collusion, at least on a small scale, is a definite possibility, and we would like to ensure that our auction mechanisms are at least somewhat collusion-resistant. Furthermore, although we might be able to prevent individual bidders from flooding the auction with low bids, it seems unlikely that we can prevent the submission of any false bids at all. In addition, there could be legitimate reasons for a buyer to submit multiple bids, for example, if his valuation suddenly increases. Handling this situation, while maintaining incentive-compatibility, would be a challenge.

References

[BEY98] Allan Borodin and Ran El-Yaniv. *Online Computation and Competitive Analysis.* Cambridge University Press, 1998.

[GHW01] Andrew Goldberg, Jason Hartline, and Andrew Wright. Competitive auctions for digital goods. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2001.

[LN00] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *Proceedings of the 2nd ACM Conference on Electronic Commerce (EC-00)*, pages 233–241, 2000.