

Online Learning in Online Auctions [★]

Avrim Blum

*Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA*¹

Vijay Kumar

Strategic Planning and Optimization Team, Amazon.com, Seattle, WA

Atri Rudra

*Department of Computer Science, University of Texas at Austin, Austin, TX*²

Felix Wu

*Computer Science Division, University of California at Berkeley, Berkeley, CA*³

Abstract

We consider the problem of revenue maximization in online auctions, that is, auctions in which bids are received and dealt with one-by-one. In this paper, we demonstrate that results from online learning can be usefully applied in this context, and we derive a new auction for digital goods that achieves a constant competitive ratio with respect to the optimal (offline) fixed price revenue. This substantially improves upon the best previously known competitive ratio for this problem of $O(\exp(\sqrt{\log \log h}))$ [4]. We also apply our techniques to the related problem of designing online *posted price* mechanisms, in which the seller declares a price for each of a series of buyers, and each buyer either accepts or rejects the good at that price. Despite the relative lack of information in this setting, we show that online learning techniques can be used to obtain results for online posted price mechanisms which are similar to those obtained for online auctions.

[★] Portions of this work appeared as an extended abstract in Proceedings of SODA'03 [5].

Email addresses: avrim@cs.cmu.edu (Avrim Blum), vijayk@amazon.com (Vijay Kumar), atri@cs.utexas.edu (Atri Rudra), felix@cs.berkeley.edu (Felix Wu).

¹ Supported in part by National Science Foundation grants CCR-0105488 and IIS-0121678.

² Work done while the author was at IBM India Research Lab, New Delhi, India.

³ Supported in part by National Science Foundation ITR grant CCR-0121555.

1 Introduction

Auctions are traditional and well-studied economic mechanisms, and economists have long studied the design of auctions intended to satisfy various goals, including that of maximizing the total revenue obtained by the auctioneer from the auction. Traditionally, however, economists have analyzed auctions under the assumption that statistical information about the participating bidders is available. Recent work in computer science has been directed toward designing auctions in the absence of such statistical assumptions, using instead a form of worst-case competitive analysis [3,4,7,10,12].

The proliferation of Internet auctions and the increasing availability of media on the Internet has prompted particular attention to the design of auctions for digital goods, that is, goods available in unlimited supply [7,10]. In this paper, we focus on such goods, though our techniques may also be useful in the case of limited supply goods. A key property of digital goods is that it will often be useful to conduct auctions of such goods over time, with bidders arriving one-by-one, rather than as a group. Hence, we are interested here in designing *online* auctions for digital goods, a problem first described by Bar-Yossef et al. [4].

In the model of Bar-Yossef et al. [4], n bidders arrive in a sequence. Each bidder i is interested in one copy of the good, and values this copy at v_i . The valuations are normalized to some range $[1, h]$, so that h is the ratio between the highest and lowest possible valuations. Bidder i places bid b_i , and the auction must then determine whether to sell the good to bidder i , and if so, at what price $s_i \leq b_i$. This is equivalent to determining a sales price s_i , such that if $s_i \leq b_i$, bidder i wins the good and pays s_i ; otherwise, bidder i does not win the good and pays nothing.

The utility of a bidder is then given by $v_i - s_i$ if bidder i wins; 0 if bidder i does not win. As in Bar-Yossef et al. [4], we are interested in auctions which are incentive-compatible, that is, auctions in which each bidder's utility is maximized by bidding truthfully and setting $b_i = v_i$. As shown in that paper, this condition is equivalent to the condition that each s_i depends only on the first $i - 1$ bids, and not on the i th bid. Hence, the auction mechanism is essentially trying to guess the i th valuation, based on the first $i - 1$ valuations.

Note that in an online auction, the sales prices s_i are not actually revealed to the bidders, since we need the bidders to declare their valuations, so that the auction can use this information in dealing with future bidders. In auctions conducted remotely over networks, however, the bidders may not trust the auctioneer to set sales prices before seeing the next bid. Buyers would clearly prefer to receive these sales prices directly and then to make a decision ac-

cordingly whether or not to purchase the good. (Buyers purchase if and only if $s_i \leq v_i$.) We call such a mechanism a *posted price* mechanism [11]. The trade-off in using such a mechanism is that in exchange for the greater trust of the buyers, the seller loses the complete information about the buyers' valuations.

As in previous papers [4,10,12], we will use competitive analysis to analyze the performance of any given auction or mechanism. That is, we are interested in the worst-case ratio (over all sequences of valuations) between the revenue of the “optimal offline” auction and the revenue of the online auction. Following previous papers [4,10], we take the optimal offline auction to be the one which optimally sets a single fixed price for all bidders. Thus, our goal is what is sometimes called “static optimality.” The revenue of the optimal fixed price auction is given by $\mathcal{F}(\bar{v}) = \max_i \{v_i n_i\}$, where $n_i = |\{j \mid v_j \geq v_i\}|$ is the number of bidders with valuation at least v_i . An online auction A with revenue $R_A(\bar{v})$ is said to be *c-competitive* if for any sequence \bar{v} , $R_A(\bar{v}) \geq \mathcal{F}(\bar{v})/c$. We take R_A to be the expected revenue if A is randomized.

In Section 2, we present an asymptotically constant-competitive online auction for digital goods. By asymptotically, we mean that our auction achieves a revenue which is a constant fraction of \mathcal{F} , but minus an additive term. (In our case, this term is $O(h \ln \ln h)$.) Hence, as \mathcal{F} becomes large, this additive term becomes negligible. Nevertheless, it is important to minimize this term, since it roughly corresponds to the size of the smallest auctions for which we can give good revenue bounds. Theorem 4 gives a general lower bound showing that our additive constant is nearly optimal: in particular, any constant-competitive algorithm must have an additive constant $\Omega(h)$.

In Section 3, we derive a similar result for the problem of designing online posted price mechanisms. (Offline posted price mechanisms have been previously studied by Hartline [11].) Such mechanisms provide much less information to the auctioneer about the bidders' valuations, but surprisingly, we are still able to obtain results very similar to those obtained in the online auction setting.

Our results are based on application of machine learning techniques to the online auction problem. Setting a single fixed price for the auction can be thought of as following the advice of a single “expert” who predicts that fixed price for every bidder. Performing well relative to the optimal fixed price is then equivalent to performing well relative to the best of these experts, a problem well-studied in learning theory [2,6,8,9,13]. The posted price setting then corresponds to a version of the “bandit” problem [2], in which the information received depends on the expert chosen at each step. Our algorithms are derived by adapting these techniques to the online auction setting.

2 Online auctions: the full information game

We use a variant of Littlestone and Warmuth’s weighted majority (**WM**) algorithm [13] given in Auer et al. [1,2]. In our context, let $X = \{x_1, \dots, x_\ell\}$ be a set of candidate fixed prices, corresponding to a set of experts. Let $r_k(\bar{v})$ be the revenue obtained by setting the fixed price x_k for the valuation sequence \bar{v} , and let $\mathcal{F}_X(\bar{v}) = \max_k r_k(\bar{v})$ be the optimal fixed price revenue on sequence \bar{v} , when restricted to fixed prices in X .

Given a parameter $\alpha \in (0, 1]$, define weights $w_k(i) = (1 + \alpha)^{r_k(v_1, \dots, v_i)/h}$. Clearly, the weights can be easily maintained using a multiplicative update. Then, for bidder i , the auction chooses $s_i \in X$ with probability

$$p_k(i) = \Pr[s_i = x_k] = \frac{w_k(i-1)}{\sum_{j=1}^{\ell} w_j(i-1)}.$$

This algorithm is shown in Figure 1.

Algorithm WM
Parameters: Reals $\alpha \in (0, 1]$ and $X \in [1, h]^\ell$.
Initialization: For each expert k , initialize $r_k() = 0$, $w_k(0) = 1$.
For each bidder $i = 1, \dots, n$:
 Set the sales price s_i to be x_k with probability $p_k(i) = \frac{w_k(i-1)}{\sum_{j=1}^{\ell} w_j(i-1)}$.
 Observe $b_i = v_i$.
 For each expert k , update $r_k(v_1, \dots, v_i)$ and $w_k(i) = (1 + \alpha)^{r_k(v_1, \dots, v_i)/h}$.

Fig. 1. **WM** in our setting

The following theorem appears in Auer et al., with the proof adapted from proofs appearing in Freund and Schapire [9] and Littlestone and Warmuth [13].

Theorem 1 [1, Theorem 3.2] *For any sequence of valuations \bar{v} , the revenue of auction **WM** is at least:*

$$R_{\mathbf{WM}}(\bar{v}) \geq (1 - \frac{\alpha}{2})\mathcal{F}_X(\bar{v}) - \frac{h \ln \ell}{\alpha}.$$

For completeness, we provide a proof here.

Proof. Let $g_k(i)$ denote the revenue gained by the k th expert from bidder i , that is, $g_k(i) = x_k$, if $v_i \geq x_k$, and $g_k(i) = 0$ otherwise. Then, $r_k(v_1, \dots, v_i) = g_k(i) + r_k(v_1, \dots, v_{i-1})$. Let $W(i) = \sum_{k=1}^{\ell} w_k(i)$ be the sum of the weights after bidder i .

The expected revenue of the auction from bidder $i + 1$ is given by

$$g_{\mathbf{WM}}(i + 1) = \frac{\sum_{k=1}^{\ell} w_k(i) g_k(i + 1)}{W(i)}.$$

We can then relate the change in $W(i)$ to the expected revenue of the auction as follows:

$$\begin{aligned} W(i + 1) &= \sum_{k=1}^{\ell} w_k(i) (1 + \alpha)^{g_k(i+1)/h} \\ &\leq \sum_{k=1}^{\ell} w_k(i) (1 + \alpha(g_k(i + 1)/h)) \\ &= W(i) + \alpha \sum_{k=1}^{\ell} w_k(i) (g_k(i + 1)/h) \\ &= W(i) (1 + \alpha(g_{\mathbf{WM}}(i + 1)/h)), \end{aligned}$$

where for the inequality, we used the fact that for $x \in [0, 1]$, $(1 + \alpha)^x \leq 1 + \alpha x$.

Since $W(0) = \ell$, we have

$$W(n) \leq \ell \cdot \prod_{i=1}^n (1 + \alpha(g_{\mathbf{WM}}(i)/h)).$$

On the other hand, the sum of the final weights is at least the value of the maximum final weight. Hence, $W(n) \geq (1 + \alpha)^{\mathcal{F}_X/h}$.

Taking logs, we have

$$\frac{\mathcal{F}_X}{h} \ln(1 + \alpha) \leq \ln \ell + \sum_{i=1}^n \ln(1 + \alpha(g_{\mathbf{WM}}(i)/h)).$$

For $x \in [0, 1]$, $x - \frac{x^2}{2} \leq \ln(1 + x) \leq x$; hence,

$$\frac{\mathcal{F}_X}{h} \left(\alpha - \frac{\alpha^2}{2} \right) \leq \ln \ell + \frac{\alpha}{h} R_{\mathbf{WM}}.$$

Rearranging this inequality yields the theorem. \square

Now let X consist of all powers of $(1 + \beta)$ between 1 and h . If we take $\alpha = \beta = \frac{\epsilon}{3}$, we get the following theorem.

Theorem 2 For any $\epsilon \in (0, 1]$, restricting to valuation sequences with $\mathcal{F}(\bar{v}) \geq \frac{24h}{\epsilon^2}(\ln \ln h + \ln(\frac{4}{\epsilon}))$, auction **WM** with $\alpha = \frac{\epsilon}{3}$ and X consisting of all powers of $(1 + \frac{\epsilon}{3})$ is $(1 + \epsilon)$ -competitive relative to the optimal fixed price revenue.

Proof. First note that $\mathcal{F}(\bar{v}) \leq (1 + \beta)\mathcal{F}_X(\bar{v})$, since rounding down to a power of $(1 + \beta)$ loses at most a factor of $(1 + \beta)$ in the revenue. From Theorem 1, we have

$$R_{\mathbf{WM}}(\bar{v}) \geq (1 - \frac{\alpha}{2})\mathcal{F}_X(\bar{v}) - \frac{h \ln \ell}{\alpha}.$$

Note that $\ln(1 + \beta) \geq \beta - \beta^2/2 = \epsilon/3 - \epsilon^2/18 \geq \epsilon/4$. Hence, by construction,

$$\frac{h \ln \ell}{\alpha} = \frac{h \ln(\frac{\ln h}{\ln(1+\beta)})}{\frac{\epsilon}{3}} \leq \frac{3h}{\epsilon}(\ln \ln h + \ln(\frac{4}{\epsilon})) \leq \frac{\epsilon}{8}\mathcal{F}(\bar{v}).$$

Thus,

$$\begin{aligned} R_{\mathbf{WM}}(\bar{v}) &\geq (1 - \frac{\epsilon}{6})\frac{\mathcal{F}(\bar{v})}{(1 + \frac{\epsilon}{3})} - \frac{\epsilon}{8}\mathcal{F}(\bar{v}) \\ &\geq (1 - \frac{\epsilon}{6} - \frac{\epsilon}{8}(1 + \frac{\epsilon}{3}))\frac{\mathcal{F}(\bar{v})}{(1 + \frac{\epsilon}{3})} \\ &\geq (1 - \frac{\epsilon}{3})\frac{\mathcal{F}(\bar{v})}{(1 + \frac{\epsilon}{3})} \geq \frac{\mathcal{F}(\bar{v})}{1 + \epsilon}. \end{aligned}$$

□

For any moderately large auction, the performance guarantee of the weighted majority auction mechanism is dramatically better than that of previous auction mechanisms. As a comparison, Bar-Yossef et al. show that their weighted buckets auction is $O(\exp(\sqrt{\log \log h}))$ -competitive [4]. However, in that case, the competitive ratio is achieved for valuation sequences with $\mathcal{F}(\bar{v}) \geq 4h$. The following theorem (Theorem 3) shows that **WM** fails on such small valuation sequences, and indeed, the theorem provides a fairly tight lower bound on the sequences for which **WM** succeeds in achieving a constant competitive ratio.

In Theorem 4, we then prove that *any* algorithm achieving a constant competitive ratio must lose an additive term $\Omega(h)$ in the revenue. (Equivalently, it is not possible to achieve a constant competitive ratio when $\mathcal{F}(\bar{v}) = o(h)$.) Thus there is an $O(\log \log h)$ gap in the additive term between the performance of **WM** (Theorem 2 above) and our general lower bound.

Theorem 3 For any function $f(h) = o(h \log \log h)$, even when restricted to valuation sequences with $\mathcal{F}(\bar{v}) \geq f(h)$, **WM** with any constant α is not

constant-competitive. Furthermore, this holds even if **WM** is allowed to begin with unequal initial weights.

Proof. We first prove the claim under the assumption that the x_k are all distinct and the initial weights are all equal (as in the algorithm described in Figure 1). In this case, note that if the competitive ratio is at most some constant c , then for every value $x \in [1, h]$, there must be some $x_k \in X$ such that $x_k \leq x \leq cx_k$. Otherwise, a sequence of bids of value x would lead to a competitive ratio more than c . Hence, $\ell \geq \log_c h = \Omega(\log h)$.

Now consider a bid sequence consisting entirely of bids of value $x_1 = 1$. If there are n bids, clearly $\mathcal{F} = n$. For $k \neq 1$, for all i , $w_k(i) = 1$, while $w_1(i) = (1 + \alpha)^{i/h}$. Hence, the expected revenue from the i th bidder is no more than $\frac{1}{\ell}(1 + \alpha)^{i/h}$. Summing over the n bidders, we get a total revenue of at most $\frac{n}{\ell}(1 + \alpha)^{n/h}$. If the competitive ratio is at most c , then we need $(1 + \alpha)^{n/h} \geq \frac{\ell}{c}$, which implies $n = \Omega(h \log \ell) = \Omega(h \log \log h)$, from which the result follows.

The above argument implicitly assumes all x_i are distinct (or, equivalently, that **WM** begins with all experts having the same weight). We can generalize the lower bound to hold even when experts begin with different weights as follows. As before, suppose the competitive ratio is at most c . Then, for any value $x \in [1, h]$, let q_x be the fraction of initial weight on experts $x_i \in [\frac{x}{2c}, x]$. Consider a sequence of n bids at the value x for which q_x is smallest. In this case, $\mathcal{F} = nx$. The online algorithm makes at most $\frac{nx}{2c}$ from experts below this window, and at most $nxq_x(1 + \alpha)^{nx/h}$ from experts inside this window. Since $q_x \leq 1/\log_{2c} h$ and since c -competitiveness implies an online revenue of at least $\frac{nx}{c}$, it must be that $(1 + \alpha)^{nx/h} \geq (\log_{2c} h)/2c$ and therefore $nx = \Omega(h \log \log h)$. Thus, the result again follows. \square

A bid sequence consisting entirely of bids of one value may seem somewhat anomalous; in particular, h does not represent the true ratio between the highest and lowest valuations, and most of the weights remain at their initial value. However, the example does not depend on these properties. To see this, one can prepend to the sequence above a set of bids, including a bid at h , such that the revenue obtained from the prefix by using any fixed price $x_i \in X$ falls in the range $[h, 2h]$. Since in the prefix $\mathcal{F} = O(h)$, for any auction, the bids in the prefix can be ordered in such a way that the auction achieves revenue at most $O(h)$ from these bids.

It is not possible to do much better using some other algorithm. We show here that *any* constant-competitive algorithm must lose an additive term $\Omega(h)$, using analysis similar to that used for one-way trading.

Theorem 4 *There is no constant-competitive algorithm for all valuation sequences with $\mathcal{F}(\bar{v}) \geq f(h)$ when $f(h) = o(h)$. Equivalently, suppose A is an on-*

line algorithm such that for all valuation sequences \bar{v} , $R_A(\bar{v}) \geq \mathcal{F}(\bar{v})/c - f(h)$, where c is a constant. Then $f(h) = \Omega(h)$.

Proof. First note that the two statements of the theorem are equivalent. In one direction, if we have an algorithm with competitive ratio c and additive term $-f(h)$, then for $\mathcal{F}(\bar{v}) \geq 2cf(h)$, the algorithm will be $2c$ -competitive. In the other direction, if we have an algorithm with competitive ratio c for $\mathcal{F}(\bar{v}) \geq f(h)$, then it is (trivially) c -competitive with an additive term $-f(h)$ on the smaller sequences. We prove the second statement below.

Let A be an online algorithm with constant competitive ratio c and additive term $-f(h)$. Let $k = 2c$ and $m = 2k^{k-1}$. We will show that $f(h) \geq h/(km)$.

Consider the very first bid, and let $\Pr[a, b]$ denote the probability that A 's sales price is in the range $[a, b]$. Suppose it is the case that $\Pr[1, h/m] \leq 1/k$. Then, if the bid comes in at h/m , the online algorithm's expected gain is at most $h/(km)$ but $\mathcal{F}(\bar{v}) = h/m$. Thus, $f(h) \geq \mathcal{F}(\bar{v})/c - R_A(\bar{v}) \geq h/(km)$. So, we can assume that $\Pr[1, h/m] > 1/k$.

In general, define the series L_t as follows: $L_0 = 0$ and $L_{t+1} = h/m + kL_t$. So, $L_{t+1} = h/m + hk/m + \dots + hk^t/m$. By definition of m , $L_k \leq h$. So, there must be some interval $(L_t, L_{t+1}] \subseteq [1, h]$ such that $\Pr(L_t, L_{t+1}] \leq 1/k$. As above, suppose the first bid comes in at L_{t+1} . In this case, the online algorithm's expected gain is at most $L_t + L_{t+1}/k$, but $\mathcal{F}(\bar{v}) = L_{t+1}$. So, $cf(h) \geq \mathcal{F}(\bar{v}) - cR_A(\bar{v}) \geq L_{t+1} - c(L_t + L_{t+1}/k) = L_{t+1}/2 - cL_t$. Plugging in the definition of L_{t+1} , this is at least $h/(2m)$, and thus $f(h) \geq h/(km)$. \square

3 Posted price mechanisms: the partial information game

As noted in Section 1, the seller using an online posted price mechanism is at a considerable disadvantage compared to a seller using an online auction, since with a posted price mechanism, the seller receives much less information about the buyers' valuations. Nevertheless, as described below, it is still possible to design an online algorithm which achieves (asymptotically) a constant competitive ratio with respect to the optimal fixed price revenue.

To do this, we use a version of the algorithm **Exp3** of Auer et al. [1]. As with an online auction, the choice of a sales price corresponds to the choice of an expert. However, in an online auction, the subsequent bid reveals exactly how well each expert would have done. In a posted price mechanism, at each step, we will know what would have happened with some, but not all, of the possible sales prices. The only sales price whose performance we are guaranteed to know is the one chosen: this corresponds to an online learning algorithm which uses

only information about the gain of the chosen expert at each step.

The algorithm **Exp3** essentially contains algorithm **WM**, described in Section 2, as a subroutine. At each step, we take the probability distribution \mathbf{p} used by **WM** and mix it with the uniform distribution to obtain a modified probability distribution $\bar{\mathbf{p}}$, which is then used to select an expert. Following each buyer's accept/reject decision, we use the information obtained about the gain of the chosen expert to formulate a simulated gain vector, which is then used to update the weights maintained by **WM**.

Figure 2 describes the algorithm **Exp3** in our setting.

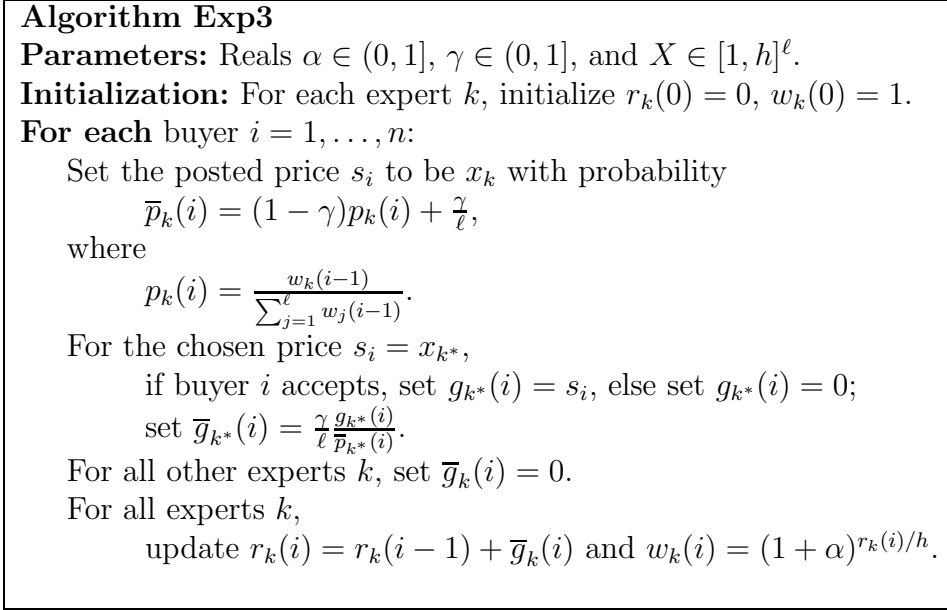


Fig. 2. **Exp3** in our setting

Theorem 4.1 in Auer et al. then becomes the following:

Theorem 5 [1, Theorem 4.1] *For any sequence of valuations \bar{v} , the revenue of auction **Exp3** is at least:*

$$R_{\text{Exp3}}(\bar{v}) \geq (1 - \gamma - \frac{\alpha}{2}) \mathcal{F}_X(\bar{v}) - \frac{h\ell \ln \ell}{\alpha\gamma}.$$

As above, let X consist of all powers of $(1+\beta)$ between 1 and h . For appropriate choices of α , β , and γ , we get the following theorem.

Theorem 6 *For any $\epsilon \in (0, 1]$, restricting to valuation sequences with $\mathcal{F}(\bar{v}) \geq \frac{2304h}{\epsilon^4} \ln h (\ln \ln h + \ln(\frac{4}{\epsilon}))$, mechanism **Exp3** with $\alpha = \frac{\epsilon}{6}$, $\gamma = \frac{\epsilon}{12}$, and X consisting of all powers of $(1 + \frac{\epsilon}{3})$ is $(1 + \epsilon)$ -competitive relative to the optimal fixed price revenue.*

Proof. As in Theorem 2, $\mathcal{F}(\bar{v}) \leq (1 + \beta)\mathcal{F}_X(\bar{v})$, and again, $\ln(1 + \beta) \geq \epsilon/4$. In this case, we have

$$\frac{h\ell \ln \ell}{\alpha\gamma} = \frac{h\left(\frac{\ln h}{\ln(1+\beta)}\right) \ln\left(\frac{\ln h}{\ln(1+\beta)}\right)}{\left(\frac{\epsilon}{8}\right)\left(\frac{\epsilon}{12}\right)} \leq \frac{288h}{\epsilon^3} \ln h(\ln \ln h + \ln(\frac{4}{\epsilon})) \leq \frac{\epsilon}{8}\mathcal{F}(\bar{v}).$$

Thus, by Theorem 5,

$$R_{\mathbf{Exp3}}(\bar{v}) \geq \left(1 - \frac{\epsilon}{12} - \frac{\epsilon}{12}\right) \frac{\mathcal{F}(\bar{v})}{\left(1 + \frac{\epsilon}{3}\right)} - \frac{\epsilon}{8}\mathcal{F}(\bar{v}) \geq \frac{\mathcal{F}(\bar{v})}{(1 + \epsilon)},$$

using the same calculations as in Theorem 2. \square

Again, we can show that this mechanism is not constant-competitive on valuation sequences with small fixed price revenue.

Theorem 7 *For any function $f(h) = o(h \log h \log \log h)$, even when restricted to valuation sequences with $\mathcal{F}(\bar{v}) \geq f(h)$, **Exp3** with any constant α is not constant-competitive.*

Proof. Suppose the competitive ratio is at most some constant c . As before, we must have $\ell = \Omega(\log h)$. Again consider a valuation sequence consisting entirely of valuations at $x_1 = 1$, and let n denote the number of buyers, so that $\mathcal{F} = n$.

For $k \neq 1$, $w_k(i) = 1$ for all i . Hence, because $r_1(i)$ is nondecreasing, $w_1(i)$, $p_1(i)$, and $\bar{p}_1(i)$ are all nondecreasing in i . Furthermore, the expected revenue from buyer i is given by $\bar{p}_1(i)$. Therefore, in order for the competitive ratio to be c , we must have $\bar{p}_1(n) \geq 1/c$.

From the definition of \bar{p} , this implies that $p_1(n) \geq 1/c$. But, $p_1(n)$ is at most $\frac{1}{2}(1 + \alpha)^{r_1(n)/h}$, so we must have $r_1(n) \geq h \log \frac{\ell}{c}$.

Recall that $r_1(n) = \sum_{i=1}^n \bar{g}_1(i)$. Furthermore, note that the expected value of $\bar{g}_1(i)$ is given by $\bar{p}_1(i)[(\gamma/\ell)(1/\bar{p}_1(i))] = \gamma/\ell$. Hence, we need $n \geq (\ell/\gamma)h \log \frac{\ell}{c} = \Omega(h\ell \log \ell) = \Omega(h \log h \log \log h)$, and the theorem follows. \square

The case of unequal initial weights can be handled analogously as in Theorem 3 above.

4 Extensions and Conclusions

Note that given any two auction mechanisms, we can achieve performance which is within a factor of two of the best of the two auctions by simply assigning probability $1/2$ to each. By combining the weighted majority and weighted buckets auctions of [4], we can achieve a constant competitive ratio for valuation sequences with large \mathcal{F} , while maintaining the $O(\exp(\sqrt{\log \log h}))$ competitive ratio for sequences with smaller \mathcal{F} .

Also note that our techniques can be applied to the limited supply case, so long as the sequence of bids can be truncated as soon as we run out of items to sell. While this is not a standard notion in competitive analysis, it does suggest that the weighted majority auction could perform well when the supply is not too small and the bids are generated in some unknown, but non-adversarial, manner. Using the standard notion of competitive ratio, Lavi and Nisan give a lower bound of $\Omega(\log h)$ for the limited supply case [12].

In this paper, we have demonstrated the power of online learning techniques in the context of online auction problems by giving a $(1 + \epsilon)$ -competitive online auction for digital goods. This auction requires valuation sequences with slightly larger, but still quite reasonable, optimal fixed price revenues. We have demonstrated that such a condition is necessary for our weighted majority-based auction. We have also devised a $(1 + \epsilon)$ -competitive online posted price mechanism under a similar assumption. This result is somewhat surprising since the amount of information available to the algorithm is much smaller in a posted-price scenario than in the standard online auction setting. In both cases, the simplicity of the underlying algorithms suggests that these mechanisms would be practical in a wide variety of settings.

References

- [1] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of the 36th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 322–331, 1995.
- [2] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.
- [3] A. Bagchi, A. Chaudhary, R. Garg, M. T. Goodrich, and V. Kumar. Seller-focused algorithms for online auctioning. In *Proceedings of the 7th International Workshop on Algorithms and Data Structures (WADS)*, volume 2125. Springer Verlag LNCS, 2001.

- [4] Z. Bar-Yossef, K. Hildrum, and F. Wu. Incentive-compatible online auctions for digital goods. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 964–970, 2002.
- [5] A. Blum, V. Kumar, A. Rudra, and F. Wu. Online learning in online auctions. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 202–204, 2003.
- [6] N. Cesa-Bianchi, Y. Freund, D. P. Helmbold, D. Haussler, R. E. Schapire, and M. K. Warmuth. How to use expert advice. *Journal of the Association for Computing Machinery*, 44(3):427–485, 1997.
- [7] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive generalized auctions. In *Proceedings of the 34th ACM Symposium on Theory of Computing (STOC)*, pages 72–81, 2002.
- [8] Y. Freund and R. E. Schapire. Game theory, on-line prediction and boosting. In *Proceedings of the 9th Annual Conference on Computational Learning Theory (COLT)*, pages 325–332, 1996.
- [9] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- [10] A. Goldberg, J. Hartline, and A. Wright. Competitive auctions and digital goods. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 735–744, 2001.
- [11] J. Hartline. Dynamic posted price mechanisms. Personal communication, 2002.
- [12] R. Lavi and N. 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.
- [13] N. Littlestone and M. K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108:212–261, 1994.