

Online Learning in Online Auctions*

Avrim Blum[†] Vijay Kumar[‡] Atri Rudra[§] Felix Wu[¶]

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 note, 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 best possible (offline) fixed price revenue. This substantially improves upon the best previously known competitive ratio [3] of $O(\exp(\sqrt{\log \log h}))$ for this problem. We apply our techniques to the related problem of online *posted price* mechanisms, where the auctioneer declares a price and a bidder only communicates his acceptance/rejection of the price. For this problem we obtain results that are (somewhat surprisingly) similar to the online auction problem.

We are primarily concerned with auctions for a single good available in unlimited supply, often described as a digital good, though our techniques may also be useful for the case of limited supply. The problem of designing online auctions for digital goods was first described by Bar-Yossef et al. [3], one of a number of recent papers interested in analyzing revenue-maximizing auctions without making statistical assumptions about the participating bidders [2, 6, 8, 10].

1 Introduction

While auctions are traditional and well-studied economic mechanisms, the popularity of internet auctions has prompted wide interest in various aspects of auctions and related mechanisms, including the question of optimizing the total revenue of an auction. A number of recent papers have addressed the design of revenue-maximizing auctions without making any statistical assumptions about the bidders who participate in the auction [2, 3, 6, 8, 10]. A particularly interesting case is that of *digital goods* [8] — goods of which infinitely many copies can be generated at no cost — considered in the online setting by Bar-Yossef et al. [3].

In the model of Bar-Yossef et al. [3], 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 the range $[1, h]$, so that h is the ratio between the highest and lowest 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 $s_i \leq b_i$.

*Portions of this work appeared as an extended abstract in Proceedings of SODA'03 [4]. This work was supported in part by National Science Foundation grants CCR-0105488 and IIS-0121678.

[†]Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA, Email: avrim@cs.cmu.edu

[‡]Strategic Planning and Optimization Team, Amazon.com, Seattle, WA, Email: vijayk@amazon.com

[§]Department of Computer Science, University of Texas at Austin, Austin, TX. This work was done while the author was at IBM India Research Lab, New Delhi, India. Email: atri@cs.utexas.edu

[¶]Computer Science Division, University of California at Berkeley, Berkeley, CA, Email: felix@cs.berkeley.edu

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. [3], 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.

As in previous papers [3, 8, 10], we will use competitive analysis to analyze the performance of any given auction. 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 [3, 8], 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 $\mathcal{F}(\bar{v}) = \max_{i \in [n]} \{v_i n_i\}$, where $n_i = |\{j \in [n] \mid v_j \geq v_i\}|$. 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 algorithm for this problem.¹ In Section 3, we derive a similar result for the related problem of online *posted-price* auctioning. In a posted-price auction [9], the auctioneer *posts* a price prior to each bidder, and the bidder communicates an acceptance or rejection of the auctioneer’s offer. Compared to the standard online auction, this mechanism provides much less information to the auctioneer about the bidders’ valuations. Surprisingly, we are still able to obtain results very similar to the full-information case. The offline version of this problem is discussed in [9].

Our results are based on application of machine learning techniques to the 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 [1, 5, 7, 11]. The posted price setting, in which the information received depends on the auctioneer’s current price, corresponds to a version of the “bandit” problem [1]. Our algorithms are derived by adapting these techniques to the online auction setting.

2 Online auction: the full information game

We use a variant of Littlestone and Warmuth’s weighted majority (WM) algorithm [11] given in Auer et al. [1]. 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} . Given a parameter $\alpha \in (0, 1]$, define weights

$$w_k(i) = (1 + \alpha)^{r_k(v_1, \dots, v_i)/h}$$

¹Specifically, our algorithm is constant-competitive, but with an additive constant that is $O(-h \ln \ln h)$. As $\mathcal{F}(\bar{v})$ gets large, this additive term becomes negligible. We also present (Theorem 2.4) general lower bounds showing that our additive constants are nearly optimal: in particular, that any constant-competitive algorithm must have an additive constant $\Omega(-h)$.

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 $w_k(i) = (1 + \alpha)^{r_k(v_1, \dots, v_i)/h}$.

Figure 1: **WM** in our setting

From Auer et al., we now have:

THEOREM 2.1. [1, Theorem 3.2] For any sequence of valuations \bar{v} ,

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

where $\mathcal{F}_X(\bar{v}) = \max_k r_k(\bar{v})$ is the optimal fixed price revenue when restricted to fixed prices in X .

For completeness, we provide the proof here.

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

Then, the expected revenue of the auction from bidder $i + 1$ is given by:

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

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

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

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

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

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

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

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

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

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

Rearranging this inequality yields the theorem. ■

Now let X contain all powers of $(1 + \beta)$ between 1 and h . Taking $\alpha = \beta = \frac{\epsilon}{3}$ yields the following:

THEOREM 2.2. *Restricting to valuation sequences with $\mathcal{F}(\bar{v}) \geq \frac{18h}{\epsilon^2} (\ln \ln h + \ln(\frac{4}{\epsilon}))$, auction WM is $(1 + \epsilon)$ -competitive relative to the optimal fixed price revenue.*

The proof follows from the theorem of Auer et al. above by analyzing the choice of parameters, and by noting 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.

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 [3]. However, in that case, the competitive ratio is achieved for valuation sequences with $\mathcal{F}(\bar{v}) \geq 4h$. The following theorem (Theorem 2.3) 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 2.4, we then prove that *any* algorithm achieving a constant competitive ratio must have an additive constant $\Omega(-h)$ (equivalently, it is not possible to achieve a constant competitive ratio for the case $\mathcal{F}(\bar{v}) = o(h)$). Thus there is an $O(\log \log h)$ gap between the performance of WM (Theorem 2.2 above) and our general lower bound.

THEOREM 2.3. *For any function $f(h) = o(h \log \log h)$, even when restricting to valuation sequences with $\mathcal{F}(\bar{v}) \geq f(h)$, WM is $\omega(1)$ -competitive. Furthermore, this holds even if we allow WM to begin with unequal initial weights.*

Proof. Let us first prove the claim under the assumption that the x_i are all distinct and the initial weights are all equal (as in the algorithm of Theorem 2.2). For this, note that if the competitive ratio is at most some constant c , then for every value $x \in [1, h]$, there must be some $x_i \in X$ such that $x_i \leq x \leq cx_i$. 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. ■

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.

Can one do much better by some other algorithm? We show here that *any* constant-competitive algorithm must have an additive term $\Omega(-h)$, using analysis similar to that used for one-way trading.

THEOREM 2.4. *There is no constant-competitive algorithm for all valuation sequences with $\mathcal{F}(\bar{v}) \geq f(h)$ when $f(h) = o(h)$. Stated another way, suppose A is an online algorithm such that for all valuation sequences \bar{v} , $R_A(\bar{v}) \geq \mathcal{F}(\bar{v})/c - f(h)$, where c is constant. Then $f(h) = \Omega(h)$.²*

²In the proof below, we will prove the second statement. This implies the first statement because if there *was* an algorithm that was constant-competitive for $f(h) = o(h)$ with no additive term, then we could just include an additive term of $-f(h)$ to make it trivially work on the smaller sequences too.

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

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 $\Pr[1, h/\alpha] \leq 1/k$. Then, if the bid comes in at h/α , the online algorithm's expected gain is at most $h/(k\alpha)$ but $\mathcal{F}(\bar{v}) = h/\alpha$. Thus, $f(h) \geq \mathcal{F}(\bar{v})/c - R_A(\bar{v}) \geq h/(k\alpha)$. So, we can assume that $\Pr[1, h/\alpha] > 1/k$.

We now argue the general case. Define the series L_t as follows: $L_0 = 0$ and $L_{t+1} = h/\alpha + kL_t$. So, $L_{t+1} = h/\alpha + hk/\alpha + \dots + hk^t/\alpha$. By definition of α , $L_k \leq h$. 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 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 $h/(2\alpha)$, and thus $f(h) \geq h/(k\alpha)$. ■

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 an asymptotically 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 about 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.

Using Theorem 4.1 in Auer et al. [1] and given an appropriate choice of parameters α , γ , and X as above, the following theorem results.

THEOREM 3.1. *There exists a constant $c(\epsilon)$ such that for all valuation sequences with $\mathcal{F}(\bar{v}) \geq ch \log h \log \log h$, mechanism **Exp3** is $(1 + \epsilon)$ -competitive relative to the optimal fixed price revenue.*

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

Algorithm Exp3**Parameters:** Reals $\alpha \in (0, 1]$, $\gamma \in (0, 1]$, and $X \in [1, h]^\ell$.**Initialization:** For each expert k , initialize $r_k(0) = 0$, $w_k(0) = 1$.**For each** buyer $i = 1, \dots, n$:Set the posted price s_i to be x_k with probability $\bar{p}_k(i) = (1 - \gamma)p_k(i) + \frac{\gamma}{\ell}$, where

$$p_k(i) = \frac{w_k(i-1)}{\sum_{j=1}^{\ell} w_j(i-1)}.$$

For the chosen price $s_i = x_{k^*}$, if buyer i accepts, set $g_{k^*}(i) = s_i$, else set $g_{k^*}(i) = 0$. Set

$$\bar{g}_{k^*}(i) = \frac{\gamma g_{k^*}(i)}{\ell \bar{p}_{k^*}(i)}.$$

For all other experts k , set $\bar{g}_k(i) = 0$.For all experts k , update $r_k(i) = r_k(i-1) + \bar{g}_k(i)$ and $w_k(i) = (1 + \alpha)^{r_k(i)/h}$.Figure 2: **Exp3** in our setting

THEOREM 3.2. *For any function $f(h) = o(h \log h)$, when restricted to valuation sequences with $\mathcal{F}(\bar{v}) \geq f(h)$, **Exp3** is $\omega(1)$ -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 $\frac{1}{\ell}(1 + \alpha)^{r_1(n)/h}$, so we must have $r_1(n) \geq h \log \frac{\ell}{c}$.

Now, let I denote the set of buyers i such that $r_1(i-1) \in [h \log \frac{\ell}{2c}, h \log \frac{\ell}{c}]$. For $i \in I$,

$$p_1(i) = \frac{w_1(i-1)}{(\ell-1) + w_1(i-1)} \geq \frac{(\ell/2c)}{2\ell} = \frac{1}{4c}$$

and hence,

$$\bar{g}_1(i) \leq \left(\frac{\gamma}{\ell}\right) \frac{1}{\bar{p}_1(i)} \leq \frac{4c\gamma}{1-\gamma} \left(\frac{1}{\ell}\right)$$

Therefore, $n \geq |I| \geq (1/\bar{g}_1(i))h(\log \frac{\ell}{c} - \log \frac{\ell}{2c}) \geq \Omega(h\ell) = \Omega(h \log h)$, and the theorem follows.

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 [3], 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 [10].

In this note, 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 auction under a similar assumption. This result is somewhat surprising since the amount of information available to the algorithm to earn from is much smaller in a posted-price scenario than in the standard online algorithm setting.

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)*, 1995.
- [2] 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 2001)*, volume 2125. Springer Verlag LNCS, 2001.
- [3] 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.
- [4] 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)*, 2003.
- [5] N. Cesa-Bianchi, Y. Freund, D. Helmbold, D. Haussler, R. Schapire, and M. Warmuth. How to use expert advice. *Journal of the ACM*, 44(3):427–485, 1997.
- [6] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive generalized auctions. In *Proceedings of the 34th ACM Symposium on Theory of Computing (STOC 02)*, 2002.
- [7] Y. Freund and R. Schapire. Game theory, on-line prediction and boosting. In *Proceedings of the 9th Annual Conference on Computational Learning Theory*, pages 325–332, 1996.
- [8] 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.
- [9] J. Hartline. Dynamic posted price mechanisms. Personal communication, 2002.
- [10] 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.
- [11] N. Littlestone and M. K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108:212–261, 1994.