Online Learning in Online Auctions

Avrim Blum * Vijay Kumar † Atri Rudra ‡ Felix Wu §

We consider here 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 which substantially improves upon the performance of previous auctions for this 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 bidders who participate in the auction [5, 6, 4, 2].

1 The Model

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 $p_i \leq b_i$. 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 $v_i - p_i$ 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 ith bid. Hence, the auction mechanism is essentially trying to guess the ith valuation, based on the first $i - 1$ valuations.

As in previous papers [3, 5, 6], we will use competitive analysis to analyze the performance of any given auction. Hence, 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, 5], we take the optimal offline auction to be the one which optimally sets a single fixed price for every bidder. The revenue of such an auction is given by $\mathcal{F}(\overline{v}) = \max_{i \in [n]} \{v_i n_i\}$, where $n_i = |\{j \in [n] \mid v_j \geq v_i\}|$. An online auction A with revenue $R_A(\overline{v})$ is said to be *c-competitive* if for any sequence $\overline{v}$, $R_A(\overline{v}) \geq \mathcal{F}(\overline{v})/c$. We take R_A to be the expected revenue if A is randomized.

2 Online Learning

The key insight connecting the online auction problem to online learning is that setting a single fixed price 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.

We use the variant of Littlestone and Warmuth's weighted majority (WM) algorithm [7] given in Auer et al. [1]. In our context, let $X = \{x_1, \ldots, x_\ell\}$ be a set of candidate fixed prices, corresponding to a set of experts. Let $r_k(\overline{v})$ be the revenue obtained by setting the fixed price x_k for the valuation sequence $\overline{v}$. Given a parameter $\alpha \in (0, 1]$, define weights

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

Clearly, the weights can be easily maintained using a multiplicative update. Then, for bidder i, the auction chooses $s_i \in X$ with probability:

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

From Auer et al., we now have:

*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

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

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

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

For completeness, we provide the proof here.

Proof. Let $g_k(i) = r_k(v_1, \ldots, v_i) - r_k(v_1, \ldots, v_{i-1})$ denote the revenue gained by the kth expert from bidder i. 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:

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

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

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

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

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

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

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

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

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

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:

THEOREM 2.2. *Restricting to valuation sequences with $\mathcal{F}(\overline{v}) \geq \frac{18h}{\epsilon^2}(\ln \ln h + \ln(\frac{4}{\epsilon}))$, 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 $\mathcal{F}(\overline{v}) \leq (1 + \beta)\mathcal{F}_X(\overline{v})$, 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 $\mathcal{F}(\overline{v}) \geq 4h$. The following theorem 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.

THEOREM 2.3. *For any function $f(h) = o(h \log \log h)$, even when restricting to valuation sequences with $\mathcal{F}(\overline{v}) \geq f(h)$, WM is $\omega(1)$-competitive.*

For the proof, first 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 ith 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 $n \geq \Omega(h \log \ell) = \Omega(h \log \log h)$, from which the theorem 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 c-competitiveness implies an online revenue of at least $\frac{nx}{c}$, the result 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 prefix 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.

3 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, 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 [6].

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. It is still open whether this condition is necessary for any constant-competitive auction.

References

[1] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert 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] Amitabha Bagchi, Amitabh Chaudhary, Rahul Garg, Michael T. Goodrich, and Vijay 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] Ziv Bar-Yossef, Kirsten Hildrum, and Felix 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] Amos Fiat, Andrew Goldberg, Jason Hartline, and Anna Karlin. Competitive generalized auctions. In *Proceedings of the 34th ACM Symposium on Theory of Computing (STOC 02)*, 2002.

[5] Andrew Goldberg, Jason Hartline, and Andrew Wright. Competitive auctions and digital goods. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 735–744, 2001.

[6] 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.

[7] Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108:212–261, 1994.