Near-Optimal Online Auctions

Avrim Blum* Jason D. Hartline†

Abstract

We consider the *online auction* problem proposed by Bar-Yossef, Hildrum, and Wu [4] in which an auctioneer is selling identical items to bidders arriving one at a time. We give an auction that achieves a constant factor of the optimal profit less an $O(h)$ additive loss term, where h is the value of the highest bid. Furthermore, this auction does not require foreknowledge of the range of bidders' valuations. On both counts, this answers open questions from [4, 5]. We further improve on the results from [5] for the *online posted-price* problem by reducing their additive loss term from $O(h \log h \log\log h)$ to $O(h \log\log h)$. Finally, we define the notion of an (offline) *attribute auction* for modeling the problem of auctioning items to consumers who are not a-priori indistinguishable. We apply our online auction solution to achieve good bounds for the attribute auction problem with 1-dimensional attributes.

1 Introduction

The *online auction* problem models the situation a seller faces when selling multiple units of an item to bidders who arrive one at a time and each desire one unit. The *unlimited supply* case is an extremal version of the problem where it is assumed that the number of units for sale exceeds the number of consumers (it is effectively infinite), e.g., a digital good or commodity item. This problem is interesting as it combines both the lack of information due to the fact that the bidders have private valuations for the good for sale (a game-theoretic issue), and the lack of information due to not knowing what bidders may arrive in the future (an online issue). The unlimited-supply online auction problem was first considered in [4] where the online auction's performance is compared with the optimal single price sale (a.k.a., the optimal static offline strategy).

To deal with the game-theoretic issues in an auction we adopt the solution concept of *truthful mechanism design*. An auction is said to be *truthful* if any bidder's optimal strategy, no matter what any of the other bidders do, is to bid their true value for the good. In this context, truthful mechanisms are exactly those that compute a price to offer each bidder independently of the bidder's bid (See, e.g., [1, 7]). Naturally, a bidder's bid is rejected if it is below the offered price. The online nature of the problem requires that the auction compute the price to offer a bidder prior to obtaining the values of any subsequent bidders. Combining the requirements of truthful mechanisms with those of online algorithms results in the following algorithmic definition of an online auction.

Definition 1. (Online Auction) *Any class of functions $f_k(\cdot)$ from $\mathbb{R}^{k-1}$ to $\mathbb{R}$ defines an* deterministic online auction *as follows. For each bidder i,*

1. $z_i \leftarrow f_i(b_1, \ldots, b_{i-1})$.
2. *If $z_i \leq b_i$ sell to bidder i at price z_i.*
3. *Otherwise, reject bidder i.*

A randomized online auction *is a distribution over deterministic online auctions.*

Let OPT denote the profit of the optimal single-price sale. For $b_{(k)}$ denoting the kth largest bid, $\mathrm{OPT} = \max_k k b_{(k)}$. Let h denote the value of the highest bid, so $\mathrm{OPT} \geq h$. It is not possible to design an online (or offline) auction that always obtains a constant fraction of h [7, 5] so instead we look to obtain an online auction that obtains profit of at least $\mathrm{OPT}/\beta - \gamma h$ on any input sequence (for constant $\beta \geq 1$ and γ as small as possible). We refer to β as the ratio and γh as the additive loss.

Prior to this work the best known online auction obtained a constant ratio with additive loss γh for $\gamma \in \Theta(\log\log h)$ and required the auction mechanism to know the range of bids in advance [5]. Our paper improves on these results by adapting and building on an expert-advice learning algorithm due to Kalai [11] and Kalai and Vempala [12], to give an auction with constant γ. Specifically, for any constant $\beta > 1$ we can obtain an expected profit of at least $\mathrm{OPT}/\beta - \Theta(h)$

*Carnegie Mellon University, Pittsburgh, PA. Email: avrim@cs.cmu.edu. This work was supported in part by the National Science Foundation under ITR grants CCR-0122581 and IIS-0312814.

†Microsoft Research, Mountain View, CA. Email hartline@microsoft.com. Part of this work was done while the author was at CMU in the ALADDIN project, supported under NSF grant CCR-0122581.

for any bid sequence. This auction also does not need to know the value of h, the highest bid, in advance. Up to constant factors, this online auction is optimal. This answers in the affirmative the outstanding open questions from [4, 5].

We also consider the *online posted-price* problem considered in [5, 13]. This problem is similar to the online auction problem except that the "bidders" are not required to make bids. Instead, the mechanism must offer each bidder a price and bidders may decide whether to accept or reject this price without informing the mechanism of their true valuation for the good. Again, the bidders will arrive one at a time and the mechanism must offer them a price prior to the subsequent bidder's arrival. The posted price mechanism may use the accept/reject responses of prior bidders in determining a price to offer future bidders.

We show how to modify the Exp3 algorithm of Auer et al. [2, 3] (and used by Blum et al. [5] for the posted-price problem) to obtain a performance bound of $\mathrm{OPT}/\beta - O(h \log\log h)$. This improves on the additive loss term in [5] of $O(h \log h \log\log h)$. The key idea is to change the exploration distribution of Exp3 to reflect the greater variance of experts at higher price levels.

In Section 6 we define the (offline) *attribute auction* problem. In an attribute auction, bidders have publicly-available attributes that distinguish them from each other. Examples of such attributes may be the bidders' zip-codes or the cost of providing them with the good or service. Attribute auctions arise as a special case of many mechanism design problems with inherent asymmetries, for example, the *multicast pricing* problem of [7]. The goal of an attribute auction is to obtain a larger profit than possible when the bidders are indistinguishable by using the attributes to perform *price discrimination*. Although we do not consider costs in this paper, this price discrimination is natural when the cost to the auctioneer of serving each bidder is different. Prior work in (offline) auctions [9, 7] explicitly assumes that the bidders are indistinguishable, making it reasonable to compare an auction's profit against the optimal single-price sale, as an auctioneer has no basis on which to charge bidders different prices. For an attribute auction, however, we would like to compare to the more difficult benchmark of the optimal pricing, OPT, obtainable by segmenting the market in some reasonable way and using a different price for each market segment.

In this paper we consider the case of single-dimensional ordered attributes, which means we can think of OPT as a piecewise-constant function, and we allow the algorithm to have an additive term that depends on the number of pieces. What we will aim for (and get) is a revenue of

$$\Omega\left(\max_{m\geq 1}[\mathrm{OPT}_m - mh]\right),$$

where OPT_m denotes the optimal revenue for an auction that is piecewise-constant with m pieces. Equivalently, we can view this as being constant-competitive with OPT, if we "charge" OPT an amount that is $O(h)$ per piece.

The way we will use our online algorithm to address attribute auctions is to view the single-dimensional attribute as a time axis, and to run an extension of our online algorithm that not only competes against the best *fixed* price, but also competes against the best strategy in hindsight that switches among a small number of prices. By achieving a bound that degrades gracefully with the number of switches, we can then get our desired bound for the attribute auction. We leave open the question of guarantees for multi-dimensional attributes.

This paper is organized as follows. In Section 2 we review the application of expert-advice learning techniques to the online auction problem. In Section 3 we give our near optimal online auction, given foreknowledge of the range of bidders bids. We remove the need for this foreknowledge in Section 4. In Section 5 we give our solution to the online posted pricing problem. Finally, in Section 6 we formally define the attribute auction and show how to adapt our solution to the online auction problem to solve the single-dimensional attribute auction problem. Conclusions and open problems are given in Section 7.

2 Combining Expert Advice

The online problem of combining expert advice has been well-studied in Computational Learning Theory [14, 8, 6, 12]. We focus here on the *decision-theoretic* version [8, 12]. In this setting, at each time t, each of k experts advocates a strategy. An algorithm must then choose the strategy of one of the experts to follow. After time t, the payoffs of the strategies of all of the experts are revealed and the algorithm obtains the payoff of the expert's strategy that it selected. It is assumed that all payoffs lie in some range (typically $[0, 1]$) known in advance. The goal of an online learning algorithm is to obtain a total payoff that is nearly as good as the payoff obtained by the best expert in hindsight.

In [5], an auction is described, parameterized by the advance knowledge that the bids are between 1 and h, that for any given $\beta > 1$ obtains profit $\mathrm{OPT}/\beta - O(h \log\log h)$. The main idea of this result is to cast the auction problem as a problem of combining expert advice. Specifically, for each price level of the form α^j

($j \in \{0, 1, \ldots, \log_\alpha h\}$ and $\alpha \approx \sqrt{\beta}$), the idea is to define an "expert" who predicts that α^j is a good single sale price. Given a new bidder i, expert j achieves a payoff of α^j if $b_i \geq \alpha^j$ and a payoff of 0 otherwise. Thus, the payoff of expert j matches the gain one would achieve by using its recommended price level and this fits into an expert-advice setting in which all payoffs lie in the range $[0, h]$. Furthermore, by definition of α, the experts' price levels are close enough together that the best expert's total gain is at most a $\sqrt{\beta}$ factor worse than the gain of the best fixed price in $[1, h]$.

One can now plug this setup into the standard *Randomized Weighted Majority*, or *Hedge*, expert-advice algorithm [14, 8]. Let us define expert j's score, s_j, after seeing the first k bidders as the profit obtained by offering price α^j to said bidders, i.e., $s_j = \alpha^j \times \left|\{i \leq k \ : \ b_i \geq \alpha^j\}\right|$. The Randomized Weighted Majority (Hedge) algorithm, parameterized by constant $\tilde{\beta} > 1$, says to weight each expert j by $\tilde{\beta}^{s_j/h}$ and pick a random expert with probability proportional to its weight. If there are N experts total and all gains are in the range $[0, h]$, then the guarantee is that the expected gain of the algorithm is at least $1/\tilde{\beta}$ times the gain of the best expert, minus an additive $O(h \log N)$ term. Plugging in $\tilde{\beta} = \sqrt{\beta}$ and $N = O(\log h)$ yields the given bound.

3 A Near-Optimal Online Auction

The auction technique we present here is based on an alternative approach to the problem of combining expert advice due to Kalai [11] and Kalai-Vempala [12], based on Hannan [10]. While their method does not improve over previous bounds for the standard expert-advice setting, we show that we can use their technique to remove the $O(\log \log h)$ term when adapted to the online auction problem.

The high-level idea of the approach of [11, 12] is that instead of picking an expert at random at each time interval, we "hallucinate" scores for each expert before time zero according to a specific probability distribution and then ever after use the deterministic go-with-the-best-expert-so-far algorithm.[1] We will first present an online auction for the case that all bids are between 1 and h. Then we will show how to modify it for the case where neither 1 nor h is known in advance. The auction is parameterized by p and α.

Definition 2. *The Hallucinated-Gains Online Auction,* HG*, is based on the scores s_j of $\log_\alpha h + 1$ experts, with expert j advocating the sale of the items at single price $\alpha^j \in [1, h]$. Score s_j will be the actual gain achieved by expert j so far plus the "hallucinated" gain made in the initialization step.*

0. *(Initialization) For each expert j, hallucinate an initial score of $s_j = k\alpha^j$ with probability $(1-p)^k p$. I.e., flip a coin with probability $1-p$ of heads until the first tails is encountered and give expert j an initial score equal to α^j times the number of heads.*
1. *When a new bidder arrives (bidder i), pick the expert, j, with highest score thus far. (We break ties arbitrarily, but consistently. For concreteness, assume we break ties in favor of experts advocating higher prices).*
2. *Offer bidder i the price α^j advocated by the chosen expert.*
3. *Update the scores of all experts that would have produced a sale: for all j such that $\alpha^j \leq b_i$, let $s_j \leftarrow s_j + \alpha^j$.*

Lemma 3.1. *Let $R = \max_j s_j$ be a random variable keeping track of the score of the best expert so far (including hallucination) as the bidders arrive. Then, the expected payoff from bidder i in* HG *is at least $(1-p)$ times the expected increase to R caused by bidder i.*

The proof follows the basic structure of the argument given by Kalai [11], except that (a) we are in a setting of gains rather than losses and (b) the experts' coins in Step 0 are not all worth the same amount (expert j's coins are worth α^j).

Proof. Imagine that at time i (after seeing the ith bid) we conceptually reflip the coins for the hallucinated gains but in the following order. Pick the expert j with the lowest score (breaking ties in favor of those advocating lower prices) and flip j's coin once. If it comes up tails, ignore this expert for the rest of the argument. If heads, add α^j to its score and re-sort the experts by score. Repeat (starting with the new lowest expert) until there is only one expert j' left that still has a coin to flip. Now, even though we are not quite done with the coin flipping, we can at this point notice that if $b_i < \alpha^{j'}$ (so expert j' gets a gain of 0 from bidder i) then expert j' must have been the leading expert prior to bidder i arriving and so the increase to R was 0 as well, and we do not care about the increase to HG. However, suppose $b_i \geq \alpha^{j'}$. Now, consider the next coin flip. If this coin comes up heads (which happens with probability $1-p$) this means that even though the score

[1] This description is assuming an "oblivious adversary" model, in which the goal is to perform well for any sequence of events determined in advance before the algorithm's randomization. This can be removed by re-randomizing at each time step, but we choose not to do that for purpose of clarity.

of j' increased, j' was the leading expert prior to bidder i arriving and our auction chose to use it. So, both R and HG increased by $\alpha^{j'}$. On the other hand, if the coin was tails, then R increased by *at most* $\alpha^{j'}$ (since j' is the new "leader") and all we can say about HG is that it increased by at least 0.

Formally, define A_j to be the event that $j = j'$ (i.e., expert j is the last expert to flip a coin in the above ordering). What we have shown is that for each j, the expected increase to HG given event A_j is at least $(1-p)$ times the expected increase to R given A_j. Thus, the expected gain of HG is at least $(1-p)$ times the expected increase to R overall. □

THEOREM 3.1. *For any constant $\beta > 1$ there is a constant γ such that the expected profit of* HG *with suitably chosen parameters α and p is at least* $\mathrm{OPT}/\beta - \gamma h$ *on inputs with bids in the interval* $[1, h]$.

Proof. Let R be the score of the leading expert in the algorithm, and let H be the Step-0 (hallucinated) portion of that score. By Lemma 3.1, our expected profit is at least $(1-p)\mathbf{E}[R-H]$. For any expert j, let d be the expected number of heads flipped in the hallucination process, i.e., $d = 1/p - 1$. We can bound the expected hallucinated gain of the leading expert by the sum of the expected hallucinated gains for all experts, H'.

$$H' = \sum_{j=0}^{\lfloor \log_\alpha h \rfloor} \alpha^j d \le \sum_{j=-\infty}^{\lfloor \log_\alpha h \rfloor} \alpha^j d = \frac{hd}{(1-1/\alpha)} \quad (3.1)$$

Of course, $\mathbf{E}[R]$ is at least OPT/α because in the worst case, the best expert had zero hallucinated gain in Step 0, and then we lose at most a factor of α due to the discretization of price levels. This gives a lower bound on the expected profit of HG of

$$(1-p)\left(\mathrm{OPT}/\alpha - \frac{(1/p-1)}{(1-1/\alpha)}h\right). \quad (3.2)$$

For $\alpha = 2$ and $p = 1/2$, this gives an expected profit of at least $\mathrm{OPT}/4 - h$. For $\alpha = \sqrt{\beta}$ and $p = 1 - \frac{1}{\sqrt{\beta}}$, this gives an expected profit of the general form desired. □

3.1 Improving the dependence on $\epsilon = \beta - 1$.

The bound (3.2) on the expected profit of HG is somewhat loose, due to bounding the maximum hallucinated gain of any expert by the sum of the hallucinated gains in the proof. In particular, if we consider $\epsilon = \beta - 1$ and look at the bound as a function of ϵ (with $\alpha = \sqrt{\beta} \approx 1 + \epsilon/2$, and $p = 1 - 1/\sqrt{\beta} \approx \epsilon/2$), then we get a bound of $\mathrm{OPT}/\beta - O(h/\epsilon^2)$. We can improve the additive term to $O(\frac{h}{\epsilon}\log\frac{1}{\epsilon})$, however, by simply performing a more careful analysis in the proof of Theorem 3.1. In particular, let S_i be the set of all experts whose price levels lie between $h/2^i$ and $h/2^{i+1}$ (for $i = 0, 1, 2, \ldots$). Each set S_i contains $O(1/\epsilon)$ experts, and thus for a given S_i, the expected maximum number of heads over all experts in S_i is $O(d\log\frac{1}{\epsilon})$. This means the expected maximum hallucinated gain over any expert in S_i is $O(\frac{dh}{2^i}\log\frac{1}{\epsilon})$. Now, summing over all sets S_i gives us $O(dh\log\frac{1}{\epsilon}) = O(\frac{h}{\epsilon}\log\frac{1}{\epsilon})$ as desired.

This additive term is nicer because it matches the dependence on ϵ of the additive term in [5]. In particular, the additive term in that result is $O(\frac{h}{\epsilon}\log(\log_{1+\epsilon} h)) = O(\frac{h}{\epsilon}\log\frac{1}{\epsilon} + \frac{h}{\epsilon}\log\log h)$.

4 Removing the need to know the range $[1, h]$ in advance

The online auction presented in the previous section, HG, as well as those in [4, 5], relies on foreknowledge of the range of bid values. Below we will show how to modify HG so that it is not necessary to know this range in advance. This modification is based on two observations. First, having a lower bound on the bid range is not necessary (from a non-computational point of view). Imagine we have experts at all powers of α less than h. These extra experts only add to the additive loss; however, the additive loss from experts at values less than 1 was already taken into account by the additive loss term in equation (3.1) of the proof of Theorem 3.1. Of course, computationally we cannot keep track of an infinite number of experts but at least conceptually this suggests the lower bound should not be necessary. Second, we can adaptively adjust the upper bound on the range by adding the "missing" experts after a new highest bid arrives. In particular, before the arrival of this new high bid, the auction actually achieves better performance without the missing experts. After the arrival we could have performed worse than the auction that had foreknowledge of the high bid; however, only by at most the value of the largest missing expert. Since each expert can only be missing once, we can charge this possible missed profit to the expert added. This gives a bound on the total possible profit missed in this fashion as the sum of the expert values. Since these values telescope to sum to $h/(1-1/\alpha)$, they just add another constant factor to the additive term. We now instantiate this intuition and make this argument more precise.

DEFINITION 3. *The Hallucinated-Gains Online Auction,* HG^+, *works identically to* HG *except for the following steps:*

0. (Initialization) Initially assume the empty range.

Offer the first bidder an arbitrary positive price.

4. *Let α^k denote the value of the current bid rounded down to the nearest power of α. Add a new expert at value α^k if one does not currently exist.*

5. *Let α^j denote the value of the current lowest expert. Add a new expert at value α^{j-1}. Also add experts at any missing values $\alpha^{j'}$ for $j' \in \{j+1, \ldots, k-1\}$.*

Give initial (hallucinated) gains to the newly-added experts as in HG*, plus credit them for gains they would have made had they been instantiated at time 0.*

THEOREM 4.1. *For any constant $\beta > 1$ there is a constant γ such that the expected profit of* HG^+ *with suitably chosen parameters α and p is at least* $\mathrm{OPT}/\beta - \gamma h$ *on any input.*

Proof. We will show that the expected profit of HG^+ is at least the expected profit of HG on $(0, h]$ minus the sum of all of HG's experts' price levels. Since the sum of those price levels is at most $h/(1 - 1/\alpha)$, our overall additive loss compared to HG is only $O(h)$ larger. Note that HG on $(0, h]$ has an infinite number of experts and has expected profit at least that given by Theorem 3.1.

To analyze HG^+, let us partition the profit made by HG into three parts: (1) profit made by following experts currently in the set used by HG^+, (2) profit made following experts above the current range used by HG^+, and (3) profit made following experts below the current range used by HG^+. The first part is easy to handle: HG^+ has at least as much probability mass on any expert in its collection as does HG, because such an expert can only be more likely to be the "leader" under HG^+ than it is under HG. So, the expected profit of HG^+ from such experts is at least as large. The second part is also easy to handle since we can charge it to the newly added expert in Step 4. In particular, α^k is the maximum profit that HG could possibly obtain from such an expert. Finally, the third part can be charged to the newly added expert in Step 5 because α^{j-1} is an upper bound on the profit obtainable by HG from experts below the current range used by HG^+.

Since we only charge experts when they are added, the total additive loss of HG^+ is at most $h/(1 - 1/\alpha) = O(h)$ more than that of HG. □

We note in passing that a similar argument to that made in Section 3.1 can be used to remove the dependence on α in the additional additive term.

5 Online Posted-Price

We now consider the online posted price selling problem [5, 13]. Here the bidders arrive one at a time and the mechanism must offer each bidder a price. However, in this scenario, the mechanism does not learn each agent's true valuation after the agent arrives. Instead, the auctioneer only learns whether the agent chose to accept or reject its offered price. That is, this corresponds to the situation faced by a shopkeeper who can post a price and see who buys and who does not, but cannot ask an exiting shopper how much they *would* have paid. In terms of the problem of learning from expert advice this corresponds to the *partial information* or *bandit* version of the problem, where the online algorithm learns only the payoff of the chosen expert at any given time, and not the potential payoff of all other experts. We will assume that each agent has a private value v_i for the good and that when offered a price $p_i \leq v_i$ then the agent will accept the offer.

DEFINITION 4. (ONLINE POSTED PRICE MECHANISM) *Any class of functions $g_k(\cdot)$ from $\{0,1\}^{k-1}$ to $\mathbb{R}$ defines an* deterministic online posted price mechanism *as follows. For each agent i,*

1. *For $j < i$, let $x_j = 1$ if agent j accepted offer $z_j = g_j(x_1, \ldots, x_{j-1})$, and 0 otherwise.*

2. *$z_i \leftarrow g_i(x_1, \ldots, x_{i-1})$.*

3. *If $z_i \leq b_i$ sell to bidder i at price z_i.*

4. *Otherwise, reject bidder i.*

A randomized online posted price mechanism *is a distribution over deterministic online posted price mechanisms.*

To solve this problem, Blum et al. [5] apply standard learning results due to Auer et al. [3] for the adversarial multi-armed bandit problem. Auer et al. [3] present an algorithm for the bandit problem called *Exp3* (for exponential-weight exploration and exploitation) that achieves a gain of $\mathrm{OPT}/\beta - O(N \log N)$, where N is the number of experts and the gains of the experts lie in the range $[0, 1]$. Using $N = O(\log h)$, and scaling the range of gains from $[0, 1]$ to $[0, h]$, gives the additive loss term in [5] of $O(h \log h \log\log h)$. We show here how this can be improved, by modifying the exploration distribution used in the Exp3 algorithm to take advantage of the structure of the posted-price problem.

THEOREM 5.1. *For any constant $\beta > 1$, we can achieve an expected profit in the online posted-price problem of at least* $\mathrm{OPT}/\beta - O(h \log\log h)$ *on any input.*

Proof Sketch: The Exp3 algorithm of [3] can be viewed as acting as an interface between the Randomized

Weighted Majority (Hedge) algorithm, which is expecting to receive a vector of gains at each time step, and the real world, which only provides a gain for the expert actually chosen. At each time step, Exp3 queries Hedge and receives a probability vector $(p_1, \ldots, p_N)$ over the N experts. It then mixes this with a uniform "exploration" distribution, producing a distribution $(q_1, \ldots, q_N)$ where $q_j = (1-\gamma)p_j + \gamma/N$, and $\gamma < 1$ is a parameter of the Exp3 algorithm. Exp3 then uses the distribution $\vec{q}$ to choose an expert j, and receives gain g_j. Finally, it provides to Hedge a "simulated" gain vector that is all-zeroes except with value g_j/q_j in the jth coordinate (so, e.g., Hedge believes it has received an expected gain of $p_j(g_j/q_j)$), and the process then repeats in the next time step.

The analysis of Exp3 is based on two properties. First, the actual gain g_j of Exp3 is at least $(1-\gamma)$ times the expected gain $p_j(g_j/q_j)$ of Hedge in its "simulated" world. Second, for each i, the expected value of the ith coordinate of the gain vector passed to Hedge is g_i (since it is g_i/q_i with probability q_i and it is 0 with probability $1-q_i$), so the expected total gain of any given expert i in the simulated world is equal to its actual total gain in the real world. This means the expected value of OPT in the simulated world is only larger than the actual value of OPT. So, we have that the gain of Exp3 is nearly as large as the expected gain of Hedge in the simulated world, which (by the guarantees of Hedge) is nearly as large as the expected value of OPT in the simulated world, which is at least as large as OPT in the real world. However, notice that the range of gain-values in the simulated world is no longer $[0, 1]$ but rather $[0, N/\gamma]$, and therefore the additive term becomes $O(N \log N)$. This is then multiplied by an extra $O(h)$ in the auction setting.

To improve Exp3 for the posted-price problem, we simply modify the exploration distribution to take advantage of the different range of gains for the different experts. Specifically, rather than giving exploration probability γ/N to each expert, we use a geometric distribution, giving the highest expert N a constant fraction $\gamma(1 - 1/\alpha)$ of the probability mass, giving expert $N-1$ a probability mass $\gamma\alpha^{-1}(1 - 1/\alpha)$, and more generally giving expert j a probability mass of $\gamma\alpha^{j-N}(1 - 1/\alpha)$. Since expert j corresponds to a price level of α^j, this ensures that $g_j/q_j = O(\alpha^j\alpha^{N-j}) = O(\alpha^N) = O(h)$. Thus, we incur only a constant-factor increase in the range of gain values, and so our additive term is only $O(h \log N) = O(h \log \log h)$. □

6 Attribute Auctions

The standard offline unlimited-supply auction problem [9] considers the problem of designing a truthful auction that performs well compared with the optimal single price mechanism, OPT (as defined in the preceding sections). An important justification for the comparison of the auction to OPT is the fact that a priori the auctioneer cannot distinguish between two bidders and therefore has no rationale for attempting to charge one bidder more than another. In this section we relax this assumption and consider the design of near-optimal auctions for the case that the bidders are not indistinguishable.

Formally, suppose that each bidder i is labeled with an attribute value $a_i \in A$. The input to the auctioneer is then the vector of attributes, $\mathbf{a} = (a_1, \ldots, a_n)$, and the vector of bidders' bids, $\mathbf{b} = (b_1, \ldots, b_n)$. The characterization of truthful mechanisms (e.g., from [7]) gives the following definition for a truthful attribute auction.

Definition 5. (Attribute Auction) *Any class of n functions $f_k(\cdot)$ from $\mathbb{R}^{n-1} \times A^n$ to $\mathbb{R}$ defines an* deterministic attribute auction *for n bidders as follows. For each bidder i,*

1. $z_i \leftarrow f_i(b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n, a_1, \ldots, a_n)$.
2. *If $z_i \leq b_i$ sell to bidder i at price z_i.*
3. *Otherwise, reject bidder i.*

A randomized attribute auction *is a distribution over deterministic attribute auctions.*

In the case that the attributes and bid values are not correlated, attributes may not aid in obtaining higher profits than the optimal single price sale. However, in the case where there is correlation, we wish to use this correlation to our advantage. In general the problem we face is first that of learning the how the bidders' values are correlated with their attributes and then that of using this learned correlation to compute prices to offer each bidder. While in general the correlations could be arbitrary, we take the intuitive model that the attributes can be used to segment the market into non-overlapping "clusters" over the range of attribute values. Specifically, we look at the case that attributes are 1-dimensional ($A = \mathbb{R}$) and look for an auction that performs well in comparison to an optimal pricing that is a piece-wise constant function over attribute values. Let OPT_m be the profit of the optimal piece-wise constant pricing having at most m pieces.[2] Given

[2] One should think of m as small compared to n. In particular, OPT_n corresponds to selling to each bidder at exactly its bid value if all bidders' attribute values are distinct.

bids in the interval $[1, h]$, we obtain an auction below that obtains an expected profit of:

$$\Omega(\max_m(\text{OPT}_m - hm)).$$

The algorithm is as follows. First, recall the online auction HG: with parameter $p = 1/2$ and $\alpha = 2$, HG obtains expected profit of at least $\text{OPT}/4 - h$. Consider now the following attribute auction:

Definition 6. *The Simulated Online Attribute Auction, SOA, works as follows:*

1. *Sort the bidders by their attribute values.*
2. *Simulate the* HG *auction (with $p = 1/2$ and $\alpha = 2$) on the ordered bidders.*
3. *Reset simulation whenever* OPT *has profit more than* $8h$.

Theorem 6.1. *SOA obtains expected profit at least* $\text{OPT}_m/16 - mh/2$ *for all* m.

Proof. Let R' be the profit of the optimal piece-wise constant pricing that changes prices only when the SOA simulation resets. Since HG has expected profit $\text{OPT}/4 - h$ on each of the segments, it is easy to see that SOA's expected profit is $R \geq R'/8$.

Now consider OPT_m. We want to show that $R' \geq \text{OPT}_m/2 - 4mh$. First of all, we can assume that OPT_m obtains profit at least $8h$ in each of its segments, since otherwise by deleting any such low-profit segments we increase the right-hand size of the desired inequality (OPT_m decreases by at most $8h$ which is paid for by decreasing m by 1). Now, let us consider a phase of the SOA algorithm. Since OPT_m obtains at least $8h$ profit in each of its segments, this phase can intersect no more than two segments of OPT_m (by definition of a phase, any middle segment would have profit less than $8h$). Now, R' uses a single price on this phase while OPT_m can use at most two prices. Thus, on this phase, R' gets at least half the profit of OPT_m. Thus, overall we have $R' \geq \text{OPT}_m/2 - 4mh$ and $R \geq \text{OPT}_m/16 - mh/2$. □

7 Conclusions and Open Problems

In this paper we showed how a natural application of expert learning algorithms can benefit from non-uniform bounds on the expert payoffs. In particular Kalai's expert algorithm and analysis allowed this in the full information case of the online auction problem and the Auer et al. algorithm and analysis allowed this in the partial information case of the online posted-price problem. These are rather general observations about non-uniform bounds on expert payoffs for the two algorithms.

Our results on the attribute auction problem suggests a number of open problems. Specifically,

1. Rather than incurring an additive cost of $O(h)$ per interval of OPT_m, can one develop an online algorithm whose additive cost is only $O(\sum_i h_i)$, where h_i is the value of the largest bid *in interval* i of OPT_m? In other words, can we be constant-competitive if we charge OPT only $O(h_i)$ in its ith interval rather than $O(h)$? While the techniques in Section 4 seem useful (they would solve the problem if we knew in advance a good way of segmenting the range of attributes) the difficulty is in determining where the algorithm's phase boundaries should be.
2. Can one achieve good bounds for d-dimensional attribute auctions for $d \geq 2$, where we allow OPT to break the space into rectangles?
3. Here is a conjectured algorithm for the problem in (2). Begin by randomly partitioning the bidders into two sets S_1 and S_2. Then, looking at all bids in S_1, find the optimal decomposition of S_1 into rectangles where we penalize OPT an amount $O(h)$ per rectangle. Finally, use this set of rectangles as prices for S_2 (and do the reverse procedure to get prices for S_1). Can this approach be shown to achieve good guarantees?

References

[1] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.

[2] Peter Auer, Nicolò 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 Symposium on Foundations of Computer Science*, pages 322–331. IEEE Computer Society Press, Los Alamitos, CA, 1995.

[3] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.

[4] Z. Bar-Yossef, K. Hildrum, and F. Wu. Incentive-compatible online auctions for digital goods. In *Proc. 13th Symp. on Discrete Alg.* ACM/SIAM, 2002.

[5] A. Blum, V. Kumar, A. Rudra, and F. Wu. Online Learning in Online Auctions. In *Proc. 14th Symposium on Discrete Algorithms*. ACM/SIAM, 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 ACM*, 44(3):427–485, 1997.

[7] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive Generalized Auctions. In *Proc. 34th ACM Symposium on the Theory of Computing*. ACM Press, New York, 2002.

[8] 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.

[9] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive Auctions and Digital Goods. In *Proc. 12th Symp. on Discrete Alg.*, pages 735–744. ACM/SIAM, 2001.

[10] J.F. Hannan. Approximation to Bayes risk in repeated play. In M. Dresher, A.W. Tucker, and P. Wolfe, editors, *Contributions to the Theory of Games*, volume III, pages 97–139. Princeton University Press, 1957.

[11] A. Kalai. Smoothed go-with-best-expert. manuscript, October 2001.

[12] A. Kalai and S. Vempala. Efficient algorithms for online decision problems. In *Proceedings of the 16th Annual Conference on Computational Learning Theory*, 2003.

[13] R. Kleinberg and T. Leighton. The Value of Knowing a Demand Curve: Bounds on Regret for Online Posted-Price Auctions. In *Proc. of the 44nd IEEE Symp. on Foundations of Computer Science*, 2003.

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