A Dynamic Learning Algorithm for Online Matching Problems with Concave Returns

Xiao Alison Chen

Department of Industrial and Systems Engineering, University of Minnesota, MN, USA chen2847@umn.edu

Zizhuo Wang

Department of Industrial and Systems Engineering, University of Minnesota, MN, USA zwang@umn.edu

We consider an online matching problem with concave returns. This problem is a generalization of the traditional online matching problem and has vast applications in online advertising. In this work, we propose a dynamic learning algorithm that achieves near-optimal performance for this problem when the inputs arrive in a random order and satisfy certain conditions. The key idea of our algorithm is to learn the input data pattern dynamically: we solve a sequence of carefully chosen partial allocation problems and use their optimal solutions to assist with the future decisions. Our analysis belongs to the primal-dual paradigm; however, the absence of linearity of the objective function and the dynamic feature of the algorithm makes our analysis quite unique. We also show through numerical experiments that our algorithm performs well for test data.

Key words: online algorithms; primal-dual; dynamic price update; random permutation model; Adwords problem

1. Introduction

In traditional optimization models, inputs are usually assumed to be known and efficient algorithms are sought to find the optimal solutions. However, in many practical cases, data does not reveal itself at the beginning. Instead, it comes in an online fashion. For example, in many revenue management problems, customers arrive sequentially and each time a customer arrives, the decision maker has to make some irrevocable decisions (e.g., what product to sell, at what prices) for this customer without knowing any of the future inputs. Such a regime is often called *online optimization*. Online optimization has gained much attention in the research community in the past few decades due to its applicability in many practical problems, and much effort has been directed toward understanding the quality of solutions that can be obtained under such settings. For an overview of the online optimization literature and its recent developments, we refer the readers to [Borodin and El-Yaniv \(1998\)](#), [Buchbinder and Naor \(2009\)](#) and [Devanur \(2011\)](#).

In this paper, we consider a special type of online optimization problem - an online matching problem. Online matching problems are considered as fundamental problems in online optimization

theory and have important applications in the online advertisement allocation problems. For a review of online matching problems, we refer the readers to [Mehta \(2012\)](#). In the problem we study, there is an underlying weighted bipartite graph $G = (I, J, E)$ with weights b_{ij} for each edge $(i, j) \in E$. The vertices in J arrive sequentially in some order, and whenever a vertex $j \in J$ arrives, the set of weights b_{ij} is revealed for all $i \in I$, $(i, j) \in E$. The decision maker then has to match j to one of its neighbors i , and a value of b_{ij} will be obtained from this matching. In our problem, the decision maker's gain from each vertex i is a function of the total matched value to this vertex, and his goal is to maximize the total gain from all vertices. Mathematically, the problem can be formulated as follows (assume $|I| = m$, $|J| = n$, and let $b_{ij} = 0$ for $(i, j) \notin E$):

$$\begin{aligned} & \text{maximize}_{\mathbf{x}} \sum_{i=1}^m M_i \left(\sum_{j=1}^n b_{ij} x_{ij} \right) \\ & \text{s.t.} \quad \sum_{i=1}^m x_{ij} \leq 1, \quad \forall j \\ & \quad \quad x_{ij} \geq 0, \quad \forall i, j, \end{aligned} \tag{1}$$

where x_{ij} denotes the fraction of vertex j that is matched to vertex i .¹ In (1), the coefficient $\mathbf{b}_j = \{b_{ij}\}_{i=1}^m$ is revealed only when vertex j arrives, and an irrevocable decision $\mathbf{x}_j = \{x_{ij}\}_{i=1}^m$ has to be made before observing the next input. For each i , $M_i(\cdot)$ is a nondecreasing concave function with $M_i(0) = 0$. In this paper, we assume that $M_i(\cdot)$ s are continuously differentiable.

As mentioned earlier, online matching problems have a very important application in the online advertisement allocation problem, which we will later refer to as the Adwords problem. In the Adwords problem, there are m advertisers (which we also call the bidders). A sequence of n keywords are searched during a fixed time horizon. Based on the relevance of the keyword, the i th bidder would bid a certain amount b_{ij} to show his advertisement on the result page of the j th keyword. The search engine's decision is to allocate each keyword to one of the m bidders (we only consider a single allocation in this paper). Note that each allocation decision can only depend on the information earlier in the arrival sequence but not on any future data. As pointed out in [Devanur and Jain \(2012\)](#), there are several practical motivations for considering a concave function of the matched bids in the Adwords problem. Among them are convex penalty costs for under-delivery in search engine-advertiser contracts, the concavity of the click-through rate in the number of allocated bids observed in empirical data and fairness considerations. In each of the situations mentioned above, one can write the objective as a concave function. We refer the readers to [Devanur and Jain \(2012\)](#) for a more thorough review of the motivations for this problem. It is worth noting that there is a special case of this problem where $M_i(x) = \min\{x, B_i\}$. In this case, one can view that the bidder has a budget B_i and the revenue from each bidder is bounded by B_i .

One important question when studying online algorithms is the assumptions on the input data. In this work, we adopt a *random permutation model*. More precisely, we assume:

¹ We allow fractional allocations in our model. However, our proposed algorithms output integer solutions. Thus all our results hold if one confines to integer solutions.

1. The total number of arrivals $n = |J|$ is known a priori.
2. The weights $\{b_{ij}\}$ can be adversarially chosen. However, the order that j arrives is uniformly distributed over all the permutations.

The random permutation model has been adopted in much recent literature in the study of online matching problems, see, e.g., [Devanur and Hayes \(2009\)](#), [Feldman et al. \(2010\)](#), [Agrawal et al. \(2014\)](#), etc. It is equivalent to saying that a set of $\mathcal{B} = \{\tilde{\mathbf{b}}_1, \tilde{\mathbf{b}}_2, \dots, \tilde{\mathbf{b}}_n\}$ is arbitrarily chosen beforehand (unknown to the decision maker). Then the arrivals $\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n$ are drawn randomly without replacement from $\mathcal{B}$. The random permutation model is an intermediate path between using a worst-case analysis and assuming each input data is drawn independently and identically distributed (i.i.d.) from a certain distribution. On one hand, compared to the worst-case analysis (see, e.g., [Mehta et al. 2005](#), [Buchbinder et al. 2007](#), [Feldman et al. 2009](#), [Devanur and Jain 2012](#)), the random permutation model is practically reasonable yet much less conservative. On the other hand, the random permutation model is much less restrictive than assuming the inputs are drawn i.i.d. from a certain distribution ([Devanur 2011](#)). Also, the assumption of the knowledge of n is necessary for any online algorithm to achieve near-optimal performance (see [Devanur and Hayes 2009](#)). Therefore, for large problems with relatively stationary inputs, the random permutation model is a good approximation and the study of such models is of practical interest. Next we define the performance measure of an algorithm under the random permutation model:

Definition 1 (c -competitiveness) *Let OPT be the optimal value for the offline problem (1). An online algorithm A is called c -competitive in the random permutation model if the expected value of the online solutions by using A is at least c times the optimal value of (1), that is*

$$\mathbb{E}_\sigma \left[\sum_{i=1}^m M_i \left(\sum_{j=1}^n b_{ij} x_{ij}(\sigma, A) \right) \right] \geq cOPT,$$

where the expectation is taken over uniformly random permutations σ of $1, \dots, n$, and $x_{ij}(\sigma, A)$ is the ij th decision made by algorithm A when the inputs arrive in order σ .

In [Devanur and Jain \(2012\)](#), the authors propose an algorithm for the online matching problem with concave returns that has a constant competitive ratio under the *worst-case model* (the constant depends on the forms of each $M_i(\cdot)$). They also show that a constant competitive ratio is the *best* possible result under that model. In this paper, we propose an algorithm under the random permutation model, which achieves *near-optimal performance* under some conditions on the input.

Our main result is stated as follows:

THEOREM 1. Fix $\epsilon \in (0, 1/2)$. There exists an algorithm (Algorithm DLA) that is $1 - \epsilon$ competitive for the online matching problem with concave returns $M_i(\cdot)$ s under the random permutation model if

$$n \geq \Omega \left(\max \left\{ \frac{\log(m/\epsilon)}{\epsilon \bar{b}^2}, \frac{m^2 \log(m^2 n/\epsilon) F(M, \eta)}{\epsilon^3 \bar{b}} \right\} \right), \quad (2)$$

where $\bar{b} = \frac{1}{n} \min_i \{ \sum_{j=1}^n b_{ij} \}$, $\eta = \frac{\min_{i,j} \{ b_{ij} | b_{ij} > 0 \}}{\max_{i,j} b_{ij}}$, and $F(M, \eta)$ is a constant that only depends on each $M_i(\cdot)$ and η .

In condition (2), $\bar{b}$ can be viewed as the average bid value of a bidder over time. Given that each bidder is at least interested in some fractions of the keywords, this average will go to a certain constant as n becomes large. Also, η can be viewed as the ratio between the value of the smallest non-zero bid and the highest bid. In practice, this is often bounded below by a constant by enforcing a reserve price and a maximum price for any single bid. The exact functional form of $F(M, \eta)$ is somewhat complicated, and is given in Proposition 1. Just to give an example, if we choose $M_i(x) = x^p$ ($0 < p < 1$), then $F(M, \eta) = \frac{2}{\eta^{(2-p)/(1-p)}}$. Therefore, condition (2) can be viewed as simply requiring the total number of inputs is large, which is often the case in practice. For example, in the Adwords problem, n is the number of keyword searches in a certain period, and for instance, Google receives more than 5 billion searches per day. Even if we focus on a specific category, the number can still be in the millions. Thus, this condition is reasonable. We note that most learning algorithms in the literature make similar requirements, see [Devanur and Hayes \(2009\)](#), [Agrawal et al. \(2014\)](#), and [Molinario and Ravi \(2014\)](#). Furthermore, as we will show in our numerical tests, our algorithm performs well even for problems with sizes that are significantly smaller than the condition requires, which validates the potential usefulness of our algorithm.

To propose an algorithm that achieves near-optimal performance, the main idea is to utilize the observed data in the allocation process. In particular, since the input data arrives in a random order, using the past input data and projecting it into the future should present a good approximation for the problem. To mathematically capture this idea, we use a primal-dual approach. We obtain the dual optimal solutions to suitably constructed optimization problems and use them to assist with future allocations. We first propose a one-time learning algorithm (OLA, see Section 2) that only solves an optimization problem once at time ϵn . By carefully examining this algorithm, we prove that it achieves near-optimal performance when the inputs satisfy certain conditions. However, the conditions are stronger than those stated in Theorem 1. To improve our algorithm, we further propose a dynamic learning algorithm (DLA, see Section 3). The dynamic learning algorithm makes better use of the observed data and updates the dual solution at a geometric pace, that is, at time ϵn , $2\epsilon n$, $4\epsilon n$ and so on. We show that these resolvings can lift the performance of the algorithm

and thus prove Theorem 1. As one will see in the proof of the DLA, the choice of the resolving points perfectly balances the trade-off between *exploration* and *exploitation*, which are the main trade-offs in such types of learning algorithms.

It is worth mentioning that a similar kind of dynamic learning algorithm has been proposed in Agrawal et al. (2014) and further studied in Wang (2012) and Molinaro and Ravi (2014). However, those works only focus on linear objectives. In our analysis, the nonlinearity of the objective function presents a non-trivial hurdle since one can no longer simply analyze the revenue generated in each time segment and add them together. In this paper, we successfully work around this hurdle by a convex duality argument. We believe that our analysis is a non-trivial extension of the previous work. Moreover, the problem solved has important applications.

The remainder of the paper is organized as follows. In Section 2, we start with a one-time learning algorithm and prove that it achieves near-optimal performance under some mild conditions on the input. The one-time learning algorithm is easy to understand and shows important insights for designing this class of learning algorithms. However, it only achieves a weaker performance than what is stated in Theorem 1. In Section 3, we propose a dynamic learning algorithm which makes better use of the data and has a stronger performance. Some numerical test results of our algorithm are presented in Section 4, which validate the strength of our algorithm. Section 5 concludes this paper.

2. One-Time Learning Algorithm

We first rewrite the offline problem (1) as follows:

$$\begin{aligned} & \text{maximize } \mathbf{x}, \mathbf{u} \quad \sum_{i=1}^m M_i(u_i) \\ \text{s.t.} \quad & \sum_{j=1}^n b_{ij} x_{ij} = u_i, \quad \forall i \\ & \sum_{i=1}^m x_{ij} \leq 1, \quad \forall j \\ & x_{ij} \geq 0, \quad \forall i, j. \end{aligned} \tag{3}$$

We define the following *dual problem*:

$$\begin{aligned} & \inf \mathbf{v}, \mathbf{y} \quad \sum_{j=1}^n y_j + \sum_{i=1}^m (M_i(v_i) - M'_i(v_i)v_i) \\ \text{s.t.} \quad & y_j \geq b_{ij} M'_i(v_i), \quad \forall i, j \\ & v_i \geq 0, \quad \forall i \\ & y_j \geq 0, \quad \forall j. \end{aligned} \tag{4}$$

Let the optimal value of (3) be P^* and the optimal value of (4) be D^* . In Devanur and Jain (2012), the authors proved the weak duality between (3) and (4). In the following lemma, we prove that in fact the strong duality holds. The proof of the lemma is relegated to the Appendix.

LEMMA 1. $P^* = D^*$. Furthermore, the objective value of any feasible solution to (4) is an upper bound of P^* .

Before we describe our algorithm, we define the following partial optimization problem:

$$\begin{aligned}
 (\mathbf{P}_\epsilon) \quad & \text{maximize } \mathbf{x}, \mathbf{u} \quad \sum_{i=1}^m M_i(u_i) \\
 \text{s.t.} \quad & \sum_{j=1}^{\epsilon n} \frac{b_{ij}}{\epsilon} x_{ij} = u_i, \quad \forall i \\
 & \sum_{i=1}^m x_{ij} \leq 1, \quad \forall j \\
 & x_{ij} \geq 0, \quad \forall i, j.
 \end{aligned} \tag{5}$$

Now we define the one-time learning algorithm as follows:

Algorithm 1 One-Time Learning Algorithm (OLA)

1. During the first ϵn arrivals, no allocation is made.
2. After observing the first ϵn arrivals, solve $(\mathbf{P}_\epsilon)$ and denote the optimal solutions by $\hat{\mathbf{x}}$ and $\hat{\mathbf{u}}$.
3. For any m dimensional vector $\mathbf{w}, \mathbf{q} \geq 0$, define

$$x_i(\mathbf{w}, \mathbf{q}) = \begin{cases} 1 & \text{if } i = \operatorname{argmax}_k \{q_k M'_k(w_k)\} \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

Here, ties among $q_k M'_k(w_k)$ are broken arbitrarily. For the $(\epsilon n + 1)$ th to the n th arrival, the allocation rule $x_{ij} = x_i(\hat{\mathbf{u}}, \mathbf{b}_j)$ is used.

Now we provide some intuition for the algorithm. The idea of the algorithm is to use the first ϵn inputs to learn an approximate $\hat{\mathbf{u}}$ and then use it to make all the future allocations based on the complementarity conditions between the primal and dual problems ((3) and (4)). Here $\hat{\mathbf{u}}$ is solved from $(\mathbf{P}_\epsilon)$ which projects the allocation in the first ϵn inputs to the entire problem. The decision rule in (6) can be explained as choosing the i with the highest product of the nominal bid value b_{ij} and the marginal contribution rate to the total projected reward $M'_i(\hat{\mathbf{u}})$. Note that a similar idea has been used to construct algorithms for an online matching problem with linear objective functions (see e.g., [Devanur and Hayes 2009](#), [Agrawal et al. 2014](#), [Molinaro and Ravi 2014](#)). However, the analyses of those algorithms all depend on the linearity of the objective function which we do not possess in this problem. Instead, an analysis with the use of concavity is required in our analysis, making it quite different from those in the prior literature. In the following, we assume without loss of generality that $\max_{i,j} b_{ij} \leq 1$ (we can always scale the inputs to make this hold). We also make a technical assumption as follows:

ASSUMPTION 1. *The inputs of the problem are in a general position. That is, for any vector $\mathbf{p} = (p_1, \dots, p_m) \neq 0$, there are at most m terms among $\operatorname{argmax}_i \{b_{ij} p_i\}$, $j = 1, \dots, n$, that are not singleton sets.*

The assumption says that we only need to break ties in (6) no more than m times. This assumption is not necessarily true for all inputs. However, as pointed out by [Devanur and Hayes \(2009\)](#) and [Agrawal et al. \(2014\)](#), one can always perturb b_{ij} by adding a random variable η_{ij} taking uniform

distribution on $[0, \eta]$ for some very small η . By doing so the assumption holds with probability one and the effect to the solution can be made arbitrarily small. Given this assumption and by the complementarity conditions, we have the following lemma, whose proof is in the Appendix.

LEMMA 2.

$$\epsilon \hat{u}_i - m \leq \sum_{j=1}^{\epsilon n} b_{ij} x_i(\hat{\mathbf{u}}, \mathbf{b}_j) \leq \epsilon \hat{u}_i + m.$$

We first prove the following proposition about the performance of the OLA, which relies on a condition of the solution to $(\mathbf{P}_\epsilon)$.

PROPOSITION 1. *For any given $\epsilon \in (0, 1/2)$, if $\min_i \hat{u}_i \geq \Omega\left(\frac{m \log(m^2 n / \epsilon)}{\epsilon^3}\right)$, then the OLA is a $1 - \epsilon$ -competitive algorithm.*

Before we prove Proposition 1, we define some notation.

- We define the optimal offline solution to (3) by $(\mathbf{x}^*, \mathbf{u}^*)$ with optimal value OPT.
- Define $\sum_{j=1}^n b_{ij} x_i(\hat{\mathbf{u}}, \mathbf{b}_j) = \bar{u}_i$, note that $\bar{u}_i$ normally does not equal $\hat{u}_i$.

We show the following lemma:

LEMMA 3. *For any given $\epsilon \in (0, 1/2)$, if $\min_i \hat{u}_i \geq \frac{12m \log(m^2 n / \epsilon)}{\epsilon^3}$, then with probability $1 - \epsilon$,*

$$(1 - \epsilon)\hat{u}_i \leq \bar{u}_i \leq (1 + \epsilon)\hat{u}_i, \quad \text{for all } i. \quad (7)$$

Proof. The proof will proceed as follows: For any fixed $\hat{\mathbf{u}}$, we define that a random sample (the first ϵn arrivals) S is *bad* for this $\hat{\mathbf{u}}$ if and only if $\hat{\mathbf{u}}$ is the optimal solution to (5) for this S , but $\bar{u}_i < (1 - \epsilon)\hat{u}_i$, or $\bar{u}_i > (1 + \epsilon)\hat{u}_i$, for some i . First, we show that the probability of a bad sample is small for every fixed $\hat{\mathbf{u}}$ (satisfying $\min_i \hat{u}_i \geq \frac{12m \log(m^2 n / \epsilon)}{\epsilon^3}$) and i . Then, we take a union bound over all distinct i and $\hat{u}_i$ s to prove the lemma.

To start with, we fix $\hat{\mathbf{u}}$ and i . Define $Y_j = b_{ij} x_i(\hat{\mathbf{u}}, \mathbf{b}_j)$. By Lemma 2 and the condition on $\hat{u}_i$, we have

$$(1 - \epsilon^2)\epsilon \hat{u}_i \leq \epsilon \hat{u}_i - m \leq \sum_{j \in S} Y_j \leq \epsilon \hat{u}_i + m \leq (1 + \epsilon^2)\epsilon \hat{u}_i.$$

Therefore, the probability of bad S is bounded by the sum of the following two terms ($N = \{1, 2, \dots, n\}$):

$$P\left(\sum_{j \in S} Y_j \leq \epsilon(1 + \epsilon^2)\hat{u}_i, \sum_{j \in N} Y_j > (1 + \epsilon)\hat{u}_i\right) + P\left(\sum_{j \in S} Y_j \geq \epsilon(1 - \epsilon^2)\hat{u}_i, \sum_{j \in N} Y_j < (1 - \epsilon)\hat{u}_i\right). \quad (8)$$

For the first term, we first define $Z_t = \frac{(1 + \epsilon)\hat{u}_i Y_t}{\sum_{j \in N} Y_j}$ and we have

$$P\left(\sum_{j \in S} Y_j \leq \epsilon(1 + \epsilon^2)\hat{u}_i, \sum_{j \in N} Y_j > (1 + \epsilon)\hat{u}_i\right) \leq P\left(\sum_{j \in S} Z_j \leq \epsilon(1 + \epsilon^2)\hat{u}_i, \sum_{j \in N} Z_j = (1 + \epsilon)\hat{u}_i\right).$$

Then we have

$$\begin{aligned}
P\left(\sum_{j \in S} Z_j \leq \epsilon(1 + \epsilon^2)\hat{u}_i, \sum_{j \in N} Z_j = (1 + \epsilon)\hat{u}_i\right) &\leq P\left(\left|\sum_{j \in S} Z_j - \epsilon \sum_{j \in N} Z_j\right| > \frac{\epsilon^2}{2}\hat{u}_i, \sum_{j \in N} Z_j = (1 + \epsilon)\hat{u}_i\right) \\
&\leq P\left(\left|\sum_{j \in S} Z_j - \epsilon \sum_{j \in N} Z_j\right| > \frac{\epsilon^2}{2}\hat{u}_i \mid \sum_{j \in N} Z_j = (1 + \epsilon)\hat{u}_i\right) \\
&\leq 2 \exp\left(-\frac{\epsilon^3 \hat{u}_i}{4(2 + \epsilon)}\right) \leq \frac{\epsilon}{2m(m^2 n)^m} \doteq \delta.
\end{aligned}$$

Here the second inequality follows from the Hoeffding-Bernstein's inequality for sampling without replacement, see Lemma 8 in the Appendix. Similarly, we can get the same result for the second term in (8), which is also bounded by δ . Therefore, the probability of a bad sample is bounded by 2δ for fixed $\hat{\mathbf{u}}$ and i .

Next, we take a union bound over all distinct $\hat{\mathbf{u}}$ s. We call $\hat{\mathbf{u}}$ and $\hat{\mathbf{u}}'$ *distinct* if and only if they result in different allocations, i.e., $x_i(\hat{\mathbf{u}}, \mathbf{b}_j) \neq x_i(\hat{\mathbf{u}}', \mathbf{b}_j)$ for some i, j . Denote $M'_i(\hat{u}_i) = v_i$. For each j , by the definition in (6), the allocation is uniquely defined by the signs of the following terms:

$$b_{ij}v_i - b_{i'j}v_{i'}, \quad \forall 1 \leq i < i' \leq m.$$

There are $m(m-1)/2$ such terms for each j . Therefore, the entire allocation profiles for all the n arrivals can be determined by the signs of no more than $m^2 n$ differences. Now we find out how many different allocation profiles can arise by choosing different v s. By Orlik and Terao (1992), the total number of different profiles for the $m^2 n$ differences can not exceed $(m^2 n)^m$. Therefore, the number of distinct $\hat{\mathbf{u}}$ s is no more than $(m^2 n)^m$. Now we take a union bound over all distinct $\hat{\mathbf{u}}$ s and $i = 1, \dots, m$, and Lemma 3 follows. $\square$

Next we show that the OLA achieves a near-optimal solution under the condition in Proposition 1. We first construct a feasible solution to (4):

$$\hat{v}_i = \hat{u}_i, \quad \hat{y}_j = \max_i \{b_{ij} M'_i(\hat{u}_i)\}.$$

By Lemma 1, $\sum_{j=1}^n \hat{y}_j + \sum_{i=1}^m (M_i(\hat{u}_i) - M'_i(\hat{u}_i)\hat{u}_i)$ is an upper bound of OPT. Thus, we have

$$\begin{aligned}
\text{OPT} - \sum_{i=1}^m M_i(\bar{u}_i) &\leq \sum_{i=1}^m (M_i(\hat{u}_i) - \hat{u}_i M'_i(\hat{u}_i)) - \sum_{i=1}^m M_i(\bar{u}_i) + \sum_{j=1}^n \hat{y}_j \\
&= \sum_{i=1}^m (M_i(\hat{u}_i) - M_i(\bar{u}_i)) + \sum_{i=1}^m (\bar{u}_i M'_i(\hat{u}_i) - \hat{u}_i M'_i(\hat{u}_i)) - \sum_{i=1}^m \bar{u}_i M'_i(\hat{u}_i) + \sum_{j=1}^n \hat{y}_j \\
&= \sum_{i=1}^m (M_i(\hat{u}_i) - M_i(\bar{u}_i) + (\bar{u}_i - \hat{u}_i) M'_i(\hat{u}_i)),
\end{aligned}$$

where the last equality is because by the allocation rule (6):

$$\sum_{j=1}^n \hat{y}_j = \sum_{i=1}^m \sum_{j=1}^n x_i(\hat{\mathbf{u}}, \mathbf{b}_j) b_{ij} M'_i(\hat{u}_i) = \sum_{i=1}^m \bar{u}_i M'_i(\hat{u}_i).$$

Now, we claim that if condition (7) holds,

$$\sum_{i=1}^m (M_i(\hat{u}_i) - M_i(\bar{u}_i) + (\bar{u}_i - \hat{u}_i) M'_i(\hat{u}_i)) \leq 2\epsilon \sum_{i=1}^m M_i(\bar{u}_i).$$

We consider the following two cases:

- Case 1: $\bar{u}_i \leq \hat{u}_i$. In this case,

$$M_i(\hat{u}_i) - M_i(\bar{u}_i) + (\bar{u}_i - \hat{u}_i) M'_i(\hat{u}_i) \leq M_i(\hat{u}_i) - M_i(\bar{u}_i) \leq \left| \frac{\hat{u}_i - \bar{u}_i}{\bar{u}_i} \right| M_i(\bar{u}_i) \leq 2\epsilon M_i(\bar{u}_i),$$

where the second inequality holds because of the concavity of $M_i(\cdot)$.

- Case 2: $\bar{u}_i > \hat{u}_i$. In this case,

$$M_i(\hat{u}_i) - M_i(\bar{u}_i) + (\bar{u}_i - \hat{u}_i) M'_i(\hat{u}_i) \leq (\bar{u}_i - \hat{u}_i) M'_i(\hat{u}_i) \leq \left| \frac{\bar{u}_i - \hat{u}_i}{\hat{u}_i} \right| M_i(\hat{u}_i) \leq \epsilon M_i(\bar{u}_i).$$

Again, the second inequality is due to the concavity of $M_i(\cdot)$.

Thus, under the condition that $\min_i \hat{u}_i \geq \frac{12m \log(m^2 n / \epsilon)}{\epsilon^3}$, with probability $1 - \epsilon$,

$$\text{OPT} - \sum_{i=1}^m M_i(\bar{u}_i) \leq 2\epsilon \sum_{i=1}^m M_i(\bar{u}_i) \leq 2\epsilon \text{OPT},$$

i.e., $\sum_{i=1}^m M_i(\bar{u}_i) \geq (1 - 2\epsilon) \text{OPT}$.

Lastly, we note that the actual allocation in our algorithm for i is $\tilde{u}_i = \sum_{j=\epsilon n+1}^n b_{ij} x_i(\hat{\mathbf{u}}, \mathbf{b}_j)$ (since we ignore the first ϵn arrivals). By Lemma 2, we have

$$\tilde{u}_i = \bar{u}_i - \sum_{j=1}^{\epsilon n} b_{ij} x_i(\hat{\mathbf{u}}, \mathbf{b}_j) \geq \bar{u}_i - \epsilon(1 + \epsilon^2) \hat{u}_i.$$

Thus when condition (7) holds, $\tilde{u}_i \geq (1 - 3\epsilon) \bar{u}_i$. Therefore,

$$\sum_{i=1}^m M_i(\tilde{\mathbf{u}}) \geq \sum_{i=1}^m M_i((1 - 3\epsilon) \bar{u}_i) \geq (1 - 3\epsilon) \sum_{i=1}^m M_i(\bar{u}_i).$$

The last inequality is due to the concavity of $M_i(\cdot)$ s and that $M_i(0) = 0$. Therefore, given $\min_i \hat{u}_i \geq \frac{12m \log(m^2 n / \epsilon)}{\epsilon^3}$, with probability $1 - \epsilon$,

$$\sum_{i=1}^m M_i(\tilde{u}_i) \geq (1 - 5\epsilon) \text{OPT}.$$

Therefore, Proposition 1 is proved. $\square$

Proposition 1 shows that the OLA is near-optimal under some conditions on $\hat{\mathbf{u}}$. However, $\hat{\mathbf{u}}$ is essentially an output of the algorithm. Although such types of conditions are not uncommon in the study of online algorithms (e.g., in the result of Devanur and Hayes 2009, Feldman et al. 2010), it is quite undesirable. In the following, we address this problem by providing a set of sufficient conditions which only depend on the input parameters (i.e., $m, n, \mathbf{b}$ s and $M(\cdot)$ s). We show that our algorithm achieves near-optimal performance under these conditions. We start with the following lemma.

LEMMA 4. *For any $C > 0$, suppose the following condition holds:*

$$n \geq \max \left\{ \frac{12 \log(m/\epsilon)}{\epsilon \bar{b}^2}, \frac{4mCF(M, \eta)}{\epsilon \bar{b}} \right\} \quad (9)$$

where $\bar{b} = \frac{1}{n} \min_i \{\sum_{j=1}^n b_{ij}\}$, $\eta = \min_{i,j} \{b_{ij} | b_{ij} > 0\}$ and $F(M, \eta)$ is such that

$$M'_i(\eta F(M, \eta)C) < \eta M'_{i'}(C), \forall i, i'.$$

Then with probability $1 - \epsilon$, $\hat{u}_i \geq C$, for all i .

The proof of Lemma 4 is relegated to the Appendix (it is proved together with Lemma 7). Now combining Proposition 1 and Lemma 4, we have the following result for the OLA:

PROPOSITION 2. *Fix any $\epsilon \in (0, 1/2)$. Suppose*

$$n \geq \max \left\{ \frac{12 \log(m/\epsilon)}{\epsilon \bar{b}^2}, \frac{4mCF(M, \eta)}{\epsilon \bar{b}} \right\} \quad (10)$$

where $\bar{b} = \frac{1}{n} \min_i \{\sum_{j=1}^n b_{ij}\}$, $\eta = \min_{i,j} \{b_{ij} | b_{ij} > 0\}$ and $F(M, \eta)$ is such that

$$M'_i(\eta F(M, \eta)C) < \eta M'_{i'}(C), \forall i, i' \quad (11)$$

with $C = \frac{12m \log(m^2 n/\epsilon)}{\epsilon^3}$. Then the OLA is $1 - \epsilon$ -competitive under the random permutation model.

Here we give some comments on the definition of $F(M, \eta)$. The definition of $F(M, \eta)$ basically ensures that we rule out the possibility that one i receives nearly all the allocation while some others receive almost none. Note that such $F(M, \eta)$ always exists and is finite if $\lim_{x \rightarrow \infty} M'_i(x) = 0$ for all i . In practice, this is usually true as there is usually an upper bound on the possible reward from each bidder i . In particular, if $M_i(\cdot) = M(\cdot)$ for all i and $\lim_{x \rightarrow \infty} M'(x) = 0$, then one can choose

$$F(M, \eta) = \frac{M'^{-1}(\eta M'(C))}{\eta},$$

where $M'^{-1}(\cdot)$ denotes the inverse function of $M'(\cdot)$. For example, if one chooses $M_i(x) = x^p$ ($0 < p < 1$), then one can further choose $F(M, \eta) \geq \frac{2}{\eta^{(2-p)/(1-p)}}$. Therefore, in most practical situations,

one can view $F(M, \eta)$ as a constant. Finally, we want to remark that the conditions in Lemma 4 (or Proposition 2) are only one set of sufficient conditions which have the nice feature of only depending on the problem inputs. In practice, one can always resort to the condition in Proposition 1 ($\min_i \hat{u}_i \geq \frac{12m \log(m^2 n / \epsilon)}{\epsilon^3}$) if they are more favorable. In addition, as we will show in our numerical tests in Section 4, our algorithm performs quite well even if some of the conditions in Lemma 4 are not satisfied. Therefore, the applicability of our algorithm could be well beyond what the conditions require.

3. Dynamic Learning Algorithm

In the previous section, we introduced the OLA that can achieve near-optimal performance. While the OLA illustrates the ideas of our approach and requires solving a convex optimization problem only once, the conditions it requires to achieve near-optimality are stricter than what we claim in Theorem 1. In this section, we propose an enhanced algorithm that lessens the conditions and thus improves the OLA.

The main idea for the enhancement is the following: In the one-time learning algorithm, we only solve a partial optimization problem once. However, it is possible that there is some error for that solution due to the random order of arrival. If we could modify the solution as we gather more data, we might be able to improve the performance of the algorithm. In the following, we introduce a dynamic learning algorithm based on this idea, which updates the *allocation policy* every time the history doubles; that is, it computes a new $\hat{\mathbf{u}}$ at time $t = \epsilon n, 2\epsilon n, 4\epsilon n, \dots$ and uses it to perform the matching for the next time period. We define the following problem:

$$\begin{aligned}
 (\mathbf{P}_\ell) \quad & \text{maximize } \mathbf{x}, \mathbf{u} \quad \sum_{i=1}^m M_i(u_i) \\
 \text{s.t.} \quad & \sum_{j=1}^{\ell} \frac{n}{\ell} b_{ij} x_{ij} = u_i, \quad \forall i \\
 & \sum_{i=1}^m x_{ij} \leq 1, \quad \forall j \\
 & x_{ij} \geq 0, \quad \forall i, j.
 \end{aligned}$$

We further define $(\mathbf{x}^\ell, \mathbf{u}^\ell)$ to be the optimal solution to $(\mathbf{P}_\ell)$.

We define the dynamic learning algorithm as follows:

Algorithm 2 Dynamic Learning Algorithm (DLA)

1. During the first ϵn arrivals, no allocation is made.
 2. For $r = 0, 1, \dots$, for $2^r n \epsilon < j \leq 2^{r+1} n \epsilon$, set $x_{ij} = x_i(\mathbf{u}^\ell, \mathbf{b}_j)$ for all i , where $\ell = \lceil 2^r n \epsilon \rceil$.
-

In the following, without loss of generality, we assume that $L = -\log_2 \epsilon$ is an integer (otherwise one can just choose a smaller ϵ and prove the same result). Define $\ell_k = 2^{k-1} \epsilon n$, $k = 1, \dots, L$, and define $\hat{\mathbf{u}}^k = \mathbf{u}^{\ell_k}$. We first prove the following proposition:

PROPOSITION 3. *If for all k , $\min_i \hat{u}_i^k \geq \Omega \left\{ \frac{m \log(m^2 n / \epsilon)}{\epsilon^2} \right\}$, then the DLA is $1 - \epsilon$ -competitive under the random permutation model.*

Before we proceed to the proof, we first define some more notation. We define:

$$\bar{u}_i^k = \sum_{j=\ell_k+1}^{\ell_{k+1}} b_{ij} x_i(\hat{\mathbf{u}}^k, \mathbf{b}_j), \quad \tilde{u}_i^k = \sum_{j=1}^n b_{ij} x_i(\hat{\mathbf{u}}^k, \mathbf{b}_j), \quad \bar{u}_i = \sum_{k=1}^L \bar{u}_i^k.$$

Note that in these definitions, $\bar{u}_i^k$ is the allocated values for i in the period $\ell_k + 1$ to ℓ_{k+1} using $\hat{\mathbf{u}}^k$, which is the actual allocation in that period. $\tilde{u}_i^k$ is the allocation for i in all periods if $\hat{\mathbf{u}}^k$ is used. $\bar{u}_i$ is the actual allocation for i during the entire algorithm. We first prove the following lemma bounding the differences between $\bar{u}_i^k$, $\tilde{u}_i^k$ and $\hat{u}_i^k$.

LEMMA 5. *If $\min_i \hat{u}_i^k \geq \frac{16m \log(m^2 n / \epsilon)}{\epsilon^2}$, then with probability $1 - \epsilon$, for all i ,*

$$\left(1 - \epsilon \sqrt{\frac{n}{\ell_k}}\right) \hat{u}_i^k \leq \frac{n}{\ell_k} \bar{u}_i^k \leq \left(1 + \epsilon \sqrt{\frac{n}{\ell_k}}\right) \hat{u}_i^k \quad (12)$$

and

$$\left(1 - \epsilon \sqrt{\frac{n}{\ell_k}}\right) \hat{u}_i^k \leq \tilde{u}_i^k \leq \left(1 + \epsilon \sqrt{\frac{n}{\ell_k}}\right) \hat{u}_i^k. \quad (13)$$

Lemma 5 shows that with high probability, $\frac{n}{\ell_k} \bar{u}_i^k$, $\tilde{u}_i^k$ and $\hat{u}_i^k$ are close to each other. In particular, when k is small, the factor $(1 \pm \epsilon \sqrt{n/\ell_k})$ is relatively loose while as k increases, the factor becomes tight. The proof of Lemma 5 is similar to that of Lemma 3 and is relegated to the Appendix.

The next lemma gives a bound on the revenue obtained by the DLA.

LEMMA 6. *If $\hat{u}_i^k \geq \frac{16m \log(m^2 n / \epsilon)}{\epsilon^2}$ for all i and k , then with probability $1 - \epsilon$,*

$$\sum_{i=1}^m M_i \left(\frac{n}{\ell_i} \bar{u}_i^k \right) \geq \left(1 - 6\epsilon \sqrt{\frac{n}{\ell_k}}\right) OPT.$$

The proof of Lemma 6 can be found in the Appendix.

Finally, we prove Proposition 3. We bound the objective value of the actual allocation. Note that the actual allocation for each i can be written as

$$\sum_{k=1}^L \bar{u}_i^k = \sum_{k=1}^L \alpha_k \frac{n}{\ell_k} \bar{u}_i^k,$$

where $\alpha_k = \frac{\ell_k}{n}$. By the property of concave functions, we have

$$\sum_{i=1}^m M_i \left(\sum_{k=1}^L \bar{u}_i^k \right) = \sum_{i=1}^m M_i \left(\sum_{k=1}^L \alpha_k \frac{n}{\ell_k} \bar{u}_i^k + \left(1 - \sum_{k=1}^L \alpha_k\right) \cdot 0 \right) \geq \sum_{i=1}^m \sum_{k=1}^L \alpha_k M_i \left(\frac{n}{\ell_k} \bar{u}_i^k \right).$$

By Lemma 6, with probability $1 - \epsilon$

$$\begin{aligned} \sum_{i=1}^m \sum_{k=1}^L \alpha_k M_i \left(\frac{n}{\ell_k} \bar{u}_i^k \right) &\geq \sum_{k=1}^L \frac{\ell_k}{n} \left(1 - 6\epsilon \sqrt{\frac{n}{\ell_k}} \right) \text{OPT} = (1 - \epsilon) \text{OPT} - 6\epsilon \sum_{k=1}^L \sqrt{\frac{\ell_k}{n}} \text{OPT} \\ &\geq (1 - 16\epsilon) \text{OPT}, \end{aligned}$$

where the last inequality is because

$$\sum_{k=1}^L \sqrt{\frac{\ell_k}{n}} = \sqrt{\frac{1}{2}} + \sqrt{\frac{1}{4}} \dots \leq 1 + \sqrt{2} \leq 2.5.$$

Therefore, Proposition 3 is proved. $\square$

Similar to Lemma 4, we have the following conditions on the input parameters such that with high probability, the conditions in Proposition 3 hold.

LEMMA 7. *For any $C > 0$, suppose the following condition holds:*

$$n \geq \max \left\{ \frac{24 \log(m/\epsilon)}{\epsilon \bar{b}^2}, \frac{4mCF(M, \eta)}{\epsilon \bar{b}} \right\} \quad (14)$$

where $\bar{b} = \frac{1}{n} \min_i \{\sum_{j=1}^n b_{ij}\}$, $\eta = \min_{i,j} \{b_{ij} | b_{ij} > 0\}$ and $F(M, \eta)$ is such that

$$M'_i(\eta F(M, \eta) C) < \eta M'_{i'}(C), \forall i, i'.$$

Then with probability $1 - \epsilon$, $\hat{u}_i^k \geq C$, for all i .

The proof of Lemma 7 is given in the Appendix. Finally, we combine Proposition 3 and Lemma 7, and Theorem 1 follows.

The same remark after Lemma 4 applies here. In particular, the conditions in Lemma 7 is only one set of sufficient conditions for our algorithm to achieve the target performance. However, one may also use the conditions in Proposition 3 if they turn out to hold in practice. In the next section, we show that the DLA works well even if the conditions in Lemma 7 are not satisfied.

4. Numerical Experiments

In this section, we report some numerical test results for our algorithms (both the OLA and the DLA). The objective is to validate the strength of our approaches and investigate the relationship between the performance of our algorithms and the input parameters.

In our numerical tests, we consider the Adwords problem. We assume there are m advertisers (bidders), n keywords arriving sequentially, and b_{ij} is the amount bidder i would like to pay to display his advertisement on keyword search j . We introduce a base problem in which we set $m = 50$, $n = 10,000$ and $M_i(x) = x^p$ with $p = 0.9$. The bidding values b_{ij} are generated in the following way:

1. Assume there are $k = 100$ categories of keywords. For each category k , there is a base valuation of bidder i , denoted by $\bar{b}_{ik}$, which is generated according to the following distribution:

$$\bar{b}_{ik} = \begin{cases} 0 & \text{with probability } 0.7 \\ U[0.2, 1] & \text{with probability } 0.3, \end{cases}$$

where $U[a, b]$ denotes a uniformly distributed random variable on $[a, b]$.

2. For each arriving keyword, we first randomly choose a category. The probability for each category i , denoted by ρ_i , is randomly chosen on the simplex $\{\rho_i | \sum_{i=1}^k \rho_i = 1, \rho_i \geq 0\}$. Then if category k_0 is chosen, the final bid value for bidder i will be $\bar{b}_{ik_0} \cdot U[0.9, 1.1]$.

Although the way b_{ij} is chosen seems arbitrary, we believe it reflects some major features of the bid values in practice. In the Adwords problem, each bidder is interested in certain categories of keywords. For example, a sport product company is interested in keywords related to sports. The $\bar{b}_{ik}$ s represent such interest levels. Then the bidder i 's actual bid on such a keyword is the base value $\bar{b}_{ik}$ multiplied by a random number, which reflects some level of idiosyncrasies of each keyword arrival. We also tested other ways to generate b_{ij} , and the test results are similar. We will report those test results in the end of this section.

To evaluate the performance of our algorithms, we introduce the notion of Relative Loss (RL) defined as follows:

$$RL = 1 - \frac{\text{Actual Revenue}}{\text{Offline Optimal Revenue}}.$$

In the numerical experiment, there is one key parameter we need to set in both of our algorithms: ϵ . In Theorem 1 and Proposition 1, we gave sufficient conditions on the inputs such that the algorithms will have expected RL less than ϵ . However, the theoretical results are asymptotic and thus may not represent the best practical choice of ϵ . In Table 1 and Figure 1, we first test both our OLA and DLA with different choices of ϵ . We have the following observations from Table 1 and Figure 1 (each number in Table 1 is the average of 100 independent runs, the standard deviations of the results are insignificant compared to the average value):

- For the DLA, choosing a smaller ϵ improves its performance. There are two reasons for that. First, choosing a smaller ϵ reduces the loss due to ignoring the first ϵn bids. Second, it increases the number of price updates which help the decision maker to refine the decision policy and achieve better performance. Therefore one should choose a smaller ϵ in the DLA.

- For the OLA, the optimal choice of ϵ is more subtle. There are two countervailing forces when one chooses a smaller ϵ . On one hand, by choosing a smaller ϵ , the loss due to the failure to allocate any bid during period 1 to ϵn becomes smaller, which benefits the algorithm. On the other hand, if ϵ is too small, the learned price may not be accurate enough which may lead to poor allocation in the remaining periods. In the test example, the optimal choice is $\epsilon = 0.02$.

ϵ	0.001	0.005	0.01	0.02	0.05	0.10
OLA	22.02%	5.17%	4.07%	3.71%	6.10%	10.26%
DLA	0.47%	0.84%	1.27%	2.15%	4.89%	9.41%

Table 1 Performance of the OLA and the DLA for Different Choices of ϵ

Figure 1 Performance of the OLA and the DLA for Different Choices of ϵ

- The DLA outperforms the OLA for all choices of ϵ .

Next we focus ourselves on the DLA. As shown in Table 1, we prefer to choose a smaller ϵ in the DLA. In the following experiments, we will choose $\epsilon = 0.001$. Next we compare the performance of the DLA to a myopic allocation method which simply allocates each incoming keyword to the bidder with the highest b_{ij} value. We also study the impact of the two parameters, n and p , on the performance of our algorithm. We generate 100 instances of the input b_{ij} and compare the average performance. The results of the average RL are shown in Table 2 (the standard deviations are shown in the parentheses) as well as in Figures 2(a) and 2(b).

n	1,000	2,000	5,000	10,000	20,000
DLA	1.29% (0.29%)	0.87% (0.26%)	0.58% (0.12%)	0.57% (0.17%)	0.57% (0.16%)
Myopic	3.07% (0.47%)	3.03% (0.61%)	2.90% (0.49%)	3.11% (0.54%)	2.96% (0.47%)
p	0.5	0.6	0.7	0.8	0.9
DLA	2.30% (0.82%)	1.87% (0.67%)	1.55% (0.56%)	0.99% (0.34%)	0.57% (0.17%)
Myopic	14.38% (2.01%)	12.46% (1.81%)	9.46% (1.38%)	6.51% (1.05%)	3.11% (0.54%)

Table 2 Performance of the DLA and the Myopic Policy

From Table 2, we can see that the DLA consistently performs better than the myopic approach. In particular, the performance of the DLA gradually improves when n increases, while the performance of the myopic approach seems to be insensitive to the size n of the problem. Moreover, even for small values of n , the performance of the DLA is still very good. This means that the DLA works

Figure 2 Performance of the DLA (the Bottom Ones) and the Myopic Policy (the Top Ones)

well even for problems whose size is much smaller than what Theorem 1 requires. For the parameter p , we can see that both the DLA and the myopic algorithm deteriorate when p decreases, but the DLA deteriorates much slower. Finally, we comment that these results are computed when we ignore the first ϵn bids. In practice, one does not need to do that and the performance of the DLA would be even better.

Finally, we repeat the above test for different setups of the inputs. We fix the parameters $m = 50$, $n = 10,000$, and $M_i(x) = x^p$ with $p = 0.9$ in the base problem and generate b_{ij} s in the following ways:

1. b_{ij} follows a normal distribution (truncated at 0 and 1). The parameters of each normal distribution (mean μ and standard deviation σ) are randomly generated from a uniform distribution on $[0, 1]$.
2. b_{ij} follows a Beta distribution. The parameters (α, β) of the Beta distribution are generated from a uniform distribution on $[0, 1]$.
3. b_{ij} follows a mixed normal and Beta distribution. That is, with probability 0.5, b_{ij} follows a truncated normal distribution with mean 0.5 and standard deviation 0.5, and with probability 0.5, b_{ij} follows a Beta distribution with $\alpha = \beta = 1/2$.

Next, we compare the performance of the DLA (choose $\epsilon = 0.001$) and the myopic algorithm. For each case, we generate 100 instances of the input b_{ij} and compare the average RL. The results are shown in Table 3.

From Table 3, we can see that the DLA outperforms the myopic approach under all the above three setups. The RL of the DLA decreases as the problem size grows, while the RL of the myopic policy is not sensitive to n . Also, as p changes, the performance of the DLA is rather stable, while the performance of the myopic algorithm varies a lot. The overall trend of the DLA and the myopic

Case 1	n	1,000	2,000	5,000	10,000	20,000
	DLA	1.56% (0.33%)	0.84% (0.18%)	0.36% (0.07%)	0.21% (0.05%)	0.14% (0.03%)
	Myopic	1.45% (0.45%)	1.37% (0.47%)	1.42% (0.45%)	1.41% (0.48%)	1.31% (0.47%)
	p	0.5	0.6	0.7	0.8	0.9
	DLA	0.20% (0.04%)	0.22% (0.06%)	0.22% (0.05%)	0.21% (0.03%)	0.21% (0.05%)
	Myopic	9.10% (2.01%)	7.98% (2.00%)	5.71% (1.32%)	3.38% (0.96%)	1.44% (0.46%)
Case 2	n	1,000	2,000	5,000	10,000	20,000
	DLA	1.19% (0.24%)	0.61% (0.12%)	0.25% (0.08%)	0.13% (0.07%)	0.12% (0.08%)
	Myopic	13.91% (2.07%)	13.66% (2.48%)	13.90% (2.07%)	13.78% (1.97%)	13.62% (2.12%)
	p	0.5	0.6	0.7	0.8	0.9
	DLA	0.20% (0.16%)	0.29% (0.20%)	0.21% (0.12%)	0.17% (0.11%)	0.13% (0.07%)
	Myopic	40.57% (6.50%)	38.30% (5.98%)	32.45% (4.71%)	24.67% (3.26%)	13.46% (2.07%)
Case 3	n	1,000	2,000	5,000	10,000	20,000
	DLA	1.50% (0.31%)	0.89% (0.34%)	0.35% (0.07%)	0.21% (0.06%)	0.14% (0.01%)
	Myopic	10.84% (2.50%)	11.22% (2.64%)	11.30% (2.44%)	10.91% (2.54%)	11.07% (2.85%)
	p	0.5	0.6	0.7	0.8	0.9
	DLA	0.21% (0.12%)	0.22% (0.12%)	0.15% (0.45%)	0.15% (0.35%)	0.21% (0.06%)
	Myopic	37.00% (7.43%)	34.02% (6.82%)	30.22% (5.89%)	20.72% (4.70%)	10.91% (2.54%)

Table 3 Performance of the DLA and the Myopic Policy

algorithm resembles that in the experiment in the beginning of this section. Finally, we observe that the DLA seems robust toward various problem setups, while the myopic approach does not.

5. Conclusion

In this paper, we propose a dynamic learning algorithm for an online matching problem with concave returns. We show that our algorithm achieves near-optimal performance when the data arrives in a random order and satisfies some conditions. The analysis is primal-dual based, however, the nonlinear objective function requires us to work around nontrivial hurdles that do not exist in previous work. Numerical experiment results show that our algorithm works well in test problems.

6. Acknowledgement

We thank the anonymous referees for their insightful comments. The research of both authors is supported by National Science Foundation under research grant CMMI-1434541.

References

- Agrawal, S., Z. Wang, Y. Ye. 2014. A dynamic near-optimal algorithm for online linear programming. *Operations Research* **62**(4) 876–890.
- Borodin, A., R. El-Yaniv. 1998. *Online computation and competitive analysis*. Cambridge University Press.
- Buchbinder, N., K. Jain, J. Naor. 2007. Online primal-dual algorithms for maximizing ad-auctions

- revenue. *ESA '07: Proceedings of the 15th Annual European Symposium on Algorithms*. 253–264.
- Buchbinder, N., J. Naor. 2009. Online primal-dual algorithms for covering and packing. *Mathematics of Operations Research* **34**(2) 270–286.
- Devanur, N. 2011. Online algorithms with stochastic input. *SIGecom Exch.* **10**(2) 40–49.
- Devanur, N., T. Hayes. 2009. The adwords problem: Online keyword matching with budgeted bidders under random permutations. *EC'09: Proceedings of the 10th ACM Conference on Electronic Commerce*. 71–78.
- Devanur, N., K. Jain. 2012. Online matching with concave returns. *STOC'12: Proceedings of the 44th Annual Symposium on Theory of Computing*. 137–144.
- Feldman, J., M. Henzinger, N. Korula, V. Mirrokni, C. Stein. 2010. Online stochastic packing applied to display ad allocation. *ESA '10: Proceedings of the 18th Annual European Symposium on Algorithms: Part I*. 182–194.
- Feldman, J., N. Korula, V. Mirrokni, S. Muthukrishnan, M. Pal. 2009. Online ad assignment with free disposal. *WINE'09: Proceedings of the 5th International Workshop on Internet and Network Economics*. 374–385.
- Mehta, A. 2012. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Sciences* **8**(4) 265–368.
- Mehta, A., A. Saberi, U. Vazirani, V. Vazirani. 2005. Adwords and generalized on-line matching. *FOCS'05: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*. 264–273.
- Molinaro, M., R. Ravi. 2014. Geometry of online packing linear programs. *Mathematics of Operations Research* **39**(1) 46–59.
- Orlik, P., H. Terao. 1992. *Arrangement of hyperplanes*. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], Springer-Verlag, Berlin.
- van der Vaart, A., J. Wellner. 1996. *Weak convergence and empirical processes: with applications to statistics (Springer Series in Statistics)*. Springer.
- Wang, Z. 2012. Dynamic learning mechanism in revenue management problems. Ph.D. thesis, Stanford University, Palo Alto.

Appendix

In our proofs, we will frequently use the following Hoeffding-Bernstein's Inequality for sampling without replacement:

LEMMA 8 (Theorem 2.14.19 in van der Vaart and Wellner 1996:). Let $u_1, u_2, \dots, u_r$ be random samples without replacement from real numbers $\{c_1, c_2, \dots, c_R\}$. Then for every $t > 0$,

$$P\left(\left|\sum_{i=1}^r u_i - r\bar{c}\right| \geq t\right) \leq 2 \exp\left(-\frac{t^2}{2r\sigma_R^2 + t\Delta_R}\right)$$

where $\Delta_R = \max_i c_i - \min_i c_i$, $\bar{c} = \frac{1}{R} \sum_{i=1}^R c_i$, and $\sigma_R^2 = \frac{1}{R} \sum_{i=1}^R (c_i - \bar{c})^2$.

Proof of Lemma 1. We first write down the Lagrangian dual of (3). By associating p_i to the first set of constraints and y_j to the second set of constraints, the Lagrangian dual of (3) is:

$$\begin{aligned} \inf_{\mathbf{p}, \mathbf{y}} \quad & \sum_{j=1}^n y_j + \sup_{u_i \geq 0} \sum_{i=1}^m (M_i(u_i) - p_i u_i) \\ \text{s.t.} \quad & y_j \geq b_{ij} p_i, \quad \forall i \\ & y_j \geq 0, \quad \forall j. \end{aligned} \tag{15}$$

Since the primal problem is convex and only has linear constraints, Slater's condition holds, thus the strong duality theorem holds and (3) and (15) have the same optimal value. Next we show that (4) and (15) are equivalent. To show this, assume the range of $M'_i(\cdot)$ on $[0, \infty)$ is $(a_i, b_i]$ or $[a_i, b_i]$ (by the assumption that $M(\cdot)$ s are continuously differentiable, it must be either one of these two forms). Now we argue that the optimal p_i must be within $[a_i, b_i]$ in (15). First we must have $p_i \geq a_i$, otherwise the term $\sup_{u_i \geq 0} \{M_i(u_i) - p_i u_i\}$ goes to infinity as u_i increases and it cannot be the optimal solution to (15). On the other hand, if $p_i > b_i$, the optimal u_i must be 0, and one can always set $p_i = b_i$ and achieves a smaller value of the objective function. Therefore, $p_i \in [a_i, b_i]$ at optimality.

Now if $p_i \in (a_i, b_i]$ at optimality, one can always find one v_i such that $M'_i(v_i) = p_i$, and that v_i must be the optimal solution to $\sup_{u_i} \{M_i(u_i) - p_i u_i\}$ (the optimal solution must be attainable in this case). Therefore, each feasible solution of (15) will correspond to a feasible solution of (4) and vice versa. The only case left now is when $p_i = a_i$ at optimality. In this case, $\sup_{u_i} \{M_i(u_i) - a_i u_i\} = \lim_{x \rightarrow \infty} \{M_i(x) - a_i x\}$. Also, we know that $\lim_{x \rightarrow \infty} M'_i(x) = a_i$, therefore, there exists a sequence of feasible solution of (4) such that the limit of the objective value equals the objective obtained when $p_i = a_i$ in (15). Therefore, the lemma is proved. $\square$

Proof of Lemma 2. To prove this lemma, it suffices to show that for each fixed i , $x_i(\hat{\mathbf{u}}, \mathbf{b}_j)$ and $\hat{x}_{ij}$ (recall that $\hat{\mathbf{x}} = \{\hat{x}_{ij}\}$ is the optimal solution to (5)) differ by no more than m terms. If this is true, then note that $\sum_{j=1}^n b_{ij} \hat{x}_{ij} = \epsilon \hat{u}_i$ and $0 \leq b_{ij} \leq 1$, the lemma holds.

To show that $x_i(\hat{u}, \mathbf{b}_j)$ and $\hat{x}_{ij}$ differ by no more than m terms, we first construct the dual problem of (5) (according to (4)):

$$\begin{aligned} \inf_{\mathbf{v}, \mathbf{y}} \quad & \sum_{j=1}^{\epsilon n} y_j + \sum_{i=1}^m (M_i(v_i) - M'_i(v_i)v_i) \\ \text{s.t.} \quad & y_j \geq \frac{b_{ij}}{\epsilon} M'_i(v_i), \quad \forall i, j \\ & v_i \geq 0, \quad \forall i \\ & y_j \geq 0, \quad \forall j. \end{aligned}$$

By Lemma 2.1, strong duality holds and thus any optimal solution should satisfy the complementarity conditions. Among them we should have $\hat{x}_{ij}(y_j - \frac{b_{ij}}{\epsilon} M'_i(v_i)) = 0$. Therefore, if there is no tie when we defined $x_i(\hat{u}, \mathbf{b}_j)$, we must have $x_i(\hat{u}, \mathbf{b}_j) = \hat{x}_{ij}$. By Assumption 1, there are no more than m ties. Thus, Lemma 2 is proved. $\square$

Proof of Lemma 5. We first prove (12). The idea is similar to the proof of the one-time learning case. For any fixed $\hat{\mathbf{u}}^k$, we define that a random sample S (a sequence of arrival) is *bad* if and only if $\hat{\mathbf{u}}^k$ is the optimal solution to $(\mathbf{P}_{\ell_k})$ but $\hat{\mathbf{u}}^k$ does not satisfy (12) for some i . First, we show that the probability of a bad sample is small for any fixed $\hat{\mathbf{u}}^k$ and fixed i . Then we take a union bound over all distinct $\hat{\mathbf{u}}^k$ s and i s to show the result.

Fix $\hat{\mathbf{u}}^k$ and i . We define $Y_j = b_{ij}x_i(\hat{\mathbf{u}}^k, \mathbf{b}_j)$. By Lemma 2 and the assumption on $\hat{u}_i^k$, we have

$$\frac{\ell_k}{n} \hat{u}_i^k - \epsilon^2 \hat{u}_i^k \leq \sum_{j=1}^{\ell_k} Y_j \leq \frac{\ell_k}{n} \hat{u}_i^k + \epsilon^2 \hat{u}_i^k.$$

Therefore, the probability of a bad sample is bounded by the following two terms:

$$\begin{aligned} & P \left(\sum_{j=1}^{\ell_k} Y_j \leq \frac{\ell_k}{n} \hat{u}_i^k + \epsilon^2 \hat{u}_i^k, \sum_{j=\ell_k+1}^{\ell_{k+1}} Y_j > \frac{\ell_k}{n} (1 + \sqrt{\frac{n}{\ell_k}} \epsilon) \hat{u}_i^k \right) \\ & + P \left(\sum_{j=1}^{\ell_k} Y_j \geq \frac{\ell_k}{n} \hat{u}_i^k - \epsilon^2 \hat{u}_i^k, \sum_{j=\ell_k+1}^{\ell_{k+1}} Y_j < \frac{\ell_k}{n} (1 - \sqrt{\frac{n}{\ell_k}} \epsilon) \hat{u}_i^k \right). \end{aligned} \quad (16)$$

For the first term, we have

$$\begin{aligned} & P \left(\sum_{j=1}^{\ell_k} Y_j \leq \frac{\ell_k}{n} \hat{u}_i^k + \epsilon^2 \hat{u}_i^k, \sum_{j=\ell_k+1}^{\ell_{k+1}} Y_j > \frac{\ell_k}{n} \left(1 + \sqrt{\frac{n}{\ell_k}} \epsilon \right) \hat{u}_i^k \right) \\ & = P \left(\sum_{j=1}^{\ell_k} Y_j \leq \frac{\ell_k}{n} \hat{u}_i^k + \epsilon^2 \hat{u}_i^k, \sum_{j=1}^{\ell_{k+1}} Y_j > \frac{\ell_k}{n} \left(2 + \sqrt{\frac{n}{\ell_k}} \epsilon \right) \hat{u}_i^k \right) \\ & \leq P \left(\left| \sum_{j=1}^{\ell_k} Y_j - \frac{1}{2} \sum_{j=1}^{\ell_{k+1}} Y_j \right| > \frac{\epsilon}{4} \sqrt{\frac{n}{\ell_k}} \frac{\ell_k}{n} \hat{u}_i^k \mid \sum_{j=1}^{\ell_{k+1}} Y_j > \frac{\ell_k}{n} \left(2 + \sqrt{\frac{n}{\ell_k}} \epsilon \right) \hat{u}_i^k \right) \\ & \leq 2 \exp \left(-\frac{\epsilon^2 \hat{u}_i^k}{16} \right) \leq \frac{\epsilon}{2m(m^2 n)^m} \doteq \delta. \end{aligned}$$

Here the second inequality follows from Lemma 8, and the third inequality is due to the condition of $\hat{u}_i^k$. Similarly, we can get the bound for the second term in (16). Therefore, the probability of a bad sample is bounded by 2δ .

Now we take union bound over all distinct $\hat{\mathbf{u}}^k$ and i . Similar to the proof of Lemma 3, we call $\mathbf{us}$ to be *distinct* if they result in different allocations. As argued earlier, there are no more than $(m^2n)^m$ distinct $\mathbf{us}$. Therefore, we know that with probability $1 - \epsilon$, (12) holds.

Next we prove (13). The idea is similar. Fix $\hat{\mathbf{u}}^k$ and i . We define $Y_j = b_{ij}x_i(\hat{\mathbf{u}}^k, \mathbf{b}_j)$. Applying Lemma 8, we get

$$\begin{aligned} & P\left(\sum_{j=1}^{\ell_k} Y_j \leq \frac{\ell_k}{n} \hat{u}_i^k + \epsilon^2 \hat{u}_i^k, \sum_{j=1}^n Y_j > \left(1 + \epsilon \sqrt{\frac{n}{\ell_k}}\right) \hat{u}_i^k\right) \\ & \leq P\left(\left|\sum_{j=1}^{\ell_k} Y_j - \frac{\ell_k}{n} \sum_{j=1}^n Y_j\right| > \frac{\epsilon}{4} \sqrt{\frac{\ell_k}{n}} \hat{u}_i^k \mid \sum_{j=1}^n Y_j > \left(1 + \epsilon \sqrt{\frac{n}{\ell_k}}\right) \hat{u}_i^k\right) \\ & \leq \exp\left(-\frac{\epsilon^2 \hat{u}_i^k}{16}\right) \doteq \delta. \end{aligned}$$

Using the same argument as above, Lemma 5 holds. $\square$

Proof of Lemma 6. The proof consists of two main steps. First we show that with probability $1 - \epsilon$, the following is true for all k :

$$\sum_{i=1}^m M_i(\tilde{u}_i^k) \geq \left(1 - 2\epsilon \sqrt{\frac{n}{\ell_k}}\right) \text{OPT}. \quad (17)$$

To show this, we follow a similar step when we prove the optimality of the one-time learning algorithm. Define

$$\hat{v}_i^k = \hat{u}_i^k \text{ and } \hat{y}_i^k = \max_i \{b_{ij} M'_i(\hat{u}_i^k)\}.$$

Since (v_i^k, y_i^k) is a feasible solution to (4), we know that

$$\sum_{j=1}^n \hat{y}_j^k + \sum_{i=1}^m (M_i(\hat{u}_i^k) - M'_i(\hat{u}_i^k) \hat{u}_i^k)$$

is an upper bound of OPT. Therefore, by using the same argument as in (12), we know that

$$\text{OPT} - \sum_{i=1}^m M_i(\tilde{u}_i^k) \leq \sum_{i=1}^m (M_i(\hat{u}_i^k) - M_i(\tilde{u}_i^k) + (\tilde{u}_i^k - \hat{u}_i^k) M'_i(\hat{u}_i^k)).$$

Now for each term above, we consider two cases. If $\hat{u}_i^k \geq \tilde{u}_i^k$, then

$$M_i(\hat{u}_i^k) - M_i(\tilde{u}_i^k) + (\tilde{u}_i^k - \hat{u}_i^k) M'_i(\hat{u}_i^k) \leq M_i(\hat{u}_i^k) - M_i(\tilde{u}_i^k) \leq \frac{M_i(\tilde{u}_i^k)}{\tilde{u}_i^k} (\hat{u}_i^k - \tilde{u}_i^k)$$

and with probability $1 - \epsilon$, this is less than $2\epsilon \sqrt{\frac{n}{\ell_k}} M_i(\tilde{u}_i^k)$; if $\hat{u}_i^k < \tilde{u}_i^k$, then

$$M_i(\hat{u}_i^k) - M_i(\tilde{u}_i^k) + (\tilde{u}_i^k - \hat{u}_i^k) M'_i(\hat{u}_i^k) \leq (\tilde{u}_i^k - \hat{u}_i^k) M'_i(\hat{u}_i^k) \leq \frac{M_i(\hat{u}_i^k)}{\hat{u}_i^k} (\tilde{u}_i^k - \hat{u}_i^k).$$

Again, with probability $1 - \epsilon$, this is less than $2\epsilon\sqrt{\frac{n}{\ell_k}}M_i(\tilde{u}_i^k)$. Therefore, with probability $1 - \epsilon$,

$$\sum_{i=1}^m (M_i(\hat{u}_i^k) - M_i(\tilde{u}_i^k) + (\tilde{u}_i^k - \hat{u}_i^k)M'_i(\hat{u}_i^k)) \leq 2\epsilon\sqrt{\frac{n}{\ell_k}}\text{OPT}.$$

Therefore, (17) is proved. Next we show that

$$\sum_{i=1}^m M_i(\tilde{u}_i^k) - \sum_{i=1}^m M_i\left(\frac{n}{\ell_k}\bar{u}_i^k\right) \leq 4\epsilon\sqrt{\frac{n}{\ell_k}}\text{OPT}.$$

To see this, by Lemma 5, we know that with probability $1 - \epsilon$,

$$\tilde{u}_i^k - \frac{n}{\ell_k}\bar{u}_i^k \leq 2\epsilon\sqrt{\frac{n}{\ell_k}}\hat{u}_i^k.$$

Therefore, for each i , we have

$$\frac{\tilde{u}_i^k - \frac{n}{\ell_k}\bar{u}_i^k}{\tilde{u}_i^k} \leq \frac{2\epsilon\sqrt{\frac{n}{\ell_k}}\hat{u}_i^k}{(1 - \epsilon\sqrt{\frac{n}{\ell_k}})\hat{u}_i^k} \leq 4\epsilon\sqrt{\frac{n}{\ell_k}}.$$

Now we analyze $M_i(\tilde{u}_i^k) - M_i(\frac{n}{\ell_k}\bar{u}_i^k)$ for each i . We only need to focus on the case when $\tilde{u}_i^k > \frac{n}{\ell_k}\bar{u}_i^k$ (otherwise the difference is less than 0). In this case, by the concavity of $M_i(\cdot)$, we have

$$M_i(\tilde{u}_i^k) - M_i\left(\frac{n}{\ell_k}\bar{u}_i^k\right) \leq \frac{M_i(\tilde{u}_i^k)}{\tilde{u}_i^k}(\tilde{u}_i^k - \frac{n}{\ell_k}\bar{u}_i^k) \leq 4\epsilon\sqrt{\frac{n}{\ell_k}}M_i(\tilde{u}_i^k).$$

Therefore, we have

$$\sum_{i=1}^m M_i(\tilde{u}_i^k) - \sum_{i=1}^m M_i\left(\frac{n}{\ell_k}\bar{u}_i^k\right) \leq 4\epsilon\sqrt{\frac{n}{\ell_k}} \sum_{i=1}^m M_i(\tilde{u}_i^k) \leq 4\epsilon\sqrt{\frac{n}{\ell_k}}\text{OPT}.$$

Together with (17), Lemma 6 holds. $\square$

Proof of Lemma 4 and Lemma 7. First, we note that Lemma 7 implies Lemma 4 (except for the constant part, which can be strengthened easily by only considering one ℓ_k in the following proof). Therefore, it suffices to prove Lemma 7.

We first prove for each k , with probability $1 - \epsilon/\log(1/\epsilon)$, $\min_i \hat{u}_i^k > C$. Then we take a union bound to prove Lemma 7. To show that for each k , with probability $1 - \epsilon/\log(1/\epsilon)$, $\min_i \hat{u}_i^k > C$, first we show that with probability $1 - \epsilon/\log(1/\epsilon)$, $\sum_{j=1}^{\ell_k} b_{ij} \geq \ell_k \bar{b}/2$ for all i . To see this, we use Lemma 8, we have for any i ,

$$P\left(\left|\sum_{j=1}^{\ell_k} b_{ij} - \frac{\ell_k}{n} \sum_{j=1}^n n_{ij}\right| \geq \ell_k \bar{b}/2\right) \leq 2\exp(-\ell_k \bar{b}^2/12) < \frac{\epsilon}{m \log(1/\epsilon)},$$

where the last inequality is due to condition (14). Next we show that given $\sum_{j=1}^{\ell_k} b_{ij} \geq \ell_k \bar{b}/2$, there cannot exist an i such that $\hat{u}_i^k < C$ in the optimal solution to the partial program $(\mathbf{P}_{\ell_k})$. We prove by contradiction. Let $K = \eta \frac{\epsilon n \bar{b} - 2\epsilon C}{2mC}$. If there exists i such that $\hat{u}_i^k < C$ in the optimal solution, then we argue that there must exist $1 \leq j \leq \ell_k$ such that

1. $j \in S_k = \{j : x_{ij} < 1, b_{ij} > \eta\}$, and
2. There exists i' such that $x_{i'j} > 0$ and $\hat{u}_{i'}^k \geq KC$.

Here these two conditions mean that there must exist a keyword j such that we allocated it (at least partially) to bidder i' whose total allocation had already exceeded KC when we could have allocated it to bidder i whose final allocation is less than C .

To see this, we note that we have proved with probability $1 - \epsilon/\log(1/\epsilon)$, $\sum_{j=1}^{\ell_k} b_{ij} \geq \ell_k \bar{b}/2$. However, by the definition of i , $u_i^k < C$, thus we also have $\sum_{j=1}^{\ell_k} b_{ij} x_{ij} \leq \ell_k C/n$. Therefore, combined with the assumption that $\max_{i,j} b_{ij} \leq 1$, there must exist at least $\ell_k \bar{b}/2 - \ell_k C/n$ j s between 1 and ℓ_k such that $x_{ij} < 1$ but $b_{ij} \geq \eta$, i.e., $|S_k| \geq \ell_k \bar{b}/2 - \ell_k C/n$.

Next we show that among $j \in S_k$, there exists at least one j such that $x_{i'j} > 0$ while $\hat{u}_{i'}^k \geq KC$ for some i' . We define $T_k = \{i : \hat{u}_i^k < KC\}$. We have

$$\sum_{i \in T_k, j \in S_k} x_{ij} \leq \frac{1}{\eta} \sum_{i \in T_k, j \in S_k} b_{ij} x_{ij} < \frac{mKC}{\eta}. \quad (18)$$

Here the second inequality is because $|T| < m$. However, we also have

$$\sum_{i,j \in S_k} x_{ij} \geq \frac{\ell_k \bar{b}}{2} - \frac{\ell_k C}{n}. \quad (19)$$

This is because $M_i(\cdot)$ s are increasing, thus each $\sum_i x_{ij}$ must equal 1 at optimality. Then, by taking the difference between (18) and (19), we have that

$$\sum_{i \notin T_k, j \in S_k} x_{ij} > \frac{\ell_k \bar{b}}{2} - \frac{\ell_k C}{n} - \frac{mKC}{\eta} \geq 0.$$

The last inequality is by the definition of K and that $\ell_k \geq \epsilon n$ for all k . Therefore, there exists $j \in S_k$ such that the bid is allocated to some i' with $\hat{u}_{i'}^k \geq KC$. We denote such j by j^* .

Finally, we consider another allocation that increases the allocation of j^* to i while decreasing the allocation to i' . The local change of the objective function at this point is:

$$M'_i(\hat{u}_i^k) b_{ij} - M'_{i'}(\hat{u}_{i'}^k) b_{i'j} \geq M'_i(C) \eta - M'_{i'}(KC) > 0$$

where the first inequality is due to the concavity of $M_i(\cdot)$ s and the last inequality is due to condition (14). However, this contradicts the assumption that the solution is optimal. Thus, Lemma 7 is proved. $\square$