

The Generalized Magician Problem under Unknown Distributions and Related Applications*

Aravind Srinivasan

University of Maryland, College Park
College Park, MD, USA
asriniv1@umd.edu

Pan Xu

New Jersey Institute of Technology
Newark, NJ, USA
pxu@njit.edu

ABSTRACT

The *Magician Problem* (MP) and its generalization, the *Generalized Magician Problem* (GMP), were introduced by Alaei *et al.* (APPROX-RANDOM 2013) and Alaei (SICOMP 2014) and have been used as powerful ingredients in online-algorithm design for many hard problems such as the k -choice prophet inequality, mechanism design in Bayesian combinatorial auctions, and the generalized assignment problem. The adversarial model here is essentially that of an oblivious adversary.

In this paper, we introduce generalizations of GMP (MP) under two different arrival settings (by making the adversary stronger): unknown independent identical distributions (UIID) and unknown adversarial distributions (UAD). Different adversary models capture a range of arrival patterns. For GMP under UIID, we show that a natural greedy algorithm Greedy is optimal. For the case of MP under UIID, we show that Greedy has an optimal performance of $1 - \frac{B^B}{B!e^B} \geq 1 - \frac{1}{\sqrt{2\pi B}}$, where B is the budget, and show an application to online B -matching with stochastic rewards. For GMP under UAD, we present a simple algorithm, which is near-optimal among all non-adaptive algorithms. We consider the simple case of MP under UAD with $B = 1$, and give an exact characterization of the respective optimal adaptive and optimal non-adaptive algorithms for any finite time horizon. We offer an example of MP under UAD on which there is a provable gap between the classical MP under adversarial order and MP under UAD even with a time horizon $T = 4$.

KEYWORDS

magician problem; online algorithms; prophet inequality

ACM Reference Format:

Aravind Srinivasan and Pan Xu. 2022. The Generalized Magician Problem under Unknown Distributions and Related Applications. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022)*, Online, May 9–13, 2022, IFAAMAS, 9 pages.

1 INTRODUCTION

Several “prophet” and “magician” problems have been studied intensively over the last several years, motivated by online problems

*Aravind Srinivasan was supported in part by NSF awards CCF-1422569 and CCF-1749864, as well as by research awards from Adobe, Amazon, and Google. Pan Xu was partially supported by NSF CRII Award IIS-1948157. We thank the reviewers for their helpful comments.

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

in E-commerce. The generalized magician problem (GMP) was introduced by Alaei *et al.* [4] to model several online problems. The formal description is as follows.

Generalized Magician Problem: Suppose we have a budget of B and that at each time¹ $t = 1, 2, \dots$, an item Y_t will arrive where Y_t is an independent random variable taking values from $[0, 1]$. Here are **Three Rules** to follow: (1) once the item Y_t arrives, its *distribution* (but not its value) is revealed to us and we need to make an instant and irrevocable decision – either to accept it or reject; (2) we can accept Y_t only when we have at least one unit budget remaining (referred to as “being safe”); (3) once the item Y_t is accepted, a realization $y_t \in [0, 1]$ of Y_t sampled from Y_t ’s distribution is revealed to us and our budget will be reduced by y_t accordingly. The adversary can choose an arbitrary sequence of items subject to the constraint that $\sum_t \mathbb{E}[Y_t] \leq B$, before the arrivals start; thus, the model is that of an *oblivious adversary* [7]. Note that the distribution of each item Y_t is unknown until its arrival and that the Y_t ’s are independent. We know B upfront and our task is to design an online algorithm to maximize the value γ such that each item will be accepted with probability at least γ . When all items are restricted to be Bernoulli random variables, GMP is reduced to the classical Magician Problem (MP), which was first introduced by Alaei [1]. The main results regarding MP and GMP are summarized as follows.

THEOREM 1.1 (ALAEI [2]). *For a given budget B , there exists an algorithm which accepts each arriving item with probability at least $1 - \frac{1}{\sqrt{B}}$ and $1 - \frac{1}{\sqrt{B+3}}$ for GMP and MP respectively.*

Here are several applications of GMP and MP presented in [1, 4]. Alaei [1] considered a generalization of prophet inequalities, where both the gambler and the prophet are allowed to pick B numbers and each to receive a reward equal to their sum, called the *B-choice prophet inequality*, which was first introduced by Hajiaghayi *et al.* [22]. These works designed a randomized strategy for the gambler invoking MP as blackbox, which achieves at least a fraction of $1 - \frac{1}{\sqrt{B+3}}$ of the reward obtained by the prophet. Furthermore, Alaei [1] presented a general framework for approximately reducing the mechanism design problem for multiple agents to single-agent subproblems in the context of Bayesian combinatorial auctions. Alaei *et al.* [4] introduced the online generalized assignment problem (GAP), which can be viewed as a generalization of Prophet-Inequality Matching (PIM) [3] with each item having a random size. Both GAP and PIM capture applications in ad allocation arising from cellular networks [5]. Alaei *et al.* [4] designed a near-optimal algorithm for the online GAP invoking GMP as a blackbox, which results in an online competitive ratio of $1 - \frac{1}{\sqrt{B}}$.

¹We use “time” and “round” interchangeably throughout this paper.

The arrival assumption considered by GMP (MP) in [1, 4] is called *adversarial order*, i.e., the whole arrival sequence is unknown and fixed. This is driven by theoretical interest whereas in practical cases, two other common arrival settings, called *known* and *unknown* distributions, have received far more attention. The basic setting is: We now have a set \mathcal{X} of items, and in each round, an item is sampled from \mathcal{X} according to a fixed but known or unknown distribution. These two arrival settings are commonly used to capture several stochastic online arrival patterns naturally emerging in the real world: keywords in the online advertising business [29], workers in online crowdsourcing markets [33, 34] and queries in online mobile advertising systems [5], just to name a few. As for known distributions, there is a long line of research in the context of online bipartite matching, see, e.g. [6, 9, 19, 21, 24, 28]. For the setting of unknown distributions, a notable special case, called *unknown i.i.d.*, refers to the scenario when the sampling distribution is assumed unknown but the same and independent over each round. There are several works considering this setting in the domain of online bipartite matching and Adwords², see, e.g. [15, 16, 25, 27]. Studies of [15, 16] introduced the unknown adversarial arrival model for the Adwords problem (they called it *adversarial stochastic input* there). It can be viewed as a generalization of unknown *i.i.d.*: in each round a keyword is sampled from a fixed but unknown distribution, which can change over time. Inspired by these works above, we introduce a similar “unknown distributions” generalization of GMP as follows: One of our applications is an improved analysis of the algorithm for online B -matching with stochastic rewards (see the full version). As seen below, this generalization also models fairness in online algorithms for stochastic bin packing. Let $[k]$ denote the set $\{1, 2, \dots, k\}$ for any positive integer k .

GMP (MP) under unknown distributions: Suppose we have T rounds and a set of items $\mathcal{X} = \{X_1, X_2, \dots, X_n\}$, where each X_i is a random variable taking values from $[0, 1]$. During each round $t \in [T]$, an item i is sampled (we also say i arrives) with probability $p_{i,t}$ for each $i \in [n]$ such that $\sum_{i \in [n]} p_{i,t} \leq 1$. In other words, with probability $1 - \sum_{i \in [n]} p_{i,t}$, no item will be sampled at t . In addition, each time when an item X_i arrives, its index i and distribution are revealed to us. Let $\mathcal{P} = \{p_{i,t} | i \in [n], t \in [T]\}$ be the arrival distributions over the T rounds, which are independent round-by-round. For a given budget B , an instance $I = (T, \mathcal{X}, \mathcal{P})$ is called feasible iff $\sum_{t \in [T]} \sum_{i \in [n]} p_{i,t} \cdot \mathbb{E}[X_i] \leq B$. Let \mathcal{I}_B be the set of all feasible instances with respect to B . Note that only B is known in advance while \mathcal{P} and \mathcal{X} are never revealed to us: the adversary can choose an arbitrary unknown feasible instance I from \mathcal{I}_B before observing our strategy. The same **Three Rules** shown in the original GMP apply here as well.

For an algorithm (strategy) ALG, let σ_i be the expected number of acceptances of item i (over the randomness of ALG and the online arrivals through the T rounds). Define the fairness achieved by ALG over an instance I as $\gamma(\text{ALG}, I) \doteq \min_{i \in [n]} \left(\frac{\sigma_i}{\sum_t p_{i,t}} \right)$, where the denominator refers to the expected number of arrivals of item i over the T rounds. Let $\gamma(\text{ALG}, B) = \inf_{I \in \mathcal{I}_B} \gamma(\text{ALG}, I)$ be the fairness achieved by ALG with respect to budget B . Our goal is to design an

²Unknown *i.i.d.* is typically studied together with another closely-related variant, called random arrival order. algorithm such that the fairness achieved is maximized. A natural

application of the above is in fair online algorithms for stochastic bin packing. As is common in Internet advertising [29], each X_i denotes a “type” of item – e.g., type of job submitted to a server, such as data-intensive, highly-parallel, low-variance runtime, etc. – before a job is accepted, we only have distributional information about it. Maximizing γ is a natural fairness objective in this type of stochastic bin-packing problem. We can define MP under unknown distributions in a same way by just adding one more constraint for the adversary: all X_i are required to be Bernoulli random variables. **Remarks on the adversarial model.** Note that our model of adversary is more powerful than that of an oblivious adversary (which will have to decide on the arrival sequence upfront) [7]. On the other hand, our adversary has less power than, say, adaptive online adversaries [7]: indeed, our approach is to strengthen the oblivious adversary by adding an ingredient of randomness (the $p_{i,t}$), the realizations of which the adversary has no control over. This is an attempt to model non-completely-adversarial, yet random and hence unpredictable, arrival models, as in the case of random item-types in the above-seen online bin-packing problem (where item sizes are also stochastic). Indeed, as shown in Example 1.2, our adversarial model is provably (strictly) stronger than the oblivious adversary of the classical GMP and MP. See, e.g. [11, 13, 31] for additional works on online algorithms under different types of adversaries.

A notable special case is when the arrival distributions are the same throughout the T rounds (not necessarily the same among the n items), i.e., for each $i \in [n]$, $p_{i,t} = p_i$ for all $t \in [T]$. We refer to this as GMP under *unknown independent identical distributions* (UIID). For the general case where the arrival distributions are allowed to change over time, we refer to it as GMP under *unknown adversarial distributions* (UAD) instead. Similarly we can define MP under UIID and UAD respectively when all items are restricted to be Bernoulli random variables. We refer to the original GMP (MP) introduced in [1, 4] as GMP (MP) under (oblivious) adversarial order or simply as GMP (MP) when the context is clear.

We can view GMP as a special case of GMP under UAD: for any instance $\{Y_1, Y_2, \dots\}$, just take $p_{i,t} = 1$ for all t ; we can verify that for a given ALG and instance I , the minimum acceptance probability γ over all items in I is exactly the fairness achieved by ALG on I . For GMP, recall that the adversary’s strategy is fixed in advance, obviously of our strategy; thus, at any time t , we know the adversary’s choices for the full distributions at times $1, 2, \dots, t-1$. This key assumption is not valid for GMP under UAD since \mathcal{P} or \mathcal{X} is never revealed to us, causing significantly more algorithmic challenges. Take the “ γ -conservative” strategy for example as shown in [1, 4]. The main idea is to adaptively compute a sequence of thresholds $\{\tau_t \geq 1\}$ at each time t and the corresponding strategy at t as follows: add the arriving item X_t with probability 1 and 0 respectively if the remaining budget R_t (at the beginning of t) satisfies $R_t > \tau_t$ or $R_t < \tau_t$; add X_t with probability β_t if $R_t = \tau_t$. The sequence of thresholds $\{\tau_t \geq 1\}$ and β_t are computed in an interesting manner in [1, 4] such that

$$\Pr[R_t > \tau_t] + \Pr[R_t = \tau_t] \cdot \beta_t = \gamma, \quad \forall t \in [T]. \quad (1)$$

Note that in GMP, the distribution of R_t can be computed from our previous strategies, together with previous observed inputs. Thus, theoretically for each given target γ , we can sequentially solve

τ_t and β_t from (1) at t . However, this idea fails in the GMP under UAD since we are not informed of \mathcal{P} or \mathcal{X} and thus the distribution of R_t is not computable just from our previous strategies and observed inputs. This is seen in the following example.

Example 1.2. Consider MP with $B = 1$. Consider the following simple algorithm ALG. **Step (1):** For time $t = 1$, accept X_1 with probability $\frac{1}{2}$; **Step (2):** Let $\mathbb{E}[X_1] = \mu_1$ and thus, we see that we are safe (i.e., we have at least one unit budget) at $t = 2$ with probability $\alpha_2 = 1 - \frac{\mu_1}{2} \geq \frac{1}{2}$. For any arriving item X_2 at $t = 2$, ALG will do the following: if we are not safe, then stop; otherwise accept X_2 with probability $\frac{1}{2\alpha_2} \leq 1$. Note that both X_1 and X_2 are accepted by ALG with probability equal to $\frac{1}{2}$; **Step (3):** For a general time $t = 3, 4, \dots$, let all previously-arrived items be Bernoulli random variables with respective means μ_1, \dots, μ_{t-1} . Suppose each arrived item at $t' < t$ is accepted with probability equal to $\frac{1}{2}$ in ALG. We can compute the probability that we are safe at t as follows:

$$\alpha_t = 1 - \frac{1}{2} \sum_{t'=1}^{t-1} \mu_{t'} \geq \frac{1}{2},$$

since $\sum_{t'=1}^{t-1} \mu_{t'} \leq 1$ by the definition of a feasible instance. ALG can continue a similar strategy as outlined in **Step (2)** at t : if safe, then accept X_t with probability $\frac{1}{2\alpha_t} \leq 1$. In this way X_t is accepted by ALG with probability again equal to $\frac{1}{2}$. We can thus verify that ALG accepts all items with probability equal to $\frac{1}{2}$ for MP with $B = 1$.

For MP under unknown distributions with $B = 1$, ALG fails. We cannot do the same as stated in **Step (2)** for MP to compute α_2 , the probability that we are safe at $t = 2$. In fact, after running **Step (1)** here, the updated value $\alpha_2 = 1 - \sum_i p_{i,1} \cdot \mu_i \cdot \frac{1}{2}$, which is not evaluable since $\{p_{i,1}\}$ and the majority of $\{\mu_i\}$ are not revealed to us (only one single μ_i is revealed during $t = 1$). We can rigorously show that no (adaptive) algorithm can achieve a fairness equal to $\frac{1}{2}$ for MP under unknown distributions. See the details in 5.2. ■

Our Contributions: We give a high-level summary of our contributions followed by the details. The first contribution is the new adversary model, which we hope is useful for non-adversarial yet random arrival issues in, say, E-commerce and cloud computing; we give a concrete example showing that this model is strictly stronger than the classical GMP and MP under adversarial order [2, 4]. This also models fairness for certain models of, e.g., online stochastic bin packing. We show a simple greedy algorithm for GMP under UIID, prove that it is optimal, and explicitly describe its (optimal) fairness; as an application, we obtain improved competitive analysis for online B -matching with stochastic rewards under known IID. For the general GMP under UAD, we present a simple non-adaptive algorithm and prove that it is near-optimal over all non-adaptive algorithms. Adaptive algorithms are harder to reason about: for the special case of budget $B = 1$, we present an exact characterization of the optimal adaptive and optimal non-adaptive algorithms. Overall, our algorithms are simple, and the optimality/hardness results are more challenging. We next present further details of our results.

First we consider the GMP under UIID and show that Algorithm 1, referred to as Greedy, is optimal (Section 3).

THEOREM 1.3. *For GMP under UIID, Greedy is optimal. Furthermore, for MP under UIID, Greedy achieves an optimal fairness $\gamma^* = 1 - \frac{B^B}{B!e^B} \geq 1 - \frac{1}{\sqrt{2\pi B}}$.*

One consequence is that we can use MP under UIID as a black-box to re-analyze the performance of the algorithm presented in [9]. Our result shows that for “online B -matching with stochastic rewards under known IID” [9], Greedy achieves an online competitive ratio of $1 - \frac{1}{\sqrt{2\pi B}}(1 + o(1))$, strictly better than the ratio of $(1 - B^{-1/2+\epsilon})$ as shown in [9]. Our analysis also implies that this improved competitive ratio is applicable to a more general online-bipartite-matching scenario where each offline vertex u has capacity at least B : see more details in the full version.

For the general case of GMP under UAD, we restrict our attention to all *non-adaptive* algorithms. Generally speaking, an algorithm ALG is called non-adaptive, if it can be characterized by $\{\beta_t \in [0, 1] | t \in [T]\}$ such that at each time t , ALG will accept the arriving item non-adaptively with probability β_t iff we are safe. Notice that Greedy is a special non-adaptive when $\beta_t = 1$ for all $t \in [T]$. Our second contribution is a near-optimal algorithm among all *non-adaptive* algorithms, denoted by $\text{NAdap}(B, \delta^*)$, for GMP under UAD (Section 4.1).

THEOREM 1.4. *For GMP under UAD with budget B , $\text{NAdap}(B, \delta^*)$ achieves a fairness at least $1 - 2\delta^*$, where $\delta^* = \frac{\sqrt{\ln B}}{\sqrt{B}}(1 + o(1))$ and $o(1)$ is a vanishing term when $B \rightarrow \infty$. Moreover, no non-adaptive algorithm can achieve a fairness better than $1 - \delta^*$ even for GMP.*

Recall that GMP is a special case of GMP under UAD. Thus Theorem 1.4 implies that $\text{NAdap}(B, \delta^*)$ achieves a near-optimal fairness among all non-adaptive algorithms for GMP under UAD.

Finally, we consider a special case of MP under UAD with $B = 1$. We offer exact characterizations for the optimal adaptive and optimal non-adaptive algorithms (Section 5.1). We consider a concrete example of MP under UAD with $T = 4$ and show that the optimal fairness achieved by any algorithm is strictly less than $\frac{1}{2}$. This contrasts with the fact that there is an algorithm which achieves a fairness of $\frac{1}{2}$ for the classical MP with $B = 1$ for all T [2].

2 RELATED WORK

Based on different arrival assumptions, online problems can be divided into the following four categories. The first is Adversarial: the arrival sequence is unknown but fixed. See, e.g., Online matching [26, 35], Adwords [10, 30]). The MP and GMP introduced in [1] and [4] fall into this class. The classical prophet inequality also assumes this arrival setting but *oblivious* adversarial. The second is random arrival order: the set of items is unknown but fixed, and the arrival sequence is a random permutation over all items. See, e.g., online matching [25, 27], Adwords [14, 20], and prophet inequality [12, 17]. The third is unknown distributions: the set of items is unknown but fixed; each round, an item is sampled from a fixed but unknown distribution. If the sampling distributions are required to be the same during each round, we refer to it as *unknown i.i.d.* (UIID) (e.g., [15, 16]); otherwise, we call it *unknown adversarial distributions* (UAD). See, e.g., [15]³. The fourth is known distributions: in each round, an item is sampled from a known distribution. Similarly, we have *known i.i.d.* (KIID) (e.g., [9, 19, 21, 23, 24, 28])

³In [15, 16], this is referred to as adversarial stochastic input.

and *known adversarial distributions* (KAD) (e.g., [3, 4, 8]), depending on if the sampling distributions are allowed to be different over time. Note that Huang and Shu [23] studied the setting of Poisson arrivals that share the essence with KIID.

Among all variants of online bipartite matching, online B -matching (also known as *Display Ads*) might be the most relevant model for GMP (MP). A typical setting is as follows: we have a bipartite graph $G = (U, V, E)$, and in each round, a vertex $v \in V$ arrives (while all of U is “offline” – i.e., already available) and we need to either discard v or assign it to one of its neighbors from U right away. Each u has a capacity of $c_u \geq B$, i.e., it can be assigned at most c_u times. The goal is to design an online allocation policy such that the total expected weight of the assignment is maximized. Feldman et al. [18] considered Display Ads under adversarial order with a “free disposal” assumption and showed an algorithm achieving an online ratio of $1 - 1/e$ when all $c_u \rightarrow \infty$. Alaei et al. [3] introduced Display Ads under known adversarial distributions as Prophet Inequality Matching, where the arrival distributions change over time. They gave an algorithm achieving an online ratio of $1 - \frac{1}{\sqrt{B+3}}$. Brubach et al. [9] introduced online B -matching under known IID with stochastic rewards, where each $e = (u, v)$ is present independently with a known probability p_e ; this models the user v ’s click-rate for the ads u . The online generalized assignment problem (GAP) introduced by Alaei et al. [4] considered all features together and can be viewed as online B -matching under known adversarial distributions such that each assignment $e = (u, v)$ has a random time-sensitive cost of $S_{u,v,t} \in [0, 1]$ for u . For the online GAP, the analysis in [4] actually implies that the standard LP-based online algorithm has an online ratio of γ , when provided with an oracle achieving a ratio of γ for GMP.

3 GMP UNDER UIID

In this section, we consider the case when all $p_{i,t} = p_i$ for each $i \in [n]$. Please see the example below which inspires us to design an optimal algorithm. Throughout this paper, we say “we are safe at t ” iff at the beginning of time t , we still have at least one unit budget left.

Example 3.1. Let $B = 1, n = 1$, i.e., there is only one item X . Assume that for each time t , the single item arrives with probability $p_t = 1$ and that X_t is Bernoulli($1/T$). Observe that the expected number of arrivals of X is T and recall – from the definition of fairness for GMP under IID (or under UAD) – that our objective is reduced to maximizing $\mathbb{E}[Z]/T$ where Z is the total number of acceptances of X .

Consider any optimal algorithm OPT and let $\gamma_t = \Pr[Z_t = 1]$ where Z_t indicates if X is accepted at t in OPT. Thus $\mathbb{E}[Z] = \sum_t \mathbb{E}[Z_t]$. Let $H_t = Z_t \cdot X_t$ indicate if we use the one unit budget B or not. Observe that $\{H_t | t \in [T]\}$ are mutually exclusive events and that $\Pr[H_t = 1] = \frac{\gamma_t}{T}$. Therefore we are safe at t with probability equal to $\Pr[\bigwedge_{t' < t} (H_{t'} = 0)] = 1 - \sum_{t' < t} \gamma_{t'}/T$. Thus, OPT can be viewed as choosing the values $\{\gamma_t | t \in [T]\}$ in order to maximize $\mathbb{E}[Z] = \sum_t \gamma_t$ subject to $\gamma_t \leq 1 - \sum_{t' < t} \frac{\gamma_{t'}}{T}$ for each t . (These constraints are valid since for each time t , the probability that X is accepted by OPT should be at most the probability that we are safe t .) Therefore we obtain the following LP:

$$\max \sum_{t \in [T]} \gamma_t : 0 \leq \gamma_1 \leq 1, 0 \leq \gamma_t \leq 1 - \sum_{t' < t} \frac{\gamma_{t'}}{T}, \forall t \geq 2.$$

We can verify that the optimal solution is $\gamma_t^* = (1 - 1/T)^{t-1}$ for each t . This suggests that OPT will accept X whenever we are safe at each time t , i.e., OPT is essentially greedy. ■

Inspired by the above example, we present a formal description of our greedy algorithm, denoted by Greedy, as follows.

Algorithm 1: Greedy

- 1 For each time t , we accept the arriving item if we are safe, i.e., we still have at least one unit budget.
-

Observe that under UIID, the constraint $\sum_t \sum_i \mathbb{E}[X_i] p_{i,t} \leq B$ is reduced to $\sum_i \mathbb{E}[X_i] p_i \leq B/T$. Greedy treats each item in a uniform way and accepts any item as long as we are safe. For Greedy, the online arrival process can be viewed in the following way: in each round $t \in [T]$, a single random variable $X = \sum_i l_i \cdot X_i$ arrives with probability 1, where l_i is a Bernoulli random variable with mean p_i indicating if X_i comes each round. Notice that (1) $\{l_i | i \in [n]\}$ are independent of $\{X_i | i \in [n]\}$ and (2) X takes values from $[0, 1]$ with $\mathbb{E}[X] = \sum_i p_i \mathbb{E}[X_i] \leq B/T$. Let Y_t indicate if we are safe at t when we run Greedy on X , and $Y = \sum_{t \in [T]} Y_t$ which denotes the (random) number of rounds after which we become unsafe.

Note that the distribution of each Y_t is completely determined by X . For a given B and T , suppose the adversary tries to minimize $\mathbb{E}[Y]/T = \sum_{t \in [T]} \mathbb{E}[Y_t]/T$ by constructing a proper random variable $X = X^*$ with $X^* \in [0, 1]$ and $\mathbb{E}[X^*] \leq B/T$. Let γ^* be the resultant optimal objective value for the adversary. We prove that:

THEOREM 3.2. *For any given B and T , Greedy will achieve a fairness equal to γ^* . Furthermore, no algorithm can achieve a fairness better than γ^* , i.e., Greedy is optimal.*

PROOF. Consider a given input $X = \{X_1, \dots, X_n\}$ and focus on an item X_i . Let Z_t indicate that X_i comes and is accepted at t in Greedy and Y_t indicate if we are safe at t when we run Greedy on X . Thus $\Pr[Z_t = 1] = \Pr[Y_t = 1] p_i$. The fairness achieved by Greedy for item i will be $\gamma_i = \sum_t \mathbb{E}[Z_t]/(T p_i) = \sum_t \mathbb{E}[Y_t]/T$. By the definition of γ^* , we see $\gamma_i \geq \gamma^*$. Thus we see that Greedy achieves a fairness at least γ^* .

Now we show that no algorithm can achieve a fairness better than γ^* . Consider an input I where X consists of one single item $X = X^*$ and X comes with probability 1 in each of the T rounds. Let Z be the number of acceptances of X in any given algorithm ALG. Since Greedy accepts X whenever we are safe, we see that the number of acceptances of X in any ALG should be no larger than that of Greedy. Thus we claim $Z \leq Y$, where Y is the number of acceptances when Greedy runs on I . Therefore, $\mathbb{E}[Z]/T \leq \mathbb{E}[Y]/T = \gamma^*$. Note that $\mathbb{E}[Z]/T$ is exactly the fairness achieved by ALG for the item X . Thus we have proven our claim. □

3.1 MP under UIID

We now consider a special case: MP under UIID, i.e., all the X_i are Bernoulli variables, and give an asymptotically tight expression for γ^* when $B \gg 1$.

THEOREM 3.3. *The fairness achieved by Greedy for MP under UIID with input B is $\gamma^* = 1 - \frac{B^B}{B!} e^{-B} \approx 1 - \frac{1}{\sqrt{2\pi B}}$.*

In the case when all X_i are Bernoulli, we see that $X = \sum_i I_i X_i$ is also a Bernoulli random variable with mean $\mathbb{E}[X] \leq B/T$. For a given B and T , the adversary will hence arrange $X^* \sim \text{Ber}(B/T)$ to ensure that $\gamma = \mathbb{E}[Y]/T$ is minimized, where $Y \leq T$ is the number of rounds after which we become unsafe. Consider the random process where we have a bin with capacity of B and where, in each round, we independently receive a ball with probability B/T . Let Y' be the number of rounds after which we exhaust the capacity. Observe that $\mathbb{E}[Y'] = T$ and $Y = \min(Y', T)$. Note that our problem is different from that raised and solved in [16, 36] for Adwords: they are concerned with $\min(Y'', B)$ where Y'' is the number of balls received after T rounds.

Notice that for each given $t \geq B$,

$$\Pr[Y' = t] = \left(\frac{B}{T}\right)^B \binom{t-1}{B-1} \left(1 - \frac{B}{T}\right)^{t-B} = \left(\frac{B}{T-B}\right)^B \binom{t-1}{B-1} \left(1 - \frac{B}{T}\right)^t.$$

Therefore, we have

$$\begin{aligned} T - \mathbb{E}[Y] &= \mathbb{E}[Y'] - \mathbb{E}[Y] \\ &= \sum_{t=T+1}^{\infty} \left(\frac{B}{T-B}\right)^B \binom{t-1}{B-1} \left(1 - \frac{B}{T}\right)^t (t-T) \\ &= \left(\frac{B}{T-B}\right)^B \left(1 - \frac{B}{T}\right)^T \sum_{\ell=1}^{\infty} \binom{T+\ell-1}{B-1} \left(1 - \frac{B}{T}\right)^\ell \ell \\ &= \frac{B^B}{B!} \left(1 - \frac{B}{T}\right)^T \frac{T(T-1) \cdots (T-B)}{(T-B)^B}. \end{aligned}$$

Thus, we have

$$\gamma = \frac{\mathbb{E}[Y]}{T} = 1 - \frac{B^B}{B!} \left(1 - \frac{B}{T}\right)^T \frac{(T-1) \cdots (T-B)}{(T-B)^B} = 1 - F(B, T). \quad (2)$$

LEMMA 3.4. *For each given $B \geq 1$, $F(B, T)$ is an increasing function of $T \geq B$ and $\lim_{T \rightarrow \infty} F(B, T) = \frac{B^B}{B!} e^{-B}$.*

PROOF. Let $T = B + k$ with $k \geq 0$. Notice that

$$F(B, T) = \frac{B^B}{B!} \left(\frac{k}{B+k}\right)^{k+1} \left(\frac{k+1}{k+B} \cdot \frac{k+2}{k+B} \cdots \frac{k+B-1}{k+B}\right),$$

which implies that $F(B, T)$ is an increasing function of $T \geq B$. When $T \rightarrow \infty$, we see that the limit of $F(B, T)$ is $\frac{B^B}{B!} e^{-B}$. \square

Therefore for each given B , the adversary will arrange an instance of MP under UIID with $T \rightarrow \infty$ such that γ is minimized. Plugging the result of Lemma 3.4 into (2) yields Theorem 3.3. By combining the results of Theorems 3.2 and 3.3, we obtain Theorem 1.3.

4 GMP UNDER UAD

In this section we consider the GMP under UAD. Recall that in this case, the arrival distributions can change in different rounds.

4.1 A near-optimal non-adaptive algorithm

First, we show that Greedy can be arbitrarily bad for GMP under UAD.

Example 4.1. Let \mathcal{X} be the disjoint union of $\{X_1, \dots, X_n\}$ and $\{X_a, X_b\}$. $\mathbb{E}[X_i] = \mu_i = 1$ for all $i \in [n]$ and $\mathbb{E}[X_a] = \mu_a = 1 - \epsilon$ and $\mathbb{E}[X_b] = \mu_b = \epsilon/(T - B)$. For $t \in [B - 1]$, $p_{i,t} = 1/n$ for each $i \in [n]$ and $p_{a,t} = p_{b,t} = 0$. For $t = B$, $p_{i,t} = 0$ for each $i \in [n]$ and $p_{b,t} = 0, p_{a,t} = 1$; for $B < t \leq T$, $p_{i,t} = 0$ for each $i \in [n]$ and $p_{a,t} = 0, p_{b,t} = 1$. Observe that $\sum_{t \in [T]} \sum_{i \in [n] \cup \{a, b\}} p_{i,t} \mu_i = B$. Now suppose we run Greedy on \mathcal{X} . First, the expected number of arrivals is $\sum_{t \in [T]} p_{b,t} = T - B$. Second, note that at the beginning of $t = B + 1$, we are safe with probability ϵ , implying that the expected number of acceptances of X_b is at most $\epsilon(T - B)$. Thus, Greedy achieves fairness at most ϵ for X_b . \blacksquare

Recall that for a given B , an instance $\mathcal{X} = \{X_1, \dots, X_T\}$ of GMP is called feasible iff each X_i lies in $[0, 1]$ with $\sum_{t \in [T]} \mathbb{E}[X_i] \leq B$. Let $\{Y_t | t \in [T]\}$ be T i.i.d. Bernoulli random variables with mean $1 - \delta$ each, where δ is a parameter. Set $H_{\mathcal{X}} = \sum_{t \in [T]} X_t \cdot Y_t$. For a given δ and B , let $F(B, \delta) = \sup_{\mathcal{X} \in \mathcal{J}_B} \Pr[H_{\mathcal{X}} > B - 1]$ where \mathcal{J}_B is the set of all feasible instances of GMP with respect to B .

Let NALG be the set of all possible non-adaptive algorithms. For an instance I of GMP and an algorithm $\text{ALG} \in \text{NALG}$, let $\gamma(\text{ALG}, I)$ be the minimum acceptance probability of ALG over all the items. We can verify that this coincides with the fairness defined for GMP under UAD. For a given B , let $\gamma^*(B) = \sup_{\text{ALG} \in \text{NALG}} \inf_{I \in \mathcal{J}_B} \gamma(\text{ALG}, I)$, which refers to the fairness achieved by the optimal non-adaptive algorithm for GMP with input B .

THEOREM 4.2. *Let $\delta^* \in [0, 1]$ be the unique solution to the equation $F(B, \delta) = \delta$. Then, $\gamma^*(B) \leq 1 - \delta^*$.*

PROOF. Observe that $F(B, \delta)$ is a non-increasing function of δ . Also, we can verify that $F(B, 0) = 1$ and $F(B, 1) = 0$, which justifies the existence of a unique solution in $[0, 1]$ for the equation $F(B, \delta) = \delta$.

Let NOPT be an optimal non-adaptive algorithm for GMP with input B . Suppose for a contradiction that the fairness achieved by NOPT is $\gamma^*(B) > 1 - \delta^*$. Let $\epsilon = (\gamma^*(B) - (1 - \delta^*))/2$. By the definition of $F(B, \delta^*)$, we see that there exists a feasible instance $\mathcal{X}' = \{X_1, X_2, \dots, X_T\}$ such that $\Pr[H_{\mathcal{X}'} > B - 1] > \delta^* - \epsilon$ where $H_{\mathcal{X}'} = \sum_{t \in [T]} X_t * Y_t$ with each $Y_t \sim \text{Ber}(1 - \delta^*)$. Consider the instance $\mathcal{X}'' = \mathcal{X}' \cup \{X_{T+1}\}$ where X_{T+1} is the last item with mean 0. Thus \mathcal{X}'' is still a feasible instance of GMP with respect to B . Suppose we run NOPT on \mathcal{X}'' and define U_T to be the usage of the budget at the end of time T . Since the last item X_{T+1} can be accepted with probability at least $\gamma^*(B)$ by NOPT, we must have $\Pr[U_T > B - 1] \leq 1 - \gamma^*(B)$.

Notice that in NOPT, each item in \mathcal{X}'' is accepted with probability at least $\gamma^*(B) > 1 - \delta^*$ non-adaptively. Consider such a non-adaptive ALG which accepts each item with probability exactly equal to $1 - \delta^*$ whenever we are safe. Let U'_T be the expected usage of budget for ALG at the end of T . We see that

$$\Pr[U'_T > B - 1] = \Pr[H_{\mathcal{X}'} > B - 1] \leq \Pr[U_T > B - 1] \leq 1 - \gamma^*(B)$$

On the other hand, we have that $\Pr[H_{\mathcal{X}'} > B - 1] > \delta^* - \epsilon$, which implies that $\epsilon > \delta^* - (1 - \gamma^*(B))$. This contradicts our assumption that $\epsilon = (\delta^* - (1 - \gamma^*(B)))/2 > 0$. \square

Since GMP is a special case of GMP under UAD, we have the following corollary.

COROLLARY 4.3. *For GMP under UAD with budget B , no non-adaptive algorithm can achieve a fairness better than $1 - \delta^*$.*

We next present a simple non-adaptive algorithm, denoted by $\text{NAdap}(B, \delta^*)$, for GMP under UAD with budget B :

Algorithm 2: $\text{NAdap}(B, \delta^*)$

1 For each time t , we accept the arriving item with probability $1 - \delta^*$ (non-adaptively) whenever we are safe.

THEOREM 4.4. *Algorithm $\text{NAdap}(B, \delta^*)$ achieves a fairness at least $(1 - \delta^*)^2 \geq 1 - 2\delta^*$ for GMP under UAD with budget B .*

PROOF. Consider a given feasible instance $I = \{T, \mathcal{X}, \mathcal{P}\}$ of GMP under UAD with budget B , where $\mathcal{X} = \{X_i | i \in [n]\}$ and $\mathcal{P} = \{p_{i,t} | i \in [n], t \in [T]\}$. Since $\text{NAdap}(B, \delta^*)$ is non-adaptive, we can view the arrival process as follows: in each round $t \in [T]$ a single random variable $X'_t = \sum_i l_{i,t} \cdot X_i$ arriving with probability 1 where $l_{i,t}$ indicates that X_i comes at t , with $\mathbb{E}[l_{i,t}] = p_{i,t}$; we accept each X'_t non-adaptively with probability $1 - \delta^*$ whenever we are safe. Let $\{Y_t | t \in [T]\}$ be *i.i.d.* Bernoulli random variables with mean $1 - \delta^*$ each. Let U_t be the usage of the budget at the end of t when we run $\text{NAdap}(B, \delta^*)$ over I ; thus, $\chi(U_t \leq B - 1)$ indicates that we are safe at time $t + 1$.

Observe that $\Pr[U_t > B - 1] = \Pr[\sum_{\ell \leq t} X'_\ell \cdot Y_\ell > B - 1]$. Also each $X'_\ell \in [0, 1]$ and $\mathbb{E}[\sum_{\ell \leq t} X'_\ell] \leq B$; thus, $\{X'_\ell | \ell \leq t\}$ itself can be viewed as a feasible instance of GMP (under adversarial) with respect to B . From the definition of $F(B, \delta^*)$, we see that for each t ,

$$\Pr[U_t > B - 1] = \Pr\left[\sum_{\ell \leq t} X'_\ell \cdot Y_\ell > B - 1\right] \leq F(B, \delta^*) = \delta^*,$$

which implies that $\Pr[U_t \leq B - 1] \geq 1 - \delta^*$. Notice that for each given X_i with $i \in [n]$, the expected number of acceptances of X_i over the T rounds is $\mathbb{E}[\sum_t \chi(U_t \leq B - 1) \cdot Y_{t+1} \cdot l_{i,t+1}] \geq (1 - \delta^*)^2 \sum_t l_{i,t}$ while its expected number of arrivals is $\sum_t l_{i,t}$. Thus from the definition of fairness, $\text{NAdap}(B, \delta^*)$ achieves a fairness at least $(1 - \delta^*)^2$ on the instance I . This completes the proof. \square

4.2 Computation of δ^*

By Theorem 4.4, we obtain a near-optimal non-adaptive algorithm for GMP under UAD whenever δ^* is known. In this section, we assume $B \gg 1$ and give an asymptotically tight form for δ^* .

THEOREM 4.5. $\delta^* = (1 + o(1))B^{-1/2}\sqrt{\ln B}$, where $o(1)$ is a vanishing term as B increases.

4.2.1 A Lower Bound on δ^* . First we show a lower bound on δ^* as follows.

LEMMA 4.6. $\delta^* \geq (1 + o(1))B^{-1/2}\sqrt{\ln B}$, where $o(1)$ is a vanishing term as B increases.

PROOF. Consider the following instance of MP with respect to B : $\mathcal{X}' = \{X_1, \dots, X_T\}$ where $\{X_t\}$ are T *i.i.d.* Bernoulli random variables with mean B/T each. Recall that $H_{\mathcal{X}} = \sum_{t \in [T]} X_t \cdot Y_t$ where $\{Y_t | t \in [T]\}$ are T *i.i.d.* Bernoulli random variables with mean $1 - \delta$ each, which are independent from $\{X_t\}$. For our instance \mathcal{X}' , we see that $H_{\mathcal{X}'}$ is a sum of T *i.i.d.* Bernoulli random variables with mean $B(1 - \delta)/T$ each. For each given B and δ , we have that $H_{\mathcal{X}'}$

follows a Poisson distribution with mean $B(1 - \delta)$, when $T \rightarrow \infty$. By applying the Berry-Esseen Theorem [32], we see that $\frac{H_{\mathcal{X}'} - B(1 - \delta)}{\sqrt{B(1 - \delta)}}$ can be approximated by $\mathcal{N}(0, 1)$ with error at most $\frac{1}{2\sqrt{B(1 - \delta)}}$. Thus we have

$$\begin{aligned} \Pr[H_{\mathcal{X}'} > B - 1] &= \Pr\left[\frac{H_{\mathcal{X}'} - B(1 - \delta)}{\sqrt{B(1 - \delta)}} \geq \frac{\sqrt{B}\delta}{\sqrt{1 - \delta}}\right] + O\left(\frac{1}{\sqrt{B}}\right) \\ &= \frac{1}{\sqrt{2\pi}} \int_{\tau}^{\infty} e^{-x^2/2} dx + O\left(\frac{1}{\sqrt{B}}\right) \\ &= \frac{1}{\sqrt{2\pi}} \frac{\exp(-\tau^2/2)}{\tau} (1 - o(1)) + O\left(\frac{1}{\sqrt{B}}\right) \left(\tau \doteq \frac{\sqrt{B}\delta}{\sqrt{1 - \delta}}\right), \end{aligned}$$

where the term $o(1)$ in the last line vanishes when $B \rightarrow \infty$ and lies between 0 and $1/\tau^2$.

Consider the equation $\Pr[H_{\mathcal{X}'} > B - 1] = \delta$, which is reduced to

$$\frac{1}{\sqrt{2\pi}} \exp\left(\frac{-B\delta^2}{2(1 - \delta)}\right) (1 - o(1)) + O(\delta) = \frac{\delta^2}{\sqrt{1 - \delta}} \sqrt{B}.$$

Solve this, we get that $\delta' = (1 + o(1))B^{-1/2}\sqrt{\ln B}$, where $o(1)$ is a vanishing term when $B \rightarrow \infty$. Notice that $F(B, \delta) = \sup_{\mathcal{X} \in \mathcal{I}_B} \Pr[H_{\mathcal{X}} > B - 1] \geq \Pr[H_{\mathcal{X}'} > B - 1]$. In other words, if we view $F(B, \delta)$ and $\Pr[H_{\mathcal{X}'} > B - 1]$ as two functions of δ for a given B , we see that both are decreasing over $\delta \in [0, 1]$ and the graph of $F(B, \delta)$ lies above that of $\Pr[H_{\mathcal{X}'} > B - 1]$. Since δ^* and δ' are respectively the intersection points of $y = \delta$ and the two functions, we claim that $\delta^* \geq \delta'$. \square

4.2.2 An Upper Bound on δ^* . Now we show an upper bound on δ^* as follows.

LEMMA 4.7. $\delta^* \leq (1 + o(1))B^{-1/2}\sqrt{\ln B}$, where $o(1)$ is a vanishing term as B increases.

In this paper we use the following form of the Chernoff bound to prove the above lemma.

Definition 4.8 (Chernoff Bound). Let X_1, \dots, X_n be n independent random variables with $0 \leq X_i \leq 1$. Let $X = X_1 + \dots + X_n$ and $\mu = \mathbb{E}[X]$. Then for any $\Delta > 0$,

$$\Pr[X \geq (1 + \Delta)\mu] \leq \exp\left(-\frac{\Delta^2}{2 + \Delta}\mu\right). \quad (3)$$

PROOF OF LEMMA 4.7. Consider a given B and a given feasible instance $\mathcal{X} \in \mathcal{I}_B$. Recall that $H_{\mathcal{X}} = \sum_{t \in [T]} X_t \cdot Y_t$ where each $X_t \in [0, 1]$ and Y_t is Bernoulli random variable with mean $1 - \delta$ with $\sum_{t \in [T]} \mathbb{E}[X_t] \leq B$. Thus we can view $H_{\mathcal{X}}$ as a sum of T independent random variable each taking value from $[0, 1]$ with total mean at most $B(1 - \delta)$. To apply the Chernoff bound shown in 4.8, we have $\mu = B(1 - \delta)$ and $\Delta = \frac{1 - 1/B}{1 - \delta} - 1$. Here we assume $\delta \gg B^{-1}$. This is allowed since B can be taken arbitrarily large.

$$\begin{aligned} \Pr[H_{\mathcal{X}} > B - 1] &\leq \exp\left(-\frac{\Delta^2}{2 + \Delta}\mu\right) \\ &= \exp\left(-\frac{B(1 - \delta)}{2 + \frac{1 - 1/B}{1 - \delta} - 1} \left(\frac{1 - 1/B}{1 - \delta} - 1\right)^2\right) \\ &= \exp\left(-\frac{B\delta^2}{2} (1 + o(1))\right), \end{aligned}$$

where $o(1)$ is a vanishing term when $B \rightarrow \infty$. Solving the equation $\exp\left(-\frac{B\delta^2}{2}(1+o(1))\right) = \delta$, we get that $\delta' = (1+o(1))B^{-1/2}\sqrt{\ln B}$. Notice that \mathcal{X} is an arbitrary feasible instance in \mathcal{I}_B , thus we claim that $F(B, \delta) \leq \exp\left(-\frac{B\delta^2}{2}(1+o(1))\right)$. Again suppose we try to view $F(B, \delta)$ and $\exp\left(-\frac{B\delta^2}{2}(1+o(1))\right)$ as two decreasing functions of $\delta \in [0, 1]$ while δ^* and δ' are respectively the intersection points of the line $y = \delta$ with the two functions. Therefore we have that $\delta^* \leq \delta'$ and we prove our claim. \square

Theorem 4.5 directly follows from the results in Lemmas 4.6 and 4.7. The combination of results in Corollary 4.3 and Theorems 4.4 and 4.5 yields that NAdap is nearly an optimal non-adaptive algorithm for GMP under UAD.

5 MP UNDER UAD WITH $B = 1$

In this section, we consider the following special case: MP under UAD with $B = 1$.

5.1 An optimal non-adaptive algorithm

Consider a given finite T . Let $\{\beta_t | t \in [T]\}$ denote an optimal non-adaptive algorithm, denoted by NOPT. In other words, NOPT will accept the arriving item at t with probability β_t whenever safe. Now we discuss how to compute $\{\beta_t | t \in [T]\}$ for a given T . For each $t \in [T]$, let \mathcal{J}_t be the set of all feasible instances of GMP with $B = 1$ and time horizon t . In other words, $\mathcal{J}_t = \{A_{t'} | 1 \leq t' \leq t\}$ where $\{A_{t'} | t' \in [t]\}$ are all independent random variables and each takes a value from $[0, 1]$ and $\sum_{t' \in [t]} \mathbb{E}[A_{t'}] \leq 1$.

LEMMA 5.1. NOPT achieves a fairness of γ for MP under UAD with $B = 1$ iff for each $1 \leq t \leq T$,

$$\gamma \leq \mathbb{E}\left[\beta_t \prod_{t' < t} (1 - \beta_{t'} A_{t'})\right],$$

for all possible $\{A_{t'} | t' < t\} \in \mathcal{J}_{t-1}$.

PROOF. We first show the necessity. Consider a given $1 \leq t \leq T$ and a given set of $\{A_{t'} | t' < t\}$ such that $\{A_{t'}\}$ are all independent random variables taking values from $[0, 1]$ with total mean at most 1. WLOG assume the union of support of $A_{t'}$ is a finite set $\mathcal{S} = \{\mu_1, \mu_2, \dots, \mu_n\}$ and for each $t' < t$, $\Pr[A_{t'} = \mu_i] = q_{i, t'}$. Consider the following instance I created by the adversary: $\mathcal{X} = \{X_1, X_2, \dots, X_n\} \cup \{X_{n+1}\}$ such that $X_i = \text{Ber}(\mu_i)$ for each $i \in [n]$, and $X_{n+1} = \text{Ber}(0)$. During each round of t' with $t' < t$, item X_i comes with probability $p_{i, t'}$ for each $i \in [n]$ and X_{n+1} comes with probability 0; during the round t , item X_{n+1} comes with probability 1 and no one else will come; no item will come after round t . Note that the sum of mean of $A_{t'}$ is at most 1 implies that instance I is feasible. We can verify that the fairness achieved by NOPT on the item X_{n+1} is exactly equal to $\mathbb{E}\left[\beta_t \prod_{t' < t} (1 - \beta_{t'} A_{t'})\right]$. The fact that NOPT achieves a fairness of γ justifies the constraint at t .

Now we show the sufficiency. Consider a given instance $I = (T, \mathcal{X}, \mathcal{P})$, where $\mathcal{X} = \{X_1, \dots, X_n\}$ and $\mathcal{P} = \{p_{i, t}\}$. Let $\mu_i = \mathbb{E}[X_i]$. For each $t \in [T]$, define a random variable A_t such that $\Pr[A_t = \mu_i] = p_{i, t}$ for each $i \in [n]$. Observe that each A_t takes values from $[0, 1]$ with $\sum_{i \in [n]} \mathbb{E}[A_t] \leq 1$. For each $t \in [T]$, we can verify that

we are safe with probability equal to $\mathbb{E}\left[\prod_{t' < t} (1 - \beta_{t'} A_{t'})\right]$. This implies that for each item X_i , the expected number of acceptances should be

$$\sum_{t \in [T]} p_{i, t} \cdot \beta_t \cdot \mathbb{E}\left[\prod_{t' < t} (1 - \beta_{t'} A_{t'})\right] \geq \gamma \sum_{t \in [T]} p_{i, t}.$$

Thus, we claim that NOPT achieves a fairness at least γ . \square

To find a maximum γ satisfying all the constraints in Lemma 5.1, we just need to solve the following maximization program.

$$\begin{aligned} \max \gamma & \\ \gamma &\leq \beta_t (1 - \beta_{t'}) \quad \forall 1 \leq t' < t \leq T \\ 0 &\leq \beta_t \leq 1 \quad \forall t \in [T] \end{aligned} \quad (4)$$

LEMMA 5.2. An optimal solution $\{\beta_t | 1 \leq t \leq T\}$ and γ to the maximization program (4) can be solved from the following equations:

$$\beta_1 = \gamma, \beta_t = \frac{\gamma}{1 - \beta_{t-1}}, \forall 1 < t < T, \gamma = 1 - \beta_{T-1}, \beta_T = 1. \quad (5)$$

PROOF. Let $\{\beta^*\}$ and γ^* be the unique solution to the program (5). Ignore the equation $\gamma = 1 - \beta_{T-1}$; each $\beta_t, t < T$ can be viewed as an function of γ , which is strictly increasing over $[0, 1]$. Thus we can get an unique solution γ^* from the last equation $\gamma = 1 - \beta_{T-1}$.

The feasibility of $\{\beta^*\}$ and γ^* to the maximization program (4) can be verified straightforwardly. Now we show the proof of optimality. Suppose the optimal value to the program (4) is $\gamma' > \gamma^*$ which is achieved on $\{\beta'_t\}$. Then we see that

$$\beta'_1 \geq \gamma', \beta'_t \geq \frac{\gamma'}{1 - \beta'_{t-1}}, \forall 1 < t < T, 1 - \beta'_{T-1} \geq \gamma'.$$

By comparing the above inequalities to the equations in program (5), we claim that $\beta'_{T-1} > \beta^*_{T-1}$, which is followed by $\gamma' \leq 1 - \beta'_{T-1} < 1 - \beta^*_{T-1} = \gamma^*$. We get a contradiction. \square

Let $\gamma^*(T)$ be the optimal value to the program (5). Numerically we can verify that

$$\gamma^*(1) = 1, \gamma^*(2) = \frac{1}{2}, \gamma^*(3) = \frac{3 - \sqrt{5}}{2} \sim 0.3819, \gamma^*(4) = \frac{1}{3}.$$

LEMMA 5.3. $\lim_{T \rightarrow \infty} \gamma^*(T) = 0$.

PROOF. Consider a given γ . Ignore the equation $\gamma = 1 - \beta_{T-1}$ and view each $\beta_t, t < T$ as an increasing function of γ . Since $\beta_{T-1} \leq 1$, we claim that $\{\beta_1, \beta_2, \dots, \beta_{T-1}\}$ must converge to a value satisfying the equation $\frac{\gamma}{1 - \beta} = \beta$ when $T \rightarrow \infty$. From the last equation, we see $\gamma = 1 - \beta$, which implies $\beta = 1, \gamma = 0$. \square

5.2 An optimal adaptive algorithm

In this section, we present an optimal adaptive algorithm, which is denoted by OPT.

LEMMA 5.4. The optimal adaptive strategy at t for the MP under UAD with $B = 1$ just needs to adapt to $\mathcal{U}_{t-1} = (\mu_1, \dots, \mu_{t-1})$ subject to being safe, where $\mu_{t'}$ is the mean of the arriving item at t' for each $t' < t$. It need not adapt to the random outcomes of previous optimal strategies and realizations of accepted items.

PROOF. Consider a finite (known) T and suppose the worst unknown but fixed input arranged by the adversary is $I = (\mathcal{X}, \mathcal{P})$, where $\mathcal{X} = \{X_1, X_2, \dots, X_n\}$ and $\mathcal{P} = \{p_{i,t} | i \in [n], t \in [T]\}$ with $\sum_{i,t} p_{i,t} \mathbb{E}[X_i] \leq 1$ and each X_i is a Bernoulli random variable.

Consider a state $\mathcal{S}_t = (t, B = 1, \mathcal{U}_t)$ with $\mathcal{U}_t = \{\mu_1, \dots, \mu_{t-1}, \mu_t\}$, i.e., we are safe at time t and all observations so far are $\mathcal{U}_t = \{\mu_1, \dots, \mu_{t-1}, \mu_t\}$ where $\mu_{t'}$ is the mean of the arriving item at $t' \in [t]$. Let PATH be the set of all paths in the decision tree of OPT ending at \mathcal{S} . Notice that we can extract the same information from each path $P \in \text{PATH}$ for the future arrivals \mathcal{X} and $\{p_{i,t'} | t' > t\}$. Thus we claim the optimal strategy should be the same if we end at \mathcal{S}_t following each path $P \in \text{PATH}$.

Let $f_t(\mathcal{U}_t)$ be the probability that OPT accepts the arrival item Y_t at t conditioning on we are safe at t and \mathcal{U}_t . Now we show further that $f_t(\mathcal{U}_t)$ has nothing to do with $\mu_t = \mathbb{E}[Y_t]$. Suppose we run OPT and at time t , we are safe and observe an arriving item $Y_t = \text{Ber}(\mu_t)$. Consider the following worst scenario for Y_t arranged by the adversary: Y_t never comes before (its index revealed to us is different from all previous) and it comes exclusively at t with probability ϵ . Thus OPT achieving a fairness of γ over Y_t is equivalent to the property that for all $\mu_t \in [0, 1]$, we have

$$\min_{I(t-1) \in \mathcal{I}_1(t-1)} \mathbb{E}_{\mathcal{U}_{t-1}} \left[\Pr[\text{OPT is safe at } t | \mathcal{U}_{t-1}] \cdot f_t(\mu_1, \dots, \mu_t) \right] \geq \gamma,$$

where (1) \mathcal{U}_{t-1} is a random vector of means of all arrival items coming during the first $t-1$ rounds whose distribution is determined by the unknown input projected into the first $t-1$ rounds, denoted by $I(t-1)$; (2) $\mathcal{I}_1(t-1)$ is the set of all feasible instances of MP under UAD with $B = 1$ and time horizon $T = t-1$. Notice that neither of the values $\Pr[\text{OPT is safe at } t | \mathcal{U}_{t-1}]$ and $\mathcal{I}_1(t-1)$ is connected to μ_t , thus we claim that $f_t(\mu_1, \mu_2, \dots, \mu_t)$ in OPT has no dependence on μ_t either. \square

Consider a finite T . From Lemma 5.4 we see that OPT can be characterized by $\{f_t | t \in [T]\}$ where for each $t \in [T]$, $f_t : [0, 1]^{t-1} \rightarrow [0, 1]$ is such that $f_t(\mathcal{U}_{t-1})$ denotes the probability that OPT accepts the arriving item at t conditioning on being safe at t and \mathcal{U}_{t-1} . The following lemma gives a sufficient and necessary condition when $\text{OPT} = \{f_t | t \in [T]\}$ achieves a fairness of γ . Recall that \mathcal{J}_t is the set of all feasible instances of GMP with $B = 1$ and a time horizon of t .

LEMMA 5.5. *An OPT parameterized by $\{f_t | t \in [T]\}$ achieves a fairness of γ for MP under UAD with $B = 1$ and a time horizon of T iff for each $1 \leq t \leq T$,*

$$\gamma \leq \mathbb{E} \left[f_t(A_1, A_2, \dots, A_{t-1}) \prod_{t' < t} \left(1 - f_{t'}(A_1, \dots, A_{t'-1}) A_{t'} \right) \right],$$

for all possible $\{A_{t'} | 1 \leq t' < t\} \in \mathcal{J}_{t-1}$.

The proof of Lemma 5.5 is similar to that of Lemma 5.1. We omit it here. For a given T , let $\widehat{\gamma}(T)$ be the maximum value satisfying all constraints in Lemma 5.5. It turns out to be much more challenging than before to compute $\widehat{\gamma}(T)$. Here we only consider the simple case $T = 4$ and show $\widehat{\gamma}(4) < \frac{1}{2}$. Note that for the classic MP with $B = 1$ for an arbitrary T , Theorem 1.1 in [2] indicates that there is an algorithm which achieves a fairness of $\frac{1}{2}$. This suggests that MP under UAD is strictly harder than the classical MP under adversarial order.

LEMMA 5.6. *For MP under UAD with $B = 1$ and $T = 4$, OPT achieves a fairness of $\widehat{\gamma}(4) < \frac{1}{2}$.*

PROOF. We prove by contradiction. Suppose $\widehat{\gamma}(4) \geq 1/2$. From Lemma 5.5, we see that the optimal choice is $f_1 = \gamma$, $f_2(x) = \frac{1}{1-\gamma x}$ for each $x \in [0, 1]$ and $f_4 = 1$. The key issue is to choose $f_3 \doteq f$. The conditions stated in Lemma 5.5 for $t = 3, 4$ are reduced to the following:

$$\begin{aligned} \mathbb{E} \left[f(X_1, X_2) (1 - \gamma(X_1 + X_2)) \right] &\geq \gamma, \\ \mathbb{E} \left[\left(1 - f(X_1, X_2) X_3 \right) (1 - \gamma(X_1 + X_2)) \right] &\geq \gamma, \end{aligned} \quad (6)$$

for all possible $\{X_1, X_2\} \in \mathcal{J}_2$ and $\{Y_1, Y_2, Y_3\} \in \mathcal{J}_3$. Consider these two concrete examples: $\{X_1, X_2\}$ are two *i.i.d.* Bernoulli random variables with mean $1/2$ while $\{Y_1, Y_2, Y_3\}$ are three *i.i.d.* Bernoulli random variables with mean $1/3$. Let $a = (f(1, 0) + f(0, 1))/2$ and $b = f(0, 0)$. After simplification, the two inequalities in (6) on these two examples with $\gamma \geq 1/2$ are reduced to the following: $a + b \geq 2$, $a + 2b \leq \frac{9}{4}$. Notice that both a and b take value in $[0, 1]$ and thus this linear system is infeasible. \square

6 CONCLUSION AND FUTURE WORK

In this paper, we have considered a generalization of GMP (MP) as introduced in [1, 4] and presented two near-optimal non-adaptive online algorithms for unknown IID and the more-general unknown adversarial distributions.

For GMP under UIID, we have proven that Greedy is optimal while the exact optimal fairness is not known yet. We conjecture that it should have the same performance as that of MP under UIID, i.e., that the adversary will arrange all items as Bernoulli random variables in the worst case. Another direct future direction is to show some hardness results for GMP under UAD regarding the optimal adaptive algorithm: is $\text{NAdap}(B, \delta^*)$ also near-optimal among all adaptive algorithms for GMP under UAD?

REFERENCES

- [1] Saeed Alaei. 2011. Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers. In *Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, 512–521.
- [2] Saeed Alaei. 2014. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM J. Comput.* 43, 2 (2014), 930–972.
- [3] Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. 2012. Online prophet-inequality matching with applications to ad allocation. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 18–35.
- [4] Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. 2013. The online stochastic generalized assignment problem. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Springer, 11–25.
- [5] Saeed Alaei, Mohammad T Hajiaghayi, Vahid Liaghat, Dan Pei, and Barna Saha. 2011. Adcell: Ad allocation in cellular networks. In *European Symposium on Algorithms*. Springer, 311–322.
- [6] Bahman Bahmani and Michael Kapralov. 2010. Improved bounds for online stochastic matching. *Algorithms-ESA 2010, 18th Annual European Symposium (2010)*, 170–181.
- [7] Shai Ben-David, Allan Borodin, Richard M. Karp, Gábor Tardos, and Avi Wigderson. 1994. On the Power of Randomization in On-Line Algorithms. *Algorithmica* 11, 1 (1994), 2–14.
- [8] Allan Borodin, Calum MacRury, and Akash Rakheja. 2021. Prophet Inequality Matching Meets Probing with Commitment. *arXiv preprint arXiv:2102.04325* (2021).
- [9] Brian Brubach, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. 2016. New Algorithms, Better Bounds, and a Novel Model for Online Stochastic Matching. In *24th Annual European Symposium on Algorithms (ESA 2016)*.
- [10] Niv Buchbinder, Kamal Jain, and Joseph Seffi Naor. 2007. Online primal-dual algorithms for maximizing ad-auctions revenue. In *European Symposium on Algorithms*. Springer, 253–264.
- [11] Nicolò Cesa-Bianchi, Ofer Dekel, and Ohad Shamir. 2013. Online Learning with Switching Costs and Other Adaptive Adversaries. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, Christopher J. C. Burges and Léon Bottou, Zoubin Ghahramani, and Kilian Q. Weinberger (Eds.), 1160–1168.
- [12] José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. 2017. Posted price mechanisms for a random stream of customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. 169–186.
- [13] Ofer Dekel, Ambuj Tewari, and Raman Arora. 2012. Online Bandit Learning against an Adaptive Adversary: from Regret to Policy Regret. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012*. icml.cc / Omnipress.
- [14] Nikhil R Devanur and Thomas P Hayes. 2009. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *Proceedings of the 10th ACM conference on Electronic commerce*. ACM, 71–78.
- [15] Nikhil R Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A Wilkens. 2011. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, 29–38.
- [16] Nikhil R Devanur, Balasubramanian Sivan, and Yossi Azar. 2012. Asymptotically optimal algorithm for stochastic adwords. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 388–404.
- [17] Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monemizadeh. 2017. Prophet secretary. *SIAM Journal on Discrete Mathematics* 31, 3 (2017), 1685–1701.
- [18] Jon Feldman, Nitish Korula, Vahab Mirrokni, S Muthukrishnan, and Martin Pál. 2009. Online ad assignment with free disposal. In *International Workshop on Internet and Network Economics*. Springer, 374–385.
- [19] Jon Feldman, Aranyak Mehta, Vahab Mirrokni, and S Muthukrishnan. 2009. Online stochastic matching: Beating 1-1/e. In *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*. IEEE, 117–126.
- [20] Gagan Goel and Aranyak Mehta. 2008. Online budgeted matching in random input models with applications to adwords. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 982–991.
- [21] Bernhard Haeupler, Vahab S Mirrokni, and Morteza Zadimoghaddam. 2011. Online stochastic weighted matching: Improved approximation algorithms. In *International Workshop on Internet and Network Economics*. Springer, 170–181.
- [22] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. 2007. Automated online mechanism design and prophet inequalities. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence*, Vol. 7. 58–65.
- [23] Zhiyi Huang and Xinkai Shu. 2021. Online Stochastic Matching, Poisson Arrivals, and the Natural Linear Program. *arXiv preprint arXiv:2103.13024* (2021).
- [24] Patrick Jaillet and Xin Lu. 2013. Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations Research* 39, 3 (2013), 624–646.
- [25] Chinmay Karande, Aranyak Mehta, and Pushkar Tripathi. 2011. Online bipartite matching with unknown distributions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*. ACM, 587–596.
- [26] Richard M Karp, Umesh V Vazirani, and Vijay V Vazirani. 1990. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*. ACM, 352–358.
- [27] Mohammad Mahdian and Qiqi Yan. 2011. Online bipartite matching with random arrivals: an approach based on strongly factor-revealing lps. In *Proceedings of the 43rd ACM Symposium on Theory of Computing*. ACM, 597–606.
- [28] Vahideh H Manshadi, Shayan Oveis Gharan, and Amin Saberi. 2012. Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research* 37, 4 (2012), 559–573.
- [29] Aranyak Mehta. 2012. Online matching and ad allocation. *Theoretical Computer Science* 8, 4 (2012), 265–368.
- [30] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. 2007. Adwords and generalized online matching. *Journal of the ACM (JACM)* 54, 5 (2007), 22.
- [31] Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. 2011. Online Learning: Stochastic, Constrained, and Smoothed Adversaries. In *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011. Proceedings of a meeting held 12-14 December 2011, Granada, Spain*, John Shawe-Taylor, Richard S. Zemel, Peter L. Bartlett, Fernando C. N. Pereira, and Kilian Q. Weinberger (Eds.), 1764–1772.
- [32] Irina Shevtsova. 2011. On the absolute constants in the Berry-Esseen type inequalities for identically distributed summands. *arXiv preprint arXiv:1111.6554* (2011).
- [33] Yaron Singer and Manas Mittal. 2013. Pricing mechanisms for crowdsourcing markets. In *Proceedings of the 22nd international conference on World Wide Web*. ACM, 1157–1166.
- [34] Adish Singla and Andreas Krause. 2013. Truthful incentives in crowdsourcing tasks using regret minimization mechanisms. In *Proceedings of the 22nd international conference on World Wide Web*. ACM, 1167–1178.
- [35] Xiaoming Sun, Jia Zhang, and Jialin Zhang. 2016. Near optimal algorithms for online weighted bipartite matching in adversary model. *Journal of Combinatorial Optimization* (2016), 1–17.
- [36] Qiqi Yan. 2011. Mechanism design via correlation gap. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 710–719.