Exploring the Tradeoff between Competitive Ratio and Variance in Online-Matching Markets⋆

Pan Xu

Department of Computer Science, New Jersey Institute of Technology
`pxu@njit.edu`

Abstract. In this paper, we propose an online-matching-based model to study the assignment problems arising in a wide range of online-matching markets, including online recommendations, ride-hailing platforms, and crowdsourcing markets. It features that each assignment can request a random set of resources and yield a random utility, and the two (cost and utility) can be arbitrarily correlated with each other. We present two linear-programming-based parameterized policies to study the tradeoff between the *competitive ratio* (CR) on the total utilities and the *variance* on the total number of matches (unweighted version). The first one (SAMP) is simply to sample an edge according to the distribution extracted from the clairvoyant optimal, while the second (ATT) features a time-adaptive attenuation framework that leads to an improvement over the state-of-the-art competitive-ratio result. We also consider the problem under a large-budget assumption and show that SAMP achieves asymptotically optimal performance in terms of competitive ratio.

1 Introduction

Online-matching models have been used to study a wide range of real-world matching markets, including ride-hailing platforms, online recommendations, and crowdsourcing markets. One of the central problems there is to design an online-matching policy such that the expected total utility is maximized over all matches. Most of the current models assume that each match requests one single deterministic resource and what is more, the profit gained through the match is typically supposed to be independent of the amount of resource requested. These assumptions make current models hard to capture many real applications perfectly. Consider the following three motivating examples.

Bundle recommendations. Consider online bundle recommendations [1]. We have a ground set $\mathcal{U}$ of all offline items to sell. Upon the arrival of an online buyer, say j, we need to select a bundle $S \subseteq \mathcal{U}$ of offline items to offer to j, and then the user j will accept and reject S with certain respective probabilities, say p and $1-p$. Assume the acceptance of bundle S will yield some profit, say $w_{j,S}$, to the platform (*e.g.*, Amazon) and each type of offline item has a limited number of copies in stock. In this case, we observe that after "matching" S with j: With probability p, we will deplete a copy of all items in S and get a profit $w_{j,S}$, and with probability $1-p$, it will incur no cost and no profit.

Display advertising. Consider display advertising business [2, 3]. We have a ground set $\mathcal{U}$ of all offline impressions (or ads). Upon the arrival of an online user of type j, one ads platform (*e.g.*, Google) will display to her a set of ads, say $S \subseteq \mathcal{U}$. Then, the user j will select a subset $S' \subseteq S$ to click, which occurs with some probability $p_{j,S'}$, and this yields profit $w_{j,S'}$ to the ads platform as a result. Assume each ad has a displaying capacity due to the budget of the advertiser. In this context, matching j with S will lead to a random consumption of budgets of ads in S and a random profit, and the two (consumption and profit) are positively correlated with each other.

Task assignment in crowdsourcing markets. Consider task assignment problem in crowdsourcing human-resource markets [4], in which we crowdsource arriving workers to complete as many tasks as possible. We have a ground set $\mathcal{U}$ of offline tasks. Upon the arrival of an online worker of type j, the platform (*e.g.*, Amazon Mechanical Turk) assigns her a set $S \subseteq \mathcal{U}$ of relevant tasks, and then the worker will select a subset

---

$S' \subseteq S$ of tasks to work according to her preferences, which occurs with a certain probability $p_{j,S'}$. Assume every task has a matching capacity reflecting the limited stock. In this context, after matching j and S, with probability $p_{j,S'}$, we will exhaust one copy of each task in S' and gain a utility of $|S'|$ (the number of tasks completed).

In all the three examples, each assignment may incur a random set of resources and yield a stochastic profit, and the cost and profit can be highly correlated with each other. In this paper, we propose a unifying model, called *Multi-Budgeted Online Stochastic Matching* (MB-OSM), to study these applications.

Multi-Budgeted Online Stochastic Matching (MB-OSM). Suppose we have a bipartite graph $G = (I, J, E)$, where I and J denote the sets of offline and online agents, respectively. We have a set of K resources, denoted by $[K] \doteq \{1, 2, \ldots, K\}$, and each resource $k \in [K]$ has a budget $B_k \in \mathbb{Z}^+$. Each edge $e \in E$ is associated with a random utility $W_e \geq 0$ and a random vector-valued cost, $\mathcal{A}_e = (A_{e,k})_{k \in [K]}$, which takes values over $\{0, 1\}^K$. We define the support of $\mathcal{A}_e$ as $\mathcal{S}_e = \{k \in [K] : \mathbb{E}[A_{e,k}] > 0\}$. In our context, $\mathcal{A}_e$ with $e = (i, j)$ captures the random set of resources requested by matching agents i and j, and $\mathcal{S}_e$ represents the set of all possible resources potentially consumed by e. Note that offline vertices in I are static, while online vertices in J arrive dynamically. Specifically, we consider a given time horizon T. For each time (or round) $t \in [T] \doteq \{1, 2, \ldots, T\}$, one single online vertex $\hat{j}$ will be sampled (called $\hat{j}$ arrives) following a known distribution $\{p_j\}$ *with replacement* such that $\Pr[\hat{j} = j] = p_j$ for each $j \in J$ with $\sum_{j \in J} p_j = 1$. Note that the sampling process is independent and identical across the T rounds. For each j, let $r_j = T \cdot p_j$, which is called the *arrival rate* of j with $\sum_{j \in J} r_j = T$. Our arrival assumption is commonly referred to as the *known identical independent distributions* (KIID).[1] Upon the arrival of an online vertex j, an immediate and irrevocable decision is required before observing the next arrival of online vertex: either reject j or match it with an offline neighbor i with $e = (i, j) \in E$. In the latter case, instant cost and utility realizations will be observed, *e.g.*, $\mathcal{A}_e = (a_{e,k})$ and $W_e = w_e$, and the budget of each resource k will get reduced by $a_{e,k}$ and we will gain a utility of w_e as a result.

Remarks on the model of MB-OSM. (1) The following information is all known as part of the input and is accessible to the algorithm: $G = (I, J, E)$, $\{B_k | k \in [K]\}$, distributions of $\{\mathcal{A}_e, W_e | e \in E\}$, $\{p_j, r_j | j \in J\}$ and T. *Additionally, we assume $T \gg 1$, and some of our results are obtained by taking $T \to \infty$*, a common practice in studying competitive ratio for theoretical online-matching models under KIID [13–15]. (2) Cost distributions among all edges ($\{\mathcal{A}_e | e \in E\}$) are independent; the same for all utility distributions of $\{W_e | e \in E\}$. For each given edge e, its cost $\mathcal{A}_e$ and utility W_e can be arbitrarily correlated, however. (3) *Throughout this paper, we assume each edge will incur at most Δ different possible resources,* i.e., $|\mathcal{S}_e| \leq \Delta$ *for every $e \in E$, where Δ is called sparsity*. The sparsity typically takes a small constant value in practice,[2] though the total number K of all resources can be huge. Fortunately, as shown later, the performance of algorithms proposed here will rely only on the sparsity Δ, regardless of K. (4) Safe policies. We say a policy (or algorithm) ALG is *safe* if ALG can make an assignment e only when it will not violate any budget constraint for any possible realizations of $\mathcal{A}_e$ (*i.e., every resource in $\mathcal{S}_e$ should have at least one unit budget remaining then*). Safe policies are required in most real-world applications. For example, we can offer a bundle to a buyer only when all items in the bundle have at least one copy in stock. *For this reason, we consider safe policies only throughout this paper.*

Remarks on sources of randomness in MB-OSM. Consider a given (randomized) policy ALG. Let $\mathcal{M}$ be a set of (random) assignments. The expected total amount of utilities obtained by ALG is defined as $\mathbb{E}[\text{ALG}] = \mathbb{E}[\sum_{e \in \mathcal{M}} W_e]$, where the expectation is taken over the following four sources of randomness: **(R1)** dynamic arrivals of online agents over the T rounds; **(R2)** randomness possibly used by the policy ALG; **(R3)** randomness in cost realizations of $\{\mathcal{A}_e\}$; and **(R4)** randomness in utility realizations of $\{W_e\}$. Our goal is to design a policy that achieves as large utilities as possible while with a variance as small as possible.

---

[1] KIID is mainly inspired from the fact that we can often learn the arrival distribution from historical logs [5–7]. It is widely adopted to study practical online-matching markets [8–10], and it is also a common setting in theoretical online-matching models [11–16].
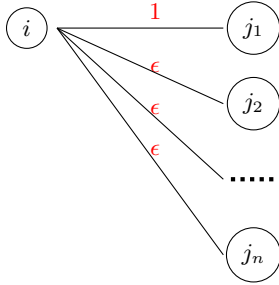
[2] This can be seen from the fact that sparsity captures the tolerance on the bundle size of online buyers, the patience on the number of ads displayed to on online user simultaneously, and working capacity among online workers in the three aforementioned applications, respectively.

1.1 Preliminaries

Throughout this paper, we set $[n] := \{1, 2, \ldots, n\}$ for any integer n.

Competitive ratio. Competitive ratio (CR) is a commonly used metric to evaluate the performance of online algorithms. Consider maximization of the total utilities in MB-OSM as studied here for example. Consider a given algorithm ALG and an instance $\mathcal{I}$ of MB-OSM. Let $\text{ALG}(\mathcal{I})$ and $\text{OPT}(\mathcal{I})$ be the expected total utilities achieved by ALG and a clairvoyant optimal OPT on $\mathcal{I}$, respectively. We say ALG achieves a competitive ratio of $\rho \in [0, 1]$ if $\text{ALG}(\mathcal{I}) \geq \rho \, \text{OPT}(\mathcal{I})$ for all possible instances $\mathcal{I}$ of MB-OSM.

Here are a few similarities and differences between ALG and OPT: (1) The expected performance of ALG and OPT are both taken over the four sources of randomness (*i.e.,* **R1-R4**); (2) ALG is required to make an immediate matching decision upon every arrival of an online vertex before the next one, while OPT enjoys the privilege of accessing the full arrival sequence of online vertices before any decisions; (3) Neither ALG nor OPT has access to realizations of the cost or utility of an edge until the edge has been added; (4) Both ALG and OPT should follow rules of safe policies, *i.e.,* an edge e can be added only when every resource in $\mathcal{S}_e$ has at least one remaining budget. The toy example below shows that the natural two heuristics, Greedy and Ranking, both achieve a competitive ratio of zero on MB-OSM.[3]

$I = \{i\}, J = \{j_1, \ldots, j_n\}, E = \{(i, j_\ell) | \ell \in [n]\};$

$K = 1, B = 1;$

$\mathcal{A}_e = 1$ with probability $1, \forall e \in E;$

$W_{(i,j_1)} = 1,$ with probability $1;$

$W_{(i,j_\ell)} = \epsilon,$ with probability $1, \forall 1 < \ell \leq n;$

$T = n, p_j = 1/n, r_j = 1, \forall j \in [n];$

Greedy = Ranking $\leq \epsilon + 1/n;$

OPT $\geq 1 - 1/\mathsf{e}.$

Fig. 1. A toy example on which Greedy and Ranking both achieve a competitive ratio of zero.

*Example 1 (*Greedy *and* Ranking *both achieve a competitive ratio of zero).* Consider such a toy example as shown in Figure 1. We have a star graph with $I = \{i\}$ and $J = \{j_1, \ldots, j_n\}$, $T = n$, and $r_j = 1$ for every $j \in J$. Thus, during each round $t \in [T]$, one single online agent $\hat{j}$ will be sampled uniformly at random with replacement such that $\Pr[\hat{j} = j_\ell] = 1/n$ for every $\ell \in [n]$. We have one single resource with a unit budget, and each edge will cost one unit resource. The edge (i, j_1) has a deterministic weight of one, while the rest have a deterministic weight of $\epsilon > 0$. Our example captures a simple instance of the classical online matching under KIID, where the offline vertex i can be viewed as the single resource with a unit matching capacity.

We can verify that (1) Greedy and Ranking reduce to the same and both achieve an expected utility of $1/n \cdot 1 + (1 - 1/n) \cdot \epsilon \leq 1/n + \epsilon$; (2) OPT (a clairvoyant optimal) achieves an expected utility of $(1 - 1/\mathsf{e}) \cdot 1 + 1/\mathsf{e} \cdot \epsilon$, where OPT will assign i to j_1 if j_1 arrives at least once (that happens with probability $1 - 1/\mathsf{e}$). By definition, Greedy and Ranking both achieve a competitive ratio no more than $(1/n + \epsilon)/(1 - 1/\mathsf{e}) \to 0$ when $n \to \infty$ and $\epsilon \to 0$. ∎

Variance analysis. In this paper, we pioneer variance analysis for online algorithms in the context of online stochastic matching under known distributions. For most classical optimization problems formulated in an offline setting where input information is fully accessible (*e.g.,* finding a maximum weighted matching in

---

[3] Note that Greedy and Ranking operate as follows: Greedy matches each arriving $j \in J$ with a neighbor i such that $e = (i, j)$ has the largest expectation of utility among all safe assignments (*i.e.,* no budget violation will be caused); Ranking first chooses a random order π over I and then matches each arriving j with a neighbor i such that i has the lowest order in π among all safe choices.

a general graph), we are allowed to run an algorithm multiple times on a given input instance. In this context, suppose we design a randomized algorithm ALG and show that its *expected* performance ($\mathbb{E}[\text{ALG}]$) is good enough. By applying de-randomization techniques like conditional expectations [17], we can get a *deterministic* version that can performance as good as $\mathbb{E}[\text{ALG}]$. Note that for online optimization problems like MB-OSM, we can run an online algorithm *only once* on a given input instance and thus, de-randomization techniques fail to work here. This highlights the importance of variance analysis in online-algorithm design, in addition to the popular competitive-ratio (CR) analysis. Observe that the CR metric reflects only the gap between an online algorithm (ALG) and a clairvoyant optimal (OPT) in terms of their *expected* performance: it shows no any guarantee on the variance or robustness of ALG.

In this paper, we focus on analyzing the variance on the total (random) number of matches instead of the total utilities achieved by any algorithm. Note that the randomness in utility realizations (**R4**) can contribute to an unbounded variance in the total utilities even for simple deterministic algorithms. Consider such a toy example as follows: There is only one single edge $e = (i, j)$ in the graph and one single round $T = 1$ where j will arrive with probability one; We have one single resource with a unit budget; $A_e = 1$ deterministically while $W_e = 1/\epsilon$ with probability ϵ and $W_e = 0$ otherwise. We can verify that the total utilities achieved by any policy adding edge e will have a variance equal to $1/\epsilon - 1$, which can be arbitrarily large. For this reason, we focus on analyzing the variance on the total number of matches, which is due to randomness sources as outlined in **R1**, **R2**, and **R3** only (see **Remarks on sources of randomness** in MB-OSM).

1.2 Main Contributions

Our contributions are summarized as follows. First, we present a canonical LP (Section 2) as the benchmark, whose optimal value proves a valid upper bound on the total expected utilities achieved by a clairvoyant optimal. Second, we design two LP-based parameterized policies to study the tradeoff between the competitive ratio (CR) on the total utilities and the variance on the total number of matches. The first one (SAMP) is simply to sample an edge according to the distribution extracted from a clairvoyant optimal, while the second (ATT) features a time-adaptive attenuation framework. In the last, we study a special case under the large-budget assumption and show that the first algorithm SAMP can achieve an asymptotically optimal CR that approaches one when budgets go infinity. Here are the details.

Theorem 1. *[Section 3] There exists a parameterized LP-based sampling algorithm* $\text{SAMP}(\alpha)$ *with* $\alpha \in [0, 1]$ *such that (i) it achieves a competitive ratio (CR) equal to* $(1 - \mathsf{e}^{-\alpha\Delta})/\Delta$ *on the total utilities with respect to the benchmark LP (1); and (ii) it achieves a variance at most* $(\alpha T)^2 \cdot \mathsf{g}\big(\min(\Delta\alpha, \eta)\big) + O(T)$ *on the total number of matches, where* $\mathsf{g}(x) := \big(1 - \mathsf{e}^{-2x} - 2x\mathsf{e}^{-x}\big)/x^2$, *and* $\eta \sim 1.126$ *is the unique maximizer of* $\mathsf{g}(x)$ *when* $x \in [0, \infty]$. *Both the CR and variance analyses are tight.*[4]

Theorem 2. *[Section 4] There exists a parameterized sampling algorithm* $\text{ATT}(\alpha)$ *with* $\alpha \in [0, 1]$ *such that (i) it achieves a competitive ratio (CR) equal to* $(1 - \mathsf{e}^{-\alpha\Delta})/\Delta$ *on the total utilities with respect to the benchmark LP (1); and (ii) it achieves a variance at most* $(\alpha T)^2 \cdot \mathsf{g}(\alpha\Delta) + O(T)$ *on the total number of matches, where* $\mathsf{g}(x) := \big(1 - \mathsf{e}^{-2x} - 2x\mathsf{e}^{-x}\big)/x^2$. *Both CR and variance analyses are tight.*

Remarks on $\text{SAMP}(\alpha)$ **and** $\text{ATT}(\alpha)$. (i) Both $\text{SAMP}(\alpha)$ and $\text{ATT}(\alpha)$ achieve a competitive ratio (CR) *equal to* $(1 - \mathsf{e}^{-\alpha\Delta})/\Delta$ on the total expected utilities. Note that $(1 - \mathsf{e}^{-\alpha\Delta})/\Delta$ is an increasing function of $\alpha \in [0, 1]$ for any given integer Δ. Meanwhile, we can verify that the upper bounds of variance for $\text{SAMP}(\alpha)$ and $\text{ATT}(\alpha)$ are both increasing functions of $\alpha \in [0, 1]$ for any given integer Δ. This suggests that when α takes a larger value, both algorithms will achieve a higher CR on the total utilities at a price of a higher variance on the total number of matches and vice versa. (ii) Recall that CR is defined based on the worst-case on which the ratio of the performance of an algorithm to that of a clairvoyant optimal gets minimized. Though SAMP and ATT achieve the same worst-case CR, ATT shows more robust than SAMP in the way that $\text{ATT}(\alpha)$ will achieve a CR equal to $(1 - \mathsf{e}^{-\alpha\Delta})/\Delta$ on *every* input instance, while $\text{SAMP}(\alpha)$ will achieve the worst-case CR on a very specialized instance as identified in Example 2. This partially explains why the

---

[4] Tightness on the analysis means we can identify an instance on which the CR (or variance) achieved by SAMP (or ATT) matches the claimed bound.

upper bound of variance of SAMP is slightly larger than that of ATT. (iii) Though ATT is conceptually more complicated than SAMP, both the CR and variance analyses of ATT turn out much simpler than those of SAMP.

We complement the above lower bounds of CR by showing some upper bound due to the benchmark LP (conditional hardness result).

Theorem 3. *[Section 5] No algorithm can achieve a competitive ratio better than $\big(1 - e^{-(\Delta - 1 + 1/\Delta)}\big)/(\Delta - 1 + 1/\Delta)$ with respect to the benchmark LP (1) with $\Delta - 1$ being a prime.*

Remarks on CR for online-matching models related to MB-OSM. (i) When $\Delta = 1$, the upper bound in Theorem 3 matches the best lower bound of $1 - 1/e$ (which arrives at $\alpha = 1$) in Theorems 1 and 2. This suggests the tightness of lower bounds on CR with respect to the current benchmark LP (1).[5] (ii) Dickerson et al. [18] considered a special case of MB-OSM when each edge is associated with a *deterministic* vector-valued cost with sparsity Δ but under known adversarial distributions, which allow the arrival distributions to change over time. They gave an upper bound (hardness result) of $1/(\Delta - 1 + 1/\Delta)$ and a lower bound of $1/(\Delta + 1)$, respectively. Note that both the lower and upper bounds are improved here: the best lower bound of SAMP and ATT (when $\alpha = 1$) satisfies $(1 - e^{-\Delta})/\Delta \geq 1/(\Delta + 1)$; see Figure 2. (iii) Brubach et al. [16] considered a variant of online stochastic matching, which can be cast as a *strict* special case of MB-OSM in the way that each edge is associated with a Bernoulli random vector-valued cost with $\Delta = 1$. From Theorems 1 and 2, both SAMP(α) and ATT(α) achieve a CR of $1 - 1/e$ with $\alpha = 1$ when $\Delta = 1$, which matches that of [16]. (iv) Kesselheim et al. [19] considered online Δ-hypergraph matching that can be cast as a special case of our model where each edge takes a deterministic cost vector with sparsity Δ. They considered the random arrival order and gave a CR of $1/(e \cdot \Delta)$, which is much worse than the best CR as stated in Theorems 1 and 2 that is equal to $(1 - e^{-\Delta})/\Delta$ when $\alpha = 1$.
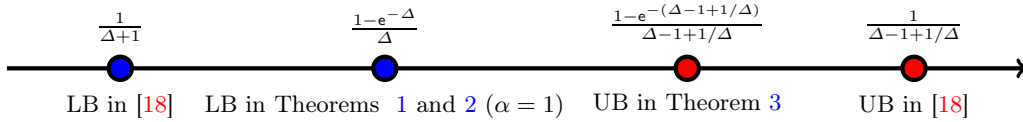
Fig. 2. Upper bounds (UB) and lower bounds (LB) on competitive ratio of MB-OSM as shown in this paper and in [18]: Note that $(1 - e^{-\Delta})/\Delta \geq \big(1 - 1/(\Delta + 1)\big)/\Delta = 1/(\Delta + 1)$, where $\Delta := \max_{e \in E} |\mathcal{S}_e|$.

In the last, we consider MB-OSM under the large-budget assumption. Let $B = \min_{k \in [K]} B_k$, which denotes the minimum budget over all resources. When all budgets are large with $B \gg 1$, we show that SAMP(α) with $\alpha = 1$ can achieve an asymptotically optimal CR with respect to the benchmark LP (1).

Theorem 4. *[Section 6] There exists an LP-based sampling algorithm* SAMP(1) *such that (i) it achieves a competitive ratio (CR) of $1 - \frac{1}{\sqrt{2\pi B}}\big(1 + o(1)\big)$ when $\Delta = 1$ and $T \gg B \gg 1$; (ii) it achieves a competitive ratio of $1 - \kappa \cdot \sqrt{\frac{\ln \Delta}{B}}(1 + o(1))$ with $\sqrt{2} < \kappa \leq 2\sqrt{2}$ when $T \gg B \gg \ln \Delta \gg 1$. The competitive ratios for cases (i) and (ii) are both asymptotically optimal with respect to the benchmark LP (1).*

Remarks on results of Theorem 4. (1) The term $o(1)$ in part (i) vanishes when $B \to \infty$, while that in part (ii) vanishes when $B = \omega(\ln \Delta)$ and $\Delta \to \infty$. (2) As mentioned before, Brubach et al. [16] considered a strictly special case of MB-OSM with $\Delta = 1$. When all budgets are large, they gave an algorithm achieving an asymptotical online ratio of $1 - B^{-1/2+\epsilon}(1 + o(1))$ for any given $\epsilon > 0$. Our result in part (i) significantly improves that. (3) The constants of the leading terms as stated in parts (i) and (ii) are asymptotically optimal with respect to the current benchmark LP (1), which are $\frac{1}{\sqrt{2\pi}}$ and κ, respectively. That is to say, *e.g.,* no policy can achieve a CR of $1 - \frac{c}{\sqrt{2\pi B}}(1 + o(1))$ with a constant $c < 1$ when $B \gg 1$ and $\Delta = 1$ if

---

[5] Tightness here means that the lower bound of CR is the best we can get based on the current benchmark LP.

compared against the optimal value of LP (1) (which proves a valid upper bound on a clairvoyant optimal). (4) The result in part (i) (*i.e.,* $1 - \frac{1}{\sqrt{2\pi B}}(1 + o(1))$) appears in multiple contexts before, including Adwords and correlation gap [20–22]. However, the analysis here is essentially different from there. Let $\text{Ber}(B/T)$ denote a Bernoulli random variable with mean B/T. For Adwords and related applications [20–22], they all care about $\mathbb{E}[\min(X, B)]/B$ when $T \to \infty$, where X is the sum of T *i.i.d.* $\text{Ber}(B/T)$s. In contrast, we need to figure out $\mathbb{E}[\min(T', T)]/T$ when $T \to \infty$, where T' is the number of copies of *i.i.d.* $\text{Ber}(B/T)$s needed such that the total sum is equal to B. Though we can draw a subtle connection between the two (see the proof in Section C.3), it is not straightforward to see the two are the same.

1.3 Main Techniques and Other Related Works

Main techniques. Overall, both the competitive-ratio (CR) and variance analyses in this paper feature a sequential identification process of the worst-scenario (WS) structure on which the exact lower bound (for CR) and the exact upper bound (for variance) are attained. For the algorithm-design part, our second parameterized policy (ATT) features a time-adaptive attenuation framework, which differs from previous time-oblivious attenuations widely used before [23–27]. Generally speaking, in a time-oblivious attenuation framework, we set a uniform attenuation target, say a given constant $\gamma = 1/2$, such that a "good" event will happen with probability exactly *equal to* γ for every online agent regardless of her arriving time. In contrast, we propose a time-adaptive attenuation framework, where we carefully craft an attenuation target function $\gamma(t)$, which is adaptive to the arriving time t of an online agent. This partially leads to an improvement on CR for MB-OSM over the previous work [18] that adopts a time-oblivious attenuation. For the variance-analysis part, we propose several balls-and-bins models to facilitate the analysis and exploit related *negative-association* properties exclusively applied to balls-and-bins models; see [28–30].

Other related works. MB-OSM falls under the family of *online packing* problems. There have been a few studies investigating CR of MB-OSM under the large-budget assumption but under the arrival setting of random arrival order (RAO), which is less restrictive than KIID as considered here; see [31–33]. For RAO, a powerful algorithm-design paradigm is called the primal-and-dual approach; see the survey book [34]. It is interesting to compare the setting studied here with those of [35] and [36], both of which considered online resource allocation with multiple budget constraints under RAO. However, they assumed that both the cost vector and utility are deterministic for each edge. In contrast, we assume each edge can have a random vector-valued cost and a random utility, and the two can be arbitrarily correlated with each other. Another difference is that they considered fractional cost in the way that each $\mathcal{A}_e \in [0, 1]^K$ (deterministic), while we assume that each $\mathcal{A}_e \in \{0, 1\}^K$ (random) here. Devanur et al. [36] gave a CR of $1 - \omega(\sqrt{\ln \Delta/B})$, which was improved to $1 - O(\sqrt{\ln \Delta/B})$ by Kesselheim et al. [35]. Note that the work of [35] has not identified any proper constant included inside the term $O(\sqrt{\ln \Delta/B})$, whereas it is one of the main focuses in this paper. The offline version of MB-OSM captures the stochastic Δ-set packing problem as a special case, which was introduced by Bansal et al. [37]. They gave a 2Δ-approximation algorithm for the stochastic Δ-set packing problem, which was improved to $\Delta + o(\Delta)$ by Brubach et al. [38] later. Baveja et al. [39] considered Δ-uniform stochastic hypergraph matching, which can be viewed as a special case of the stochastic Δ-set packing problem. They gave two approximation algorithms that achieve a ratio of $\Delta + 1/2$ and a ratio of $\Delta + \epsilon$ for any given $\epsilon > 0$, respectively.

There is a large body of research works that have studied budgeted online resource allocation in an *online learning* setting, where distributions of utility and/or cost associated with assignments are unknown. In that context, a common practice is to formulate the problem as one of the renowned Multi-armed bandit variants [40] and then conduct regret analysis, showing the expected total regret (defined as the gap in the total utility achieved by a given policy and a prophet optimal) is upper bounded by a certain function of the total time horizon [41–44]. A few recent works investigate the potential tradeoff between variance and regret in online learning; see, *e.g.,* [45, 46]. In particular, Vakili et al. [46] introduced and analyzed the performance of several risk-averse policies in both bandit and full information settings under the metric of mean-variance [47].

Glossary of notations. We offer a glossary of notations used throughout this paper; see Table 1.

Table 1. A glossary of notations used throughout this paper.

$[n]$	Set of integers $\{1, 2, \ldots, n\}$ for any generic integer n.
$G = (I, J, E)$	Input compatibility graph where I and J are sets of offline and online vertices.
K	Toal number of resources.
B_k	Budget on the resource $k \in [K]$.
$\mathcal{A}_e = (A_{e,k})$	Random vector-valued cost on edge e with $\mathbb{E}[A_{e,k}] = a_{e,k}$.
$\mathcal{S}_e$	Support of the cost of edge e, $i.e.$, $\mathcal{S}_e = \{k : \mathbb{E}[A_{e,k}] > 0\}$.
Δ	Sparsity defined as the largest size of edge cost support, $i.e.$, $\Delta = \max_{e \in E} |\mathcal{S}_e|$.
T	Total number of online rounds (times).
p_j	Probability that online vertex j arrives during each round.
r_j	Expected arrival rate of online vertex j with $r_j = T \cdot p_j$.
W_e	Random (non-negative) utility on edge e with $\mathbb{E}[W_e] = w_e$.
e (Italic)	Edge or assignment $e \in E$.
e (Non-italic)	Natural base taking the value around 2.718.

2 Benchmark LP

For an edge $e \in E$, let x_e be the expected number of times that edge e is added in a clairvoyant optimal. For each vertex j (i), let E_j (E_i) be the set of relevant edges incident to j (i). Let $\mathbb{E}[A_{e,k}] = a_{e,k}$ and $\mathbb{E}[W_e] = w_e$ for each $e \in E$ and $k \in [K]$. Our benchmark LP is formally stated as follows.

$$\max \quad \sum_{e \in E} w_e x_e \tag{1}$$

$$\sum_{e \in E_j} x_e \leq r_j \qquad \forall j \in J \tag{2}$$

$$\sum_{e \in E} a_{e,k} \cdot x_e \leq B_k \quad \forall k \in [K] \tag{3}$$

$$0 \leq x_e \qquad \forall e \in E. \tag{4}$$

Throughout this paper, we refer to the LP above simply as LP (1).

Lemma 1. *The optimal value of* LP (1) *is a valid upper bound on the total expected utilities achieved by a clairvoyant optimal.*

Proof. For each given edge $e = (i, j)$, let X_e be the random number of times that e is added in a clairvoyant optimal (denoted by OPT) with $x_e = \mathbb{E}[X_e]$. We try to justfy that $\{x_e | e \in E\}$ satisfy all constraints in the LP above.

Note that X_e can take values larger than 1 due to the potential multiple arrivals of the online vertex j. However, the total number of edges added with respect to j should be no larger than that of arrivals of j during the online phase, say R_j. Thus, $\sum_{e \in E_j} X_j \leq R_j$ holds with probability one. Taking expectation on both sides, we get $\sum_{e \in E_j} x_j \leq \mathbb{E}[R_j] = r_j$, which leads to Constaint (2). Since we consider safe policies, we are sure that no budget could get vioalted throughout the online process. Thus, for each given resource $k \in [K]$, the event $\sum_{e \in E} A_{e,k} \cdot X_e \leq B_k$ occurs with probability 1. Taking expectation on both sides, we have

$$\mathbb{E}\Big[\sum_{e \in E} A_{e,k} \cdot X_e\Big] = \sum_{e \in E} \mathbb{E}[A_{e,k} \cdot X_e] = \sum_{e \in E} \mathbb{E}[X_e] \cdot \mathbb{E}[A_{e,k}] = \sum_{e \in E} x_e \cdot a_{e,k} \leq B_k, \tag{5}$$

which yields Constaint (3). Note that since we focus on safe policies, the random realization of $A_{e,k}$ and that of W_e both should be independent of when OPT adds the edge e. This is why equalities in (5) hold. Moreover, the total expected utilities gained by OPT is

$$\mathbb{E}\Big[\sum_{e \in E} X_e \cdot W_e\Big] = \sum_{e \in E} \mathbb{E}[X_e] \cdot \mathbb{E}[W_e] = \sum_{e \in E} x_e \cdot w_e,$$

which is consistent with the objective function. Therefore, we conclude that the optimal value to LP (1) should be a valid upper bound for the expected performance of OPT. □

3 An LP-based Sampling Algorithm

Recall that E_j and E_j denote the set of edges incident to j and i, respectively. Our first parameterized algorithm SAMP is formally stated as follows. Note that Step (6) in SAMP(α) is valid since $\sum_{e \in E_j} \alpha x_e^*/r_j \leq$

Algorithm 1: An LP-based sampling algorithm SAMP(α) with $\alpha \in [0, 1]$.

1 **Offline Phase**:
2 Solve LP (1) and let $\{x_e^*|e \in E\}$ be an optimal solution.
3 **Online Phase**:
4 **for** $t = 1, \ldots, T$ **do**
5   Let an online vertex j arrive at time t.
6   Sample an edge $e \in E_j$ with probability $\alpha \cdot x_e^*/r_j$.
7   **if** e *is safe (*i.e.*, there is at least one unit budget of each resource in $\mathcal{S}_e$);* **then**
8   Make the edge e.
9   **else**
10   Reject j.

$\sum_{e \in E_j} x_e^*/r_j \leq 1$ due to Constraint (2) of LP (1).

3.1 Competitive-ratio (CR) analysis for SAMP(α)

In this section, we prove the first part of Theorem 1, which states as follows.

Theorem 5. SAMP(α) *achieves a competitive ratio equal to* $(1 - e^{-\alpha\Delta})/\Delta$ *with respect to* LP (1).

For SAMP(α), we can re-interpret the online phase as an edge-arriving process such that during each round $t \in [T]$, one edge $e = (i, j) \in E$ will arrive with replacement (*i.e.*, j arrives and e gets sampled) with a probability $(r_j/T) \cdot (\alpha x_e^*/r_j) = \alpha x_e^*/T$. Note that $\sum_{e \in E} \alpha x_e^*/T = \sum_{j \in J} \sum_{e \in E_j} \alpha x_e^*/T \leq \sum_j r_j/T = 1$, where the last inequality is due to Constraint (2) in LP (1) and $\alpha \in [0, 1]$. Let $\mathbf{Z} = (Z_k)_{k \in [K]} \in \{0, 1\}^K$ be the random consumption of resources involved in each round of SAMP(α) when all resources are abundant. By definition, we have $\mathbf{Z} = \mathcal{A}_e$ with probability $\alpha x_e^*/T$ for each $e \in E$. Note that $\mathbb{E}[Z_k] = \sum_{e \in E} \mathbb{E}[A_{e,k}] \cdot (\alpha x_e^*/T) = \sum_{e \in E} a_{e,k} \cdot (\alpha x_e^*/T) \leq \alpha B_k/T$, where the last inequality follows from Constraint (3) in LP (1). The proof below for Theorem 5 will need two lemmas, namely, Lemmas 4 and 5, which proofs are deferred to Appendix.

Proof. Consider a given edge e and a given time $t \in [T]$. For each $k \in \mathcal{S}_e$, let $U_{k,t}$ denote the total random cost of resource k at (the beginning of) t in SAMP(α). Let $\mathsf{SF}_{e,t} = \bigwedge_{k \in \mathcal{S}_e} (U_{k,t} \leq B_k - 1)$ denote the event that e is safe at t. WLOG assume $\Delta = |\mathcal{S}_e|$. Recall that $\mathbf{Z} = (Z_k)_k$ denotes the random consumption of resources involved in each single round of SAMP(α) when all resources have remaining budgets. For each $k \in [K]$, let $\{Z_{k,t'}|1 \leq t' < t\}$ be $t - 1$ *i.i.d.* copies of Z_k. Therefore,

$$\Pr[\mathsf{SF}_{e,t}] = \Pr\Big[\bigwedge_{k \in \mathcal{S}_e} (U_{k,t} \leq B_k - 1)\Big] \geq \Pr\Big[\bigwedge_{k \in \mathcal{S}_e} \big(\sum_{1 \leq t' < t} Z_{k,t'} \leq B_k - 1\big)\Big]. \tag{6}$$

Observe that for each $k \in \mathcal{S}_e$, $\mathbb{E}[Z_k] \leq \alpha B_k/T$. By Lemma 4, we see that the right-hand-side value in Inequality (6) gets minimized when $\mathbf{Z}(e) := (Z_k)_{k \in \mathcal{S}_e}$ are negatively correlated according to the following distribution (denoted by $\mathcal{D}^*$): with probality $\alpha B_k/T$, $\mathbf{Z}(e) = \mathbf{1}(Z_k)$ for each $k \in \mathcal{S}_e$, where $\mathbf{1}(Z_e)$ refers to the

standard basis vector with the only entry being 1 at the position of Z_k, and with probability $1 - \sum_{k \in \mathcal{S}_e} \alpha B_k / T$, $\mathbf{Z}(e) = \mathbf{0}$ (a zero vector of length Δ). Thus,

$$\Pr[\mathsf{SF}_{e,t}] \geq 1 - \Pr_{\mathbf{Z}(e) \sim \mathcal{D}^*} \Big[\exists k \in \mathcal{S}_e : \sum_{1 \leq t' < t} Z_{k,t'} \geq B_k \Big]. \tag{7}$$

The value of $\Pr_{\mathbf{Z}(e) \sim \mathcal{D}^*} \Big[\exists k \in \mathcal{S}_e : \sum_{1 \leq t' < t} Z_{k,t'} \geq B_k \Big]$ can be interpreted via the following Balls-and-Bins model: There are $t - 1$ balls and $\Delta = |\mathcal{S}_e|$ bins; each ball will be thrown independently and it will land in bin $k \in \mathcal{S}_e$ with a probability $\alpha B_k / T$ and land in none of them with probability $1 - \sum_{k \in \mathcal{S}_e} \alpha B_k / T$. The value $\Pr_{\mathbf{Z}(e) \sim \mathcal{D}^*} \Big[\exists k \in \mathcal{S}_e : \sum_{1 \leq t' < t} Z_{k,t'} \geq B_k \Big]$ then represents the probability that there exists at least one bin $k \in \mathcal{S}_e$ with at least B_k balls in the end. We claim that $\Pr_{\mathbf{Z}(e) \sim \mathcal{D}^*} \Big[\exists k \in \mathcal{S}_e : \sum_{1 \leq t' < t} Z_{k,t'} \geq B_k \Big]$ gets maximized when each B_k takes a value to maximize $\Pr[\sum_{1 \leq t' < t} Z_{k,t'} \geq B_k]$: note that the change of B_k only affects the probability of each ball falling into the bin k and the threshold B_k associated with bin k (and this has nothing to do with the rest of the bins). By Lemma 5, we see that B_k should take a value of 1 for all $k \in \mathcal{S}_e$. Therefore, in the worst case (when $\Pr[\mathsf{SF}_{e,t}]$ gets minimized), we have (i) $B_k = 1$ for $k \in \mathcal{S}_e$ and (ii) with probability α / T, $\mathbf{Z}(e) = \mathbf{1}(Z_k)$ for $k \in \mathcal{S}_e$, and with probability $1 - \Delta \alpha / T$, all $\mathbf{Z}(e) = \mathbf{0}$. Thus,

$$\Pr[\mathsf{SF}_{e,t}] = \Pr \Big[\bigwedge_{k \in \mathcal{S}_e} \big(U_{k,t} \leq B_k - 1 \big) \Big] \geq \Pr \Big[\bigwedge_{k \in \mathcal{S}_e} \big(\sum_{1 \leq t' < t} Z_{k,t'} \leq B_k - 1 \big) \Big]$$

$$= \Pr \Big[\bigwedge_{k \in \mathcal{S}_e} \big(\sum_{1 \leq t' < t} Z_{k,t'} \leq 0 \big) \Big] = \Pr \Big[\bigwedge_{k \in \mathcal{S}_e, t' < t} Z_{k,t'} = 0 \Big] = \Big(1 - \frac{\Delta \alpha}{T} \Big)^{t-1}.$$

Let M_e be the total (random) utilities gained on edge $e = (i, j)$ in SAMP(α). For each $t \in [T]$, let X_t indicate if j arrives at t and $Y_{e,t}$ indicate if e is sampled at t. Thus, we have

$$\mathbb{E}[M_e] = \sum_{t=1}^{T} \mathbb{E}[X_t \cdot Y_{e,t} \cdot \mathsf{SF}_{e,t} \cdot W_e] = \sum_{t=1}^{T} \frac{r_j}{T} \cdot \frac{\alpha x_e^*}{r_j} \cdot \Big(1 - \frac{\Delta \alpha}{T} \Big)^{t-1} \cdot w_e$$

$$\geq x_e^* \cdot w_e \cdot \sum_{t=1}^{T} \frac{\alpha}{T} \Big(1 - \frac{\Delta \alpha}{T} \Big)^{t-1} \geq x_e^* w_e \frac{1 - \mathsf{e}^{-\Delta \alpha}}{\Delta}.$$

Thus, we have $\mathbb{E}[M_e] \geq x_e^* \cdot w_e \cdot (1 - \mathsf{e}^{-\Delta \alpha}) / \Delta$. By the linearity of expectation, we claim that the total expected utilities of SAMP(α) should satisfy

$$\mathbb{E}[\text{SAMP}(\alpha)] \geq \sum_{e \in E} \mathbb{E}[M_e] \geq \sum_{e \in E} x_e^* \cdot w_e \cdot (1 - \mathsf{e}^{-\Delta \alpha}) / \Delta = \text{LP}(1) \cdot (1 - \mathsf{e}^{-\Delta \alpha}) / \Delta \geq \text{OPT} \cdot (1 - \mathsf{e}^{-\Delta \alpha}) / \Delta,$$

where LP (1) denotes the optimal value of the benchmark LP (1) and OPT the total expected utilities achieved by a clairvoyant optimal, and the last inequality above follows from Lemma 1. Thus, we claim that SAMP(α) achieves a competitive ratio of at least $(1 - \mathsf{e}^{-\Delta \alpha}) / \Delta$. The analysis above actually suggests the tightness of the CR-analysis of SAMP(α) with respect to LP (1). For completeness, we present an explicit CR worst-case structure on Example 2. □

Example 2 (The CR worst-case structure of SAMP(α)*).* Consider such an instance of MB-OSM as follows: $|I| = |J| = |E| = 1$, $K = \Delta$, $B_k = 1$ for all $k \in [K]$. The single edge e has such a cost distribution: $\mathcal{A}_e = \mathbf{1}_k$ with probability $1/T$ for each $k = 1, 2, \ldots, K$ and $\mathcal{A}_e = \mathbf{0}$ with probability $1 - K/T$, where $\mathbf{1}_k$ denotes the kth standard basis vector. Also, the edge e has a deterministic unit utility $W_e = 1$, and $p_j = 1, r_j = T$. We can verify that (i) $x_e^* = T$ in the benchmark LP (1) with an optimal value of T; (ii) SAMP(α) gets an expected total utilities equal to $T \cdot (1 - \mathsf{e}^{-\alpha \Delta}) / \Delta$. Thus, we conclude that SAMP(α) achieves a CR of no more than $(1 - \mathsf{e}^{-\alpha \Delta}) / \Delta$ with respect to LP (1). ∎

3.2 Variance analysis for SAMP(α) with $\Delta = 1$

To better expose our techniques, we start with a simple case of $\Delta = 1$ here and then go to the general case of Δ in Section 3.3.

Note that when $\Delta = 1$, each edge consumes one single resource only.[6] For each resource $k \in [K]$, let $E_k = \{e \in E : \mathcal{S}_e = \{k\}\}$, which denotes the subset of edges whose cost involves the single resource k. For notation convenience, we use a_e to denote $a_{e,k} = \mathbb{E}[A_{e,k}]$ for any $e \in E_k$. The online process of SAMP(α) can be re-interpreted via the following balls-and-bins model.

An auxiliary balls-and-bins model for variance analysis of SAMP(α). We treat each edge as a ball and each edge in E_k is labeled with type $k \in [K]$, and there are K bins and each bin k corresponds to resource k. There are T rounds and during each round $t \in [T]$, we sample a ball $e = (i, j) \in E$ with probability $(r_j/T) \cdot (\alpha x_e^*/r_j) = \alpha x_e^*/T$ and put it into bin k if $e \in E_k$. Thus, during each round, a ball will be added into bin k with probability $\sum_{e \in E_k} \alpha x_e^*/T := q_k$. Note that $\sum_{k \in [K]} q_k = \sum_{e \in E} \alpha x_e^*/T = \sum_{j \in J} \sum_{e \in E_j} \alpha x_e^*/T \leq \sum_{j \in J} \alpha r_j/T = \alpha$. Each bin has a capacity B_k, and each ball $e \in E_k$ is associated with a Bernoulli random variable of mean a_e, denoted by $\mathrm{Ber}(a_e)$. Each time after a ball e is added into bin k, the capacity of bin k gets reduced by one with probability a_e and remains unchanged otherwise. ∎

For each $k \in [K]$, let Y_k be the total (random) number of balls added into bin k by the time when either the capacity B_k is reached or at the end of $t = T$ (whichever comes first). In our context, $\sum_{k \in [K]} Y_k := Y$ captures the exact total number of edges made in SAMP(α), and we aim to upper bound $\mathsf{Var}[Y]$. Note that there are three sources of randomness in Y: (**R1**) the dynamic arrivals of online vetices in J; (**R2**) the random sampling choices of edges made by SAMP(α), and (**R3**) the random cost realization of $\mathcal{A}_e$.

Theorem 6. $\mathsf{Var}[Y] \leq T^2 \Big(1 - \mathrm{e}^{-2\alpha} - 2\alpha \mathrm{e}^{-\alpha} + O(1/T)\Big).$

Proof. Recall that in each round $t \in [T]$, a ball will be added to bin k with probability $q_k = \sum_{e \in E_k} \alpha x_e^*/T$. Let Y_k' be the sum of T *i.i.d.* copies of $\mathrm{Ber}(q_k)$, which denotes the total (random) number of balls added into bin k at the end of time T if ignoring the capacity. Observe that (i) $\{Y_k' | k \in [K]\}$ are negatively associated [28]; (2) Each Y_k can be viewed as a non-decreasing function of Y_k'. Thus, we claim that $\{Y_k\}$ are also negatively associated [29]. Therefore, by the work of [30], we have $\mathsf{Var}[Y] = \mathsf{Var}[\sum_{k \in [K]} Y_k] \leq \sum_{k \in [K]} \mathsf{Var}[Y_k]$.

Focus on a given bin k. Let $B = B_k$ and $\sum_{e \in E_k} x_e^* \cdot a_e = B \cdot \eta$ with $\eta \in (0, 1]$. For each time $t \in [T]$, let $X_t = 1$ indicate that one ball $e \in E_k$ arrives at t and $H_t = 1$ indicates that bin k has at least one capacity at (the beginning of) t. Thus, $Y_k = \sum_{t=1}^{T} X_t \cdot H_t$. Observe that (i) $\{X_t | t \in [T]\}$ are T *i.i.d.* Bernoulli random variables each with mean $q_k = \sum_{e \in E_k} \alpha x_e^*/T$; (ii) Assuming bin k has at least one capacity at the beginning of some round, bin k will have one capacity reduced in that round with probability $\sum_{e \in E_k} (\alpha x_e^*/T) \cdot a_e = \alpha \cdot \eta \cdot B/T$. Therefore, $\mathbb{E}[H_t] = \Pr[\mathrm{Ber}(\alpha\eta B/T)^{(t-1)} \leq B - 1]$, where $\mathrm{Ber}(\alpha\eta B/T)^{(t-1)}$ denotes the sum of $t - 1$ *i.i.d.* Bernoulli random variables each with mean $\alpha\eta B/T$. Observe that $\mathbb{E}[X_t \cdot H_t] = \mathbb{E}[X_t] \cdot \mathbb{E}[H_t]$ since H_t is independent of X_t for each $t \in [T]$.

$$
\begin{aligned}
\mathsf{Var}[Y_k] &= \mathsf{Var}\Big[\sum_{t=1}^{T} X_t \cdot H_t\Big] = \mathbb{E}\Big[\Big(\sum_{t=1}^{T} X_t \cdot H_t\Big)^2\Big] - \Big(\mathbb{E}\Big[\sum_{t=1}^{T} X_t \cdot H_t\Big]\Big)^2 \\
&= \mathbb{E}\Big[\sum_{t=1}^{T} X_t \cdot H_t + 2\sum_{1 \leq t' < t \leq T} X_{t'}' \cdot X_t \cdot H_{t'} \cdot H_t\Big] - \Big(q_k \cdot \sum_{t=1}^{T} \mathbb{E}[H_t]\Big)^2 \\
&= q_k \sum_{t=1}^{T} \mathbb{E}[H_t] + 2q_k^2 \sum_{1 \leq t' < t \leq T} \mathbb{E}[H_t | X_{t'} = 1] - \Big(q_k \cdot \sum_{t=1}^{T} \mathbb{E}[H_t]\Big)^2
\end{aligned}
$$

---

[6] Observe that MB-OSM with $\Delta = 1$ captures the classical online stochastic matching under KIID as a strictly special case [11–16], when each edge $e = (i, j)$ consumes one single resource of the offline vertex i.

$$\leq q_k \sum_{t=1}^{T} \mathbb{E}[H_t] + q_k^2 \cdot \sum_{t=1}^{T} 2(t-1) \cdot \mathbb{E}[H_t] - \left(q_k \cdot \sum_{t=1}^{T} \mathbb{E}[H_t] \right)^2 \tag{8}$$

$$= q_k \sum_{t=1}^{T} \mathbb{E}[H_t] + q_k^2 \left(\sum_{t=1}^{T} 2(t-1) \cdot \mathbb{E}[H_t] - \left(\sum_{t=1}^{T} \mathbb{E}[H_t] \right)^2 \right)$$

Inequality (8) follows from the fact that $\mathbb{E}[H_t] = \mathbb{E}[H_t|X_{t'} = 1]\Pr[X_{t'} = 1] + \mathbb{E}[H_t|X_{t'} = 0]\Pr[X_{t'} = 0] \geq \mathbb{E}[H_t|X_{t'} = 1]$ since $\mathbb{E}[H_t|X_{t'} = 0] \geq \mathbb{E}[H_t|X_{t'} = 1]$. Thus,

$$\mathsf{Var}[Y] \leq \sum_{k=1}^{K} \mathsf{Var}[Y_k] \tag{9}$$

$$\leq \sum_{k=1}^{K} q_k \sum_{t=1}^{T} \mathbb{E}[H_t] + \sum_{k=1}^{K} q_k^2 \left(\sum_{t=1}^{T} 2(t-1) \cdot \mathbb{E}[H_t] - \left(\sum_{t=1}^{T} \mathbb{E}[H_t] \right)^2 \right) \tag{10}$$

$$\leq \alpha T + \alpha^2 \cdot T^2 \cdot \left(1 - \mathsf{e}^{-2\alpha\eta} - 2\alpha\mathsf{e}^{-\alpha\eta} + O(1/T) \right)/(\alpha\eta)^2 \tag{11}$$

$$\leq T^2 \left(1 - \mathsf{e}^{-2\alpha} - 2\alpha\mathsf{e}^{-\alpha} + O(1/T) \right). \tag{12}$$

Inequality (11) follows from Inequality $\sum_{k=1}^{K} q_k^2 \leq \alpha^2$ due to $\sum_{k=1}^{K} q_k \leq \alpha$ and Lemma 6 (the proof is deferred to Appendix); Inequality (12) is due to that the function $\left(1 - \mathsf{e}^{-2\alpha\eta} - 2\alpha\mathsf{e}^{-\alpha\eta} \right)/(\alpha\eta)^2$ is increasing when $\eta \in (0, 1]$ for any given $\alpha \in [0, 1]$. □

The proof above suggests that the upper bound of variance stated in Theorem 6 can be *tight*. Observe that Inequality (9) becomes tight when $K = 1$; Inequalities (10), (11) will be asymptotically tight (after ignoring terms of $O(T)$) when $K = 1$, $q_k = \alpha$, and $B = 1$ (due to Lemma 6); and Inequality (12) gets tight when $\eta = 1$. This reveals the following variance worst-case structure for SAMP(α) with $\Delta = 1$.

Example 3 (A variance worst-case structure of SAMP(α) *with* $\Delta = 1$*).* Consider such an instance of MB-OSM with $\Delta = 1$ as follows: $|I| = |J| = |E| = 1$, $K = 1, B = 1$. In other words, there is one single edge and one single resource with a unit budget. The single edge e has a Bernoulli random cost: $\mathcal{A}_e = \text{Ber}(1/T)$ with mean $1/T$. We have $p_j = 1, r_j = T$. We can verify that (1) $x_e^* = T$ in the benchmark LP (1); (2) SAMP(α) samples edge e with probability α in each round $t \in [T]$; (3) The total (random) number of edges made in SAMP(α) is equal to $\min(\text{Ge}(\alpha/T), T)$, where $\text{Ge}(\alpha/T)$ denotes a Geometric random variable of mean α/T; (4) $\mathsf{Var}[\min(\text{Ge}(\alpha/T), T)] = T^2 \left(1 - \mathsf{e}^{-2\alpha} - 2\alpha\mathsf{e}^{-\alpha} + O(1/T) \right)$. ■

3.3 Variance analysis for SAMP(α) with general Δ

Recall that in the competitive-ratio analysis of SAMP(α), we re-interpret the vertex-arriving process as an edge-arriving process such that in each round $t \in [T]$, one edge is sampled (called e *arrives*) with replacement with probability $(r_j/T) \cdot (\alpha x_e^*/r_j) = \alpha x_e^*/T := q_e$ with $\sum_{e \in E} q_e \leq \alpha$. For each edge $e \in E$ and $t \in [T]$, let $X_{e,t} = 1$ indicate that e arrives at t and $\mathsf{SF}_{e,t} = 1$ indicate that e is safe at t, *i.e.*, all resources in $\mathcal{S}_e$ have at least one unit budget at (the beginning) of t. Let $Y = \sum_{e \in E, t \in [T]} X_{e,t} \cdot \mathsf{SF}_{e,t}$, which represents the total (random) number of edges made in SAMP(α). Thus, our goal is to upper bound $\mathsf{Var}[Y]$. Observe that for each $e = (i, j)$ $\mathbb{E}[X_{e,t}] = q_e$ for every $t \in [T]$.

Theorem 7.
$$\mathsf{Var}[Y] \leq (\alpha T)^2 \cdot \mathsf{g}\big(\min(\Delta\alpha, \eta) \big) + O(T),$$

where $\mathsf{g}(x) := \left(1 - \mathsf{e}^{-2x} - 2x\mathsf{e}^{-x} \right)/x^2$, *and* $\eta \sim 1.126$ *is the unique maximizer of* $\mathsf{g}(x)$ *when* $x \in [0, \infty]$ *with* $\mathsf{g}'(x) = 0$.

Proof. Note that

$$\mathsf{Var}[Y] = \mathsf{Var}\Big[\sum_{e\in E, t\in[T]} X_{e,t}\cdot\mathsf{SF}_{e,t}\Big]$$

$$= \mathbb{E}\Big[\Big(\sum_{e\in E,t\in[T]} X_{e,t}\cdot\mathsf{SF}_{e,t}\Big)^2\Big] - \Big(\mathbb{E}\Big[\sum_{e\in E,t\in[T]} X_{e,t}\cdot\mathsf{SF}_{e,t}\Big]\Big)^2 := F_1 - F_2,$$

where $F_1 = \mathbb{E}\Big[\Big(\sum_{e\in E,t\in[T]} X_{e,t}\cdot\mathsf{SF}_{e,t}\Big)^2\Big]$ and $F_2 = \Big(\mathbb{E}\Big[\sum_{e\in E,t\in[T]} X_{e,t}\cdot\mathsf{SF}_{e,t}\Big]\Big)^2$. Let $p_{e,t} = \mathbb{E}[\mathsf{SF}_{e,t}]$.

$$F_1 = \underbrace{\sum_{e\in E,t\in[T]} \mathbb{E}[(X_{e,t}\cdot\mathsf{SF}_{e,t})^2]}_{F_{1,a}} + \underbrace{2\sum_{e\in E}\sum_{1\le t'<t\le T} \mathbb{E}[X_{e,t'}]\cdot\mathbb{E}[X_{e,t}]\cdot\mathbb{E}[\mathsf{SF}_{e,t}|X_{e,t'}=1]}_{F_{1,b}} +$$

$$+ \underbrace{2\sum_{t\in[T]}\sum_{e\ne e'} \mathbb{E}[X_{e,t}\cdot X_{e',t}\cdot\mathsf{SF}_{e,t}\cdot\mathsf{SF}_{e',t}]}_{F_{1,c}} + \underbrace{2\sum_{e'\ne e}\sum_{t'<t} \mathbb{E}[X_{e',t'}\cdot\mathsf{SF}_{e',t'}\cdot X_{e,t}\cdot\mathsf{SF}_{e,t}]}_{F_{1,d}}$$

$$= F_{1,a} + F_{1,b} + F_{1,c} + F_{1,d}.$$

We upper bound the four parts of F_1 one by one as follows. Let $\sum_{e\in E} q_e = \alpha' \le \alpha$.

$$F_{1,a} = \sum_{e\in E,t\in[T]} \mathbb{E}[X_{e,t}\cdot\mathsf{SF}_{e,t}] = \sum_{e\in E,t\in[T]} q_e\cdot p_{e,t} \le \alpha'\cdot T,$$

$$F_{1,b} \le 2\sum_{e\in E} q_e^2 \sum_{t\in[T]} (t-1)\cdot p_{e,t},$$

$$F_{1,c} = 0,$$

$$F_{1,d} = 2\sum_{e'\ne e} \mathbb{E}[X_{e',t'}\cdot X_{e,t}] \sum_{1\le t'<t} \mathbb{E}[\mathsf{SF}_{e',t'}\cdot\mathsf{SF}_{e,t}|X_{e',t'}=X_{e,t}=1]$$

$$\le 2\sum_{e\in E} q_e \sum_{e'\ne e} q_{e'} \sum_{t\in[T]}\sum_{1\le t'<t} \mathbb{E}[\mathsf{SF}_{e,t}|X_{e',t'}=1] \le 2\sum_{e\in E} q_e(\alpha'-q_e)\sum_{t\in[T]}(t-1)\cdot(p_{e,t}+O(1/T)).$$

The analyses for $F_{1,a}$ and $F_{1,b}$ are similar to the previous case of $\Delta = 1$; For $F_{1,c}$, note that in each round, SAMP(α) samples at most one single edge and thus, $X_{e,t}\cdot X_{e',t} = 0$ for all t; As for $F_{1,d}$, observe that in every single round t' with $t' < t$, any arrival of e' with $e' \ne e$ will have a positive impact on the event $\mathsf{SF}_{e,t}$ no more than the fact that no edge arrives in the round t'; the latter will have a negligible impact on $\mathsf{SF}_{e,t}$ in terms of at most an additive term of $O(1/T)$. Thus,

$$\mathsf{Var}[Y] = F_1 - F_2$$

$$\le \alpha'\cdot T + 2\sum_{e\in E} q_e^2 \sum_{t\in[T]}(t-1)\cdot p_{e,t} + 2\sum_{e\in E} q_e(\alpha'-q_e)\sum_{t\in[T]}(t-1)\cdot p_{e,t} - \Big(\mathbb{E}\Big[\sum_{e,t} X_{e,t}\cdot\mathsf{SF}_{e,t}\Big]\Big)^2$$

$$= \alpha'\cdot T + \alpha'\sum_{e\in E} q_e \sum_{t\in[T]} 2(t-1)\cdot p_{e,t} - \Big(\sum_{e\in E,t\in[T]} q_e\cdot p_{e,t}\Big)^2. \tag{13}$$

The analysis above for $F_{1,d}$ suggests that for any given $\{q_e, p_{e,t}|e\in E, t\in[T]\}$, the variance-WS arrives when all edges share the same support in their random cost distribution, *i.e.*, $\{\mathcal{S}_e|e\in E\}$ are all the same. For this case, we can simply assume there is one single meta edge e^* that arrives with probability $\alpha' = \sum_e q_e$ in each round by compressing all the edges into one single edge and together with their cost distributions[7]. For each given $t\in[T]$, let $p_t = \Pr[\wedge_{k\in\mathcal{S}_{e^*}}(\mathsf{SF}_{k,t} = 1)]$ be the probability that e^* is safe at t. Note that $p_t \ge (1 - \Delta\alpha/T)^{t-1} \sim e^{-\Delta\alpha(t-1)/T}$ following the worst-case competitive-ratio analysis of SAMP(α) (see Example 2). Meanwhile, $p_t \le 1$, which can be asymptotically tight in the case when $|\mathcal{S}_{e^*}| = 1$,

---

[7] The resulting cost distribution of $\mathcal{A}_{e^*}$ can be simply expressed as $\mathcal{A}_{e^*} = \mathcal{A}_e$ with probability q_e/α'.

$\sum_{e \in E} a_{e,k} \cdot x_e^* = B_k \cdot \epsilon$ for the single resource $k \in \mathcal{S}_{e^*}$. In this case, $p_t \geq (1 - \alpha\epsilon/T)^T \sim e^{-\alpha\epsilon} \geq 1 - \alpha\epsilon$ for all $t \in [T]$. Following this argument, we can simplify the expression on (13) as follows

$$\mathsf{Var}[Y] \leq \alpha' \cdot T + \alpha'^2 \sum_{t \in [T]} 2(t-1) \cdot p_t - \Big(\sum_{t \in [T]} \alpha' \cdot p_t \Big)^2$$

$$= \alpha' \cdot T + \alpha'^2 \Big[\sum_{t \in [T]} 2(t-1) \cdot p_t - \Big(\sum_{t \in [T]} p_t \Big)^2 \Big] \leq \alpha \cdot T + \alpha^2 \Big[\sum_{t \in [T]} 2(t-1) \cdot p_t - \Big(\sum_{t \in [T]} p_t \Big)^2 \Big].$$

The analysis on $\Big[\sum_{t \in [T]} 2(t-1) \cdot p_t - \Big(\sum_{t \in [T]} p_t \Big)^2 \Big]$ can be reduced to the previous case of $\Delta = 1$ with an updated value of $\boldsymbol{\alpha} \in [0, \Delta\alpha]$ since $p_t \in [e^{-\Delta\alpha(t-1)/T}, 1]$. By Lemma 6, we have

$$\mathsf{Var}[Y] \leq \alpha \cdot T + \alpha^2 \cdot T^2 \cdot \max_{\boldsymbol{\alpha} \in [0,\Delta\alpha]} \mathsf{g}(\boldsymbol{\alpha}) = (\alpha T)^2 \cdot \mathsf{g}\big(\min(\Delta\alpha, \eta) \big) + O(T).$$

The equality above follows from the fact that $\mathsf{g}(x) := \big(1 - e^{-2x} - 2xe^{-x} \big)/x^2$ first increases when $x \in [0, \eta]$ and then decreases when $x \in [\eta, \infty]$, where $\eta \sim 1.126$ is the unique solution of $\mathsf{g}'(x) = 0$ over $x \in [0, \infty]$. $\qquad \square$

4 A Sampling Algorithm with Time-adaptive Attenuations

Here are the details of simulation-based attenuation in our context. Suppose an online vertex $j \in J$ arrives at t. Consider a given edge $e = (i, j)$. We say e is *safe* at (the beginning of) t iff each resource $k \in \mathcal{S}_e$ has at least one unit budget at t. By applying Monte-Carlo simulations (*i.e.*, simulating the online algorithm up to time t), we can get a sharp estimate of the probability that each e is safe at t (denoted by $\beta_{e,t}$). Suppose we can show that $\beta_{e,t} \geq \gamma_t = (1 - \alpha\Delta/T)^{t-1}$ for all $t \in [T]$, where γ_t is a time-adaptive targeted value we aim to achieve. By generating an auxiliary Bernoulli random variable $Z_{e,t}$ with mean $\gamma_t/\beta_{e,t}$ and adding $Z_{e,t} = 1$ as an extra condition for e to be "safe" at t, we can reduce the probability that each e is "safe" at t to be *equal* to our target γ_t. The formal statement of our algorithm is as follows. Recall that E_j is the set of edges with respect to online vertex $j \in J$.

Algorithm 2: A sampling algorithm $\mathrm{ATT}(\alpha)$ with time-adaptive attenuations with $\alpha \in [0, 1]$.

1 Offline Phase:
2 Solve LP (1) and let $\{x_e^*\}$ be an optimal solution.
   /* In Offline Phase, we aim to compute a sharp estimate of the probability that each edge e is safe at (the beginning of) t, denoted by $\beta_{e,t}$. */
3 *Initialization*: When $t = 1$, set $\beta_{e,t} = 1$ for all $e \in E$.
4 for $t = 2, 3, \ldots, T$ **do**
5 | By simulating Step (8) to Step (14) of Online Phase for all the rounds $t' = 1, 2, \ldots, t-1$, we get a sharp estimate of $\beta_{e,t}$, the probability that e is safe at t.

6 Online Phase:
7 for $t = 1, \ldots, T$ **do**
8 | Let an online vertex j arrive at time t.
9 | Sample an edge $e \in E_j$ with probability $\alpha x_e^*/r_j$.
10 | Generate an auxiliary Bernoulli random variable $Z_{e,t}$ with mean $\gamma_t/\beta_{e,t}$, where $\gamma_t = (1 - \alpha\Delta/T)^{t-1}$.
11 | **if** *e is safe and* $Z_{e,t} = 1$ **then**
12 | | Match the edge e.
13 | **else**
14 | | Reject j.

Throughout this section, we use $\mathsf{SF}_{e,t} = 1$ to indicate that e is "safe" at t, which means that both e is safe at t and $Z_{e,t} = 1$ in Step (11). By the nature of ATT, we see $\mathbb{E}[\mathsf{SF}_{e,t}] = \gamma_t$ for all $e \in E$ and $t \in [T]$. Note

that the condition of e is "safe" at t is stricter than that of e is safe at t; the latter simply means that every resource $k \in \mathcal{S}_e$ has at least one unit budget at t.

Lemma 2. $\beta_{e,t} \geq \gamma_t = (1 - \alpha\Delta/T)^{t-1}$ for all $e \in E, t \in [T]$.

Proof. We prove the lemma by induction on $t \in [T]$. Consider the base case when $t = 1$. We have $\beta_{e,t} = 1 = \gamma_t$ for all $e \in E$. Now consider a given $\bar{t} > 1$ and assume $\beta_{e,t} \geq \gamma_t$ and $\mathbb{E}[\mathsf{SF}_{e,t}] = \gamma_t$ for all $t < \bar{t}$ and all $e \in E$. We show that $\beta_{e,\bar{t}} \geq \gamma_{\bar{t}}$ for all $e \in E$. Consider a given $\bar{e}$ and a given $k \in \mathcal{S}_{\bar{e}}$. We can re-interpret the online vertex-arriving process as edge-arriving alternatively such that each edge $e = (i,j)$ arrives (*i.e.,* j arrives and e gets sampled) with probability $(r_j/T) \cdot (\alpha x_e^*/T) = \alpha x_e^*/T := q_e$ with $\sum_{e \in E} q_e \leq \alpha$. Let $U_{k,\bar{t}}$ be the (random) number of consumptions of resource k during the previous $\bar{t} - 1$ rounds. For each $t < \bar{t}$, $e \in E$, let $X_{e,t} = 1$ indicate that e arrives at t. Recall that $\mathsf{SF}_{e,t}$ be the event that both e is safe at t and $Z_{e,t} = 1$ in Step (11). Thus, we have

$$\mathbb{E}[U_{k,\bar{t}}] = \mathbb{E}\Big[\sum_{t < \bar{t}} \sum_{e \in E} X_{e,t} \cdot \mathsf{SF}_{e,t} \cdot A_{e,k}\Big] = \sum_{t < \bar{t}} \sum_{e \in E} (\alpha x_e^*/T) \cdot (1 - \alpha\Delta/T)^{t-1} \cdot a_{e,k}$$

$$= \sum_{e \in E} (x_e^* \cdot a_{e,k}) \sum_{t < \bar{t}} (1 - \alpha\Delta/T)^{t-1} \cdot (\alpha/T) \leq B_k \cdot (1/\Delta) \cdot \left(1 - \left(1 - \alpha\Delta/T\right)^{\bar{t}-1}\right). \quad (14)$$

The inequality (14) is due to Constraint (3) of LP (1). Therefore, the probability that $\bar{e}$ is safe at $\bar{t}$ before attenuations should be

$$\beta_{\bar{e},\bar{t}} = \Pr\Big[\bigwedge_{k \in \mathcal{S}_{\bar{e}}} \left(U_{k,\bar{t}} \leq B_k - 1\right)\Big] \geq 1 - \sum_{k \in \mathcal{S}_{\bar{e}}} \Pr[U_{k,\bar{t}} \geq B_k] \quad (15)$$

$$\geq 1 - \sum_{k \in \mathcal{S}_{\bar{e}}} \frac{E[U_{k,\bar{t}}]}{B_k} \geq \left(1 - \frac{\alpha\Delta}{T}\right)^{\bar{t}-1} = \gamma_{\bar{t}}. \quad (16)$$

Inequality (15) is due to union bound; the first Inequality on (16) follows from Markov's inequality while the second is due to Inequality (14) and the fact $|\mathcal{S}_{\bar{e}}| \leq \Delta$. Therefore, we complete the proof on $\bar{t}$ and finish the induction. $\square$

Proof of the competitive-ratio result of $\mathrm{ATT}(\alpha)$ **in Theorem 2.**

Proof. Consider a given edge $e = (i,j)$. Let M_e be the total (random) utilities gained on e in $\mathrm{ATT}(\alpha)$. For each $t \in [T]$, let $X_t = 1$ indicate that e arrives at t. Thus, we have

$$\mathbb{E}[M_e] = \sum_{t=1}^{T} \mathbb{E}[X_t \cdot \mathsf{SF}_{e,t} \cdot W_e] = \sum_{t=1}^{T} (\alpha x_e^*/T) \cdot \left(1 - \frac{\alpha\Delta}{T}\right)^{t-1} \cdot w_e = (x_e^* w_e) \cdot \frac{1 - e^{-\alpha\Delta}}{\Delta}.$$

The last equality above is obtained by taking $T \to \infty$. Thus, by the linearity of expectation, we claim that $\mathrm{ATT}(\alpha)$ achieves an expected sum of utilities equal to $\mathbb{E}[\mathrm{ATT}(\alpha)] = \sum_{e \in E} (x_e^* w_e) \cdot (1 - e^{-\alpha\Delta})/\Delta$, which is a factor of $(1 - e^{-\alpha\Delta})/\Delta$ of the optimal value of the benchmark LP (1). Thus, we establish the competitive ratio of $\mathrm{ATT}(\alpha)$. $\square$

4.1 Variance analysis for ATT(α)

We re-interpret the online vertex-arriving process as edge-arriving alternatively such that each edge $e = (i,j)$ arrives (*i.e.,* j arrives and e gets sampled) with probability $(r_j/T) \cdot (\alpha x_e^*/T) = \alpha x_e^*/T := q_e$ with $\sum_{e \in E} q_e \leq \alpha$. For each edge $e \in E$ and $t \in [T]$, let $X_{e,t} = 1$ indicate that e arrives at t. Recall that $\mathsf{SF}_{e,t} = 1$ indicate that both e is safe at t (all resources in $\mathcal{S}_e$ have at least one unit budget at t) and $Z_{e,t} = 1$ in Step (11). Let $Y = \sum_{e \in E, t \in [T]} X_{e,t} \cdot \mathsf{SF}_{e,t}$, which represents the total (random) number of edges made in $\mathrm{ATT}(\alpha)$. Thus, our goal is to upper bound $\mathsf{Var}[Y]$. Observe that for each $e = (i,j)$ and $t \in [T]$, $\mathbb{E}[X_{e,t}] = q_e$ and $\mathbb{E}[\mathsf{SF}_{e,t}] = (1 - \alpha\Delta/T)^{t-1}$. Let $\alpha' = \sum_{e \in E} q_e \leq \alpha$.

Theorem 8.
$$\mathsf{Var}[Y] \le (\alpha T)^2 \cdot \mathsf{g}(\alpha\Delta) + O(T),$$

where $\mathsf{g}(x) := \left(1 - \mathsf{e}^{-2x} - 2x\mathsf{e}^{-x}\right)/x^2$.

Proof. Note that

$$\mu := \mathbb{E}[Y] = \sum_{e\in E, t\in[T]} \mathbb{E}[X_{e,t} \cdot \mathsf{SF}_{e,t}] = \sum_{e\in E, t\in[T]} q_e \cdot \left(1 - \frac{\alpha\Delta}{T}\right)^{t-1} = (\alpha'T) \cdot \frac{1 - \mathsf{e}^{-\alpha\Delta}}{\alpha\Delta}.$$

Thus,

$$\mathsf{Var}[Y] = \mathsf{Var}\Big[\sum_{e\in E, t\in[T]} X_{e,t} \cdot \mathsf{SF}_{e,t}\Big] = \mathbb{E}\Big[\Big(\sum_{e\in E, t\in[T]} X_{e,t} \cdot \mathsf{SF}_{e,t}\Big)^2\Big] - \Big(\mathbb{E}\Big[\sum_{e\in E, t\in[T]} X_{e,t} \cdot \mathsf{SF}_{e,t}\Big]\Big)^2$$

$$:= F_{1,a} + F_{1,b} + F_{1,c} + F_{1,d} - \mu^2,$$

where

$$F_{1,a} = \sum_{e\in E, t\in[T]} \mathbb{E}[(X_{e,t} \cdot \mathsf{SF}_{e,t})^2] = \mu,$$

$$F_{1,b} = 2\sum_{e\in E} \sum_{1\le t'<t\le T} \mathbb{E}[X_{e,t'}] \cdot \mathbb{E}[X_{e,t}] \cdot \mathbb{E}[\mathsf{SF}_{e,t}|X_{e,t'}=1] \le 2\sum_{e\in E} q_e^2 \sum_{t\in[T]} (t-1) \cdot \left(1 - \frac{\alpha\Delta}{T}\right)^{t-1},$$

$$F_{1,c} = 2\sum_{t\in[T]} \sum_{e\neq e'} \mathbb{E}[X_{e,t} \cdot X_{e',t} \cdot \mathsf{SF}_{e,t} \cdot \mathsf{SF}_{e',t}] = 0,$$

$$F_{1,d} = 2\sum_{e'\neq e} \sum_{t'<t} \mathbb{E}[X_{e',t'} \cdot \mathsf{SF}_{e',t'} \cdot X_{e,t} \cdot \mathsf{SF}_{e,t}]$$

$$\le 2\sum_{e\in E} q_e(\alpha' - q_e) \sum_{t\in[T]} (t-1) \cdot \left(\left(1 - \frac{\alpha\Delta}{T}\right)^{t-1} + O(1/T)\right).$$

Summarizing all analyses above, we have

$$\mathsf{Var}[Y] \le \mu - \mu^2 + 2\alpha'^2 \sum_{t\in[T]} (t-1) \cdot \left(\left(1 - \frac{\alpha\Delta}{T}\right)^{t-1} + O(1/T)\right)$$

$$\le \mu - \mu^2 + 2\alpha^2 \sum_{t\in[T]} (t-1) \cdot \left(\left(1 - \frac{\alpha\Delta}{T}\right)^{t-1} + O(1/T)\right) = (\alpha T)^2 \cdot \mathsf{g}(\alpha\Delta) + O(T),$$

where $\mathsf{g}(x) := \left(1 - \mathsf{e}^{-2x} - 2x\mathsf{e}^{-x}\right)/x^2$. $\qquad\blacksquare$

5   Proof of the hardness result in Theorem 3

Example 4. Consider the projective plane $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ of order $\Delta - 1$ with $\Delta - 1$ being a prime [48]. $\mathcal{H}$ is a hypergraph such that (1) $\mathcal{H}$ is Δ-uniform, Δ-regular and intersecting; (2) $|\mathcal{V}| = |\mathcal{E}| = \Delta^2 - \Delta + 1$; and (3) the natural canonical LP on $\mathcal{H}$ has an optimal value of $\Delta - 1 + 1/\Delta$. Now based on $\mathcal{H}$, we construct an instance of MB-OSM as follows.

Let $G = (I, J, E)$ be a star graph with $|I| = 1$, $|J| = |E| = T$, and $r_j = 1$ for all $j \in J$. Set $K = \Delta^2 - \Delta + 1$, where each offline resource corresponds to one hypervertex. Set $B_k = 1$ for all $k \in [K]$. For each hyperedge $f \in \mathcal{E}$, we create $T/(\Delta^2 - \Delta + 1)$ copies of an edge $e \in E$ such that for each edge e, (1) with probability $p = (\Delta - 1 + 1/\Delta)/T$, $W_e = 1$ and $\mathcal{A}_e = \mathbf{a}_f$, where $\mathbf{a}_f \in \{0,1\}^{\Delta^2 - \Delta + 1}$ is the canonical representation of f; (2) with probability $1 - p$, $W_e = 0$ and $\mathcal{A}_e = \mathbf{0}$, where $\mathbf{0}$ is the zero vector of dimension $\Delta^2 - \Delta + 1$. Thus, we have in total $|E| = (T/(\Delta^2 - \Delta + 1)) \cdot (\Delta^2 - \Delta + 1) = T$ edges. Also, we see that every edge $e \in E$ has a support of size Δ, *i.e.,* the number of non-zero entries in $\mathcal{A}_e$. $\qquad\blacksquare$

Lemma 3. LP-(1) *has an optimal value at least* $\Delta - 1 + 1/\Delta$ *on Example* 4.

Proof. For each edge e, let y_e be the sum of x_e over its $T/(\Delta^2 - \Delta + 1)$ copies. Consider such a solution that $x_e = 1$, $y_e = T/(\Delta^2 - \Delta + 1)$ for all e. We show $\{x_e\}$ is feasible to LP-(1). First, $\{x_e\}$ is feasible to Constraint (2). Second, $\{x_e\}$ is feasible to Constraint (3), since for each k,

$$\sum_{e \in E} x_e \cdot \mathbb{E}[A_{e,k}] = |\{e \in E : \mathbb{E}[A_{e,k}] > 0\}| \cdot p = \Delta \cdot \frac{T}{\Delta^2 - \Delta + 1} \cdot \frac{\Delta - 1 + 1/\Delta}{T} = 1.$$

Note that for each resource (or hypervertex) k, there are exactly Δ different edges e such that each e has $T/(\Delta^2 - \Delta + 1)$ copies and each has a non-zero $\mathbb{E}[A_{e,k}]$ with $\mathbb{E}[A_{e,k}] = p$. Therefore, we claim that LP-(1) has an optimal value at least $\sum_{e \in E} x_e \cdot p = T \cdot p = \Delta - 1 + 1/\Delta$. □

Now we start to prove Theorem 3.

Proof. Consider Example 4. Note that all the T edges are intersecting with each other. We see that the expected total utilities achieved by any online algorithm should be no larger than $1 - (1 - p)^T = 1 - e^{-(\Delta - 1 + 1/\Delta)}$. Thus, we claim that the resulting competitive ratio with respect to LP (1) should be no larger than $\left(1 - e^{-(\Delta - 1 + 1/\Delta)}\right)/(\Delta - 1 + 1/\Delta)$. □

6 Proof of Theorem 4

In this section, we show SAMP(α) with $\alpha = 1$ achieves a CR approaching one under the large-budget assumption. WLOG assume that $B_k = B$ for all $k \in [K]$. *Throughout this section, we refer to* SAMP(1) *as* SAMP *for simplicity.*

In the heart of the CR analysis of SAMP, we need the following key result.

Theorem 9. *Assume* $B = \omega(\ln \Delta)$ *and* $\Delta, B \ll T$. *Let* $\mathbf{U} = (U_k | k \in [\Delta]) \sim \mathcal{D}$, *where* $\mathcal{D}$ *is a random distribution such that (1)* $\mathbf{U} \in \{0,1\}^\Delta$; *and (2)* $\mathbb{E}[U_k] \le B/T$ *for every* $k \in [\Delta] := \{1, 2, \ldots, \Delta\}$. *For each* $t \in [T]$, *let* $\mathbf{U}^{(t)}$ *be the sum of* $t - 1$ *i.i.d. copies of* $\mathbf{U}$. *Suppose* $\mathcal{D}^*$ *is an optimal distribution satisfying the two conditions such that* $\sum_{t \in [T]} \Pr_{\mathbf{U} \sim \mathcal{D}^*} \left[\max_{k \in [\Delta]} U_k^{(t)} \le B - 1 \right]$ *gets minimized, where* $U_k^{(t)}$ *is the kth entry of* $\mathbf{U}^{(t)}$. *We have that under* $\mathcal{D}^*$,

$$\lim_{T \to \infty} \frac{1}{T} \sum_{t=1}^T \Pr_{\mathbf{U} \sim \mathcal{D}^*} \left[\max_{k \in [\Delta]} U_k^{(t)} \le B - 1 \right] = 1 - \kappa \sqrt{\frac{\ln \Delta}{B}} (1 + o(1)), \tag{17}$$

where $\sqrt{2} < \kappa \le 2\sqrt{2}$, *and* $o(1)$ *is a vanishing term when* $\Delta \to \infty$.

Now we show how the theorem above partially implies our main Theorem 4.

Proof. Consider a given edge $e \in E$. WLOG assume $\mathcal{S}_e = [\Delta] = \{1, 2, \ldots, \Delta\}$. Let $\mathbf{U} = (U_k | k \in [\Delta])$ be the random cost of resources in $[\Delta]$ during each round when all budgets remain. Following the same analysis in the proof of Theorem 5, we have $\mathbb{E}[U_k] \le B/T$ for each $k \in [\Delta]$. For each given $t \in [T]$, let $\mathcal{U}^{(t)} \in \mathbf{Z}^\Delta$ be the total cost of resources in $[\Delta]$ during the first $(t - 1)$ rounds when running SAMP. Observe that (1) e is safe at t iff the cost of each resource is no larger than $B - 1$, *i.e.,* $\mathcal{U}^{(t)} \le (B - 1) \cdot \mathbf{1}$ (entry-wisely and $\mathbf{1}$ is the vector of ones); (2) $\mathcal{U}^{(t)} \le \mathbf{U}^{(t)}$ (entry-wisely), where $\mathbf{U}^{(t)}$ is the sum of $t - 1$ *i.i.d.* copies of $\mathbf{U}$. The first is valid since we consider safe policies only, while the second follows that some edge might be unsafe during some previous time $t' < t$. Let M_e be the total (random) utilities gained on e in SAMP. For each $t \in [T]$, let

16

$X_e = 1$ indicate that e arrives at t (*i.e.,* the online vertex arrives and e gets sampled in SAMP) and $\mathsf{SF}_t = 1$ indicate that e is safe at t. Thus, we have

$$\mathbb{E}[M_e] = \sum_{t=1}^{T} \mathbb{E}[X_e \cdot \mathsf{SF}_t \cdot W_e] = x_e^* \cdot w_e \cdot \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}[\mathsf{SF}_t] = x_e^* \cdot w_e \cdot \frac{1}{T} \sum_{t=1}^{T} \Pr[\mathcal{U}^{(t)} \le (B-1) \cdot \mathbf{1}]$$

$$\ge x_e^* \cdot w_e \cdot \frac{1}{T} \sum_{t=1}^{T} \Pr[\mathbf{U}^{(t)} \le (B-1) \cdot \mathbf{1}] \ge x_e^* \cdot w_e \cdot \left(1 - \kappa \sqrt{\frac{\ln \Delta}{B}}(1 + o(1))\right).$$

The last inequality follows from Theorem 9. By linearity of expectation, we claim that SAMP achieves a CR at least $1 - \kappa \sqrt{\frac{\ln \Delta}{B}}(1 + o(1))$ with $\kappa \in (\sqrt{2}, 2\sqrt{2}]$. The asymptotically optimality of SAMP can be seen from Theorems 10 and 11 for cases of general Δ and $\Delta = 1$, respectively. Thus, we are done. □

The harness results of Theorem 4 follows from the two theorems below. Due to the space limit, we defer the full proofs of Theorems 9, 10 and 11 to Appendix; see Section C.

Theorem 10. *No algorithm can achieve a CR asymptotically better than* $1 - \kappa \sqrt{\frac{\ln \Delta}{B}}(1 + o(1))$ *when* $T \gg B = \omega(\ln \Delta)$ *and* $\Delta \gg 1$*, where* κ *is the same value as stated in Theorem 9.*

Theorem 11. *No algorithm can achieve a CR asymptotically better than* $1 - \frac{1}{\sqrt{2\pi B}}(1 + o(1))$ *when* $T \gg B \gg 1$ *and* $\Delta = 1$*.*

7 Conclusions and Future work

In this paper, we proposed a model of multi-budgeted online stochastic matching to study assignment problems existing in a wide range of online-matching markets, including online recommendations, rideshares, and crowdsourcing markets. The model features correlated stochastic cost and utility for each assignment. We presented two LP-based parameterized algorithms and analyzed their performance in detail under competitive ratio and variance.

Our work opens a few research directions. The most urgent one is to close the competitive-ratio gap shown in Figure 2. The second is to identify the exact value of κ stated in Theorem 4 under the large-budget assumption. Right now, we can only get an upper and lower bound for it with a gap of $\sqrt{2}$. Additionally, it will be interesting to see if the results in this paper can be generalized to the case of fractional cost values. In other words, can we get similar results if we assume each assignment takes a vector-valued cost from $[0,1]^K$ instead of $\{0,1\}^K$? Lastly, we would like to see if the current variance-variance techniques can be generalized to cope with a more ambitious goal, *i.e.,* upper bounding the variance on the total utilities achieved. By previous analyses, we perhaps need to assume utilities on all assignments are upper bounded by some parameter, which is expected to play a critical role in the variance analysis.

Bibliography

[1] Tao Zhu, Patrick Harrington, Junjun Li, and Lei Tang. Bundle recommendation in ecommerce. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*, pages 657–666, 2014.

[2] Han Zhu, Junqi Jin, Chang Tan, Fei Pan, Yifan Zeng, Han Li, and Kun Gai. Optimized cost per click in taobao display advertising. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 2191–2200, 2017.

[3] Zeinab Abbassi, Aditya Bhaskara, and Vishal Misra. Optimizing display advertising in online social networks. In *Proceedings of the 24th International Conference on World Wide Web*, pages 1–11, 2015.

[4] Chien-Ju Ho and Jennifer Wortman Vaughan. Online task assignment in crowdsourcing markets. In *Twenty-sixth AAAI conference on artificial intelligence*, 2012.

[5] Huaxiu Yao, Fei Wu, Jintao Ke, Xianfeng Tang, Yitian Jia, Siyu Lu, Pinghua Gong, Jieping Ye, and Zhenhui Li. Deep multi-view spatial-temporal network for taxi demand prediction. AAAI '18, pages 2588–2595, 2018.

[6] Yaguang Li, Kun Fu, Zheng Wang, Cyrus Shahabi, Jieping Ye, and Yan Liu. Multi-task representation learning for travel time estimation. KDD '18, pages 1695–1704, 2018.

[7] Zheng Wang, Kun Fu, and Jieping Ye. Learning to estimate the travel time. KDD '18, pages 858–866, 2018.

[8] Boming Zhao, Pan Xu, Yexuan Shi, Yongxin Tong, Zimu Zhou, and Yuxiang Zeng. Preference-aware task assignment in on-demand taxi dispatching: An online stable matching approach. AAAI '19, 2019.

[9] John P Dickerson, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. Assigning tasks to workers based on historical data: Online task assignment with two-sided arrivals. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 318–326. International Foundation for Autonomous Agents and Multiagent Systems, 2018.

[10] Elaheh Fata, Will Ma, and David Simchi-Levi. Multi-stage and multi-customer assortment optimization with inventory constraints. *Available at SSRN 3443109*, 2019.

[11] Jon Feldman, Aranyak Mehta, Vahab S. Mirrokni, and S. Muthukrishnan. Online stochastic matching: Beating 1-1/e. In *50th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '09, pages 117–126, 2009.

[12] Zhiyi Huang and Xinkai Shu. Online stochastic matching, poisson arrivals, and the natural linear program. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 682–693, 2021.

[13] Bernhard Haeupler, Vahab S. Mirrokni, and Morteza Zadimoghaddam. Online stochastic weighted matching: Improved approximation algorithms. In *Internet and Network Economics - 7th International Workshop*, WINE '11, pages 170–181, 2011.

[14] Vahideh H Manshadi, Shayan Oveis Gharan, and Amin Saberi. Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research*, 37(4), 2012.

[15] Patrick Jaillet and Xin Lu. Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations Research*, 39(3), 2013.

[16] Brian Brubach, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. New algorithms, better bounds, and a novel model for online stochastic matching. In *24th Annual European Symposium on Algorithms (ESA 2016)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.

[17] Prabhakar Raghavan. Probabilistic construction of deterministic algorithms: approximating packing integer programs. *Journal of Computer and System Sciences*, 37(2):130–143, 1988.

[18] John P Dickerson, Karthik Abinav Sankararaman, Kanthi Kiran Sarpatwar, Aravind Srinivasan, Kun-Lung Wu, and Pan Xu. Online resource allocation with matching constraints. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1681–1689. International Foundation for Autonomous Agents and Multiagent Systems, 2019.

[19] Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *European Symposium on Algorithms*, pages 589–600. Springer, 2013.

[20] Nikhil R Devanur, Balasubramanian Sivan, and Yossi Azar. Asymptotically optimal algorithm for stochastic adwords. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 388–404. ACM, 2012.

[21] Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. Online prophet-inequality matching with applications to ad allocation. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 18–35, 2012.

[22] Qiqi Yan. Mechanism design via correlation gap. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 710–719. Society for Industrial and Applied Mathematics, 2011.

[23] Will Ma. Improvements and generalizations of stochastic knapsack and multi-armed bandit approximation algorithms. In *SODA*, 2014.

[24] Marek Adamczyk, Fabrizio Grandoni, and Joydeep Mukherjee. Improved approximation algorithms for stochastic matching. In *ESA*. 2015.

[25] Brian Brubach, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. Attenuate locally, win globally: An attenuation-based framework for online stochastic matching with timeouts. *Algorithmica*, 82(1):64–87, 2020.

[26] Yiding Feng, Rad Niazadeh, and Amin Saberi. Linear programming based online policies for real-time assortment of reusable resources. *Available at SSRN 3421227*, 2019.

[27] John P. Dickerson, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. Allocation problems in ride-sharing platforms: Online matching with offline reusable resources. AAAI '18, pages 1007–1014, 2018.

[28] Devdatt P Dubhashi and Desh Ranjan. Balls and bins: A study in negative dependence. *BRICS Report Series*, 3(25), 1996.

[29] Kumar Joag-Dev and Frank Proschan. Negative association of random variables with applications. *The Annals of Statistics*, pages 286–295, 1983.

[30] Qi-Man Shao. A comparison theorem on moment inequalities between negatively associated and independent random variables. *Journal of Theoretical Probability*, 13(2):343–356, 2000.

[31] Shipra Agrawal, Zizhuo Wang, and Yinyu Ye. A dynamic near-optimal algorithm for online linear programming. *Operations Research*, 62(4):876–890, 2014.

[32] Shipra Agrawal and Nikhil R Devanur. Fast algorithms for online stochastic convex programming. In *SODA*, 2014.

[33] Niv Buchbinder and Joseph Naor. Online primal-dual algorithms for covering and packing. *Mathematics of Operations Research*, 34(2):270–286, 2009.

[34] Niv Buchbinder, Joseph Seffi Naor, et al. The design of competitive online algorithms via a primal–dual approach. *Foundations and Trends® in Theoretical Computer Science*, 3(2–3):93–263, 2009.

[35] Thomas Kesselheim, Andreas Tönnis, Klaus Radke, and Berthold Vöcking. Primal beats dual on online packing lps in the random-order model. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '14, page 303–312, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450327107. https://doi.org/10.1145/2591796.2591810. URL https://doi.org/10.1145/2591796.2591810.

[36] Nikhil R. Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A. Wilkens. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. *J. ACM*, 66(1), January 2019. ISSN 0004-5411. https://doi.org/10.1145/3284177. URL https://doi.org/10.1145/3284177.

[37] Nikhil Bansal, Anupam Gupta, Jian Li, Julián Mestre, Viswanath Nagarajan, and Atri Rudra. When lp is the cure for your matching woes: Improved bounds for stochastic matchings. *Algorithmica*, 63(4):733–762, 2012.

[38] Brian Brubach, Karthik A. Sankararaman, Aravind Srinivasan, and Pan Xu. Algorithms to approximate column-sparse packing problems. *ACM Trans. Algorithms*, 16(1), November 2019. ISSN 1549-6325. https://doi.org/10.1145/3355400. URL https://doi.org/10.1145/3355400.

[39] Alok Baveja, Amit Chavan, Andrei Nikiforov, Aravind Srinivasan, and Pan Xu. Improved bounds in stochastic matching and optimization. *Algorithmica*, 80(11):3225–3252, 2018.

[40] Aleksandrs Slivkins et al. Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning*, 12(1-2):1–286, 2019.

[41] Huasen Wu, Rayadurgam Srikant, Xin Liu, and Chong Jiang. Algorithms with logarithmic or sublinear regret for constrained contextual bandits. *Advances in Neural Information Processing Systems*, 28, 2015.

[42] Santiago R Balseiro and Yonatan Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 65(9):3952–3968, 2019.

[43] Santiago R Balseiro, Haihao Lu, and Vahab Mirrokni. The best of many worlds: Dual mirror descent for online allocation problems. *Operations Research*, 2022.

[44] Negin Golrezaei, Patrick Jaillet, Jason Cheuk Nam Liang, and Vahab Mirrokni. Bidding and pricing in budget and roi constrained markets. *arXiv preprint arXiv:2107.07725*, 2021.

[45] Dirk van der Hoeven, Nikita Zhivotovskiy, and Nicolò Cesa-Bianchi. A regret-variance trade-off in online learning. *arXiv preprint arXiv:2206.02656*, 2022.

[46] Sattar Vakili, Alexis Boukouvalas, and Qing Zhao. Decision variance in risk-averse online learning. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 2738–2744. IEEE, 2019.

[47] Marc C Steinbach. Markowitz revisited: Mean-variance models in financial portfolio analysis. *SIAM review*, 43(1):31–85, 2001.

[48] Yuk Hei Chan and Lap Chi Lau. On linear and semidefinite programming relaxations for hypergraph matching. *Mathematical programming*, 135(1-2):123–148, 2012.

[49] José A Adell and P Jodrá. The median of the poisson distribution. *Metrika*, 61(3):337–346, 2005.

[50] Clément Canonne. A short note on poisson tail bounds. `http://www.cs.columbia.edu/~ccanonne/files/misc/2017-p`, 2020. Accessed: 2020-02-01.

[51] Eric V Slud. Distribution inequalities for the binomial law. *The Annals of Probability*, pages 404–412, 1977.

[52] P Borjesson and C-E Sundberg. Simple approximations of the error function q (x) for communications applications. *IEEE Transactions on Communications*, 27(3):639–643, 1979.

[53] Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press, 2017.

A   Missing proofs in Section 3.1

Recall that $\mathbf{Z}(e) = (Z_k)_{k \in \mathcal{S}_e}$ is a random binary vector of length Δ with $\mathbb{E}[Z_k] \leq \alpha B_k / T$ for each $k \in \mathcal{S}_e$. Consider a given $t \in [T]$. Let $Z_k^{(t)}$ be the sum of t *i.i.d.* copies of Z_k for each k. Let $\mathcal{D}^*$ be such a distribution on $\mathbf{Z}(e)$ that the value $\Pr_{\mathbf{Z}(e) \sim \mathcal{D}^*}\left[\exists k \in \mathcal{S}_e, Z_k^{(t)} \geq B_k\right]$ gets maximized.

Lemma 4. $\mathcal{D}^*$ *can always be realized at such a configuration that* $\mathbf{Z}(e) = \mathbf{1}(Z_k)$ *with probability* $\alpha B_k / T$ *for each* $k \in \mathcal{S}_e$ *and* $\mathbf{Z}(e) = \mathbf{0}$ *with probability* $1 - \sum_{k \in \mathcal{S}_e} \alpha B_k / T$*, where* $\mathbf{1}(Z_k)$ *refers to the standard basis vector with the only entry being one at the position of* Z_k*, and* $\mathbf{0}$ *refers to a zero vector of length* $|\mathcal{S}_e| = \Delta$*.*

Recall that we assume $T \gg \max_k B_k \geq 1$ and thus, $1 - \sum_{k \in \mathcal{S}_e} \alpha B_k / T \geq 0$.

Proof. We can interpret the value $\Pr_{\mathbf{Z}(e) \sim \mathcal{D}^*}\left[\exists k \in \mathcal{S}_e, Z_k^{(t)} \geq B_k\right]$ via the following balls-and-bins model: We have $|\mathcal{S}_e| = \Delta$ bins and t balls; during each time $t' \in [t]$, we put a ball into one or multiple bins according to a certain randomized strategy $\mathcal{D}^*$ such that each bin $k \in \mathcal{S}_e$ will receive a ball with a marginal probability no more than $\alpha B_k / T$; we repeat the process independently for t times and we aim to figure out the strategy $\mathcal{D}^*$ to maximize the chance that at least one bin $k \in \mathcal{S}_e$ will get at least B_k balls in the end.

Suppose under $\mathcal{D}^*$, with some probability $p \in [0, 1]$, we will put a ball into two different bins, say k and k'. Consider such a twisted version $\mathcal{D}$ of $\mathcal{D}^*$ that we simply modify $\mathcal{D}^*$ by putting a ball into bin k and bin k' each with a probability p and keep all the rest of $\mathcal{D}^*$. We can verify that strategy $\mathcal{D}$ is still feasible in the way that all bins will get a ball with a marginal probability the same as before. Consider the two materialization trees of $\mathcal{D}^*$ and $\mathcal{D}$ where each has a depth of t. We say a path is successful if it ends with that at least a bin $i \in \mathcal{S}_e$ has at least B_i balls. We can construct an *injective* mapping between all successful paths on the tree of $\mathcal{D}^*$ and those on tree of $\mathcal{D}$. This suggests that the chance of success under strategy $\mathcal{D}^*$ should be no larger than that under $\mathcal{D}$. Keeping on arguing in the way above, we get our claim. □

Lemma 5. *Let* $Z^{(t)}$ *be the sum of* t *i.i.d. Bernoulli random variables each has a mean of* $\alpha B / T$ *with* $\alpha \in [0, 1]$*,* $1 \leq B \ll T$*, and* $1 \leq t \leq T$*. We have that* $\Pr[Z^{(t)} \geq B]$ *get maximized when* $B = 1$*.*

Proof. It suffices to show that $\Pr[Z^{(t)} \leq B - 1]$ will get minimized at $B = 1$. Note that $\mathbb{E}[Z^{(t)}] = \alpha t B / T := \mu$.

When $\mu < B - 1$, by Lemma 11 of [39], we have $\Pr[Z^{(t)} \leq B - 1] \geq \Pr[\mathrm{Pois}(\mu) \leq B - 1]$, where $\mathrm{Pois}(\mu)$ denotes a Poisson random variable with mean μ. When $\mu \geq B - 1$, we can view that $Z^{(t)} \sim \mathrm{Pois}(\mu)$ since $1 \leq B \ll T$ (here we can treat B is a given constant since $T \gg B$). Thus, we claim that $\Pr[Z^{(t)} \leq B - 1] \geq \Pr[\mathrm{Pois}(\mu) \leq B - 1]$ with $\mu = \alpha t B / T \leq B$ [8]. Now we show that for any given $\tau := \alpha t / T \in [0, 1]$, $f(B) := \Pr[\mathrm{Pois}(\tau B) \leq B - 1]$ is an increasing function on $B = 1, 2, \ldots$. Observe that

$$f(B) \cdot e^{\tau B + \tau} = e^\tau \sum_{\ell=0}^{B-1} (\tau B)^\ell / \ell! = \left(\sum_{\ell=0}^{\infty} \tau^\ell / \ell!\right) \cdot \left(\sum_{\ell=0}^{B-1} (\tau B)^\ell / \ell!\right) := g_B(\tau) \tag{18}$$

$$f(B+1) \cdot e^{\tau B + \tau} = \sum_{\ell=0}^{B} (\tau B + \tau)^\ell / \ell! := g_{B+1}(\tau) \tag{19}$$

Suppose we regard the right-hand-side of Equations (18) and (19) both as polynomials of τ, denoted by $g_B(\tau)$ and $g_{B+1}(\tau)$, respectively. We can verify the two share the same coefficients of τ^ℓ for all $0 \leq \ell \leq B - 1$ and further more, we can verify that $g_B(\tau) - g_{B+1}(\tau) = \tau^B (\sum_{\ell=1}^{\infty} c_\ell \tau^\ell - B^B / B!)$ with $c_\ell > 0$. Observe that

$$f(B) \leq f(B+1) \Leftrightarrow f(B) \cdot e^{\tau B + \tau} \leq f(B+1) \cdot e^{\tau B + \tau} \Leftrightarrow g_B(\tau) - g_{B+1}(\tau) \leq 0 \Leftrightarrow \sum_{\ell=1}^{\infty} c_\ell \tau^\ell - B^B / B! \leq 0.$$

---

[8] More precisely, we have $\Pr[Z^{(t)} \leq B - 1] \geq \Pr[\mathrm{Pois}(\mu) \leq B - 1] - O(1/T)$. We ignore terms of $O(1/T)$ since $T \gg B \geq 1$.

21

Thus, it would suffice to show the case when $\tau = 1$. In this case, $f(B) = \Pr[\text{Pois}(B) \leq B - 1]$ is increasing on B, which follows from Lemma 1 of [49].

Summarizing all analyses above, we claim $\Pr[Z^{(t)} \leq B - 1] \geq \Pr[\text{Pois}(\alpha t B/T) \leq B - 1] \geq \Pr[\text{Pois}(\alpha t/T) \leq 0] = e^{-\alpha t/T}$. Observe that when $B = 1$, $\Pr[Z^{(t)} \leq B - 1] = (1 - \alpha/T)^t \sim e^{-\alpha t/T}$ when $T \gg 1$. Therefore, we establish that $\Pr[Z^{(t)} \leq B - 1]$ gets minimized at $B = 1$. $\qquad\square$

B Missing proofs in Section 3.2

Lemma 6. *Let $H_t = 1$ indicate that $\text{Ber}(\alpha B/T)^{(t-1)} \leq B - 1$ and $H_t = 0$ otherwise, where $\text{Ber}(\alpha B/T)^{(t-1)}$ is the sum of $t-1$ i.i.d. Bernoulli random variables each with mean $\alpha B/T$. We have that for any given $\alpha > 0$,*

$$F_\alpha(B) := \sum_{t=1}^{T} 2(t-1) \cdot \mathbb{E}[H_t] - \Big(\sum_{t=1}^{T} \mathbb{E}[H_t] \Big)^2 \leq \frac{T^2}{\alpha^2} \Big(1 - e^{-2\alpha} - 2\alpha e^{-\alpha} + O(1/T) \Big),$$

where the inequality above becomes tight when $B = 1$.

Proof. For notation convenience, let $Z^{(t-1)} := \text{Ber}(\alpha B/T)^{(t-1)}$, which denotes the sum of $t-1$ *i.i.d.* Bernoulli random variables each with mean $\alpha B/T$. Let $p_t = \mathbb{E}[H_t] = \Pr[Z^{(t-1)} \leq B-1]$. Consider a given t and suppose we try to view $F_\alpha(B)$ as a function of p_t and all $\{p_{t'}|t' \neq t\}$ are constants. We can verify that

$$F_\alpha(B) = -p_t^2 + 2p_t \Big(t - 1 - \sum_{t' \neq t} p_{t'} \Big) + C := f(p_t),$$

where C is a constant (a function of $\{p_{t'}|t' \neq t\}$). Note that by Lemma 5, $p_t = \Pr[Z^{(t-1)} \leq B - 1]$ gets minimized at $B = 1$ and thus, $p_t \geq (1 - \alpha/T)^{t-1}$. Therefore, we can verify that $\sum_{t' \neq t} p_{t'} \geq \Omega(T)$, which suggests that for any given $t \leq \Omega(T)$, $f(p_t)$ is decreasing when $p_t \in [0, 1]$. In the proof of Lemma 5, we show $p_t = \Pr[Z^{(t-1)} \leq B - 1] \geq \Pr[\text{Pois}((t-1)\alpha B/T) \leq B - 1] := \ell_t$ for any $t \in [T]$. This implies that $\text{Var}[H]$ will get non-decreased if we replace p_t with ℓ_t for any $t \leq (1 - 1/e)T$. Note that $p_t = \ell_t$ for any $t \geq \Omega(T)$.[9] Thus,

$$F_\alpha(B) \leq \sum_{t=1}^{T} 2(t-1) \cdot \ell_t - \Big(\sum_{t=1}^{T} \ell_t \Big)^2 \tag{20}$$

$$\leq \sum_{t=0}^{T-1} 2t \cdot e^{-\alpha Bt/T} \sum_{\ell=0}^{B-1} \frac{(\alpha Bt/T)^\ell}{\ell!} - \Big(\sum_{t=0}^{T-1} e^{-\alpha Bt/T} \sum_{\ell=0}^{B-1} \frac{(\alpha Bt/T)^\ell}{\ell!} \Big)^2 \tag{21}$$

$$= \sum_{\ell=0}^{B-1} T^2 \sum_{t=0}^{T-1} \frac{1}{T} \frac{2t}{T} e^{-\alpha Bt/T} \frac{(\alpha Bt/T)^\ell}{\ell!} - \Big(\sum_{\ell=0}^{B-1} T \sum_{t=0}^{T-1} \frac{1}{T} e^{-\alpha Bt/T} \frac{(\alpha Bt/T)^\ell}{\ell!} \Big)^2 \tag{22}$$

$$= \sum_{\ell=0}^{B-1} T^2 \Big(\int_0^1 2\zeta \cdot e^{-\alpha B\zeta} \frac{(\alpha B\zeta)^\ell}{\ell!} d\zeta + O(1/T) \Big) - \Big[\sum_{\ell=0}^{B-1} T \Big(\int_0^1 e^{-\alpha B\zeta} \frac{(\alpha B\zeta)^\ell}{\ell!} d\zeta + O(1/T) \Big) \Big]^2 \tag{23}$$

$$= T^2 \Big[\int_0^1 2\zeta \cdot \Pr\left[\text{Pois}(\alpha B\zeta) \leq B - 1\right] d\zeta - \Big(\int_0^1 \Pr\left[\text{Pois}(\alpha B\zeta) \leq B - 1\right] d\zeta \Big)^2 \Big] + o(T^2). \tag{24}$$

We show that the expression on (24) gets maximized at $B = 1$ any given constant $\alpha > 0$ (Actually, we can show it is strictly increasing when $B = 1, 2, \dots$ for any given $\alpha > 0$). For the convenience of exposition, we

---

[9] For analysis convenience, we ignore all $O(1/T)$ terms involved here since $T \gg B \geq 1$.

just show the case when $\alpha = 1$ [10].

$$f(B) := \int_0^1 2\zeta \cdot \Pr\left[\text{Pois}(B\zeta) \leq B - 1\right] d\zeta - \left(\int_0^1 \Pr\left[\text{Pois}(B\zeta) \leq B - 1\right] d\zeta\right)^2 \tag{25}$$

$$\leq \int_0^1 2\zeta d\zeta - \left(1 - \int_0^1 \Pr\left[\text{Pois}(B\zeta) \geq B\right] d\zeta\right)^2 \tag{26}$$

$$\leq 1 - \left(1 - \int_0^1 e^{-\zeta^2 B/2} d\zeta\right)^2 = 1 - \left(1 - \frac{1}{\sqrt{B}} \int_0^{\sqrt{B}} e^{-\zeta^2/2} d\zeta\right)^2 \tag{27}$$

$$\leq 1 - \left(1 - \frac{1}{\sqrt{B}} \int_0^\infty e^{-\zeta^2/2} d\zeta\right)^2 = \sqrt{\frac{\pi}{2B}}\left(2 + \sqrt{\frac{\pi}{2B}}\right), \tag{28}$$

where Inequality (27) is due to the upper-tail bound of a Poisson random variable [50]. We can manually verify that $f(B)$ is strictly decreasing when $B = 1, 2, \ldots, 1000$ with $f(1) = 1 - e^{-2} - 2e^{-1} \sim 0.1289$, and $f(B) \leq 0.08$ when $B > 1000$ by the upper bound shown above. Thus, we claim that $f(B)$ gets maximized at $B = 1$.

Substituting $B = 1$ back to Inequality (24), we have

$$F_\alpha(B) \leq T^2 \left[\int_0^1 2\zeta \cdot \Pr\left[\text{Pois}(\alpha\zeta) \leq 0\right] d\zeta - \left(\int_0^1 \Pr\left[\text{Pois}(\alpha\zeta) \leq 0\right] d\zeta\right)^2\right] + o(T^2) \tag{29}$$

$$= T^2\left(1 - e^{-2\alpha} - 2\alpha e^{-\alpha}\right)/\alpha^2 + o(T^2). \tag{30}$$

$\square$

C   Proofs of Theorems 9, 10, and 11

C.1   Proof of Theorem 9

From Lemma 4, we see that the optimal distribution $\mathcal{D}^*$ stated in Theorem 9 can always be realized at such a configuration that $\mathbf{U} = \mathbf{1}_k$ with probability B/T for each $k \in [\Delta]$ and $\mathbf{U} = \mathbf{0}$ with probability $1 - \Delta \cdot B/T$, where $\mathbf{1}_k$ denotes the kth standard basis vector. Under this setting, the question presented in Theorem 9 can be formulated as a Balls-and-Bins problem as follows.

A Balls-and-Bins Problem (BBP). Let Δ, B and T are three integers with $\Delta \leq B \ll T$. Suppose we have Δ bins and T rounds. In each round $t \in [T]$, we throw a ball such that it will hit each bin with probability B/T and it will miss all the bins with probability $1 - \Delta B/T$. Let $\mathbf{U} = (U_k | k \in [\Delta])$ be the random vector taking values from $\{0, 1\}^\Delta$, which captures the numbers of balls falling into the Δ bins in one single round. Observe that $\mathbf{U}$ follows the exact distribution of $\mathcal{D}^*$ as stated in Lemma 4. Let $\widehat{T} \leq T$ and $T' \leq T$ be the last rounds such that every bin has at most $B - 1$ balls before and after the arrival of ball, respectively. We see that $T' \leq \widehat{T} \leq T' + 1$, and

$$\lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^T \Pr_{\mathbf{U}\sim\mathcal{D}^*}\left[\max_{k\in[\Delta]} U_k^{(t)} \leq B - 1\right] = \lim_{T\to\infty} \frac{\mathbb{E}[\widehat{T}]}{T} = \lim_{T\to\infty} \frac{\mathbb{E}[T']}{T} \doteq \text{Ratio}.$$

An upper bound on Ratio.

Lemma 7. Ratio $\leq 1 - (\sqrt{2} - \epsilon)\sqrt{\frac{\ln \Delta}{B}}(1 + o(1))$ *for any given* $\epsilon > 0$*, where* $o(1)$ *is a vanishing term when* Δ *approaches infinity.*

---

[10] For a general case $\alpha > 0$, we can reduce to the case $\alpha = 1$ by treating $Z^{(t-1)}$ as $t - 1$ i.i.d. Ber(B/T') with $T' = T/\alpha$.

To prove the above lemma, we need Slud's Inequality [51] stated as follows.

Lemma 8 (Slud's inequality [51]). *Let $\{X_i | i \in [n]\}$ be n i.i.d. Bernoulli random variables with $\mathbb{E}[X_i] = p$ and let $k \leq n$ be a given integer. If either (a) $p \leq 1/4$ and $np \leq k$, or (b) $np \leq k \leq n(1-p)$, then*

$$\Pr\Big[\sum_{i=1}^{n} X_i \leq k - 1\Big] \leq \Phi\Big(\frac{k - np}{\sqrt{np(1-p)}}\Big),$$

where Φ is the cdf of a standard normal distribution.

Here is an explanation for the above inequality. When n is large, we can approximation the distribribution of $\sum_i X_i$ by $\mathcal{N}(np, np(1-p))$. Thus, we see that $\Pr[\sum_i X_i \leq k - 1] \sim \Phi((k - 1 - npt)/\sqrt{np(1-p)})$. Now we start to prove Lemma 7.

Proof. For each time $t \in [T]$ and bin $k \in [\Delta]$, let $U_k^{(t)}$ be the number of balls in the ith bin at the end of time t. Thus, we have

$$\mathbb{E}[T'] = \sum_{t=1}^{T} \Pr\Big[\bigwedge_{i \in [\Delta]} (U_k^{(t)} \leq B - 1)\Big] \leq \sum_{t=1}^{T} \prod_{i \in [\Delta]} \Pr[U_k^{(t)} \leq B - 1] \tag{31}$$

$$\leq B - 1 + \sum_{t=B}^{T} \Phi^{\Delta}\Big(\frac{B - tB/T}{\sqrt{t(B/T)(1 - B/T)}}\Big). \tag{32}$$

The first inequality (31) is due to the fact that $\{U_k^{(t)} | i \in [\Delta]\}$ are *negatively associated* [29]. The second inequality (32) is due to Slud's inequality. Observe that when $t \leq B - 1$, $\Pr[U_k^{(t)} \leq B - 1] = 1$ for every $i \in [\Delta]$. When $t \geq B$, we can view each $U_k^{(t)}$ as a sum of t *i.i.d.* Bernoulli random variable with mean B/T, and verify that Condition (a) stated in Lemma 8 is satisfied since $B/T \leq 1/4$ and $tB/T \leq B$. Therefore,

$$\begin{aligned}
\mathsf{Ratio} &= \lim_{T \to \infty} \frac{\mathbb{E}[T']}{T} \\
&\leq \lim_{T \to \infty} \Big(\frac{B - 1}{T} + \frac{1}{T} \sum_{t=B}^{T} \Phi^{\Delta}\Big(\frac{B - tB/T}{\sqrt{t(B/T)(1 - B/T)}}\Big)\Big) = \int_0^1 dz \Phi^{\Delta}\Big(\frac{\sqrt{B}(1 - z)}{\sqrt{z}}\Big) \\
&= \int_0^{1-\delta} dz \Phi^{\Delta}\Big(\frac{\sqrt{B}(1 - z)}{\sqrt{z}}\Big) + \int_{1-\delta}^1 dz \Phi^{\Delta}\Big(\frac{\sqrt{B}(1 - z)}{\sqrt{z}}\Big) \\
&\leq (1 - \delta) + \delta \cdot \Phi^{\Delta}\big(\sqrt{B}\delta/\sqrt{1 - \delta}\big). \tag{33}
\end{aligned}$$

Inequality (33) follows from the fact that $f(z) \doteq \Phi^{\Delta}\big(\sqrt{B}(1 - z)/\sqrt{z}\big)$ is an non-increasing function over $z \in [0, 1]$. Set $\delta = c\sqrt{\ln \Delta/B}$ with $c = \sqrt{2} - \epsilon$. In the following, we prove that $\Phi^{\Delta}\big(\sqrt{B}\delta/\sqrt{1 - \delta}\big) = o(1)$ when $\Delta \to \infty$ for any given $\epsilon > 0$. This yields our main result.

Let $Q(x) = 1 - \Phi(x)$. Observe that when $\delta = c\sqrt{\ln \Delta/B}$ with $c = \sqrt{2} - \epsilon$,

$$\Phi^{\Delta}\big(\sqrt{B}\delta/\sqrt{1 - \delta}\big) = \Phi^{\Delta}\big(c\sqrt{\ln \Delta}/\sqrt{1 - \delta}\big) \tag{34}$$

$$= \Big(1 - Q\big(c\sqrt{\ln \Delta}/\sqrt{1 - \delta}\big)\Big)^{\Delta} \tag{35}$$

$$\leq \Big(1 - \frac{e^{-c^2 \ln \Delta/(2(1-\delta))}}{\sqrt{2\pi}\big(2c\sqrt{\ln \Delta}\big)}\Big)^{\Delta} \tag{36}$$

$$\leq \exp\Big(-\Delta^{1 - \frac{c^2}{2(1-\delta)}}\frac{1}{2\sqrt{2\pi} \cdot c \cdot \sqrt{\ln \Delta}}\Big) = o(1). \tag{37}$$

Inequality (36) follows from the fact that $Q(x) \geq \phi(x)/(x + 1/x)$ for $x > 0$ with $\phi(x)$ being the pdf of the standard norm distribution [52]. Thus, we claim that $\Phi^{\Delta}\left(\sqrt{B}\delta/\sqrt{1-\delta}\right)$ is a vanishing term when $\Delta \to \infty$ with $\delta = c\sqrt{\ln \Delta/B} = o(1)$ since $B = \omega(\ln \Delta)$ and $c < \sqrt{2}$. $\qquad\square$

A lower bound on Ratio.

Lemma 9. Ratio $\geq 1 - 2\sqrt{2}\sqrt{\frac{\ln \Delta}{B}}(1 + o(1))$ *for any given* $\epsilon > 0$*, where* $o(1)$ *is a vanishing term when* Δ *approaches infinity.*

To prove the above lemma, we need the following key lemma, which appears as Lemma 5.10 on page 110 of the book [53].

Lemma 10 ([53]). *Consider a balls-and-bins model with m balls and n bins. Let $Y_i^{(m)}$ be the number of balls in the ith bin at the end of m rounds. let $\{Z_i^{(m)}|i \in [n]\}$ be n i.i.d. Poisson random variables each with mean m/n. Let $f(x_1, \ldots, x_n)$ be a nonnegative function such that $\mathbb{E}[f(Y_1^{(m)}, \ldots, Y_n^{(m)})]$ is either monotonically increasing or monotonically decreasing in m. Then*

$$\mathbb{E}[f(Y_1^{(m)}, \ldots, Y_n^{(m)})] \leq 2\mathbb{E}[f(Z_1^{(m)}, \ldots, Z_n^{(m)})].$$

Now we start to prove Lemma 9.

Proof. Recall that Ratio $= \lim_{T \to \infty} \mathbb{E}[T']/T$, where T' is the last round when all bins have at most $B - 1$ balls.

$$\mathbb{E}[T'] = \sum_{t=1}^{T} \Pr[\max_k U_k^{(t)} \leq B - 1] = \sum_{t=1}^{T} \left(1 - \Pr[\max_i U_k^{(t)} \geq B]\right) = T - \sum_{t=1}^{T} \Pr[\max_k U_k^{(t)} \geq B].$$

Therefore,

$$\mathbb{E}[T']/T = 1 - \frac{1}{T} \sum_{t=1}^{T} \Pr[\max_k U_k^{(t)} \geq B] = 1 - \frac{1}{T} \sum_{t=B}^{T} \Pr[\max_k U_k^{(t)} \geq B] \tag{38}$$

$$\geq 1 - \frac{2}{T} \sum_{t=B}^{T} \Pr[\max_k Z_k^{(t)} \geq B] \quad \left(\{Z_k^{(t)}\} \text{ are } \Delta \text{ i.i.d. Poisson r.v.s each with mean } tB/T\right) \tag{39}$$

$$= 1 - \frac{2}{T} \sum_{t=B}^{T} \left[1 - \left(\Pr[Z_k^{(t)} \leq B - 1]\right)^{\Delta}\right] \tag{40}$$

$$= -1 + O(B/T) + \frac{2}{T} \sum_{t=1}^{T} \left(\Pr\left[\text{Pois}(tB/T) \leq B - 1\right]\right)^{\Delta}. \tag{41}$$

Inequality (39) is due to Lemma 10. In our case, we have $n = T/B$ and $m = t$, and $f\left(Y_1^{(m)}, \ldots, Y_n^{(m)}\right) \doteq \mathbf{1}\left(U_k^{(t)} \geq B, \forall 1 \leq k \leq \Delta\right)$, which is an indicator function showing if one of the first Δ bins has at least B balls by the round of t. We can verify that $\mathbb{E}[f(Y_1^{(t)}, \ldots, Y_n^{(t)})]$ is monotonically increasing in t under our definition.

Thus, Inequality (41) suggests that

$$\text{Ratio} = \lim_{T \to \infty} \frac{\mathbb{E}[T']}{T} \geq -1 + 2 \int_0^1 dx \left(\Pr\left[\text{Pois}(xB) \leq B - 1\right]\right)^{\Delta} \doteq 1 - 2F, \tag{42}$$

25

where $F = 1 - \int_0^1 dx \left(1 - \Pr[\mathrm{Pois}(xB) \geq B]\right)^\Delta$. We try to upper bound F as follows. Note that,

$$F \leq \int_0^1 dx \left[1 - \left(1 - \exp\left(\frac{-B}{2}x^2\right)\right)^\Delta\right] \tag{43}$$

$$= \sqrt{\frac{2}{B}} \int_0^{\sqrt{\frac{B}{2}}} d\theta \left[1 - \left(1 - \mathrm{e}^{-\theta^2}\right)^\Delta\right] \leq \sqrt{\frac{2}{B}} \int_0^\infty d\theta \left[1 - \left(1 - \mathrm{e}^{-\theta^2}\right)^\Delta\right] \tag{44}$$

$$\leq \sqrt{\frac{2}{B}} \left[\sqrt{\ln \Delta} + \int_{\sqrt{\ln \Delta}}^\infty d\theta \left[1 - \left(1 - \mathrm{e}^{-\theta^2}\right)^\Delta\right]\right] \tag{45}$$

$$\leq \sqrt{\frac{2}{B}} \left[\sqrt{\ln \Delta} + \Delta \int_{\sqrt{\ln \Delta}}^\infty d\theta \, \mathrm{e}^{-\theta^2}\right] \tag{46}$$

$$\leq \sqrt{\frac{2}{B}} \left[\sqrt{\ln \Delta} + \Delta \cdot \frac{1}{2\sqrt{\ln \Delta}} \cdot \mathrm{e}^{-(\sqrt{\ln \Delta})^2}\right] = \sqrt{2}\sqrt{\frac{\ln \Delta}{B}}(1 + o(1)). \tag{47}$$

Inequality (43) is due to the upper tail bound of a Poisson random variable [50]. Inequality (46) is due to the fact that $f(\theta) \doteq \left[1 - \left(1 - \mathrm{e}^{-\theta^2}\right)^\Delta\right] \leq \Delta \cdot \mathrm{e}^{-\theta^2}$. Inequality (47) follows from the fact that $Q(x) = 1 - \Phi(x) \leq \phi(x)/x$ for any $x > 0$ with $\phi(x)$ being the density function of the standard normal distribution [52].

Plugging the result of Inequality (47) back to (42), we get that

$$\mathsf{Ratio} = \lim_{T \to \infty} \frac{\mathbb{E}[T']}{T} \geq 1 - 2\sqrt{2}\sqrt{\frac{\ln \Delta}{B}}(1 + o(1)).$$

$\square$

C.2   Proof of Theorem 10

Based on the structure shown in Lemma 4, we construct the worst scenario of SAMP as follows.

Example 5. Consider such a simple case that $|I| = |J| = 1$ and $q_j = 1$ and $r_j = T$ where T is the total number of online rounds. In other words, there is one single online agent which arrives with probability one during each round $t \in [T]$, and there is one single edge, denoted by e. Assume $W_e = 1$ with probability one. Let $K = \Delta$ and $\mathcal{A}_e = \mathbf{1}_k$ with probability B/T for each $k \in [\Delta]$ and $\mathcal{A}_e = \mathbf{0}$ with probability $1 - \frac{\Delta B}{T}$, where $\mathbf{1}_k \in \{0,1\}^\Delta$ refers to the kth standard basis vector. Thus, by definition we have $\mathcal{S}_e = [\Delta]$ with $|\mathcal{S}_e| = \Delta$. Set $B_k = B$ for all $k \in [K]$. $\square$

Lemma 11. *No algorithm can achieve an online competitive ratio better than $1 - \kappa\sqrt{\frac{\ln \Delta}{B}}(1 + o(1))$ on Example 5, where κ is the value as stated in Theorem 9 and $o(1)$ vanishes when $B = \omega(\ln \Delta)$ and $\Delta \to \infty$.*

Proof. We can verify that the optimal value to LP-(1) is T since $x_e = T$ is the unique optimal solution. Note that there is one single edge e, which arrives during each round with probability one. Thus, the expected total utilities achieved by any policy should be no greater than that by Greedy, which will add the edge e whenever e is safe. Observe that e is safe iff the consumption of each resource k is no larger than $B - 1$. Let $T' \leq T$ be the last round such that the consumption of every resource $k \in [K]$ is no larger than $B - 1$ (before the arrival of online agent). We see that Greedy $\leq \mathbb{E}[T']$, which suggests that the CR of any algorithm will never beat $\mathbb{E}[T']/T$ with respect to LP-(1). Note that if applying Greedy to Example 5, the random resource consumption can be captured by the Bins-and-Balls model in Section C.1. Particularly, T' is exactly equal to the last round that all bins have at most $B - 1$ balls. Therefore, we claim that any policy should achieve a CR no larger than $1 - \kappa\sqrt{\frac{\ln \Delta}{B}}(1 + o(1))$ on Example 5. $\square$

C.3   Proof of Theorem 11

Lemma 12. *When $\Delta = 1$, $\lim_{T\to\infty} \mathbb{E}[\widehat{T}]/T = 1 - \frac{1}{\sqrt{2\pi B}}(1 + o(1))$, where $o(1)$ is a vanishing term when $B \to \infty$.*

Proof. Consider a special case of BPP stated in Section C.1 with $\Delta = 1$. In this case, there is only one bin and during each round $t \in [T]$, it will get a ball with probability B/T and nothing otherwise. Recall that $\widehat{T} \leq T$ is the last round such that the bin has at most $B - 1$ balls before the arrival of ball. We can recast $\widehat{T}$ as the smallest integer $n \leq T$ such that $\sum_{i=1}^{n} X_i = B$ or $\widehat{T} = T$ where $\{X_i\}$ are *i.i.d.* Bernoulli random variables each with mean B/T. Let $X = \sum_{i=1}^{\widehat{T}} X_i$. Note that $\widehat{T}$ qualified as a *stopping* time. By Wald's equation, we have $\mathbb{E}[X] = \mathbb{E}[\widehat{T}] \cdot \mathbb{E}[X_i]$. Thus,

$$\lim_{T\to\infty} \frac{\mathbb{E}[\widehat{T}]}{T} = \lim_{T\to\infty} \frac{\mathbb{E}[X]}{\mathbb{E}[X_i]} \frac{1}{T} = 1 - \frac{1}{\sqrt{2\pi B}}(1 + o(1)).$$

The last equality is due to the fact that $\mathbb{E}[X] = B\left(1 - \frac{1}{\sqrt{2\pi B}}\left(1 + o(1)\right)\right)$, which result appears in multiple contexts before such as Adwords and online B-matching problems [20, 21], and correlation gap [22]. □

Proof of Theorem 11. Lemma 12 suggests that when $\Delta = 1$, the limit on Equation (17) in Theorem 9 will be $1 - \frac{1}{\sqrt{2\pi B}}(1 + o(1))$. Following the same procedure of applying Theorem 9 to prove the general results ($\Delta \gg 1$) of Theorem 4 as shown in Section 6, we can show that SAMP will achieve a CR of $1 - \frac{1}{\sqrt{2\pi B}}(1+o(1))$ when $\Delta = 1$. The asymptotic optimality can be seen on an example similar to Example 5: we just need to set up one single resource, and the single edge e has $\mathcal{A}_e = 1$ with probability B/T and 0 otherwise. All the rest proofs are the same as shown in Lemma 11. □