

---

# Maillard Sampling: Boltzmann Exploration Done Optimally

---

Jie Bian

University of Arizona

Kwang-Sung Jun

University of Arizona

## Abstract

The PhD thesis of Maillard (2013) presents a rather obscure algorithm for the  $K$ -armed bandit problem. This less-known algorithm, which we call Maillard sampling (MS), computes the probability of choosing each arm in a *closed form*, which is not true for Thompson sampling, a widely-adopted bandit algorithm in the industry. This means that the bandit-logged data from running MS can be readily used for counterfactual evaluation, unlike Thompson sampling. Motivated by such merit, we revisit MS and perform an improved analysis to show that it achieves both the asymptotical optimality and  $\sqrt{KT \log T}$  minimax regret bound where  $T$  is the time horizon, which matches the known bounds for asymptotically optimal UCB. We then propose a variant of MS called MS<sup>+</sup> that improves its minimax bound to  $\sqrt{KT \log K}$ . MS<sup>+</sup> can also be tuned to be aggressive (i.e., less exploration) without losing the asymptotic optimality, a unique feature unavailable from existing bandit algorithms. Our numerical evaluation shows the effectiveness of MS<sup>+</sup>.

## 1 INTRODUCTION

The  $K$ -armed bandit problem (Thompson, 1933; Lattimore and Szepesvári, 2020) poses a unique challenge of balancing between exploration and exploitation in sequential decision-making tasks. Researchers have studied this problem and propose algorithms extensively since its widely adopted commercial applications including online news recommendations (Li et al., 2010). In this problem, the learner is given a set of  $K$  arms where each arm  $i$  has an associated reward distribution  $\nu_i$  with an unknown mean  $\mu_i \in \mathbb{R}$ . At each time  $t$ , the learner

chooses an arm  $I_t \in [K] := \{1, \dots, K\}$  and then receives a reward  $y_t \sim \nu_{I_t}$  associated with the chosen arm  $I_t$ . In this paper, we focus on the assumption that  $y_t - \mu_{I_t}$  is  $\sigma^2$ -sub-Gaussian, i.e.,  $\mathbb{E}[\exp(\lambda \eta_t)] \leq \exp(\lambda^2 \sigma^2 / 2)$ , and that the algorithm knows  $\sigma^2$ , which is standard.<sup>1</sup> The standard performance measure of a bandit algorithm is cumulative (pseudo-)regret, or simply *regret*, defined as follows:

$$\text{Reg}_T = \sum_t^T \max_{a \in [K]} \mu_a - \mu_{I_t} \quad (1)$$

where  $T$  is the time horizon.

Among many algorithms with regret guarantees, randomized algorithms such as Thompson sampling have attracted significant attention due to their high performance and the fact that the actions are chosen randomly (Chapelle and Li, 2011). The latter property is useful for off-policy evaluation where one has access to a set of chosen arms and their rewards as a result of deploying a bandit algorithm  $A$  and desires to estimate “How much reward would I have collected, had I used a different algorithm  $B$ ?” (Precup, 2000; Li et al., 2011; 2015). Solutions for such a question are attractive since they enable the evaluation of alternative algorithms without deploying thus are an active area of research. However, all the existing methods with good regret guarantees require that the probability value with which the pulled arm was chosen is available precisely; e.g., inverse propensity scoring (Horvitz and Thompson, 1952) and doubly robust estimators (Robins and Rotnitzky, 1995). Thompson sampling, however, does not maintain explicit probability values, and there is no closed-form expression for computing them. While one can simulate the sampling rule from Thompson sampling and estimate the probability, this is computationally expensive especially for web-scale deployments.

Surprisingly, there exists a novel randomized algorithm in the PhD thesis of Maillard (2013) that computes the probability of choosing each arm in a closed form,

---

Proceedings of the 25<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

---

<sup>1</sup>In many cases, practitioners have access to a reasonable value of  $\sigma^2$ . For example, clicks ( $\{0, 1\}$ ) or scaled 5-star ratings ( $\{1, 2, 3, 4, 5\}$ ) are bounded, and if a reward is bounded by  $[a, b]$  then it is  $\frac{(a-b)^2}{4}$ -sub-Gaussian.

which we call *Maillard sampling* (MS). At time step  $t$ , MS for  $\sigma^2 = 1$  chooses arm  $a$  with probability

$$p_{t,a} \propto \exp\left(-\frac{N_{t-1,a} \widehat{\Delta}_{t-1,a}^2}{2}\right)$$

where  $N_{t-1,i}$  is the number of times arm  $i$  has been pulled up to (and including) time step  $t-1$ ,  $\widehat{\Delta}_{t-1,i} := \max_j \hat{\mu}_{t-1,j} - \hat{\mu}_{t-1,i}$  is the empirical gap, and  $\hat{\mu}_{t-1,i}$  is the empirical mean up to time step  $t-1$ . The author shows a regret bound of MS that is asymptotically optimal, although the bound does not satisfy the sub-UCB-ness and implies a minimax regret of  $O(\sqrt{KT}^{3/4})$  that is substantially higher than  $\sqrt{KT \log T}$  achieved by existing algorithms like UCB1 (Auer et al., 2002) (see Section 2 for the definitions of the optimalities). We provide a detailed comparison of MS and other algorithms in Section 3.

MS is a lesser-known type of bandit algorithm that can be viewed as a correction to the popular Boltzmann exploration (BE) (Sutton, 1990; Kaelbling et al., 1996), which chooses arm  $i$  with  $p_{t,i} \propto \exp(\eta_t \hat{\mu}_{t-1,i})$  for some step size  $\eta_t$ . That is, since BE can be rewritten as  $p_{t,i} \propto \exp(-\eta_t \widehat{\Delta}_{t-1,a})$ , one can obtain MS by squaring the empirical gap and taking arm-specific step size of  $N_{t-1,i}/2$  in place of  $\eta_t$ . For practitioners, we recommend that they use Eq. (4) as a replacement to BE.

Motivated by the simplicity of MS and its friendliness to the off-policy evaluation, we revisit MS and make two main contributions. First, we streamline the proof of its theoretical guarantee<sup>2</sup> and show that MS achieves not only asymptotically optimality but also the sub-UCB-ness and a minimax regret bound of  $O(\sqrt{KT \log(T)})$ , matching the guarantees of asymptotically optimal UCB Lattimore and Szepesvári (2020, Section 8). Second, we propose a new algorithm called Maillard Sampling<sup>+</sup> (MS<sup>+</sup>). Besides all the optimalities that MS possesses, MS<sup>+</sup> further enjoys a minimax regret bound of  $O(\sqrt{KT \log(K)})$ . We present the algorithms and the regret bounds of MS and MS<sup>+</sup> in Section 4 and 5 respectively.

MS<sup>+</sup> further possesses a unique, practically relevant property. In the design of MS<sup>+</sup>, we allow an exploitation parameter that we call the ‘booster’, which encourages exploitation without breaking theoretical guarantees. Increasing the booster parameter only affects a lower-order term in the regret bound. This means that even for a very aggressive booster, which is likely to incur linear regret for some time in the beginning, the algorithm has the ability to recover from it in a finite time. This is in stark contrast to the common heuristic of reducing the confidence width of UCB (Zhang et al., 2016) or

posterior variance of Thompson sampling (Chapelle and Li, 2011), which breaks the regret guarantee and risks suffering a linear regret from which one may never recover. The practical implication is that, given domain knowledge on the noise level, one can tune the booster parameter to perform well without such a risk.

In section 6, we perform an empirical evaluation of MS<sup>+</sup>, which achieves either the best or on par regret among the anytime algorithms. Our evaluation also shows the effect of the booster parameter, which indeed allows us to reduce the expected regret without risking a linear regret at the price of increasing variance. Finally, we conclude our paper in Section 7 with future research directions enabled by our paper.

## 2 PRELIMINARIES

**Notations.** Without loss of generality, we assume  $\mu_1 \geq \mu_2 \geq \dots \geq \mu_K$ . Define the gap  $\Delta_a := \mu_1 - \mu_a$ ,  $\forall i \in [K] := \{1, \dots, K\}$ .

**Regret optimality.** For the expected regret  $\mathbb{E} \text{Reg}_T$ , there are multiple optimality criteria. An algorithm is called asymptotically optimal if it satisfies the following (Lai and Robbins, 1985; Burnetas and Katehakis, 1996):

$$\limsup_{T \rightarrow \infty} \frac{\mathbb{E} \text{Reg}_T}{\log(T)} = \sum_{i: \Delta_i > 0} \frac{2\sigma^2}{\Delta_i}. \quad (2)$$

The minimax optimal regret is the smallest achievable worst-case regret with respect to the number of arms  $K$  and the time horizon  $T$ . The minimax optimal regret of the bandit problem is  $\Theta(\sqrt{KT})$  (Audibert et al., 2009), meaning that there exists an algorithm with the minimax regret of  $O(\sqrt{KT})$  and a problem instance where any algorithm must suffer regret of  $\Omega(\sqrt{KT})$ . We say that an algorithm has a minimax ratio of  $f(K, T)$  if it has a minimax regret bound of  $O(\sqrt{KT} f(K, T))$ . Finally, an algorithm is said to be sub-UCB if  $\mathbb{E} \text{Reg}_T \leq C_1 \sum_{i \in [K]} \Delta_i + C_2 \sum_{i: \Delta_i > 0} \frac{\sigma^2}{\Delta_i} \log(T)$  for absolute constants  $C_1, C_2 > 0$  (Lattimore and Szepesvári, 2020). While sub-UCB is satisfied by most existing algorithms including UCB1 (Auer et al., 2002), some algorithms like MOSS (Audibert et al., 2009) are so aggressive on certain instances that they are not sub-UCB; see Lattimore and Szepesvári (2020) for details.

An algorithm  $\mathcal{A}$  is said to be anytime if it does not take in the time horizon  $T$  whereas  $\mathcal{A}$  is fixed budget if it requires  $T$  as input and enjoys a regret bound only for the time horizon  $T$ .

## 3 RELATED WORK

We summarize existing  $K$ -armed bandit algorithms for sub-Gaussian rewards in Table 1. In this table, TS-SG

<sup>2</sup>In fact, we were not able to verify the correctness of the original proof. The original author has confirmed that the proof is not clear.

	Sub-UCB	Asympt. optimality	Minimax ratio	Anytime	Closed-form probability
UCB1 (Auer et al., 2002)	✓	✗	$\sqrt{\log T}$	✓	N/A
AOUCB (Lattimore and Szepesvári, 2020, Sec. 8)	✓	✓	$\sqrt{\log T}$	✓	N/A
TS-SG (Agrawal and Goyal, 2017, Alg. 2)	✓	✗	$\sqrt{\log K}$	✓	✗
MOSS (Lattimore and Szepesvári, 2020, Sec. 9)	✗	✗	1	✗	N/A
UCB <sup>+</sup> (Lattimore, 2015, Alg. 2)	✓	✗*	$\sqrt{\log K}$	✗	N/A
OCUCB (Lattimore, 2015, Alg. 1)	✓	✗*	1	✗	N/A
MS (this work; Algorithm 1)	✓	✓	$\sqrt{\log T}$	✓	✓
MS <sup>+</sup> (this work; Algorithm 2)	✓	✓	$\sqrt{\log K}$	✓	✓

Table 1: Comparison of bandit algorithms for sub-Gaussian rewards. Our algorithms enjoy the best guarantee except for the minimax ratio. See Section 3 for details. The mark ✗\* means that the optimality is not explicitly reported but might actually be achieved.

stands for Thompson sampling for sub-Gaussian rewards, which is originally designed for bounded rewards ( $y_t \in [0, 1]$ ) yet trivially enjoys the same guarantee for  $(1/2)^2$ -sub-Gaussian rewards. We emphasize that TS-SG is different from TS for Gaussian rewards (Korda et al., 2013) as TS-SG inflates the posterior variance by a factor of 4, which is why it does not have the asymptotic optimality. Note that there are much more algorithms if we restrict to Gaussian rewards. We refer to Lattimore and Szepesvári (2020, Table 2) for a more comprehensive summary.

Note that we only report results available in the literature in Table 1. For example, while MOSS and UCB<sup>+</sup> are fixed-budget algorithms, one should be able to modify these algorithms and prove *anytime* guarantees using existing techniques. Such a modification comes with a cost of inflating the width of the confidence bound, which likely degrades the practical performance. Similarly, one should be able to modify AdaUCB (Lattimore, 2018), which achieves all the possible optimality criteria in the literature for *Gaussian* rewards, to allow *sub-Gaussian* rewards. However, we emphasize that the focus of this paper is the (re-)introduction of a new type of algorithm with desirable properties and good regret guarantees, leaving further improvements in regret bounds for future work.

Finally, MS is closely related to many existing algorithms including Thompson sampling, we review in detail in Appendix A.

After the submission of the paper, the authors of Honda and Takemura (2015) have informed us that MS is in fact an instantiation of MED (Honda and Takemura, 2011) for Gaussian rewards and that at that time (‘pre-Thompson-sampling’ era) randomized algorithms were not preferred by reviewers, which necessitated deterministic versions of MED such as DMED (Honda and Takemura, 2010) and IMED (Honda and Takemura, 2015). However, the analysis of MED makes a lot of assumptions and is inherently an asymptotic argument. In contrast, our analysis of MS is finite time, and shows

---

### Algorithm 1 Maillard Sampling

---

**Input:**  $K \geq 2, \sigma^2 > 0$ .  
**for**  $t = 1, 2, \dots$  **do**  
    **if**  $t \leq K$  **then**  
        Pull the arm  $I_t = t$  and observe reward  $y_t$ .  
    **else**  
         $p_{t,a} \propto \exp(-\frac{1}{2\sigma^2} N_{t-1,a} \widehat{\Delta}_{t-1,a}^2)$   
        Pull the arm  $I_t \sim p_t$  and observe reward  $y_t$ .  
    **end if**  
**end for**

---

that it achieves a near-optimal minimax regret for the first time.

## 4 MAILLARD SAMPLING

We now formally introduce Maillard sampling (MS), which originally appeared in Maillard (2013, Figure 1.8). Algorithm 1 describes the pseudocode of MS with an explicit dependence on the sub-Gaussian parameter  $\sigma^2$ . MS has a very simple closed-form computation of sampling probability that directly uses the estimated gap  $\widehat{\Delta}_{t,a}$ . Interestingly, the sampling probability  $p_{t,a}$  is proportional to the sub-Gaussian tail at deviation  $\widehat{\Delta}_{t-1,a}$ . Such a form of probability has a close connection to Thompson sampling as we discuss in Section A.

**Theorem 1.** (Maillard, 2013, Theorem 1.10) *Maillard sampling (Algorithm 1) with  $\sigma^2 = 1/4$  satisfies that,  $\forall T \geq 1$  and  $c > 0$ ,*

$$\mathbb{E} \text{Reg}_T \leq \sum_{a:\Delta_a > 0} \left( \frac{(1+c)^2}{2\Delta_a} \log(T) + O\left(\frac{(1+c)^4}{c^4 \Delta_a^3} K\right) \right).$$

Theorem 1 implies that MS is asymptotically optimal by choosing  $c = 1/o(\log^{1/4}(T))$ . Specifically, such a value of  $c$  makes  $(1+c)^2$  asymptotically 1 and  $\frac{(1+c)^4}{c^4}$  asymptotically  $o(\log(T))$ . However, Theorem 1 does not imply that MS is sub-UCB due to the second term with a larger dependence on  $1/\Delta_a$ . Furthermore, using the standard technique such as Lattimore and Szepesvári (2020, Theorem 7.2), one can show that

Theorem 1 implies a minimax regret bound of  $\mathbb{E} \text{Reg}_T = O(\sqrt{KT}^{3/4})$ .

Altogether, Theorem 1 implies that MS has a worse regret upper bound than basic UCB algorithms such as AOUCB (asymptotically optimal UCB) (Cappé et al., 2013; Lattimore and Szepesvári, 2020), which is not only asymptotically optimal but also sub-UCB along with minimax regret of  $O(\sqrt{KT} \log(T))$ . Is this an intrinsic limitation of MS or merely a consequence of a loose analysis?

Our main result below shows that it is indeed a consequence of a loose analysis and that MS enjoys the same regret guarantees as AOUCB.

**Theorem 2.** *MS satisfies that,  $\forall T \geq 1$  and  $c > 0$ ,*

$$\mathbb{E} \text{Reg}_T \leq \sum_{a \in [K]: \Delta_a > 0} \left( \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{\sigma^2}\right)}{\Delta_a} + O\left(\Delta_a \vee \left(\frac{(1+c)^2\sigma^2}{c^2\Delta_a} \ln \frac{(1+c)^2\sigma^2}{c^2\Delta_a^2}\right)\right) \right).$$

Theorem 2 implies that MS is sub-UCB (note  $\ln(T\Delta_a^2) + \ln(\frac{1}{\Delta_a^2}) = \ln(T)$ ). Setting  $c = 1/o(\ln^{1/2}(T))$ , we achieve the asymptotic optimality in the same way described above. Finally, we show the minimax regret bound in the following corollary as a consequence of Theorem 2. Hereafter, all the proofs are in the appendix unless noted otherwise.

**Corollary 3.** *MS satisfies that that,  $\forall T \geq 1$ ,*

$$\mathbb{E} \text{Reg}_T \leq O(\sigma\sqrt{KT \ln(T)}).$$

#### 4.1 Proof of Theorem 2

Let  $T \geq K$ . Let  $N_{t,a}$  be the number of times arm  $a$  has been pulled up to (and including) time step  $t$ . Using  $\mathbb{E} \text{Reg}_T = \sum_{a: \Delta_a > 0} \Delta_a \mathbb{E}[N_{T,a}]$ , it suffices to bound  $\mathbb{E}[N_{T,a}]$ . Fix  $a \in [K]$  such that  $\Delta_a > 0$  and define

$$u := \left\lceil \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{2\sigma^2}\right)}{\Delta_a^2} \right\rceil.$$

Then,

$$\begin{aligned} \mathbb{E}[N_{T,a}] &\leq u + \mathbb{E}\left[\sum_{t=u+1}^T \mathbb{1}\{I_t = a \text{ and } N_{t-1,a} > u\}\right] \\ &= u + \mathbb{E}\left[\sum_{t=u}^{T-1} \mathbb{1}\{I_{t+1} = a \text{ and } N_{t,a} > u\}\right]. \end{aligned}$$

Let  $\hat{\mu}_{t,\max} = \max_a \hat{\mu}_{t,a}$ . Let  $\varepsilon > 0$  be a variable to be tuned later. Recall that  $\Delta_a = \mu_1 - \mu_a$  and  $\widehat{\Delta}_{t,a} = \hat{\mu}_{t,\max} - \hat{\mu}_{t,a}$ . We further split the second term as

follows:

$$\begin{aligned} &\mathbb{E}\left[\sum_{t=u}^{T-1} \mathbb{1}\{I_{t+1} = a \text{ and } N_{t,a} > u\}\right] \\ &= \mathbb{E}\sum_{t=u}^{T-1} \mathbb{1}\left\{I_{t+1} = a, N_{t,a} > u, \hat{\mu}_{t,\max} \geq \mu_1 - \varepsilon, \widehat{\Delta}_{t,a} > \frac{\Delta_a}{1+c}\right\} \\ &+ \mathbb{E}\sum_{t=u}^{T-1} \mathbb{1}\left\{I_{t+1} = a, N_{t,a} > u, \hat{\mu}_{t,\max} \geq \mu_1 - \varepsilon, \widehat{\Delta}_{t,a} \leq \frac{\Delta_a}{1+c}\right\} \\ &+ \mathbb{E}\sum_{t=u}^{T-1} \mathbb{1}\left\{I_{t+1} = a, N_{t,a} > u, \hat{\mu}_{t,\max} < \mu_1 - \varepsilon\right\}. \end{aligned}$$

We call the three terms above (F1), (F2), and (F3) respectively.

The term (F1) is a desirable event where both the empirically best mean is close to  $\mu_1$  and the gap estimation of arm  $a$  is good enough. Then,

$$\begin{aligned} &\text{(F1)} \\ &\leq \sum_{t=u}^{T-1} \mathbb{P}\left(A_{t+1} = a \mid N_{t,a} > u, \widehat{\Delta}_{t,a} > \frac{\Delta_a}{1+c}\right) \\ &\quad \times \mathbb{P}\left(N_{t,a} > u, \widehat{\Delta}_{t,a} > \frac{\Delta_a}{1+c}\right) \\ &\leq \sum_{t=u}^{T-1} \exp\left(-\frac{1}{2\sigma^2} u \left(\frac{\Delta_a}{1+c}\right)^2\right) \cdot 1 \leq \frac{2\sigma^2}{\Delta_a^2}. \end{aligned}$$

For the term (F2), it is easy to see that

$$\begin{aligned} &\{\hat{\mu}_{t,\max} \geq \mu_1 - \varepsilon\} \cap \left\{\widehat{\Delta}_{t,a} \leq \frac{1}{1+c} \Delta_a\right\} \\ &\subseteq \left\{\hat{\mu}_{t,a} - \mu_a \geq \frac{c}{1+c} \Delta_a - \varepsilon\right\}. \end{aligned}$$

Let  $\tau_a(k)$  be the time step  $t$  such that the arm  $a$  was pulled at  $t$  and the number of arm pulls of arm  $a$  becomes  $k$  at the end of the time step  $t$ ; i.e.,  $\tau_a(k) = \min\{t \geq 1 : N_{t,a} = k\}$ . We use the shortcut  $\hat{\mu}_{a,(k)} := \hat{\mu}_{\tau_a(k),a}$  to denote the sample mean of arm  $a$  after pulling it for the  $k$ -th time, . Thus,

$$\begin{aligned} &\mathbb{E}\sum_{t=u}^{T-1} \mathbb{1}\left\{I_{t+1} = a, N_{t,a} > u, \hat{\mu}_{t,\max} \geq \mu_1 - \varepsilon, \widehat{\Delta}_{t,a} \leq \frac{\Delta_a}{1+c}\right\} \\ &\leq \sum_{t=u}^{T-1} \mathbb{P}\left(I_{t+1} = a, \hat{\mu}_{t,a} - \mu_a \geq \frac{c\Delta_a}{1+c} - \varepsilon\right) \\ &\leq \mathbb{E}\left[\sum_{k=1}^{\infty} \sum_{t=\tau_1(k)}^{\tau_1(k+1)-1} \mathbb{1}\{I_{t+1} = a\} \cdot \mathbb{1}\left\{\hat{\mu}_{t,a} - \mu_a \geq \frac{c\Delta_a}{1+c} - \varepsilon\right\}\right] \\ &= \mathbb{E}\left[\sum_{k=1}^{\infty} \mathbb{1}\left\{\hat{\mu}_{a,(k)} - \mu_a \geq \frac{c\Delta_a}{1+c} - \varepsilon\right\} \sum_{t=\tau_1(k)}^{\tau_1(k+1)-1} \mathbb{1}\{I_{t+1} = a\}\right] \\ &= \mathbb{E}\left[\sum_{k=1}^{\infty} \mathbb{1}\left\{\hat{\mu}_{a,(k)} - \mu_a \geq \frac{c\Delta_a}{1+c} - \varepsilon\right\}\right] \\ &= \sum_{k \geq 1} \mathbb{P}\left(\hat{\mu}_{a,(k)} - \mu_a \geq \frac{c\Delta_a}{1+c} - \varepsilon\right) \quad (u \geq k) \end{aligned}$$

$$\begin{aligned} &\leq \sum_{k \geq 1} \exp\left(-\frac{k}{2\sigma^2} \left(\frac{c\Delta_a}{1+c} - \varepsilon\right)^2\right) \\ &= \frac{1}{\exp\left(\frac{1}{2\sigma^2} \left(\frac{c\Delta_a}{1+c} - \varepsilon\right)^2\right) - 1} \leq \frac{2\sigma^2}{\left(\frac{c\Delta_a}{1+c} - \varepsilon\right)^2}. \end{aligned}$$

The term (F3) deals with the case where the best empirical mean is much lower than the true best mean  $\mu^*$ . This part of the proof is the most nontrivial part where we essentially show that there exists a sufficient amount of probability assigned to the true best arm so that eventually we can recover from the case where the empirical mean of the best arm is actually not empirically the best. Let  $\mathcal{F}_t$  be the  $\sigma$ -algebra generated by  $(A_1, R_1, \dots, A_t, R_t)$ . We define the shortcuts  $\mathbb{P}_t(\mathcal{E}) := \mathbb{P}(\mathcal{E} | \mathcal{F}_t)$  and  $\mathbb{E}_t[X] := \mathbb{E}[X | \mathcal{F}_t]$ .

We first relate the probability of choosing the arm  $a$  with that of the true best arm. If  $\hat{\mu}_{t,\max} < \mu_1 - \varepsilon$ , then

$$\begin{aligned} &\mathbb{P}_t(I_{t+1} = a) \\ &= \frac{\exp(-\frac{1}{2\sigma^2} N_{t,a} \cdot \widehat{\Delta}_{t,a}^2)}{\exp(-\frac{1}{2\sigma^2} N_{t,1} \cdot \widehat{\Delta}_{t,1}^2)} \mathbb{P}_t(I_{t+1} = 1) \\ &\leq \frac{1}{\exp(-\frac{1}{2\sigma^2} N_{t,1} \cdot (\mu_1 - \varepsilon - \hat{\mu}_{t,1})_+^2)} \mathbb{P}_t(I_{t+1} = 1). \end{aligned} \quad (3)$$

Let  $\tau_1(k)$  be the time step after which the number of arm pulls of just becomes  $k$ ; i.e.,  $\tau_1(k) = \min\{t \geq 1 : N_{t,1} = k\}$ . Let us use the shortcut  $\hat{\mu}_{1,(k)} := \hat{\mu}_{\tau_1(k),1}$ . Define  $B_t := \mathbb{1}\{\hat{\mu}_{t,\max} < \mu_1 - \varepsilon\}$ , which is at most  $B_{t,1} = \mathbb{1}\{\hat{\mu}_{t,1} < \mu_1 - \varepsilon\}$ . Let  $A_{t+1,i} = \mathbb{1}\{I_{t+1} = i\}$ ,  $\forall i \in [K]$ . With this, we can bound (F3) as follows:

$$\begin{aligned} &\mathbb{E} \sum_{t=u}^{T-1} \mathbb{1}\{I_{t+1} = a, N_{t,a} > u, \hat{\mu}_{t,\max} < \mu_1 - \varepsilon\} \\ &\leq \mathbb{E} \left[ \sum_{t=u}^{T-1} \mathbb{E}_t [B_t A_{t+1,a}] \right] \\ &= \mathbb{E} \left[ \sum_{t=u}^{T-1} B_t \mathbb{E}_t [A_{t+1,a}] \right] \\ &\leq \mathbb{E} \left[ \sum_{t=u}^{T-1} B_t \exp\left(\frac{N_{t,1}(\mu_1 - \varepsilon - \hat{\mu}_{t,1})^2}{2\sigma^2}\right) \mathbb{E}_t [A_{t+1,1}] \right] \\ &\quad \text{(by (3))} \\ &= \mathbb{E} \left[ \sum_{t=u}^{T-1} \mathbb{E}_t \left[ B_t \exp\left(\frac{N_{t,1}(\mu_1 - \varepsilon - \hat{\mu}_{t,1})^2}{2\sigma^2}\right) A_{t+1,1} \right] \right] \\ &= \mathbb{E} \left[ \sum_{t=u}^{T-1} B_t \exp\left(\frac{N_{t,1}(\mu_1 - \varepsilon - \hat{\mu}_{t,1})^2}{2\sigma^2}\right) \cdot A_{t+1,1} \right] \\ &\leq \mathbb{E} \left[ \sum_{t=u}^{T-1} B_{t,1} \exp\left(\frac{N_{t,1}(\mu_1 - \varepsilon - \hat{\mu}_{t,1})^2}{2\sigma^2}\right) \cdot A_{t+1,1} \right] \\ &= \mathbb{E} \left[ \sum_{k=1}^{\infty} \sum_{t=\tau_1(k)+1}^{\tau_1(k+1)} B_{t,1} \exp\left(\frac{N_{t,1}(\mu_1 - \varepsilon - \hat{\mu}_{1,(k)})^2}{2\sigma^2}\right) A_{t+1,1} \right]. \end{aligned}$$

Using the trivial fact that  $\sum_{t=\tau_1(k)+1}^{\tau_1(k+1)} A_{t+1,1} = 1$  and with a new notation  $\bar{X}_k := \mu_1 - \hat{\mu}_{1,(k)}$ , we have that

$$(F3) \leq \mathbb{E} \left[ \sum_{k=1}^{\infty} \mathbb{1}\{\bar{X}_k > \varepsilon\} \cdot \exp\left(\frac{1}{2\sigma^2} k \cdot (\bar{X}_k - \varepsilon)^2\right) \right].$$

To evaluate this integral, we perform peeling to upper bound the above by

$$\begin{aligned} &\mathbb{E} \left[ \sum_{k=1}^{\infty} \sum_{q=1}^{\infty} \mathbb{1}\left\{ \mu_1 - \hat{\mu}_{1,(k)} \in \left( \varepsilon + \frac{q-1}{2}\varepsilon, \varepsilon + \frac{q}{2}\varepsilon \right) \right\} \right. \\ &\quad \left. \times \exp\left(\frac{1}{2\sigma^2} k \cdot \left(\frac{q\varepsilon}{2}\right)^2\right) \right] \\ &= \sum_{k=1}^{\infty} \sum_{q=1}^{\infty} \mathbb{P} \left( \mu_1 - \hat{\mu}_{1,(k)} \in \left( \varepsilon + \frac{q-1}{2}\varepsilon, \varepsilon + \frac{q}{2}\varepsilon \right) \right) \\ &\quad \times \exp\left(\frac{1}{2\sigma^2} k \cdot \left(\frac{q\varepsilon}{2}\right)^2\right) \\ &\leq \sum_{k=1}^{\infty} \sum_{q=1}^{\infty} \exp\left(-\frac{1}{2\sigma^2} k \left(\frac{q+1}{2}\varepsilon\right)^2\right) \cdot \exp\left(\frac{1}{2\sigma^2} k \cdot \left(\frac{q\varepsilon}{2}\right)^2\right) \\ &= \sum_{k=1}^{\infty} \sum_{q=1}^{\infty} \exp\left(-\frac{1}{4\sigma^2} k \varepsilon^2 \left(q + \frac{1}{2}\right)\right) \\ &= \sum_{q=1}^{\infty} \frac{1}{\exp\left(\frac{1}{4\sigma^2} \left(q + \frac{1}{2}\right) \varepsilon^2\right) - 1} \\ &\stackrel{(a)}{\leq} \frac{4\sigma^2}{\varepsilon^2} \sum_{q=1}^{\bar{q}-1} \frac{1}{q + \frac{1}{2}} + \sum_{q=\bar{q}}^{\infty} \frac{1}{\exp\left(\frac{1}{4\sigma^2} \left(q + \frac{1}{2}\right) \varepsilon^2\right) - 1} \end{aligned}$$

where (a) introduces a free variable  $\bar{q} \in \{1, 2, \dots\}$  and uses  $e^z \geq z + 1$ . Note that as long as  $\bar{q} = \Theta(1)$ , this is tight because  $X \in (0, 1)$  then  $1 + X \leq e^X \leq (e-1)X + 1$ . It remains to bound the above and tune  $\bar{q}$ . Leaving the algebra to the appendix, we obtain

$$(F3) = O\left(\frac{\sigma^2}{\varepsilon^2} (1 + \ln(1 + \frac{\sigma^2}{\varepsilon^2}))\right).$$

Altogether, choosing  $\varepsilon = \frac{c\Delta_a}{2(1+c)}$ , we have that

$$\begin{aligned} \mathbb{E}[N_{T,a}] &\leq u + (F1) + (F2) + (F3) \\ &\leq u + O\left(\frac{(1+c)^2\sigma^2}{c^2\Delta_a^2} (1 + \ln(1 + \frac{(1+c)^2\sigma^2}{c^2\Delta_a^2}))\right), \end{aligned}$$

which concludes the proof.

## 5 MAILLARD SAMPLING PLUS (MS<sup>+</sup>)

We now propose an improved version of MS described in Algorithm 2, which we call MS<sup>+</sup>. MS<sup>+</sup> differs from MS in two aspects. First, the probability of being sampled is inflated by a small amount, and the empirical best arms are also inflated but in a slightly different way. This helps achieving the minimax ratio of  $\sqrt{\log(K)}$ , which is better than  $\sqrt{\log(T)}$  of MS. Second, we have

---

**Algorithm 2** Maillard Sampling<sup>+</sup> (MS<sup>+</sup>)
 

---

**Input:**  $B \geq 1, 0 \leq C \leq 1, 0 < D \leq 1, \sigma^2 > 0$ 
**for**  $t = 1, 2, \dots$  **do**

   **if**  $t \leq K$  **then**

     Pull the arm  $I_t = t$  and observe reward  $y_t$ .
 
   **else**

     For all  $a \in [K]$  with  $\hat{\mu}_{t-1,a} = \max_b \hat{\mu}_{t-1,b}$ ,
 

$$\hat{p}_{t,a} = B \cdot (1 + C \ln(1 + \ln(\frac{t}{N_{t-1,a}}))) .$$

For other arms,

$$\begin{aligned} \hat{p}_{t,a} = & \exp\left(-\frac{1}{2\sigma^2} N_{t-1,a} \widehat{\Delta}_{a,t-1}^2\right) \\ & + \ln\left(1 + \frac{D}{2\sigma^2} N_{t-1,a} \widehat{\Delta}_{a,t-1}^2\right) . \end{aligned}$$

     Compute  $p_{t,a} = \frac{\hat{p}_{t,a}}{\sum_b \hat{p}_{t,b}}$ .
 
     Pull the arm  $I_t \sim p_t$  and observe reward  $y_t$ .
 
   **end if**
**end for**


---

introduced a number of parameters  $B$ ,  $C$ , and  $D$ . We remark that while the parameters  $C$  and  $D$  are less important in practice, which we set to  $C = D = 0.01$  in our experiments later, the parameter  $B$ , which we call the ‘booster’, affects the performance a lot. We recommend  $B = 4$  to be safe, but depending on the situation one may want to set it to  $B = 8$  or even a higher value; see our experiments in Section 6 for more discussion.

MS<sup>+</sup> enjoys the following regret bound.

**Theorem 4.** *MS<sup>+</sup> (Algorithm 2) satisfies that, for every  $T \geq 1$  and  $c > 0$ ,*

$$\begin{aligned} \mathbb{E} \text{Reg}_T \leq & \sum_{a: \Delta_a > 0} \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{2\sigma^2} (1 + \ln(1 + \frac{T\Delta_a^2}{2\sigma^2}))\right)}{\Delta_a} \\ & + O\left(B \frac{\sigma^2(1+c)^2 \ln(1 + \ln(\frac{T\Delta_a^2}{2\sigma^2}))}{c^2 \Delta_a} + \Delta_a\right) \end{aligned}$$

where we omit the dependence on  $C$  and  $D$  for brevity.

*Proof.* The key difference from the proof of MS is the term (F3) therein that resulted in  $O(\frac{\sigma^2}{\Delta^2} \ln(\frac{\sigma^2}{\Delta^2}))$ . Inspecting how we obtain the minimax regret bound, one can see that we would obtain the minimax ratio of  $\sqrt{\log(K)}$  if we could bound (F3) by  $O(\frac{\sigma^2}{\Delta^2})$ . Indeed, this was the driving force behind the design of MS<sup>+</sup>. The inflated sampling probability helps us to obtain such a bound. We refer to our appendix for the full proof.  $\square$

Just like MS, MS<sup>+</sup> is sub-UCB, one can obtain an asymptotic optimality by setting  $c = 1/o(\ln^{1/2}(T))$ . One can see that the dependence on  $B$  is on the lower order term only. This means that however large  $B$  is, for large enough  $T$ , the leading term will dominate the regret.

The following corollary shows that we can obtain a minimax bound of  $\sqrt{KT \log(K)}$ , which is an improvement of MS.

**Corollary 5.** *MS<sup>+</sup> with absolute constants  $B, C$ , and  $D$  satisfies that,  $\forall T \geq 1$ ,*

$$\mathbb{E} \text{Reg}_T \leq O(\sigma \sqrt{KT \ln(K)}) .$$

We remark that MS<sup>+</sup> enjoys the same minimax bound with  $C = 0$ , and the parameter  $C$  can be viewed as another way to encourage exploitation that goes away with a large enough  $t$ , which has a milder effect compared to  $B$ .

## 6 EXPERIMENTS

In this section, we numerically investigate the effects of the parameter  $B$  of MSP and evaluate the performance of MSP with existing bandit algorithms for sub-Gaussian rewards. Throughout, we set  $C = D = 0.01$  for MSP, which seems to work well overall.

**The effect of the booster parameter  $B$ .** We consider a Gaussian arm set with variance 1. There are 10 arms with mean rewards that decrease linearly:  $\{0.9, 0.8, \dots, 0.0\}$ . We run MS<sup>+</sup> 200 times with  $B \in \{2, 4, 8, \dots, 128\}$ . Figure 1 plots the average regret and its standard deviation. The average regret goes down up to  $B = 16$  and then starts to increase. The variance starts to increase from  $B = 8$  and then soon increases dramatically. This shows the tradeoff between the expected regret and the variance of the regret. Note, however, that even with a very aggressive booster  $\log_2(B) = 7$ , the average regret is not dramatically higher than the best and certainly far from suffering a linear regret. This is not surprising because MSP still enjoys a valid regret bound; in fact, even with  $B = 128$ , MSP is asymptotically optimal. Such aggressive exploitation of the empirical best arm without losing theoretical guarantee (and without linear regret) is the unique characteristic of MSP that is not found in other existing algorithms.

**Methods.** Since there are so many bandit algorithms for various reward settings, we limit the set of baseline methods to those that have a regret guarantee for sub-Gaussian rewards. For MSP, we use  $B \in \{4, 8, 16\}$ , which we denote by MS<sup>+</sup>4, MS<sup>+</sup>8, and MS<sup>+</sup>16 respectively. We consider the following baseline methods:

- TS: Thompson sampling for Gaussian rewards (Ko-

Noise	Gaps	$K$	$\mu_{1:K}$
Gaussian	Linear	$\in \{10, 100\}$	$\mu_i = 1 - \frac{i}{K}$
Gaussian	Equal	$\in \{2, 10, 100\}$	$\mu_1 = 1, \mu_{2:K} = .5$
Bernoulli	Linear	$\in \{10, 100\}$	$\mu_i = 1 - \frac{i}{K}$
Bernoulli	Equal	$\in \{2, 10, 100\}$	$\mu_1 = 0.1, \mu_{2:K} = .05$

Table 2: Synthetic bandit problems for evaluation.

rda et al., 2013). This is the only baseline that does not have a regret guarantee for sub-Gaussian rewards, which we include as a comparison point.

- TS-SG: Gaussian Thompson sampling for bounded rewards by Agrawal and Goyal (2017) that is (trivially) adapted for sub-Gaussian rewards. The theoretical guarantees for bounded rewards transfer to the sub-Gaussian rewards naturally. The key difference is that this algorithm inflates the posterior variance by a factor of 4 for technical reasons.
- AOUCB: Asymptotically optimal UCB (Lattimore and Szepesvári, 2020).
- UCB<sup>+</sup>: Almost optimally confidence UCB (a.k.a. UCB plus) (Lattimore, 2015). This is a fixed-budget algorithm (more comments below).
- OCUCB: Optimally confidence UCB (Lattimore, 2015), a fixed-budget algorithm.
- MOSS: MOSS (Audibert et al., 2009), a fixed-budget algorithm.

Note that the comparison will not be entirely fair since the fixed-budget algorithms use the extra information of the time horizon  $T$  to adjust the amount of exploration. Note that it should be possible to adjust UCB<sup>+</sup> and MOSS to not require  $T$  as input such as Ménard and Garivier (2017), but this will inflate the confidence width, likely increasing regret in practice. On the other hand, it is not clear to us how to make OCUCB anytime. All the algorithms were implemented so it takes in the sub-Gaussian parameter  $\sigma^2 = 1/4$ .

**Synthetic bandit problems.** We test the algorithms above with the bandit problems summarized in Table 2. For Bernoulli problems, our test algorithms with the sub-Gaussian parameter of  $\sigma^2 = 1/4$ , which is the tightest one without losing their regret guarantees. The intention of running Bernoulli problems is not to claim that the tested algorithms will excel at them but to observe how they behave under non-Gaussian rewards.

**Results.** We run evaluation of the algorithms on every bandit problem in Table 2 with  $T = 20,000$  and 200 trials and report the resulting average regrets and its standard deviations in Figure 2 and 3. Note that the number of trials is high enough to make the average regret a good estimate of the expected regret except for a few cases with a very large variance. For Gaussian environments, the overall winner is OCUCB, which is




Figure 1: The effect of tuning  $B$  for MSP. The error bars are standard deviation over 200 trials.

not surprising given that its regret bounds are much better than the rest and that it is a fixed-budget algorithm. Among anytime algorithms, MS<sup>+</sup>8, MS<sup>+</sup>16, and TS are all comparable with respect to the average regret while MS<sup>+</sup>8 and MS<sup>+</sup>16 often come with a larger variance. This is still an encouraging result against TS since TS is optimized for Gaussian rewards; we expect that optimizing MS<sup>+</sup> for Gaussian rewards may make MS<sup>+</sup> be comparable to TS in both expected regret and the variance. For the remaining methods, MS<sup>+</sup>4 has smaller variances than other MS<sup>+</sup>'s, but its average regrets are mostly behind them, and both TS-SG and AOUCB are much worse than the rest. We remark that while the variance of MS<sup>+</sup>16 seems detrimental in Figure 2(a), it is still far from suffering a linear regret; the maximum regret of MS<sup>+</sup>16 over 200 trials was 450.5 while the largest possible regret of any algorithm is  $20,000 \cdot 0.5 = 10,000$ .

For Bernoulli instances, we still observe the same trend except for minor differences: (i) the relative performance of TS is not as good as Gaussian environments due to the non-Gaussian nature of the rewards, and (ii) MS<sup>+</sup>16 consistently has the smallest average regret among anytime algorithms. The second observation is due to the fact that the actual variances are small. This confirms the benefit of MS<sup>+</sup>. When we have domain knowledge or experience from similar tasks that the variances tend to be small, we can tune the booster parameter to enjoy a better performance without worrying about suffering a linear regret when such information is not true. In contrast, existing tuning techniques of UCB or TS (Zhang et al., 2016) break their regret guarantees and risk suffering from a linear regret.

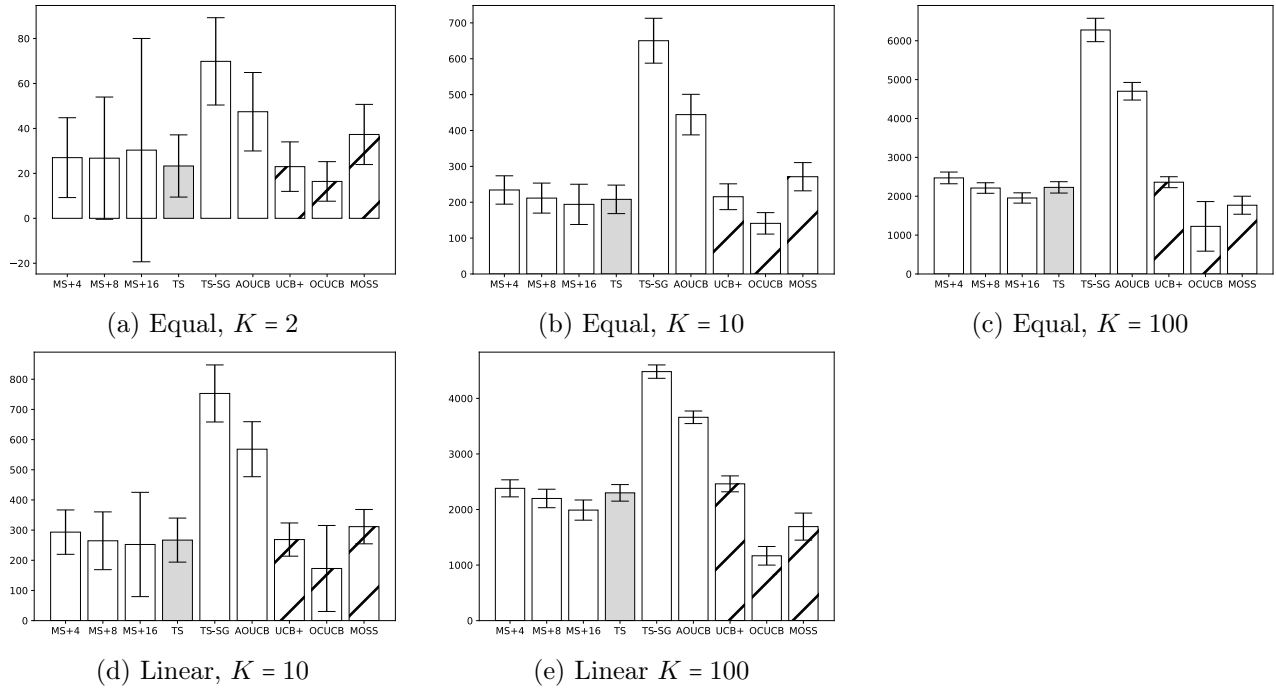


Figure 2: Experimental results for Gaussian environments from Table 2 with  $T = 20,000$  and 200 trials. We show the average regrets as bars and the standard deviations with error bars. The shaded bar means that the method does not have a regret guarantee for sub-Gaussian rewards. The hatched bar means that the method is fixed-budget, which requires an extra information of  $T$ , making the comparison not exactly fair.

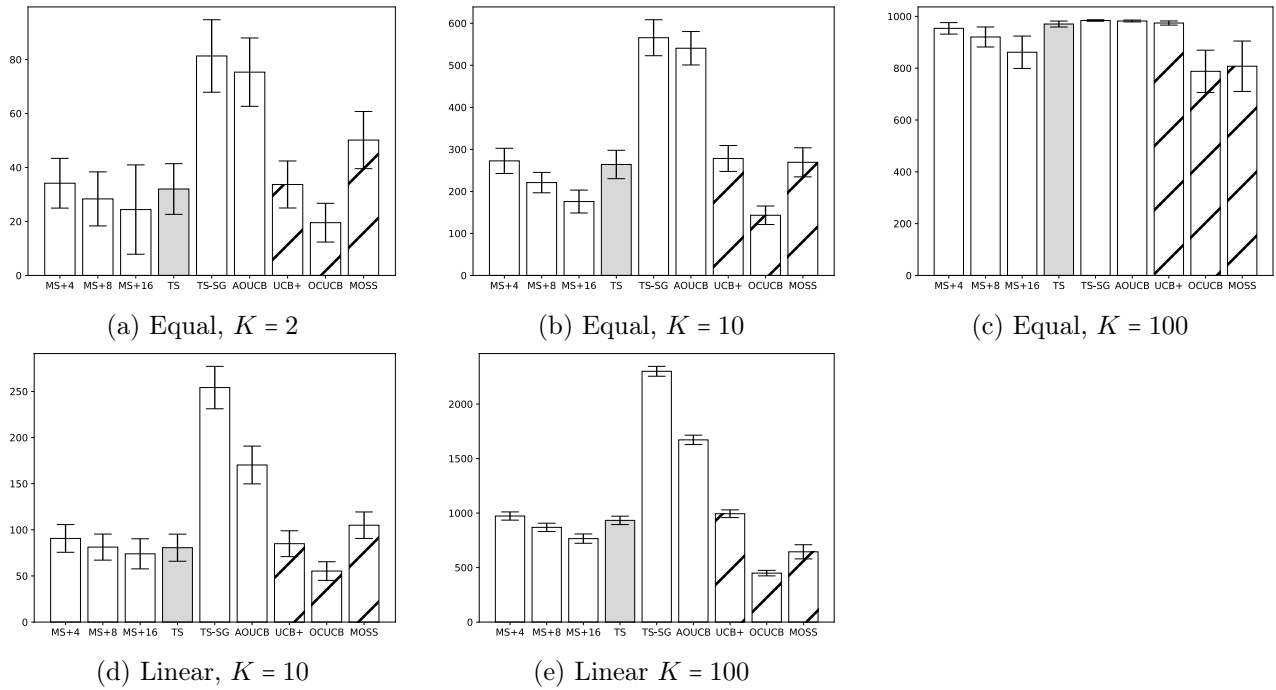


Figure 3: Experimental results for Bernoulli environments from Table 2.



## 7 CONCLUSION

We have revisited Maillard sampling, improved its regret bound, and proposed a variant that enjoys a better minimax regret bound. As Boltzmann exploration is popular in reinforcement learning for its simplicity yet is provably suboptimal (Cesa-Bianchi et al., 2017), we highly recommend that the practitioners try the following simplified version of MS as a replacement for Boltzmann exploration:

$$p_{t,a} \propto B^{\mathbb{1}\{a \in \arg \max_b \hat{\mu}_{t-1,b}\}} \exp\left(-\frac{\hat{\Delta}_{t-1,a}^2}{2\sigma^2}\right) \quad (4)$$

where both  $B \geq 1$  is to be tuned and  $\sigma^2$  can also be tuned when the sub-Gaussian parameter is unknown. This version is simple enough to cause no friction for a quick implementation yet enjoys good guarantees in bandit problems.

Our work opens up numerous exciting research directions. First, we can further investigate MS or even improve its mathematical properties. For example, while MS<sup>+</sup> allowed us to achieve the minimax ratio  $\sqrt{\log(K)}$ , it is not clear whether MS can achieve the same guarantee without further modifications. The same issue happens with Thompson sampling – the minimax ratio of  $\sqrt{\log(K)}$  of Thompson sampling is achieved only after inflating the posterior probability (Agrawal and Goyal, 2017, Theorem 1.3). Also, we conjecture that MS is better than UCB under the delayed-reward setting since this is known to be true for randomized algorithms such as Thompson sampling (Chapelle and Li, 2011); see also Dudík et al. (2011, Section 6). It would be interesting to verify the conjecture theoretically and empirically. Second, It would be interesting to further improve MS to match the regret guarantees of AdaUCB, which has the best-known regret bound in the literature. Third, it would be interesting to generalize MS to the exponential family of distributions or to the linear bandit setting. A natural candidate for the former is to take the MED algorithm (Honda and Takemura, 2011), which is a generalization of MS as mentioned in related work, and perform finite time analyses. Finally, while MS can be tuned to enjoy small expected regret, it tends to have a large variance. It would be interesting to study the fundamental tradeoff between the expected pseudo regret and its variance.

### Acknowledgements

The authors would like to thank Tor Lattimore and Junya Honda for their valuable feedback and Nikos Karampatziakis for sharing Maillard sampling in social media.

### References

- S. Agrawal and N. Goyal. Further optimal regret bounds for thompson sampling. In *Artificial intelligence and statistics*, pages 99–107, 2013.
- S. Agrawal and N. Goyal. Near-Optimal Regret Bounds for Thompson Sampling. *Journal of the ACM*, 64(5): 1–24, 2017.
- J.-Y. Audibert, S. Bubeck, and Others. Minimax Policies for Adversarial and Stochastic Bandits. In *Proceedings of the Conference on Learning Theory (COLT)*, 2009.
- P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning*, 47(2–3):235–256, 2002.
- A. N. Burnetas and M. N. Katehakis. Optimal adaptive policies for sequential allocation problems. *Advances in Applied Mathematics*, 17(2):122–142, 1996.
- O. Cappé, A. Garivier, O.-A. Maillard, R. Munos, and G. Stoltz. Kullback-Leibler upper confidence bounds for optimal sequential allocation. *The Annals of Statistics*, pages 1516–1541, 2013.
- N. Cesa-Bianchi, C. Gentile, G. Lugosi, and G. Neu. Boltzmann exploration done right. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- O. Chapelle and L. Li. An Empirical Evaluation of Thompson Sampling. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2249–2257, 2011.
- M. Dudík, D. J. Hsu, S. Kale, N. Karampatziakis, J. Langford, L. Reyzin, and T. Zhang. Efficient Optimal Learning for Contextual Bandits. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 169–178, 2011.
- J. Honda and A. Takemura. An Asymptotically Optimal Bandit Algorithm for Bounded Support Models. In *COLT*, pages 67–79. Citeseer, 2010.
- J. Honda and A. Takemura. An asymptotically optimal policy for finite support models in the multiarmed bandit problem. *Machine Learning*, 85(3):361–391, 2011.
- J. Honda and A. Takemura. Non-asymptotic analysis of a new bandit algorithm for semi-bounded rewards. *J. Mach. Learn. Res.*, 16:3721–3756, 2015.
- D. G. Horvitz and D. J. Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260):663–685, 1952.
- L. P. Kaelbling, M. L. Littman, and A. W. Moore. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285, 1996.

- E. Kaufmann, N. Korda, and R. Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In *Proceedings of the international conference on Algorithmic Learning Theory (ALT)*, pages 199–213, 2012.
- N. Korda, E. Kaufmann, and R. Munos. Thompson Sampling for 1-Dimensional Exponential Family Bandits. In *Advances in Neural Information Processing Systems*, volume 26, 2013.
- B. Kveton, C. Szepesvari, M. Ghavamzadeh, and C. Boutilier. Perturbed-history exploration in stochastic multi-armed bandits. *arXiv preprint arXiv:1902.10089*, 2019a.
- B. Kveton, C. Szepesvari, S. Vaswani, Z. Wen, T. Lattimore, and M. Ghavamzadeh. Garbage in, reward out: Bootstrapping exploration in multi-armed bandits. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 3601–3610. PMLR, 2019b.
- T. L. Lai and H. Robbins. Asymptotically Efficient Adaptive Allocation Rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- T. Lattimore. Optimally confident UCB: Improved regret for finite-armed bandits. *arXiv preprint arXiv:1507.07880v3*, 2015.
- T. Lattimore. Refining the Confidence Level for Optimistic Bandit Strategies. *Journal of Machine Learning Research*, 19(20):1–32, 2018.
- T. Lattimore and C. Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020.
- L. Li, W. Chu, J. Langford, and R. E. Schapire. A Contextual-Bandit Approach to Personalized News Article Recommendation. *Proceedings of the International Conference on World Wide Web (WWW)*, pages 661–670, 2010.
- L. Li, W. Chu, J. Langford, and X. Wang. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *Proceedings of the ACM international conference on Web search and data mining*, pages 297–306, 2011.
- L. Li, R. Munos, and C. Szepesvári. Toward minimax off-policy value estimation. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 608–616, 2015.
- A. Locatelli, M. Gutzeit, and A. Carpentier. An optimal algorithm for the Thresholding Bandit Problem. In *Proceedings of The International Conference on Machine Learning (ICML)*, pages 1690–1698, 2016.
- O.-A. Maillard. *APPRENTISSAGE SÉQUENTIEL: Bandits, Statistique et Renforcement*. PhD thesis, Université des Sciences et Technologie de Lille-Lille I, 2013.
- P. Ménard and A. Garivier. A minimax and asymptotically optimal algorithm for stochastic bandits. In *Proceedings of the international conference on Algorithmic Learning Theory (ALT)*, pages 223–237, 2017.
- D. Precup. Eligibility traces for off-policy policy evaluation. *Proceedings of The International Conference on Machine Learning (ICML)*, pages 759–766, 2000.
- J. M. Robins and A. Rotnitzky. Semiparametric efficiency in multivariate regression models with missing data. *Journal of the American Statistical Association*, 90(429):122–129, 1995.
- R. S. Sutton. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 216–224. Elsevier, 1990.
- W. R. Thompson. On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples. *Biometrika*, 25(3/4):285, 1933.
- L. Zhang, T. Yang, R. Jin, Y. Xiao, and Z.-h. Zhou. Online Stochastic Linear Optimization under One-bit Feedback. In *Proceedings of the International Conference on Machine Learning (ICML)*, volume 48, pages 392–401, 2016.

---

# Supplementary Material: Maillard Sampling: Boltzmann Exploration Done Optimally

---

## Table of Contents

---

<b>A Related Algorithms</b>	<b>11</b>
<b>B Proofs for Maillard Sampling</b>	<b>12</b>
B.1 Proof of Theorem 2	12
B.2 Proof of Corollary 3	13
<b>C Proofs for Maillard Sampling<sup>+</sup></b>	<b>14</b>
C.1 Proof of Theorem 4	14
C.2 Proof of Corollary 5	18

---

## A Related Algorithms

Beside Boltzmann exploration, there are a number of algorithms that shares a similar shape. Among deterministic algorithms, IMED is the most similar one. While IMED only has regret guarantees when rewards have (semi-)bounded support, IMED for the case of Gaussian that is reported in [Lattimore \(2018\)](#) has a striking similarity to MS:

$$I_t = \arg \min_{i \in [K]} \frac{N_{t-1,i}}{2} \widehat{\Delta}_{t-1,i}^2 + \log(N_{t-1,i}) .$$

Except for the second term, the first term is exactly the negative of the exponent of MS. In a similar vein, [Locatelli et al. \(2016\)](#) propose an algorithm called APT for thresholding bandits. While their problem is a pure exploration, their algorithm selects an arm by

$$I_t = \arg \min_{i \in [K]} \sqrt{N_{t-1,i}} \widehat{\Delta}'_{t-1,i}$$

where  $\widehat{\Delta}'_{t-1,i} = |\widehat{\mu}_{t-1,i} - \tau| + \varepsilon$  where both  $\tau$  and  $\varepsilon$  are given from the problem setup. Except for the slight differences in the definition of  $\widehat{\Delta}'_{t-1,i}$ , APT's index can be interpreted as the negative of MS's exponent. Note, however, their estimated gap  $\widehat{\Delta}'_{t-1,i}$  is relative to a given problem parameter since the target reward level is given in thresholding bandit problems. Thus, proof techniques from APT do not help analyzing MS that involves the gap based on the empirical best arm. On the other hand, problem-dependent regret bounds necessarily include gaps  $\Delta_i = \max_{j \in [K]} \mu_j - \mu_i$ , which means that design of algorithms would be easier if we knew  $\Delta_i$  since we can try to sample just enough number of times required by the lower bound. Our results on MS reveals that one can directly use  $\widehat{\Delta}_i$  in the algorithm and perform a tight finite-time analysis of it. We speculate that our regret analysis scheme can be useful for other similar bandit problems such as pure exploration.

Among many randomized stochastic bandit algorithms ([Kveton et al., 2019b;a](#); [Thompson, 1933](#); [Agrawal and Goyal, 2013](#); [Kaufmann et al., 2012](#); [Korda et al., 2013](#)), Thompson sampling (TS) is the most popular one due to its simplicity and practical performance ([Chapelle and Li, 2011](#)). While TS's sampling probabilities do not have a closed-form expression, the special case of  $K = 2$  arms reveals a connection between TS for Gaussian rewards and MS. Suppose arm 1 has a higher mean reward than arm 2, and the reward noise is Gaussian with variance 1. Let us use the Gaussian TS algorithm by [Korda et al. \(2013\)](#) whose posterior samples  $X_1 \sim \mathcal{N}(\widehat{\mu}_1, 1/T_1)$  and  $X_2 \sim \mathcal{N}(\widehat{\mu}_2, 1/T_2)$  where  $T_1$  and  $T_2$  are the pull counts of each arm. Then, the probability of sampling arm 2 is

$\mathbb{P}(Z := X_2 - X_1 \geq 0)$ . One can show that  $Z \sim \mathcal{N}(-\widehat{\Delta} := \hat{\mu}_2 - \hat{\mu}_1, \frac{1}{T_1} + \frac{1}{T_2})$ . We can show that, with  $\omega^2 := \frac{1}{T_1} + \frac{1}{T_2}$ ,

$$\begin{aligned} p_{t,2} &= \mathbb{P}(Z \geq 0) = \mathbb{P}\left(\frac{Z - (-\widehat{\Delta})}{\omega} \geq \frac{\widehat{\Delta}}{\omega}\right) \\ &= 1 - \Phi\left(\frac{\widehat{\Delta}}{\omega}\right) \\ &= \frac{1}{2}\left(1 - \operatorname{erf}\left(\frac{\widehat{\Delta}}{\sqrt{2}\omega}\right)\right) \\ &= \frac{1}{2}\operatorname{erfc}\left(\frac{\widehat{\Delta}}{\sqrt{2}\omega}\right) \\ &\leq \frac{1}{2} \frac{1}{\frac{\widehat{\Delta}}{\omega\sqrt{2}}\sqrt{\pi}} \exp\left(-\frac{\widehat{\Delta}^2}{2\omega^2}\right) \end{aligned}$$

After a large enough iterations, note that  $\omega^2$  will be close to  $1/T_2$  as most arm pulls will be on arm 1. Then, except for the factor outside  $\exp()$  above, the sampling probability almost coincides with MS. Specifically, TS pulls arm 2 with probability  $\frac{1}{2} \exp\left(-\frac{T_2\widehat{\Delta}^2}{2}\right)$  roughly speaking and MS pulls arm 2 with probability like  $\exp\left(-\frac{T_2\widehat{\Delta}^2}{2}\right)/(1 + \exp\left(-\frac{T_2\widehat{\Delta}^2}{2}\right))$ , which is asymptotically  $\exp\left(-\frac{T_2\widehat{\Delta}^2}{2}\right)$ . Note that TS has  $\operatorname{erfc}(x)$  that decays slightly faster than  $\exp(-x^2)$ . While this looks more efficient (i.e., less exploration), such a difference is rooted at the Gaussian reward assumption as opposed to the intrinsic benefit of the algorithmic framework. The MS algorithm presented in our paper is for sub-Gaussian rewards that include a much larger family of distributions, so MS must explore a bit more than Gaussian TS to maintain valid regret guarantees. Designing MS for Gaussian rewards is likely to result in a tighter exploration and possibly a contender to TS.

## B Proofs for Maillard Sampling

### B.1 Proof of Theorem 2

We finish the proof by expanding on the term (F3). Recall that (F3) is bounded by the following:

$$\frac{4\sigma^2}{\varepsilon^2} \sum_{q=1}^{\bar{q}-1} \frac{1}{q + \frac{1}{2}} + \sum_{q=\bar{q}}^{\infty} \frac{1}{\exp\left(\frac{1}{4\sigma^2}\left(q + \frac{1}{2}\right)\varepsilon^2\right) - 1}.$$

For the second term, note that, for  $q \geq \bar{q}$ , we have  $\exp\left(\frac{1}{4\sigma^2}\left(q + \frac{1}{2}\right)\varepsilon^2\right) \geq \exp\left(\frac{1}{4\sigma^2}\left(\bar{q} + \frac{1}{2}\right)\varepsilon^2\right) =: \eta$ . Using the fact that  $X \geq \eta \implies \frac{1}{X-1} \leq \frac{1}{1-\eta^{-1}} \frac{1}{X}$ ,

$$\begin{aligned} \sum_{q=\bar{q}}^{\infty} \frac{1}{\exp\left(\frac{1}{4\sigma^2}\left(q + \frac{1}{2}\right)\varepsilon^2\right) - 1} &= \frac{1}{1 - \eta^{-1}} \sum_{q=\bar{q}}^{\infty} \frac{1}{\exp\left(\frac{1}{4\sigma^2}\left(q + \frac{1}{2}\right)\varepsilon^2\right)} = \frac{1}{1 - \eta^{-1}} \cdot \frac{\exp\left(-\frac{1}{4\sigma^2}\left(\bar{q} + \frac{1}{2}\right)\varepsilon^2\right)}{1 - \exp\left(-\frac{1}{4\sigma^2}\varepsilon^2\right)} \\ &= \frac{1}{1 - \eta^{-1}} \cdot \frac{\exp\left(-\frac{1}{4\sigma^2}\left(\bar{q} - \frac{1}{2}\right)\varepsilon^2\right)}{\exp\left(\frac{1}{4\sigma^2}\varepsilon^2\right) - 1} \\ &\stackrel{(a)}{\leq} 2 \cdot \frac{\exp\left(-\frac{1}{8\sigma^2}\varepsilon^2\right)}{\frac{1}{4\sigma^2}\varepsilon^2} \leq \frac{8\sigma^2}{\varepsilon^2} \end{aligned}$$

where (a) is by  $e^z \geq z + 1$  and the choice of  $\bar{q} := \left\lceil 1 \vee \left(\frac{4\sigma^2}{\varepsilon^2} - \frac{1}{2}\right) \right\rceil$  that ensures  $\eta \geq e$  and  $\bar{q} \geq 1$ .

For the first term,

$$\begin{aligned} \frac{4\sigma^2}{\varepsilon^2} \sum_{q=1}^{\bar{q}-1} \frac{1}{q + \frac{1}{2}} &\leq \frac{4\sigma^2}{\varepsilon^2} \sum_{q=1}^{\bar{q}-1} \frac{1}{q} \leq \frac{4\sigma^2}{\varepsilon^2} \left(1 + \int_{q=1}^{\bar{q}-1} \frac{1}{q} dq\right) \\ &\leq \frac{4\sigma^2}{\varepsilon^2} (1 + \ln(\bar{q} - 1)) \end{aligned}$$

$$\leq \frac{4\sigma^2}{\varepsilon^2} \left( 1 + \ln\left(\frac{4\sigma^2}{\varepsilon^2} + \frac{1}{2}\right) \right).$$

Hence

$$\begin{aligned} \text{(F3)} &\leq \frac{8\sigma^2}{\varepsilon^2} + \frac{4\sigma^2}{\varepsilon^2} \left( 1 + \ln\left(\frac{4\sigma^2}{\varepsilon^2} + \frac{1}{2}\right) \right) \\ &= \frac{4\sigma^2}{\varepsilon^2} \left( 3 + \ln\left(\frac{4\sigma^2}{\varepsilon^2} + \frac{1}{2}\right) \right) \\ &= O\left(\frac{\sigma^2}{\varepsilon^2} \left( 1 + \ln\left(1 + \frac{\sigma^2}{\varepsilon^2}\right) \right)\right). \end{aligned}$$

Altogether with the choice of  $u := \left\lceil \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{2\sigma^2}\right)}{\Delta_a^2} \right\rceil$ ,  $\varepsilon = \frac{c\Delta_a}{2(1+c)}$

$$\begin{aligned} \mathbb{E}[N_{T,a}] &\leq u + \text{(F1)} + \text{(F2)} + \text{(F3)} \\ &= \left\lceil \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{2\sigma^2}\right)}{\Delta_a^2} \right\rceil + \frac{2\sigma^2}{\Delta_a^2} + \frac{2\sigma^2}{\left(\frac{c\Delta_a}{1+c} - \varepsilon\right)^2} + O\left(\frac{\sigma^2}{\varepsilon^2} \left( 1 + \ln\left(1 + \frac{\sigma^2}{\varepsilon^2}\right) \right)\right) \\ &= \left\lceil \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{2\sigma^2}\right)}{\Delta_a^2} \right\rceil + \frac{2\sigma^2}{\Delta_a^2} + \frac{8\sigma^2(1+c)^2}{c^2\Delta_a^2} + O\left(\frac{4(1+c)^2\sigma^2}{c^2\Delta_a^2} \left( 1 + \ln\left(1 + \frac{4(1+c)^2\sigma^2}{c^2\Delta_a^2}\right) \right)\right) \\ &= \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{\sigma^2}\right)}{\Delta_a^2} + O\left(1 \vee \left(\frac{(1+c)^2\sigma^2}{c^2\Delta_a^2} \ln\left(1 + \frac{(1+c)^2\sigma^2}{c^2\Delta_a^2}\right)\right)\right). \end{aligned}$$

Hence

$$\begin{aligned} \mathbb{E} \text{Reg}_T &= \sum_{a \in [K]: \Delta_a > 0} \Delta_a \cdot \mathbb{E}[N_{T,a}] \\ &\leq \sum_{a \in [K]: \Delta_a > 0} \left( \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{\sigma^2}\right)}{\Delta_a} + O\left(\Delta_a \vee \left(\frac{(1+c)^2\sigma^2}{c^2\Delta_a} \ln\left(1 + \frac{(1+c)^2\sigma^2}{c^2\Delta_a^2}\right)\right)\right) \right). \end{aligned}$$

## B.2 Proof of Corollary 3

*Proof.* From Section B.1 we have: MS satisfies that,  $\forall T \geq 1$  and  $c > 0$ ,

$$\mathbb{E}[N_{T,a}] \leq \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{\sigma^2}\right)}{\Delta_a^2} + O\left(1 \vee \left(\frac{(1+c)^2\sigma^2}{c^2\Delta_a^2} \ln\left(1 + \frac{(1+c)^2\sigma^2}{c^2\Delta_a^2}\right)\right)\right).$$

Therefore,

$$\begin{aligned} \mathbb{E} \text{Reg}_T &= \sum_{a \in [K]: \Delta_a > 0} \Delta_a \mathbb{E}[N_{T,a}] \\ &= \sum_{a \in [K]: \Delta_a < \Delta} \Delta_a \mathbb{E}[N_{T,a}] + \sum_{a \in [K]: \Delta_a \geq \Delta} \Delta_a \mathbb{E}[N_{T,a}] \\ &< T\Delta + K \left( \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta^2}{\sigma^2}\right)}{\Delta} + O\left(\frac{(1+c)^2\sigma^2}{c^2\Delta} \ln\left(1 + \frac{(1+c)^2\sigma^2}{c^2\Delta^2}\right)\right) \right) + \sum_{a \in [K]: \Delta_a > 0} O(\Delta_a). \end{aligned}$$

Then by choosing  $\Delta = \Theta(\sigma\sqrt{\frac{K \ln(T)}{T}})$ , the upper bound will become

$$\begin{aligned} \mathbb{E} \text{Reg}_T &< O(\sigma\sqrt{KT \ln(T)}) + O(\sigma\sqrt{KT \ln(T)}) + \sum_{a \in [K]: \Delta_a > 0} O(\Delta_a) \\ &= O(\sigma\sqrt{KT \ln(T)}). \end{aligned}$$

□

## C Proofs for Maillard Sampling<sup>+</sup>

### C.1 Proof of Theorem 4

*Proof.* Just like the proof of MS, it suffices to upper bound  $\mathbb{E}[N_a]$  for arm  $a$  such that  $\Delta_a > 0$ . Let  $T \geq K$ . First, using the definition of the initialization step, for all  $u > 0$  we have

$$\begin{aligned} \mathbb{E}[N_{T,a}] &= \mathbb{E}[N_{K,a}] + \mathbb{E}[N_{T,a} - N_{K,a}] \\ &\leq 1 + v + \mathbb{E} \left[ \sum_{t=K+v+1}^T \mathbb{1}\{I_t = a \text{ and } N_{t-1,a} > v\} \right] \\ &= 1 + v + \mathbb{E} \left[ \sum_{t=K+v}^{T-1} \mathbb{1}\{I_{t+1} = a \text{ and } N_{t,a} > v\} \right] \\ &= u + \mathbb{E} \left[ \sum_{t=K-1+u}^{T-1} \mathbb{1}\{I_{t+1} = a \text{ and } N_{t,a} \geq u\} \right]. \end{aligned}$$

Let  $\hat{\mu}_{t,\max} = \max_a \hat{\mu}_{t,a}$ . Let  $c > 0$ . Then, we consider the following branching conditions under  $\{I_{t+1} = a, N_t(a) \geq u\}$ :

- (F1)  $\hat{\mu}_{t,\max} \geq \mu_1 - \varepsilon, \widehat{\Delta}_{t,a} > \frac{\Delta_a}{1+c}$
- (F2)  $\hat{\mu}_{t,\max} \geq \mu_1 - \varepsilon, \widehat{\Delta}_{t,a} \leq \frac{\Delta_a}{1+c}$
- (F3)  $\hat{\mu}_{t,\max} < \mu_1 - \varepsilon$ .

For case (F1), by definition of  $p_{t+1}$  and  $u := \left\lceil \frac{2\sigma^2(1+c)^2 \ln(\frac{T\Delta_a^2}{2\sigma^2}(1+\ln(1+\frac{T\Delta_a^2}{2\sigma^2})))}{\Delta_a^2} \right\rceil$ , we have

$$\begin{aligned} &\mathbb{E} \sum_{t=K-1+u}^{T-1} \mathbb{1}\left\{I_{t+1} = a, N_{t,a} \geq u, \hat{\mu}_{t,\max} \geq \mu_1 - \varepsilon, \widehat{\Delta}_{t,a} > \frac{\Delta_a}{1+c}\right\} \\ &\leq \sum_{t=K-1+u}^{T-1} \mathbb{P}\left(I_{t+1} = a \mid N_{t,a} \geq u, \widehat{\Delta}_{t,a} > \frac{\Delta_a}{1+c}\right) \mathbb{P}\left(N_{t,a} \geq u, \widehat{\Delta}_{t,a} > \frac{\Delta_a}{1+c}\right) \\ &\stackrel{(a)}{\leq} \sum_{t=K-1+u}^{T-1} \frac{1}{B} \exp\left(-\frac{1}{2\sigma^2} u \left(\frac{\Delta_a}{1+c}\right)^2\right) \cdot \left(1 + \frac{D}{2\sigma^2} u \left(\frac{\Delta_a}{1+c}\right)^2\right) \cdot 1 \end{aligned}$$

where (a) is by the fact that for  $D \in (0, 1]$ ,  $(1 + Dx) \cdot \exp(-x)$  is decreasing in  $x > 0$ . Let  $A = \frac{2\sigma^2}{\Delta_a^2}$ . Then,  $u \geq (1+c)^2 A \ln(\frac{T}{A}(1+\ln(1+\frac{T}{A})))$ . The last line above is then bounded by

$$\begin{aligned} &\frac{T}{B} \exp\left(-\frac{1}{A(1+c)^2} u\right) \cdot \left(1 + \frac{D}{A(1+c)^2} u\right) \cdot 1 \\ &\leq \frac{1}{B} \frac{A}{1+\ln(1+\frac{T}{A})} \cdot \left(1 + D \ln\left(\frac{T}{A}(1+\ln(1+\frac{T}{A}))\right)\right) \\ &\leq \frac{1}{B} \cdot A \cdot 2D \\ &= \frac{2D}{B} \cdot \frac{2\sigma^2}{\Delta_a^2}. \end{aligned}$$

For case (F2), note that

$$\left\{\hat{\mu}_{t,\max} \geq \mu_1 - \varepsilon\right\} \cap \left\{\widehat{\Delta}_{t,a} \leq \frac{1}{1+c} \Delta_a\right\} \subseteq \left\{\hat{\mu}_{t,a} - \mu_a \geq \frac{c}{1+c} \Delta_a - \varepsilon\right\}.$$

Let  $\tau_a(k)$  be the time step  $t$  such that the arm  $a$  was pulled at  $t$  and the number of arm pulls of arm  $a$  becomes  $k$  at the end of the time step  $t$ ; i.e.,  $\tau_a(k) = \min\{t \geq 1 : N_{t,a} = k\}$ . We use the shortcut  $\hat{\mu}_{a,(k)} := \hat{\mu}_{\tau_a(k),a}$  to denote the sample mean of arm  $a$  after pulling it for the  $k$ -th time, . Thus,

$$\mathbb{E} \sum_{t=K-1+u}^{T-1} \mathbb{1}\left\{I_{t+1} = a, N_{t,a} \geq u, \hat{\mu}_{t,\max} \geq \mu_1 - \varepsilon, \widehat{\Delta}_{t,a} \leq \frac{\Delta_a}{1+c}\right\}$$

$$\begin{aligned}
 &\leq \sum_{t=K-1+u}^{T-1} \mathbb{P} \left( I_{t+1} = a, \hat{\mu}_{t,a} - \mu_a \geq \frac{c\Delta_a}{1+c} - \varepsilon \right) \\
 &\leq \mathbb{E} \left[ \sum_{k=1}^{\infty} \sum_{t=\tau_1(k)}^{\tau_1(k+1)-1} \mathbb{1}\{I_{t+1} = a\} \cdot \mathbb{1} \left\{ \hat{\mu}_{t,a} - \mu_a \geq \frac{c\Delta_a}{1+c} - \varepsilon \right\} \right] \\
 &= \mathbb{E} \left[ \sum_{k=1}^{\infty} \mathbb{1} \left\{ \hat{\mu}_{a,(k)} - \mu_a \geq \frac{c\Delta_a}{1+c} - \varepsilon \right\} \sum_{t=\tau_1(k)}^{\tau_1(k+1)-1} \mathbb{1}\{I_{t+1} = a\} \right] \\
 &= \mathbb{E} \left[ \sum_{k=1}^{\infty} \mathbb{1} \left\{ \hat{\mu}_{a,(k)} - \mu_a \geq \frac{c\Delta_a}{1+c} - \varepsilon \right\} \right] \\
 &= \sum_{k \geq 1} \mathbb{P} \left( \hat{\mu}_{a,(k)} - \mu_a \geq \frac{c\Delta_a}{1+c} - \varepsilon \right) \quad (N_{t,a} \geq 1, \forall t \geq K, \text{ since } u > 1) \\
 &\leq \sum_{k \geq 1} \exp \left( -\frac{1}{2\sigma^2} k \left( \frac{c\Delta_a}{1+c} - \varepsilon \right)^2 \right) = \frac{1}{\exp \left( \frac{1}{2\sigma^2} \left( \frac{c\Delta_a}{1+c} - \varepsilon \right)^2 \right) - 1} \leq \frac{2\sigma^2}{\left( \frac{c\Delta_a}{1+c} - \varepsilon \right)^2}.
 \end{aligned}$$

We now consider case (F3). This is the case where the best arm behaves badly. Let  $\mathcal{F}_t$  be the  $\sigma$ -algebra generated by  $(A_1, R_1, \dots, A_t, R_t)$ . We define the shortcut  $\mathbb{P}_t(\mathcal{E}) := \mathbb{P}(\mathcal{E} \mid \mathcal{F}_t)$  and similarly  $\mathbb{E}_t[X] = \mathbb{E}[X \mid \mathcal{F}_t]$ .

Recall  $B \geq 1$  and  $C, D \in (0, 1]$ . Let  $Z = \sum_a \hat{\rho}_{t+1,a}$ . Note that, for every  $a \in [K]$ , the following is true regardless of whether  $a$  is the empirical best arm or not:

$$\frac{1}{Z} \exp \left( -\frac{1}{2\sigma^2} N_{t,a} \widehat{\Delta}_{t,a}^2 + \ln \left( 1 + \frac{D}{2\sigma^2} N_{t,a} \widehat{\Delta}_{t,a}^2 \right) \right) \leq \mathbb{P}_t(I_{t+1} = a) \leq \frac{B(1 + C \ln(1 + \ln(\frac{t}{N_{t,a}})))}{Z}.$$

This helps us relate the probability of pulling arm  $a$  with arm 1. Thus, under the event  $\hat{\mu}_{t,\max} < \mu_1 - \varepsilon$ , we have

$$\mathbb{P}_t(I_{t+1} = a) \leq \frac{B \cdot (1 + C \ln(1 + \ln(\frac{t}{N_{t,a}})))}{\exp(-\frac{1}{2\sigma^2} N_{t,1} \cdot (\mu_1 - \varepsilon - \hat{\mu}_{t,1})_+^2) \cdot (1 + \frac{D}{2\sigma^2} N_{t,1} \cdot (\mu_1 - \varepsilon - \hat{\mu}_{t,1})_+^2)} \mathbb{P}_t(I_{t+1} = 1). \quad (5)$$

Let  $\tau_1(k)$  be the time step  $t$  such that the arm 1 was pulled at  $t$  and the number of arm pulls of arm 1 becomes  $k$  at the end of the time step  $t$ ; i.e.,  $\tau_1(k) = \min \{t \geq 1 : N_{t,1} = k\}$ . We use the shortcut  $\hat{\mu}_{1,(k)} := \hat{\mu}_{\tau_1(k),1}$ . With this, we can bound (A) as follows:

$$\begin{aligned}
 &\mathbb{E} \sum_{t=K-1+u}^{T-1} \mathbb{1}\{I_{t+1} = a, N_{t,a} \geq u, \hat{\mu}_{t,\max} < \mu_1 - \varepsilon\} \\
 &\leq \mathbb{E} \left[ \sum_{t=K-1+u}^{T-1} \mathbb{E}_t \left[ \mathbb{1}\{I_{t+1} = a, \hat{\mu}_{t,\max} < \mu_1 - \varepsilon\} \right] \right] \\
 &= \mathbb{E} \left[ \sum_{t=K-1+u}^{T-1} \mathbb{1}\{\hat{\mu}_{t,\max} < \mu_1 - \varepsilon\} \mathbb{E}_t \left[ \mathbb{1}\{I_{t+1} = a\} \right] \right] \\
 &\leq \mathbb{E} \left[ \sum_{t=K-1+u}^{T-1} \mathbb{1}\{\hat{\mu}_{t,\max} < \mu_1 - \varepsilon\} \cdot \frac{\exp(\frac{1}{2\sigma^2} N_{t,1} \cdot (\mu_1 - \varepsilon - \hat{\mu}_{t,1})^2)}{1 + \frac{D}{2\sigma^2} \cdot N_{t,1} \cdot (\mu_1 - \varepsilon - \hat{\mu}_{t,1})^2} \cdot \left( B \cdot (1 + C \ln(1 + \ln(\frac{t}{N_{t,a}}))) \right) \mathbb{E}_t \left[ \mathbb{1}\{I_{t+1} = 1\} \right] \right] \\
 &= \mathbb{E} \left[ \sum_{t=K-1+u}^{T-1} \mathbb{E}_t \left[ \mathbb{1}\{\hat{\mu}_{t,\max} < \mu_1 - \varepsilon\} \cdot \frac{\exp(\frac{1}{2\sigma^2} N_{t,1} \cdot (\mu_1 - \varepsilon - \hat{\mu}_{t,1})^2)}{1 + \frac{D}{2\sigma^2} \cdot N_{t,1} \cdot (\mu_1 - \varepsilon - \hat{\mu}_{t,1})^2} \cdot \left( B \cdot (1 + C \ln(1 + \ln(\frac{t}{N_{t,a}}))) \right) \cdot \mathbb{1}\{I_{t+1} = 1\} \right] \right] \\
 &= \mathbb{E} \left[ \sum_{t=K-1+u}^{T-1} \mathbb{1}\{\hat{\mu}_{t,\max} < \mu_1 - \varepsilon\} \cdot \frac{\exp(\frac{1}{2\sigma^2} N_{t,1} \cdot (\mu_1 - \varepsilon - \hat{\mu}_{t,1})^2)}{1 + \frac{D}{2\sigma^2} \cdot N_{t,1} \cdot (\mu_1 - \varepsilon - \hat{\mu}_{t,1})^2} \cdot \left( B \cdot (1 + C \ln(1 + \ln(\frac{t}{N_{t,a}}))) \right) \cdot \mathbb{1}\{I_{t+1} = 1\} \right] \\
 &\leq \mathbb{E} \left[ \sum_{t=K-1+u}^{T-1} \mathbb{1}\{\hat{\mu}_{t,1} < \mu_1 - \varepsilon\} \cdot \frac{\exp(\frac{1}{2\sigma^2} N_{t,1} \cdot (\mu_1 - \varepsilon - \hat{\mu}_{t,1})^2)}{1 + \frac{D}{2\sigma^2} \cdot N_{t,1} \cdot (\mu_1 - \varepsilon - \hat{\mu}_{t,1})^2} \cdot \left( B \cdot (1 + C \ln(1 + \ln(\frac{t}{N_{t,a}}))) \right) \cdot \mathbb{1}\{I_{t+1} = 1\} \right]
 \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E} \left[ \sum_{k=1}^{\infty} \sum_{t=\tau_1(k)}^{\tau_1(k+1)-1} \mathbb{1}\{\hat{\mu}_{1,(k)} < \mu_1 - \varepsilon\} \cdot \frac{\exp\left(\frac{1}{2\sigma^2}k \cdot (\mu_1 - \varepsilon - \hat{\mu}_{1,(k)})^2\right)}{1 + \frac{D}{2\sigma^2} \cdot k \cdot (\mu^* - \varepsilon - \hat{\mu}_{1,(k)})^2} \cdot \left(B \cdot (1 + C \ln(1 + \ln(\frac{T}{u})))\right) \cdot \mathbb{1}\{I_{t+1} = 1\} \right] \\
 &\leq \mathbb{E} \left[ \sum_{k=1}^{\infty} \mathbb{1}\{\hat{\mu}_{1,(k)} < \mu_1 - \varepsilon\} \cdot \frac{\exp\left(\frac{1}{2\sigma^2}k \cdot (\mu_1 - \varepsilon - \hat{\mu}_{1,(k)})^2\right)}{1 + \frac{D}{2\sigma^2} \cdot k \cdot (\mu_1 - \varepsilon - \hat{\mu}_{1,(k)})^2} \cdot \left(B \cdot (1 + C \ln(1 + \ln(\frac{T}{u})))\right)^{\tau_1(k+1)-1} \sum_{t=\tau_1(k)}^{\tau_1(k+1)-1} \mathbb{1}\{I_{t+1} = 1\} \right] \\
 &= \mathbb{E} \left[ \sum_{k=1}^{\infty} \mathbb{1}\{\hat{\mu}_{1,(k)} < \mu_1 - \varepsilon\} \cdot \frac{\exp\left(\frac{1}{2\sigma^2}k \cdot (\mu_1 - \varepsilon - \hat{\mu}_{1,(k)})^2\right)}{1 + \frac{D}{2\sigma^2} \cdot k \cdot (\mu_1 - \varepsilon - \hat{\mu}_{1,(k)})^2} \cdot \left(B \cdot (1 + C \ln(1 + \ln(\frac{T}{u})))\right) \right] \\
 &= \mathbb{E} \left[ \sum_{k=1}^{\infty} \mathbb{1}\{\bar{X}_k > \varepsilon\} \cdot \frac{\exp\left(\frac{1}{2\sigma^2}k \cdot (\bar{X}_k - \varepsilon)^2\right)}{1 + \frac{D}{2\sigma^2}k \cdot (\bar{X}_k - \varepsilon)^2} \cdot B \cdot (1 + C \ln(1 + \ln(\frac{T}{u}))) \right]
 \end{aligned}$$

where the last line is by defining  $\bar{X}_k := \mu_1 - \hat{\mu}_{1,(k)}$ .

We bound the following term:

$$\begin{aligned}
 &\mathbb{E} \left[ \sum_{k=1}^{\infty} \mathbb{1}\{\bar{X}_k > \varepsilon\} \cdot \frac{\exp\left(\frac{1}{2\sigma^2}k \cdot (\bar{X}_k - \varepsilon)^2\right)}{1 + \frac{D}{2\sigma^2}k \cdot (\bar{X}_k - \varepsilon)^2} \right] \\
 &\leq \mathbb{E} \left[ \sum_{k=1}^{\infty} \sum_{q=1}^{\infty} \mathbb{1}\left\{\mu_1 - \hat{\mu}_{1,(k)} \in \left(\varepsilon + \frac{q-1}{2}\varepsilon, \varepsilon + \frac{q}{2}\varepsilon\right)\right\} \cdot \frac{\exp\left(\frac{1}{2\sigma^2}k \cdot \left(\frac{q\varepsilon}{2}\right)^2\right)}{1 + \frac{D}{2\sigma^2}k \cdot \left(\frac{q\varepsilon}{2}\right)^2} \right] \\
 &= \sum_{k=1}^{\infty} \sum_{q=1}^{\infty} \mathbb{P}\left(\mu_1 - \hat{\mu}_{1,(k)} \in \left(\varepsilon + \frac{q-1}{2}\varepsilon, \varepsilon + \frac{q}{2}\varepsilon\right)\right) \cdot \frac{\exp\left(\frac{1}{2\sigma^2}k \cdot \left(\frac{q\varepsilon}{2}\right)^2\right)}{1 + \frac{D}{2\sigma^2}k \cdot \left(\frac{q\varepsilon}{2}\right)^2} \\
 &\leq \sum_{k=1}^{\infty} \sum_{q=1}^{\infty} \exp\left(-\frac{1}{2\sigma^2}k \cdot \left(\frac{q+1}{2}\varepsilon\right)^2\right) \cdot \frac{\exp\left(\frac{1}{2\sigma^2}k \cdot \left(\frac{q\varepsilon}{2}\right)^2\right)}{1 + \frac{D}{2\sigma^2}k \cdot \left(\frac{q\varepsilon}{2}\right)^2} \\
 &= \sum_{k=1}^{\infty} \sum_{q=1}^{\infty} \frac{\exp\left(-\frac{1}{4\sigma^2}(q + \frac{1}{2})k\varepsilon^2\right)}{1 + \frac{Dkq^2\varepsilon^2}{8\sigma^2}} \\
 &< \sum_{k=1}^{\infty} \sum_{q=1}^{\infty} \frac{\exp\left(-\frac{1}{4\sigma^2}(q + \frac{1}{2})k\varepsilon^2\right)}{1 \vee \frac{Dkq^2\varepsilon^2}{8\sigma^2}} \\
 &= \sum_{q=1}^{\infty} \left( \sum_{k=1}^{k < \frac{8\sigma^2}{Dq^2\varepsilon^2}} \exp\left(-\frac{1}{4\sigma^2}(q + \frac{1}{2})k\varepsilon^2\right) + \sum_{k \geq \frac{8\sigma^2}{Dq^2\varepsilon^2}} \frac{8\sigma^2 \exp\left(-\frac{1}{4\sigma^2}(q + \frac{1}{2})k\varepsilon^2\right)}{Dkq^2\varepsilon^2} \right).
 \end{aligned}$$

For the first term, let us define  $Q = \frac{8\sigma^2}{D\varepsilon^2}$  and  $R = \frac{\varepsilon^2}{4\sigma^2}$ . Then,

$$\begin{aligned}
 \sum_{q=1}^{\infty} \sum_{k=1}^{k < \frac{Q}{q^2}} \exp\left(-R(q + \frac{1}{2})k\right) &\leq \sum_q \frac{1 - \exp\left(-R(q + \frac{1}{2})\frac{Q}{q^2}\right)}{\exp\left(R(q + \frac{1}{2})\right) - 1} \\
 &\stackrel{(a)}{\leq} Q \sum_q \frac{1}{q^2} \\
 &= 2Q
 \end{aligned}$$

where (a) is by  $e^z \geq z + 1$ .

For the second term,

$$\begin{aligned}
 \sum_{q=1}^{\infty} \sum_{k \geq Q/q^2} \frac{Q}{kq^2} \exp\left(-R(q + \frac{1}{2})k\right) &\leq \underbrace{\sum_{q=1}^{\infty} \sum_{k \in [Q/q^2, Q/q]} \frac{Q}{kq^2} \exp\left(-R(q + \frac{1}{2})k\right)}_{=: G_1} + \underbrace{\sum_{q=1}^{\infty} \sum_{k > Q/q} \frac{Q}{kq^2} \exp\left(-R(q + \frac{1}{2})k\right)}_{=: G_2}.
 \end{aligned}$$



We bound  $G_1$  as follows:

$$\begin{aligned}
 G_1 &\leq \sum_{q=1}^{\infty} \sum_{k \in [Q/q^2, Q/q]} \frac{Q}{kq^2} \\
 &\stackrel{(a)}{\leq} \sum_{q=1}^{\infty} \frac{Q}{q^2} \left( 1 + \ln\left(\frac{Q/q}{Q/q^2}\right) \right) \\
 &= \sum_{q=1}^{\infty} \frac{Q}{q^2} (1 + \ln(q)) \\
 &= Q \left( \frac{\pi^2}{6} + \sum_{q=1}^{\infty} \frac{\ln(q)}{q^2} \right) \\
 &\leq 3Q
 \end{aligned}$$

where (a) is by the fact that

$$\sum_{i=a}^b f(i) \leq \max_{x \in [a, b]} f(x) + \int_a^b f(x) dx . \tag{6}$$

We bound  $G_2$  as follows:

$$\begin{aligned}
 G_2 &\leq \sum_{q=1}^{\infty} \sum_{k > Q/q} \frac{Q}{kq^2} \exp\left(-R\left(q + \frac{1}{2}\right)k\right) \\
 &\leq \sum_{q=1}^{\infty} \frac{1}{q} \sum_{k > Q/q} \exp\left(-R\left(q + \frac{1}{2}\right)k\right) \\
 &\leq \sum_{q=1}^{\infty} \frac{1}{q} \frac{\exp(-R(q + \frac{1}{2})((1 \vee \frac{Q}{q}) - 1))}{\exp(R(q + \frac{1}{2})) - 1} \tag{note } k \geq 1 \\
 &\leq \sum_{q=1}^{\infty} \frac{1}{q} \frac{1}{R(q + \frac{1}{2})} \tag{by } e^z \geq z - 1 \\
 &\leq \frac{2}{R} .
 \end{aligned}$$

Thus,

$$\mathbb{E} \left[ \sum_{k=1}^{\infty} \mathbb{1}\{\bar{X}_k > \varepsilon\} \cdot \frac{\exp\left(\frac{1}{2\sigma^2} k \cdot (\bar{X}_k - \varepsilon)^2\right)}{1 + \frac{D}{2\sigma^2} k \cdot (\bar{X}_k - \varepsilon)^2} \right] \leq \left(\frac{40}{D} + 8\right) \frac{\sigma^2}{\varepsilon^2} .$$

Hence

$$\left( B \cdot \left( 1 + C \ln\left(1 + \ln\left(\frac{T}{u}\right)\right) \right) \right) \cdot \mathbb{E} \left[ \sum_{k=1}^{\infty} \sum_{q=1}^{\infty} \mathbb{1}\left\{ \bar{X}_k \in \left( \varepsilon + \frac{q-1}{2}\varepsilon, \varepsilon + \frac{q}{2}\varepsilon \right] \right\} \cdot \frac{\exp\left(\frac{1}{2\sigma^2} k \cdot \left(\frac{q\varepsilon}{2}\right)^2\right)}{1 + \frac{D}{2\sigma^2} \cdot k \cdot \left(\frac{q\varepsilon}{2}\right)^2} \right] .$$

is upper bounded by

$$\begin{aligned}
 &\left( B \cdot \left( 1 + C \ln\left(1 + \ln\left(\frac{T}{u}\right)\right) \right) \right) \cdot \left(\frac{40}{D} + 8\right) \frac{\sigma^2}{\varepsilon^2} \\
 &< \left( B \cdot \left( 1 + C \ln\left(1 + \ln\left(\frac{T}{\frac{2\sigma^2(1+c)^2 \ln(T\Delta_a^2/(2\sigma^2))}}{\Delta_a^2}}\right)\right) \right) \cdot \left(\frac{40}{D} + 8\right) \frac{\sigma^2}{\varepsilon^2} \\
 &< O\left( B \cdot \left( 1 + C \ln\left(1 + \ln\left(\frac{T\Delta_a^2}{\sigma^2}\right)\right) \right) \right) \cdot \left(\frac{40}{D} + 8\right) \frac{\sigma^2}{\varepsilon^2}
 \end{aligned}$$

where the last inequality is by the safe assumption of  $\ln(T\Delta_a^2/(2\sigma^2)) \geq 1$  since otherwise we have a trivial bound

on  $u$  as follows:  $\frac{2\sigma^2}{\Delta_a^2}e > T \geq u$ . Altogether, we choose  $\varepsilon := \frac{c\Delta_a}{2(1+c)}$  and show that

$$\begin{aligned} \mathbb{E}[N_{t,a}] &\leq \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{2\sigma^2}\left(1 + \ln\left(1 + \frac{T\Delta_a^2}{2\sigma^2}\right)\right)\right)}{\Delta_a^2} + \frac{(1+2D)\sigma^2}{B\Delta_a^2} + \frac{2\sigma^2}{\left(\frac{c\Delta_a}{1+c} - \varepsilon\right)^2} \\ &\quad + O\left(B \cdot \left(1 + C\left(\ln\left(1 + \ln\left(\frac{T\Delta_a^2}{\sigma^2}\right)\right)\right)\right)\right) \cdot \left(\frac{40}{D} + 8\right) \frac{\sigma^2}{\varepsilon^2} + O(1) \\ &\leq \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{2\sigma^2}\left(1 + \ln\left(1 + \frac{T\Delta_a^2}{2\sigma^2}\right)\right)\right)}{\Delta_a^2} + O\left(1 + \left(\frac{BC}{D} + \frac{1}{B} + \frac{(1+c)^2}{c^2}\right) \cdot \frac{\sigma^2 \ln\left(1 + \ln\left(\frac{T\Delta_a^2}{\sigma^2}\right)\right)}{\Delta_a^2}\right). \end{aligned}$$

In terms of the upper bound of the regret, for every  $T \geq 1$  and  $c > 0$ ,

$$\begin{aligned} \mathbb{E} \text{Reg}_T &= \sum_{a \in [K]: \Delta_a > 0} \Delta_a \cdot \mathbb{E}[N_{t,a}] \\ &\leq \sum_{a \in [K]: \Delta_a > 0} \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{2\sigma^2}\left(1 + \ln\left(1 + \frac{T\Delta_a^2}{2\sigma^2}\right)\right)\right)}{\Delta_a} + O\left(\Delta_a + \left(\frac{BC}{D} + \frac{1}{B} + \frac{(1+c)^2}{c^2}\right) \frac{\sigma^2 \ln\left(1 + \ln\left(\frac{T\Delta_a^2}{\sigma^2}\right)\right)}{\Delta_a}\right) \\ &\leq \sum_{a: \Delta_a > 0} \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{2\sigma^2}\left(1 + \ln\left(1 + \frac{T\Delta_a^2}{2\sigma^2}\right)\right)\right)}{\Delta_a} + O\left(\Delta_a + \left(\frac{BC}{D} + \frac{1}{B} + \frac{(1+c)^2}{c^2}\right) \frac{\sigma^2 \ln\left(1 + \ln\left(\frac{T\Delta_a^2}{\sigma^2}\right)\right)}{\Delta_a}\right) \end{aligned}$$

where we omit the dependence on  $C$  for brevity.

To verify the asymptotic optimality, let us choose  $c = \ln^{-1/4}(T)$  and take  $B$ ,  $C$ , and  $D$  as absolute constants. Then,

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{\mathbb{E}[N_{t,a}]}{\ln(T)} &\leq \lim_{T \rightarrow \infty} \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{2\sigma^2}\left(1 + \ln\left(1 + \frac{T\Delta_a^2}{2\sigma^2}\right)\right)\right)}{\Delta_a^2 \ln(T)} + \lim_{T \rightarrow \infty} \frac{O\left(1 \vee \left(\frac{\sigma^2(1+c)^2 \ln\left(1 + \ln\left(\frac{T\Delta_a^2}{\sigma^2}\right)\right)}{c^2 \Delta_a^2}\right)\right)}{\ln(T)} \\ &= \lim_{T \rightarrow \infty} \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta_a^2}{2\sigma^2}\left(1 + \ln\left(1 + \frac{T\Delta_a^2}{2\sigma^2}\right)\right)\right)}{\Delta_a^2 \ln(T)} \\ &= \frac{2\sigma^2}{\Delta_a^2}. \end{aligned}$$

This shows the asymptotic optimality of  $\text{MS}^+$ . □

## C.2 Proof of Corollary 5

*Proof.* Let us set  $B$ ,  $C$  and  $D$  as absolute constants. From the Proof of  $\text{MS}^+$  in Appendix C, in terms of the upper bound of the regret,

$$\begin{aligned} \mathbb{E} \text{Reg}_T &= \sum_{a \in [K]: \Delta_a > 0} \Delta_a \cdot \mathbb{E}[N_{t,a}] \\ &= \sum_{a: \Delta_a < \Delta} \Delta_a \cdot \mathbb{E}[N_{t,a}] + \sum_{a: \Delta_a \geq \Delta} \Delta_a \cdot \mathbb{E}[N_{t,a}] \\ &< T\Delta + K \left( \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta^2}{2\sigma^2}\left(1 + \ln\left(1 + \frac{T\Delta^2}{2\sigma^2}\right)\right)\right)}{\Delta} + O\left(\sigma^2 \frac{(1+c)^2 \ln\left(1 + \ln\left(\frac{T\Delta^2}{\sigma^2}\right)\right)}{c^2 \Delta}\right) \right) + \sum_{a: \Delta_a \geq \Delta} O(\Delta_a). \end{aligned}$$

By choosing  $\Delta = \Theta\left(\sigma \sqrt{\frac{K \ln(K)}{T}}\right)$ , we have:

$$\mathbb{E} \text{Reg}_T = O\left(\sigma \sqrt{KT \ln(K)}\right) + K \left( \frac{2\sigma^2(1+c)^2 \ln\left(\frac{T\Delta^2}{2\sigma^2}\left(1 + \ln\left(1 + \frac{T\Delta^2}{2\sigma^2}\right)\right)\right)}{\Delta} + O\left(\sigma^2 \frac{(1+c)^2 \ln\left(1 + \ln\left(\frac{T\Delta^2}{\sigma^2}\right)\right)}{c^2 \Delta}\right) \right)$$

$$\begin{aligned} &+ O\left(\sum_{a \in [K]} \Delta_a\right) \\ &= O\left(\sigma\sqrt{KT\log(K)}\right). \end{aligned}$$

Hence this algorithm achieves both  $\sqrt{\ln(K)}$  minimax ratio and the asymptotically optimality.  $\square$