

---

# Rotting Infinitely Many-Armed Bandits

---

Jung-hun Kim<sup>1</sup> Milan Vojnović<sup>2</sup> Se-Young Yun<sup>1</sup>

## Abstract

We consider the infinitely many-armed bandit problem with rotting rewards, where the mean reward of an arm decreases at each pull of the arm according to an arbitrary trend with maximum rotting rate  $\varrho = o(1)$ . We show that this learning problem has an  $\Omega(\max\{\varrho^{1/3}T, \sqrt{T}\})$  worst-case regret lower bound where  $T$  is the time horizon. We show that a matching upper bound  $\tilde{O}(\max\{\varrho^{1/3}T, \sqrt{T}\})$ , up to a poly-logarithmic factor, can be achieved by an algorithm that uses a UCB index for each arm and a threshold value to decide whether to continue pulling an arm or remove the arm from further consideration, when the algorithm knows the value of the maximum rotting rate  $\varrho$ . We also show that an  $\tilde{O}(\max\{\varrho^{1/3}T, T^{3/4}\})$  regret upper bound can be achieved by an algorithm that does not know the value of  $\varrho$ , by using an adaptive UCB index along with an adaptive threshold value.

## 1. Introduction

We consider a fundamental sequential learning problem in which an agent must play one option at a time from an infinite set of options with non-stationary reward distributions, where the mean reward of an option decreases at each play of this option. This is naturally studied as the infinitely many-armed bandit problem with rotting rewards. The assumption of infinitely many arms models practical situations when there is a finite but large number of arms relative to the number of available experiments. There is an abundance of applications in which one must choose from a large set of options with rotting rewards, e.g. online advertising where arms correspond to ads and rewards decrease over exposures of an ad to a user, content recommendation systems where arms correspond to media items and rotting

arises because of user boredom when watching the same content, and clinical trials where the efficacy of a medicine may decrease because of drug tolerance when a patient takes the same medicine several times. While there has been a lot of work on multi-armed bandits with a finite number of arms with stationary or non-stationary rewards, and an infinite number of arms with stationary rewards, not much seems to be known for the case of infinitely many arms with non-stationary rewards.

In this paper we make first steps to understand the fundamental limits of sequential learning for infinite number of arms whose mean rewards decrease with the number of pulls—the case commonly referred to as the *rotting bandits*. Our focus is on rotting trends where the mean reward of an arm decreases arbitrary for at most a fixed amount  $\varrho$  at each pull of this arm. The initial mean rewards of arms are assumed to be independent and identically distributed according to uniform distribution on  $[0, 1]$ . The objective is to find a policy that minimizes the expected cumulative regret over a time horizon of  $T$  time steps with respect to playing the best arm.

We show that the worst-case regret for this problem is lower bounded by  $\Omega(\max\{\varrho^{1/3}T, \sqrt{T}\})$ , where  $\varrho$  is the maximum rotting rate, and show that this lower bound is tight up to a poly-logarithmic factor. This reveals that the rotting trend starts to have an effect on regret precisely at the threshold  $\varrho = \Theta(1/T^{3/2})$ . Our result implies that the rotting bandit problem with infinitely many arms is harder than for the stationary rewards case, as in the latter case the regret lower bound is  $\Omega(\sqrt{T})$  (Wang et al., 2009). This stands in stark contrast to the case of finite  $K$  arms in which case it is known that  $\tilde{O}(\sqrt{KT})$  can be achieved for the rotting case (Seznec et al., 2019), which matches the optimal bound in the stationary case (Auer et al., 2002b) up to a poly-logarithmic factor.

In the case of infinitely many arms with stationary rewards, it is not possible to explore all arms to find an optimal arm, hence, it is required to find a near-optimal arm; contrast this with the case of finitely many arms, where all arms must be explored to identify an optimal arm. Further, when we consider rotting rewards, the learner must keep exploring new arms because a near-optimal arm may become suboptimal as it is being pulled. Based on this fact, we design algo-

---

<sup>1</sup>Korea Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea <sup>2</sup>London School of Economics, London, UK. Correspondence to: Se-Young Yun <yunseyoung@kaist.ac.kr>, Milan Vojnović <m.vojnovic@lse.ac.uk>.

gorithms for the rotting infinitely many-armed bandit problem to achieve tight regret bounds. We summarize our contributions in more details in what follows.

### 1.1. Summary of our contributions

We show an  $\Omega(\max\{\varrho^{1/3}T, \sqrt{T}\})$  worst-case regret lower bound for the rotting rested bandit case with maximum rotting rate  $\varrho = o(1)$ . This regret lower bound matches the regret lower bound  $\Omega(\sqrt{T})$  that is known to hold for the case of stationary rewards, when rotting is sufficiently small—precisely when  $\varrho = O(1/T^{3/2})$ . Otherwise, when  $\varrho = \omega(1/T^{3/2})$ , the regret lower bound becomes worse than for the stationary case.

We show that an  $\tilde{O}(\max\{\varrho^{1/3}T, \sqrt{T}\})$  regret can be achieved by an algorithm when the maximum rotting rate  $\varrho$  is known to the algorithm. This algorithm uses a UCB index to decide whether to continue pulling an arm or remove the arm from further consideration and switch to exploring a new arm by comparing the index with a threshold. This threshold is set to account for rotting of rewards.

We further show that an  $\tilde{O}(\max\{\varrho^{1/3}T, T^{3/4}\})$  regret can be achieved by an algorithm that does not know the value of the maximum rotting rate  $\varrho$ . This algorithm uses an adaptive UCB index and an adaptive threshold value to compare the UCB index of an arm with the threshold to decide whether to continue pulling this arm or remove the arm from further consideration. This upper bound matches the lower bound up to poly-logarithmic factors when the rotting rate  $\varrho$  is sufficiently large, i.e. when  $\varrho = \Omega(1/T^{3/4})$ .

We present results of numerical experiments for randomly generated problem instances of rotting infinitely many-armed bandits. These results validate the insights derived from our theoretical results.

### 1.2. Related work

The work on multi-armed bandits can be distinguished with respect to two criteria, first whether the number of arms is finite or infinite, and second whether rewards of arms are stationary or non-stationary. For the case of non-stationary rewards, we can further distinguish rested from restless multi-armed bandit problems — in the former case, an arm’s distribution of reward may change only when the arm is pulled, while in the latter case, it may change at each time step. Our work falls in the category of multi-armed bandit problems with infinitely many non-stationary rested arms.

The case of finitely many arms with stationary rewards has been studied by many, following on [Lai & Robbins \(1985\)](#); [Auer et al. \(2002a\)](#). There exist algorithms having  $\tilde{O}(\sqrt{KT})$  worst-case regret, where  $K$  is the number of arms, and this matches the lower bound  $\Omega(\sqrt{KT})$  up to a poly-logarithmic factor ([Auer et al., 2002b](#); [Slivkins, 2019](#)).

We next discuss the case of finitely many arms with non-stationary rewards. The non-stationarity in rewards can be quantified by the number of abrupt changes or a variation budget, which is referred to as *abrupt-changing* and *slow-varying* environments, respectively. The non-stationary environments were studied by [Auer et al. \(2002b\)](#); [Garivier & Moulines \(2011\)](#); [Besbes et al. \(2014\)](#); [Auer et al. \(2019\)](#) in which proposed algorithms are based on a strategy of adapting current state rapidly and fading old history memory (e.g. sliding window, discount factor, and restarting). In addition to this, non-stationary environments were studied under various assumptions, e.g. contextual bandits and MDPs ([Cheung et al., 2019](#); [Chen et al., 2019](#); [Zhao et al., 2020](#); [Russac et al., 2019](#); [Cheung et al., 2020](#)), mortal bandits where arms have a stochastic lifetime ([Chakrabarti et al., 2008](#); [Tracà et al., 2020](#)), and bandits where arm rewards evolve according to a continuous-time stochastic process ([Slivkins & Upfal, 2008](#)).

The multi-armed bandit problem with a finite number of arms, where each arm’s mean reward decays with the number of pulls of this arm, was first studied by [Komiyama & Qin \(2014\)](#); [Heidari et al. \(2016\)](#); [Bouneffouf & Féraud \(2016\)](#); [Levine et al. \(2017\)](#). Following [Levine et al. \(2017\)](#), this problem is referred to as *rotting bandits* problem. [Levine et al. \(2017\)](#) showed that a sliding-window algorithm has a  $\tilde{O}(K^{1/3}T^{2/3})$  regret in a non-parametric rested rotting setting where the only assumption is that mean rewards are positive and non-increasing in the number of pulls. The non-parametric rotting bandit problem, allowing mean rewards to be negative with bounded decay, was subsequently studied by [Seznec et al. \(2019\)](#), showing an algorithm that has an  $\tilde{O}(\sqrt{KT})$  problem instance independent bound. [Seznec et al. \(2020\)](#) showed that a single algorithm, an adaptive-window UCB index policy, achieves near-optimal regret for both rested and restless rotting bandits. In this paper, we follow the non-parametric rested rotting setting where mean rewards can only decrease with bounded decrements.

We next discuss the case of infinitely many arms with stationary rewards. [Berry et al. \(1997\)](#); [Bonald & Proutière \(2013\)](#) proposed algorithms with asymptotically optimal regret  $O(\sqrt{T})$  for the case of arms with Bernoulli rewards and independent mean values according to uniform distribution on  $[0, 1]$ . [Wang et al. \(2009\)](#) studied the case where the mean reward distribution has support on  $[0, \mu^*]$  with  $\mu^* \leq 1$ , and for each arm  $a$  the distribution of mean reward  $\mu(a)$  is such that  $\mathbb{P}(\mu(a) \geq \mu^* - z) = \Theta(z^\beta)$ , for some  $\beta > 0$ . [Carpentier & Valko \(2015\)](#) studied the same problem but focused on simple regret, defined as the instantaneous regret at time step  $T$ . [Bayati et al. \(2020\)](#) showed that a subsampled UCB algorithm (SSUCB) that samples  $\Theta(\sqrt{T})$  arms and executes UCB only on this subset of arms has  $\tilde{O}(\sqrt{T})$  regret under 1-sub-Gaussian rewards with mean rewards according to uniform distribution on  $[0, 1]$ . In this

setting, for mean reward distributions such that there is a large enough number of near-optimal arms, an algorithm may find a near-optimal arm by exploring a restricted number of arms. There also exist several works dealing with infinitely many arms under structured reward functions such as contextual linear bandits (Abbasi-Yadkori et al., 2011) and Lipschitz bandits (Bubeck et al., 2011). In this paper, however, we focus on infinitely many arms under a mean reward distribution, where the structured-reward assumptions may not hold because of rotting.

Our work is different from the work discussed in this section in that we consider the case of infinitely many arms with non-stationary rotting arms. In the case of rotting bandits with a finite number of arms, as we mentioned, Seznec et al. (2019; 2020) achieves worst-case regret bound  $\tilde{O}(\sqrt{KT})$  which matches the near-optimal regret in the stationary stochastic setting. This result indicates that the rotting in the finitely many arms setting is not a harder problem than in the stationary rewards setting. However, in the setting of infinitely many arms, rotting of rewards makes the problem harder than in the stationary rewards case. This is because the value of the optimal mean reward is not decreasing as the arms are being pulled as there are infinitely many near-optimal arms, which requires an additional exploration to recurrently search for a new optimal arm outside of the set of already pulled arms. Our algorithms are different from previously-proposed algorithms for the case of infinitely many arms with stationary rewards in keeping to explore new arms to find a near-optimal arm over time because of rotting rewards. In more details, our algorithms use UCB policies to decide whether to continue pulling an arm or remove the arm from further consideration and explore a new arm, by comparing its UCB index with a threshold which is adjusted by using the rotting rate or an estimated value of the rotting rate.

## 2. Problem formulation

We consider a non-stationary bandit problem with infinitely many arms where the reward distributions of arms vary over time. We consider the case when the mean reward of an arm may decrease only when this arm is pulled by an agent that uses a policy  $\pi$ , which is referred to as the *rested rotting bandit* setting. Let  $\mathcal{A}$  be an infinite set of arms,  $\mu_t(a)$  be the mean reward of arm  $a$  at time  $t$  before pulling an arm at time  $t$ , and  $n_t(a)$  be the number of times arm  $a \in \mathcal{A}$  is pulled by  $\pi$  before time  $t$ . Also, denote by  $r_t$  the stochastic reward gained by pulling arm  $a_t^\pi$  at time  $t$ . Let  $r_t = \mu_t(a_t^\pi) + \eta_t$  where  $\eta_t$  is a noise term with a 1-sub-Gaussian distribution. We assume that initial mean rewards  $\{\mu_1(a)\}_{a \in \mathcal{A}}$  are i.i.d. random variables with uniform distribution on  $[0, 1]$ . The rotting of arms is defined as follows. Given a rotting rate  $0 \leq \varrho_t \leq \varrho$  at time  $t \geq 1$  with maximum rotting rate  $\varrho = o(1)$ , the mean reward of the selected arm at time  $t$

changes as follows

$$\mu_{t+1}(a_t) = \mu_t(a_t) - \varrho_t$$

whereas the mean rewards of other arms remain unchanged. The mean reward of every arm  $a \in \mathcal{A}$  at time  $t > 1$  can be represented as follows. With  $0 \leq \varrho_s \leq \varrho$  for all time steps  $0 < s < t$ ,

$$\mu_t(a) = \mu_1(a) - \sum_{s=1}^{t-1} \varrho_s \mathbb{1}(a_s = a).$$

The objective is to find a policy that minimizes the expected cumulative regret over a time horizon of  $T$  time steps, which for a given policy  $\pi$  is defined as follows

$$\mathbb{E}[R^\pi(T)] = \mathbb{E} \left[ \sum_{t=1}^T (1 - \mu_t(a_t^\pi)) \right].$$

In the regret definition, we use that the mean reward of the optimal arm at any time  $t$  is equal to 1. This is because there is an infinite number of arms in  $\mathcal{A}$  with i.i.d. mean rewards according to uniform distribution on  $[0, 1]$ , so that there always exist sufficiently many arms whose mean rewards are close enough to 1. In what follows, ‘selecting an arm’ means that a policy chooses an arm in  $\mathcal{A}$  before playing it and ‘pulling an arm’ means that the policy plays a selected arm and receives a reward.

## 3. Regret lower bound

We first discuss two different regimes for regret depending on the value of the maximum rotting rate  $\varrho$ . When  $\varrho \leq 1/T^{3/2}$ , the mean reward of any arm over the time horizon of  $T$  time steps changes for at most  $\varrho T \leq 1/\sqrt{T}$ . Therefore, for any arm, there can be at most a gap of  $1/\sqrt{T}$  between the initial mean reward and the mean reward after  $T$  time steps, which causes an additional regret of at most  $\sqrt{T}$  over the horizon of  $T$  time steps to the case of stationary arms (i.e. when  $\varrho = 0$ ). In Theorem 3 in Wang et al. (2009), the optimal regret for the stationary case, with uniform distribution of mean rewards of arms, is shown to be of the order  $\sqrt{T}$ . Therefore, the extra regret of  $\sqrt{T}$  from the rotting with  $\varrho \leq 1/T^{3/2}$  does not affect the order of the regret. When  $\varrho > 1/T^{3/2}$ , we expect that the regret lower bound may be different than for the stationary case.

By analyzing the regret lower bound for the specific case with  $\varrho_t = \varrho$  for all time steps  $t > 0$ , we provide a lower bound for the worst-case regret with respect to arbitrary rotting as given in the following theorem.

**Theorem 3.1.** *For the rotting infinitely many-armed bandit problem, there exist rotting rates  $0 \leq \varrho_t \leq \varrho = o(1)$  for all time steps  $t > 0$  such that any policy  $\pi$  has the regret over*

a time horizon of  $T$  time steps such that

$$\mathbb{E}[R^\pi(T)] = \Omega(\max\{\varrho^{1/3}T, \sqrt{T}\}).$$

From the result of the theorem, when the rotting is small enough, i.e. precisely when  $\varrho \leq 1/T^{3/2}$ , the lower bound corresponds to  $\Omega(\sqrt{T})$ , which is known to hold when rewards are stationary. Otherwise, when the rotting is sufficiently large, then the lower bound is  $\Omega(\varrho^{1/3}T)$ . For example, when  $\varrho = 1/T^\gamma$  for some  $\gamma > 0$ , we have the lower bound  $\Omega(\sqrt{T})$ , if  $\gamma \geq 3/2$  (small rotting case), and, otherwise (large rotting case), we have  $\Omega(T^{1-\gamma/3})$ . We can observe that  $\varrho = \Theta(1/T^{3/2})$  is a transition point at which the lower bound switches from  $\Omega(\sqrt{T})$  to  $\Omega(\varrho^{1/3}T)$ .

*Proof sketch.* Here we present a proof sketch of the theorem with the full version of the proof provided in Appendix A.2.

We assume that  $\varrho_t = \varrho = o(1)$  for all time steps  $t > 0$ . When  $\varrho = O(1/T^{3/2})$ , the lower bound for the stationary case of the order  $\sqrt{T}$  (Wang et al., 2009) is tight enough for the non-stationary case. This is because we only need to pay an extra regret of at most of order  $\sqrt{T}$  for small  $\varrho$ . Therefore, when  $\varrho = O(1/T^{3/2})$ , we have

$$\mathbb{E}[R^\pi(T)] = \Omega(\sqrt{T}). \quad (1)$$

We note that even though the mean rewards are rotting in our setting, we can easily obtain (1) by following the same proof steps of Theorem 3 in Wang et al. (2009). For the sake of completeness, we provide a proof in the Appendix A.2.

When  $\varrho = \omega(1/T^{3/2})$ , however, the lower bound of the stationary case is not tight enough. Here we provide the proof of the lower bound  $\Omega(\varrho^{1/3}T)$  for the case when  $\varrho = \omega(1/T^{3/2})$ . For showing the lower bound, we will classify each arm to be either bad or good or else according to the definition given shortly. To distinguish bad and good arms, we use two thresholds  $1 - c$  and  $1 - \delta$ , respectively, where  $c$  and  $\delta$  are such that  $0 < 1 - c < 1 - \delta < 1$ ,  $\delta = \varrho^{1/3}$ , and  $c$  is a constant. An arm  $a$  is said to be a *bad arm* if  $\mu_1(a) \leq 1 - c$ , and is said to be a *good arm* if  $\mu_1(a) > 1 - \delta$ .

Let  $N_T$  be the number of distinct selected good arms until time step  $T$ . We separately consider two cases when  $N_T < m$  and  $N_T \geq m$ , where  $m = \lceil (1/2)T\varrho^{2/3} \rceil$ , and show that each case has  $\Omega(\varrho^{1/3}T)$  as the regret lower bound. The main ideas for each case are outlined as follows. When the number of selected good arms is relatively small ( $N_T < m$ ), any policy  $\pi$  must pull arms with mean rewards less than  $1 - \delta$  at least  $T/2$  time steps until  $T$ , amounting to at least  $\delta$  regret for each pull (gap between 1 and mean reward of a pulled arm). Therefore, the regret is lower bounded by  $\Omega(\delta(T/2)) = \Omega(\varrho^{1/3}T)$ . When the number of selected good arms is relatively large ( $N_T \geq m$ ), we can show

---

**Algorithm 1** UCB-Threshold Policy (UCB-TP)
 

---

Given:  $T, \delta, \mathcal{A}$ ; Initialize:  $\mathcal{A}' \leftarrow \mathcal{A}$   
 Select an arm  $a \in \mathcal{A}'$   
 Pull arm  $a$  and get reward  $r_1$   
**for**  $t = 2, \dots, T$  **do**  
   Update the initial mean reward estimator  $\tilde{\mu}_t^\varrho(a)$   
   **if**  $\text{UCB}_t(a) \geq 1 - \delta$  **then**  
     Pull arm  $a$  and get reward  $r_t$   
   **else**  
      $\mathcal{A}' \leftarrow \mathcal{A}' / \{a\}$   
     Select an arm  $a \in \mathcal{A}'$   
     Pull arm  $a$  and get reward  $r_t$   
   **end if**  
**end for**

---

that any policy  $\pi$  is likely to select at least of order  $\varrho^{1/3}T$  number of distinct bad arms until  $T$ . From the fact that the selected bad arms are pulled at least once and each pull adds a constant regret of value at least  $c$ , the regret is shown to be lower bounded by  $\Omega(c\varrho^{1/3}T) = \Omega(\varrho^{1/3}T)$ . Therefore, when  $\varrho = \omega(1/T^{3/2})$ , we can obtain

$$\begin{aligned} \mathbb{E}[R^\pi(T)] &= \mathbb{E}[R^\pi(T)\mathbb{1}(N_T < m) + R^\pi(T)\mathbb{1}(N_T \geq m)] \\ &= \Omega(\varrho^{1/3}T). \end{aligned} \quad (2)$$

Finally, from (1) and (2), we have  $\mathbb{E}[R^\pi(T)] = \Omega(\max\{\varrho^{1/3}T, \sqrt{T}\})$ .  $\square$

## 4. Algorithms and regret upper bounds

In this section, we first present an algorithm for the rested rotting bandit problem with infinitely many arms for the case when the algorithm knows the value of the maximum rotting rate. We show a regret upper bound of the algorithm that matches the regret lower bound in Theorem 3.1 up to a poly-logarithmic factor. Second, we present an algorithm that does not know the maximum rotting rate and show a regret upper bound that matches the regret lower bound up to a poly-logarithmic factor, when the maximum rotting rate is large enough.

### 4.1. An algorithm knowing maximum rotting rate

We present an algorithm which requires knowledge of the maximum rotting rate in Algorithm 1. The algorithm selects an arm and pulls this arm as long as the arm is tested to be a good arm, by using a test comparing an upper confidence bound of this arm with a threshold value. Specifically, if  $a$  is the selected arm at time step  $t$ , the algorithm computes an estimator  $\tilde{\mu}_t^\varrho(a)$  of the initial mean reward of arm  $a$  and uses this estimator to compute an estimator of the mean reward of arm  $a$  at time step  $t$ , considering the worst-case rotting



rate  $\varrho$  for the estimators. Comparing the upper confidence bound for the mean reward with the threshold  $1 - \delta$ , the algorithm tests whether the arm is a good arm. If the arm is tested to be a good arm, then the algorithm continues to pull this arm. Otherwise, it discards the arm and selects a new one, and repeats the procedure described above until time horizon  $T$ . We consider Algorithm 1 with the initial mean reward estimator defined as

$$\tilde{\mu}_t^o(a) := \frac{\sum_{s=1}^{t-1} (r_s + \varrho n_s(a)) \mathbb{1}(a_s = a)}{n_t(a)}$$

and the upper confidence bound term defined as

$$\text{UCB}_t(a) := \tilde{\mu}_t^o(a) - \varrho n_t(a) + \sqrt{8 \log(T)/n_t(a)}.$$

Note that when  $\varrho_s = \varrho$  for all  $0 < s < t$ ,  $\tilde{\mu}_t^o(a)$  is an unbiased estimator of the initial mean reward  $\mu_1(a)$  of arm  $a$  and  $\tilde{\mu}_t^o(a) - \varrho n_t(a)$  is an unbiased estimator of the mean reward  $\mu_t(a)$  of arm  $a$  at time step  $t$ . The upper confidence bound  $\text{UCB}_t(a)$  follows the standard definition of an upper confidence bound. By designing the estimators to deal with the maximum rotting rate  $\varrho$ , for any arbitrary  $\rho_t \leq \varrho$  for all time steps  $t > 0$ , we show that it has a near-optimal worst-case regret upper bound in the following theorem.

**Theorem 4.1.** *For the policy  $\pi$  defined by Algorithm 1 with  $\delta = \max\{\varrho^{1/3}, 1/\sqrt{T}\}$ , and  $\varrho = o(1)$ , the regret satisfies*

$$\mathbb{E}[R^\pi(T)] = \tilde{O}(\max\{\varrho^{1/3}T, \sqrt{T}\}).$$

*Proof sketch.* Here we present a proof sketch of the theorem with a full version of the proof provided in Appendix A.3.

Observe that initial mean rewards of selected arms are i.i.d. random variables with uniform distribution on  $[0, 1]$ . We first define arms to be good or bad arms depending on the initial mean rewards. We assume  $\varrho = o(1)$  and set  $\delta = \max\{\varrho^{1/3}, 1/\sqrt{T}\}$ . Then we define an arm  $a$  to be a good arm if  $\Delta(a) \leq \delta/2$ , where  $\Delta(a) = 1 - \mu_1(a)$ , and otherwise,  $a$  is a bad arm. Because of rotting, initially good arm may become bad by pulling the arm several times. Therefore, the policy  $\pi$  may select several good arms over the entire time and we analyze the regret over time episodes defined by the selections of good arms. Given the policy, we refer to the period starting from selecting the  $i - 1$ -st good arm until selecting the  $i$ -th good arm as the  $i$ -th episode. Because of the uniform distribution of mean rewards with small  $\delta/2$ , it is likely to have several consecutive selected bad arms in each episode.

We do an episodic analysis. We first analyze the expected regret per episode and multiply it by the expected number of episodes in  $T$ . However, this proof strategy has an issue that the regret of each episode and the number of episodes




Figure 1: Episodes in a time line.

in  $T$  are not independent. To resolve the issue, we fix the number of episodes to  $m^G$  and analyze the regret not for  $T$  but for  $m^G$  episodes. Note that  $m^G$  is a fixed value, and the time after  $m^G$  episodes can exceed  $T$ . For obtaining a regret bound, we set  $m^G$  so that the total number of time steps for  $m^G$  episodes is larger than  $T$  with high probability and thus  $R^\pi(T) \leq R_{m^G}^\pi$  where  $R_{m^G}^\pi$  is the regret accumulated for  $m^G$  episodes. For the regret analysis, we denote by  $m_i^B$  the number of selections of distinct bad arms in the  $i$ -th episode. See Figure 1 for an illustration of the episodes in a time line.

In what follows, we provide an overview of the regret analysis by considering two separate cases depending on the value of the maximum rotting rate  $\varrho$ ; one for a large rotting case, and the other for a small rotting case, which we may interpret as a near-stationary case. For the analysis, we use  $R_i^G$  to denote the regret accumulated by pulling the good arm in the  $i$ -th episode, and  $R_{i,j}^B$  to denote the regret accumulated by pulling the  $j$ -th bad arm in the  $i$ -th episode.

**Case of large rotting**  $\varrho = \omega(1/T^{3/2})$ : We first show that by setting  $m^G = \lceil 2T\varrho^{2/3} \rceil$ , we have  $R^\pi(T) \leq R_{m^G}^\pi$ . If the policy selects a good arm  $a$ , where  $\mu_1(a) \geq 1 - \delta/2$ , then it must pull the arm at least  $\delta/(2\varrho)$  times, with high probability, to decrease the mean reward below the threshold  $1 - \delta$ . Then the total number of time steps for  $m^G$  episodes is at least  $(\delta/(2\varrho))m^G = (1/(2\varrho^{2/3}))\lceil 2T\varrho^{2/3} \rceil \geq T$ . This implies that  $R^\pi(T) \leq R_{m^G}^\pi$ . We next provide a bound for  $\mathbb{E}[R^\pi(T)]$  using  $\mathbb{E}[R_{m^G}^\pi]$ . For bounding  $\mathbb{E}[R_{m^G}^\pi]$ , we can show that for any  $i \in [m^G]$  and  $j \in [m_i^B]$ , we have

$$\mathbb{E}[R_i^G] = \tilde{O}\left(\frac{1}{\varrho^{1/3}}\right) \text{ and } \mathbb{E}[R_{i,j}^B] = \tilde{O}(1). \quad (3)$$

Observe that  $m_1^B, \dots, m_{m^G}^B$  are i.i.d. random variables with geometric distribution with parameter  $\delta/2$ . Therefore, for any non-negative integer  $k$ , we have  $\mathbb{P}(m_i^B = k) = (1 - \delta/2)^k \delta/2$  and  $\mathbb{E}[m_i^B] = (2/\delta) - 1$  for all  $i \in [m^G]$ . We have set  $\delta = \max\{\varrho^{1/3}, 1/\sqrt{T}\}$ . Then when  $\varrho = \omega(1/T^{3/2})$ , with  $m^G = \lceil 2T\varrho^{2/3} \rceil$  and from (3),  $\mathbb{E}[m_i^B] = 2/\delta - 1$ , and  $\delta = \varrho^{1/3}$ , we have

$$\begin{aligned} \mathbb{E}[R^\pi(T)] &= O(\mathbb{E}[R_{m^G}^\pi]) \\ &= O\left(\mathbb{E}\left[\sum_{i \in [m^G]} \left(R_i^G + \sum_{j \in [m_i^B]} R_{i,j}^B\right)\right]\right) \\ &= \tilde{O}\left(\left(\frac{1}{\varrho^{1/3}} + \mathbb{E}[m_i^B]\right) m^G\right) = \tilde{O}\left(\varrho^{1/3}T\right). \end{aligned} \quad (4)$$

**Case of small rotting**  $\varrho = O(1/T^{3/2})$ : By setting  $m^{\mathcal{G}} = C$  for some large constant  $C > 0$ , we can show that  $R^\pi(T) \leq R_{m^{\mathcal{G}}}^\pi$ . This is because if the policy selects a good arm  $a$ , then it must pull the arm for at least order  $T$  times with high probability. This is because the small rotting case is a near-stationary setting so that the policy can pull a good arm for a large amount of time steps. For bounding  $\mathbb{E}[R_{m^{\mathcal{G}}}^\pi]$ , we can show that for any  $i \in [m^{\mathcal{G}}]$  and  $j \in [m_i^{\mathcal{B}}]$ ,

$$\mathbb{E}[R_i^{\mathcal{G}}] = O(\sqrt{T}) \text{ and } \mathbb{E}[R_{i,j}^{\mathcal{B}}] = \tilde{O}(1). \quad (5)$$

Then when  $\varrho = O(1/T^{3/2})$ , with  $m^{\mathcal{G}} = C$  and from (5),  $\mathbb{E}[m_i^{\mathcal{B}}] = 2/\delta - 1$ , and  $\delta = \Theta(1/\sqrt{T})$ , we have

$$\begin{aligned} \mathbb{E}[R^\pi(T)] &= O(\mathbb{E}[R_{m^{\mathcal{G}}}^\pi]) \\ &= \tilde{O}\left(\left(\sqrt{T} + \mathbb{E}[m_i^{\mathcal{B}}]\right) m^{\mathcal{G}}\right) = \tilde{O}(\sqrt{T}). \end{aligned} \quad (6)$$

We note that in the small rotting case, the policy achieves  $\tilde{O}(\sqrt{T})$  which matches a near-optimal bound for the stationary setting (Wang et al., 2009).

**Putting the pieces together:** From (4) and (6), by taking  $\varrho = o(1)$  and  $\delta = \max\{\varrho^{1/3}, 1/\sqrt{T}\}$ , it follows

$$\mathbb{E}[R^\pi(T)] = \tilde{O}(\max\{\varrho^{1/3}T, \sqrt{T}\}).$$

□

## 4.2. An algorithm not knowing maximum rotting rate

In this section we present an algorithm which does not require information about the maximum rotting rate  $\varrho$  defined in Algorithm 2, and provide a regret upper bound for this algorithm. The algorithm adopts the strategy of hierarchical bandit algorithms similar to BOB (bandit-over-bandit) (Cheung et al., 2019). It consists of a master algorithm and several base algorithms, where EXP3 (Auer et al., 2002b) is used for the master algorithm whose goal is to find a near-optimal base algorithm, and for each base algorithm, a UCB index policy similar to that in Algorithm 1 is used with a candidate rotting rate and an adaptive threshold to decide whether to continue pulling currently selected arm. In Algorithm 2, the time horizon of  $T$  time steps is partitioned into blocks of  $H$  time steps. Before starting each block, the master algorithm selects a rotting estimator  $\tilde{\beta}$  of the unknown maximum rotting rate  $\varrho$  from a set of candidate values denoted by  $\mathcal{B}$ . Then it runs a base algorithm over  $H$  time steps which decides whether to continue pulling the selected arm based on a UCB index and a threshold tuned using the selected  $\tilde{\beta}$ . By utilizing the obtained rewards over the block as a feedback for the decision of the master algorithm, it updates the master to find a near-optimal base and repeats the procedure described above until time horizon  $T$ .

---

### Algorithm 2 Adaptive UCB-Threshold Policy (AUCB-TP)

---

Given:  $T, H, \mathcal{B}, \mathcal{A}, \alpha$   
 Initialize:  $\mathcal{A}' \leftarrow \mathcal{A}, w(\beta) \leftarrow 1$  for  $\beta \in \mathcal{B}$   
**for**  $i = 1, 2, \dots, \lceil T/H \rceil$  **do**  
     Select an arm  $a \in \mathcal{A}'$   
     Pull arm  $a$  and get reward  $r_{(i-1)H+1}$   
      $p(\beta) \leftarrow (1 - \alpha) \frac{w(\beta)}{\sum_{k \in \mathcal{B}} w(k)} + \alpha \frac{1}{B}$  for  $\beta \in \mathcal{B}$   
     Select  $\tilde{\beta} \leftarrow \beta$  with probability  $p(\beta)$  for  $\beta \in \mathcal{B}$   
      $\delta \leftarrow \tilde{\beta}^{1/3}$   
     **for**  $t = (i-1)H + 2, \dots, i \cdot H \wedge T$  **do**  
         **if**  $\text{UCB}_{i,t}(a, \tilde{\beta}) \geq 1 - \delta$  **then**  
             Pull arm  $a$  and get reward  $r_t$   
         **else**  
              $\mathcal{A}' \leftarrow \mathcal{A}' \setminus \{a\}$   
             Select an arm  $a \in \mathcal{A}'$   
             Pull arm  $a$  and get reward  $r_t$   
         **end if**  
     **end for**  
      $w(\tilde{\beta})$   
      $\leftarrow w(\tilde{\beta}) \exp\left(\frac{\alpha}{Bp(\tilde{\beta})} \left(\frac{1}{2} + \frac{\sum_{t=(i-1)H}^{i \cdot H \wedge T} r_t}{186H \log T + 4\sqrt{H \log T}}\right)\right)$   
**end for**

---

We note that the term  $1/2 + \sum_{t=(i-1)H}^{i \cdot H \wedge T} r_t / (186H \log T + 4\sqrt{H \log T})$  in updating  $w(\tilde{\beta})$  in Algorithm 2 is for re-scaling and translating rewards, which makes the rewards lie in  $[0, 1]$  with high probability. Also by optimizing the block size  $H$ , we can control regrets induced from the master and a base. By increasing  $H$ , the regret induced from the master increases and the regret induced from a base decreases. Those facts are shown later in the poof of Theorem 4.2.

In what follows, we define the inputs  $\mathcal{B}$  and  $\alpha$ , and the upper confidence bound index  $\text{UCB}_{i,t}(a, \beta)$  for  $\beta \in \mathcal{B}$ .  $\mathcal{B}$  contains candidate values of  $\beta$  to optimize  $\text{UCB}_{i,t}(a, \beta)$  and the threshold parameter  $\delta$ . We find that the optimal base parameter  $\beta^\dagger \in \mathcal{B}$  is when  $\beta^\dagger = \max\{1/H^{3/2}, \varrho\}$  including a clipped domain for the optimal threshold value  $\delta$  as in Theorem 4.1. This implies that the optimized  $\beta^\dagger \geq 1/H^{3/2} = 2^{-3/2 \log_2 H}$ . Also from  $\varrho = o(1)$ ,  $\beta^\dagger \leq 1/8$ . Therefore, we set

$$\mathcal{B} = \{2^{-3}, 2^{-4}, \dots, 2^{-\lceil (3/2) \log_2 H \rceil}\},$$

in which the cardinality of the set is restricted by  $O(\log H)$  which does not hurt the regret from EXP3 up to a logarithmic factor. Let  $B = |\mathcal{B}|$ . Then we set  $\alpha = \min\{1, \sqrt{B \log B / ((e-1)\lceil T/H \rceil)}\}$ , which is used to guarantee a least selection probability for each base. Let  $n_{i,t}(a)$  be the number of times that arm  $a \in \mathcal{A}$  is pulled by the algorithm from time step  $(i-1)H + 1$  before time step

t. Let  $\tilde{\mu}_{i,t}^o(a, \beta)$  be defined as

$$\tilde{\mu}_{i,t}^o(a, \beta) := \frac{\sum_{s=(i-1)H+1}^{t-1} (r_s + \beta n_{i,s}(a)) \mathbb{1}(a_s = a)}{n_{i,t}(a)}$$

and the  $\text{UCB}_{i,t}(a, \beta)$  index be defined as

$$\text{UCB}_{i,t}(a, \beta) := \tilde{\mu}_{i,t}^o(a, \beta) - \beta n_{i,t}(a) + \sqrt{8 \log(H)/n_{i,t}(a)}.$$

We provide a worst-case regret upper bound for Algorithm 2 in the following theorem.

**Theorem 4.2.** *With  $\varrho = o(1)$ , for the policy  $\pi$  defined by Algorithm 2 with  $H = \lceil T^{1/2} \rceil$ , the regret satisfies*

$$\mathbb{E}[R^\pi(T)] = \tilde{O}(\max\{\varrho^{1/3}T, T^{3/4}\}).$$

*Proof sketch.* Here we present a proof sketch of the theorem with a full version of the proof given in Appendix A.4.

The policy  $\pi$  consists of two strategies: EXP3 for the master and UCB-Threshold Policy (Algorithm 1) for bases. We can decompose the regret into two parts: regret incurred by playing a base with  $\beta \in \mathcal{B}$  over each block of  $H$  time steps and regret incurred due to the master trying to find a near-optimal base parameter in  $\mathcal{B}$ . In what follows, we define the regret decomposition formally. Let  $\pi_i(\beta)$  for  $\beta \in \mathcal{B}$  denote the base policy with  $\beta$  for time steps between  $(i-1)H+1$  and  $i \cdot H \wedge T$ . Denote by  $a_t^{\pi_i(\beta)}$  the pulled arm at time step  $t$  by policy  $\pi_i(\beta)$ . Then, for  $\beta^\dagger \in \mathcal{B}$ , which is set later for a near-optimal base, we have

$$\begin{aligned} \mathbb{E}[R^\pi(T)] &= \mathbb{E} \left[ \sum_{t=1}^T 1 - \sum_{i=1}^{\lceil T/H \rceil} \sum_{t=(i-1)H+1}^{i \cdot H \wedge T} \mu_t(a_t^\pi) \right] \\ &= \mathbb{E}[R_1^\pi(T)] + \mathbb{E}[R_2^\pi(T)], \end{aligned} \quad (7)$$

where

$$\begin{aligned} R_1^\pi(T) &= \sum_{t=1}^T 1 - \sum_{i=1}^{\lceil T/H \rceil} \sum_{t=(i-1)H+1}^{i \cdot H \wedge T} \mu_t(a_t^{\pi_i(\beta^\dagger)}), \\ R_2^\pi(T) &= \sum_{i=1}^{\lceil T/H \rceil} \sum_{t=(i-1)H+1}^{i \cdot H \wedge T} \left( \mu_t(a_t^{\pi_i(\beta^\dagger)}) - \mu_t(a_t^\pi) \right). \end{aligned}$$

Note that  $R_1^\pi(T)$  accounts for the regret caused by the near-optimal base algorithm  $\pi_i(\beta^\dagger)$  against the optimal mean reward and  $R_2^\pi(T)$  accounts for the regret caused by the master algorithm by selecting a base with  $\beta \in \mathcal{B}$  at every block against the base with  $\beta^\dagger$  as illustrated in Figure 2. We set  $\beta^\dagger$  to be the smallest value in  $\mathcal{B}$  which is larger than  $\max\{\varrho, 1/H^{3/2}\}$ . Then the base has the threshold parameter  $\beta^{\dagger 1/3}$  of order  $\max\{\varrho^{1/3}, 1/\sqrt{H}\}$  which coincides with




Figure 2: Regret decomposition for Algorithm 2.

the optimal threshold parameter of Algorithm 1 by replacing  $T$  with  $H$ . We note that the policy  $\pi$  does not require knowing  $\beta^\dagger$  and it is defined only for the proof.

In what follows, we provide upper bounds for each regret component. We first provide an upper bound for  $\mathbb{E}[R_1^\pi(T)]$  by following the proof steps in Theorem 4.1. We can easily find that regret of the base with  $\beta^\dagger$  for each block of size  $H$  that has the same regret bound as in Theorem 4.1 by replacing  $T$  with  $H$  amounting to  $\tilde{O}(\max\{\varrho^{1/3}H, \sqrt{H}\})$ . Then by adding the regret for  $\lceil T/H \rceil$  number of blocks, we have

$$\begin{aligned} \mathbb{E}[R_1^\pi(T)] &= \tilde{O}((T/H) \max\{\varrho^{1/3}H, \sqrt{H}\}) \\ &= \tilde{O}(\max\{\varrho^{1/3}T, T/\sqrt{H}\}). \end{aligned} \quad (8)$$

Then we provide an upper bound for  $\mathbb{E}[R_2^\pi(T)]$  using a regret bound for EXP3 in (Auer et al., 2002b). The EXP3 in policy  $\pi$  selects a base in  $\mathcal{B}$  before starting a block and gets feedback at the end of the block and repeats this over  $\lceil T/H \rceil$  blocks. Therefore, EXP3 in  $\pi$  can be thought to be run for  $\lceil T/H \rceil$  decision rounds and the number of decision options for each round is  $B$ . Let  $Q$  be an upper bound for the absolute sum of rewards for any block with length  $H$  with high probability. Then from Corollary 3.2 in (Auer et al., 2002b), we can show that

$$\mathbb{E}[R_2^\pi(T)] = \tilde{O}(Q\sqrt{B(T/H)}). \quad (9)$$

By considering that mean rewards may become negative because of rotting and using a Chernoff's bound, we show that with high probability

$$\left| \sum_{t=(i-1)H+1}^{i \cdot H \wedge T} r_t \right| \leq 93H \log T + 2\sqrt{H \log T}. \quad (10)$$

Then with  $B = O(\log T)$ , from (9) and (10), we have

$$\begin{aligned} \mathbb{E}[R_2^\pi(T)] &= \tilde{O} \left( H \log(T) \sqrt{B(T/H)} \right) \\ &= \tilde{O} \left( \sqrt{HT} \right). \end{aligned} \quad (11)$$

Finally, from (7), (8), and (11), with  $H = \lceil T^{1/2} \rceil$ , we have

$$\begin{aligned} \mathbb{E}[R^\pi(T)] &= \tilde{O}(\max\{\varrho^{1/3}T, T/\sqrt{H}\} + \sqrt{HT}) \\ &= \tilde{O}(\max\{\varrho^{1/3}T, T^{3/4}\}). \end{aligned}$$

This concludes the proof.  $\square$

The regret bound for Algorithm 2 in Theorem 4.2 is larger than or equal to that for Algorithm 1 in Theorem 4.1. This is because the master in Algorithm 2 needs to learn the unknown maximum rotting rate to find a near-optimal base algorithm, which produces extra regret. In the following remarks, we discuss the region of the maximum rotting rate  $\varrho$  for which Algorithm 2 achieves the near-optimal regret and discuss the computation and memory efficiency of our algorithms.

*Remark 4.3.* When  $\varrho = \Omega(1/T^{3/4})$ , Algorithm 2 achieves the optimal regret bound  $\tilde{O}(\varrho^{1/3}T)$  up to a poly-logarithmic factor. This is because when  $\varrho = \Omega(1/T^{3/4})$ , the additional regret from learning the maximum rotting rate is negligible compared with the regret from the rotting of rewards. It is an open problem to achieve the optimal regret bound for any value of  $\varrho$ , without knowing the value of this parameter.

*Remark 4.4.* Note that we can achieve the same regret bounds as in Theorems 4.1 and 4.2 for the worst-case rotting by replacing our UCB index with adaptive-window UCB index in (Sez nec et al., 2020) which is known to achieve near-optimal regret for the case of rotting with a finite number of arms. However, computing the adaptive-window UCB over horizon  $T$  has high computation cost,  $O(T^2)$  and  $O(TH)$ , and memory space  $O(T)$  and  $O(H)$ , in Algorithms 1 and 2, respectively, for optimizing the window size. On the other hand, our proposed UCB index has a lower computation cost of  $O(T)$  and requires only  $O(1)$  memory space using a simple trick for updating mean estimators in online computation settings (Welford, 1962).

## 5. Numerical experiments

In this section we present results of our numerical experiments using synthetic data in the rotting setting with infinitely many arms<sup>1</sup>. To the best of our knowledge, there are no previously-proposed algorithms for the setting of rotting with infinitely many arms. We compare the performance of Algorithms 1 and 2 with SSUCB (Bayati et al., 2020), which was proposed for infinitely many arms with stationary rewards, and is known to have a near-optimal regret for stationary sub-Gaussian reward distributions. From the theoretical results in Section 3 and Section 4, we expect that Algorithm 1 and SSUCB would have similar performance when the maximum rotting rate is sufficiently small, which may be regarded as a nearly stationary case, and that both Algorithm 1 and Algorithm 2 would outperform SSUCB when the rotting rate is sufficiently large. We expect that Algorithm 2 would not be competitive in the nearly stationary case because of the extra regret from the rotting rate estimation. It requires the rotting rate to be sufficiently large

<sup>1</sup>Our code is available at [https://github.com/junghunkim7786/rotting\\_infinite\\_armed\\_bandits](https://github.com/junghunkim7786/rotting_infinite_armed_bandits)




Figure 3: Performance of Algorithms 1 and 2, and SSUCB: (a) regret versus  $\varrho^{1/3}$  for fixed time horizon  $T$ , and (b,c,d) regret versus time horizon  $T$  with rotting rates  $\varrho = 1/T^{1/2}$  (b),  $1/T^{3/5}$  (c), and  $1/T^{3/2}$  (d).

to have the near-optimal regret bound. We will confirm these insights by our numerical results in what follows.

In our experiments, we consider the case of identical rotting with  $\varrho_t = \varrho$  for all  $t \in [T]$ . We generate initial mean rewards of arms by sampling from uniform distribution on  $[0, 1]$ . In each time step, stochastic reward from pulling an arm has a Gaussian noise with mean zero and variance 1. We repeat each experiment 10 times and compute confidence intervals for confidence probability 0.95.

We first investigate the performance of algorithms for varied rotting rate  $\varrho$  and fixed time horizon  $T$ . We ran experiments for rotting rate  $\varrho$  set to  $1/T, 1/T^{0.9}, 1/T^{0.8}, \dots, 1/T^{0.3}$ , and measured the expected regret for the time horizon  $T = 10^6$ . In Figure 3 (a), we can confirm that our algorithms show more robust performance than SSUCB for various rotting rates with linearly increasing regret with respect to  $\varrho^{1/3}$  for large rotting, which matches Theorems 4.1 and 4.2. We observe that for large rotting cases, our algorithms have similar performance and both outperform SSUCB. For sufficiently small rotting, we can observe that all the three algorithms have comparable performance while Algorithm 2 performs slightly worse. These results conform to the insights derived from our theoretical analysis.

We next investigate the performance of algorithms versus time horizon  $T$  and rotting rate  $\varrho$  depending on  $T$ . We ran experiments for time horizon  $T$  taking values  $1, 1 \times 10^5, 2 \times 10^5, \dots, 10^6$ , and measured expected regret of each case. We set the rotting rate  $\varrho$  to  $1/T^{1/2}, 1/T^{3/5}$  and  $1/T^{3/2}$ . Note that the case  $1/T^{3/2}$  may be considered as a



nearly-stationary case, because in this case the regret lower bound is  $\Omega(\sqrt{T})$  from Theorem 3.1. In Figure 3 (b) and (c), corresponding to large rotting rates, we observe that Algorithms 1 and 2 have similar performance and outperform SSUCB. The gaps between regrets of our algorithms and SSUCB become smaller by decreasing the rotting rate from  $1/T^{1/2}$  to  $1/T^{3/5}$ . This is because SSUCB is designed for the case of stationary rewards and has a near-optimal regret in the stationary case. In the near-stationary case, in Figure 3 (d), we observe that Algorithm 1 has best performance and, as expected, Algorithm 2 has worst performance.

## 6. Conclusion

In this paper we studied the infinitely many-armed bandit problem with rested rotting rewards. We provided a regret lower bound and proposed an algorithm which achieves a near-optimal regret, when the the maximum rotting rate is known to the algorithm. We also proposed an algorithm which does not require knowledge of the rotting rate and we showed that it achieves a near-optimal regret for any large enough rotting rate. In future work, it may be of interest to relax the assumption of uniform distribution for initial mean rewards.

## References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24, 2011.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002a.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002b.
- Auer, P., Gajane, P., and Ortner, R. Adaptively tracking the best bandit arm with an unknown number of distribution changes. In *Conference on Learning Theory*, pp. 138–158. PMLR, 2019.
- Bayati, M., Hamidi, N., Johari, R., and Khosravi, K. Unreasonable effectiveness of greedy algorithms in multi-armed bandit with many arms. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1713–1723. Curran Associates, Inc., 2020.
- Berry, D. A., Chen, R. W., Zame, A., Heath, D. C., and Shepp, L. A. Bandit problems with infinitely many arms. *The Annals of Statistics*, 25(5):2103 – 2116, 1997.
- Besbes, O., Gur, Y., and Zeevi, A. Stochastic multi-armed-bandit problem with non-stationary rewards. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1, NIPS’14*, pp. 199–207, Cambridge, MA, USA, 2014. MIT Press.
- Bonald, T. and Proutière, A. Two-target algorithms for infinite-armed bandits with bernoulli rewards. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS’13*, pp. 2184–2192, Red Hook, NY, USA, 2013. Curran Associates Inc.
- Bouneffouf, D. and Féraud, R. Multi-armed bandit problem with known trend. *Neurocomputing*, 205:16–21, 2016.
- Brown, D. G. How I wasted too long finding a concentration inequality for sums of geometric variables. *Found at <https://cs.uwaterloo.ca/~brown/g/negbin.pdf>*, 8(4), 2011.
- Bubeck, S., Stoltz, G., and Yu, J. Y. Lipschitz bandits without the lipschitz constant. In *International Conference on Algorithmic Learning Theory*, pp. 144–158. Springer, 2011.
- Carpentier, A. and Valko, M. Simple regret for infinitely many armed bandits. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37, ICML’15*, pp. 1133–1141. JMLR.org, 2015.
- Chakrabarti, D., Kumar, R., Radlinski, F., and Upfal, E. Mortal multi-armed bandits. In *Proceedings of the 21st International Conference on Neural Information Processing Systems, NIPS’08*, pp. 273–280, Red Hook, NY, USA, 2008. Curran Associates Inc.
- Chen, Y., Lee, C.-W., Luo, H., and Wei, C.-Y. A new algorithm for non-stationary contextual bandits: Efficient, optimal and parameter-free. In *Conference on Learning Theory*, pp. 696–726. PMLR, 2019.
- Cheung, W. C., Simchi-Levi, D., and Zhu, R. Learning to optimize under non-stationarity. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1079–1087. PMLR, 2019.
- Cheung, W. C., Simchi-Levi, D., and Zhu, R. Reinforcement learning for non-stationary markov decision processes: The blessing of (more) optimism. In *International Conference on Machine Learning*, pp. 1843–1854. PMLR, 2020.
- Garivier, A. and Moulines, E. On upper-confidence bound policies for switching bandit problems. In Kivinen, J., Szepesvári, C., Ukkonen, E., and Zeugmann, T. (eds.), *Algorithmic Learning Theory*, pp. 174–188, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg. ISBN 978-3-642-24412-4.

- Heidari, H., Kearns, M., and Roth, A. Tight policy regret bounds for improving and decaying bandits. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, IJCAI'16, pp. 1562–1570, 2016.
- Komiyama, J. and Qin, T. Time-decaying bandits for non-stationary systems. In Liu, T.-Y., Qi, Q., and Ye, Y. (eds.), *Web and Internet Economics*, pp. 460–466, Cham, 2014. Springer International Publishing.
- Lai, T. and Robbins, H. Asymptotically efficient adaptive allocation rules. *Adv. Appl. Math.*, 6(1):4–22, mar 1985.
- Levine, N., Crammer, K., and Mannor, S. Rotting bandits. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pp. 3077–3086, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- Rigollet, P. and Hütter, J.-C. High dimensional statistics. *Lecture notes for course 18S997*, 813:814, 2015.
- Russac, Y., Vernade, C., and Cappé, O. Weighted linear bandits for non-stationary environments. *arXiv preprint arXiv:1909.09146*, 2019.
- Seznec, J., Locatelli, A., Carpentier, A., Lazaric, A., and Valko, M. Rotting bandits are no harder than stochastic ones. In Chaudhuri, K. and Sugiyama, M. (eds.), *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 2564–2572. PMLR, 16–18 Apr 2019.
- Seznec, J., Menard, P., Lazaric, A., and Valko, M. A single algorithm for both restless and rested rotting bandits. In Chiappa, S. and Calandra, R. (eds.), *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 3784–3794. PMLR, 26–28 Aug 2020.
- Slivkins, A. Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning*, 12(1-2):1–286, 2019.
- Slivkins, A. and Upfal, E. Adapting to a changing environment: the brownian restless bandits. In *21st Annual Conference on Learning Theory - COLT 2008, Helsinki, Finland, July 9-12, 2008*, pp. 343–354, 2008.
- Tracà, S., Rudin, C., and Yan, W. Reducing exploration of dying arms in mortal bandits. In Adams, R. P. and Gogate, V. (eds.), *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115 of *Proceedings of Machine Learning Research*, pp. 156–163. PMLR, 22–25 Jul 2020.
- Tsun, A. Probability & statistics with applications to computing, 2020. URL [https://www.alextsun.com/files/Prob\\_Stat\\_for\\_CS\\_Book.pdf](https://www.alextsun.com/files/Prob_Stat_for_CS_Book.pdf).
- Wang, Y., Audibert, J.-y., and Munos, R. Algorithms for infinitely many-armed bandits. In Koller, D., Schuurmans, D., Bengio, Y., and Bottou, L. (eds.), *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc., 2009.
- Welford, B. Note on a method for calculating corrected sums of squares and products. *Technometrics*, 4(3):419–420, 1962.
- Zhao, P., Zhang, L., Jiang, Y., and Zhou, Z.-H. A simple approach for non-stationary linear bandits. In *International Conference on Artificial Intelligence and Statistics*, pp. 746–755. PMLR, 2020.

## A. Appendix

### A.1. Preliminaries

Here we state some known concentration inequalities that we use in our proofs.

**Lemma A.1** (Theorem 6.2.35 in Tsun (2020)). *Let  $X_1, \dots, X_n$  be identical independent Bernoulli random variables. Then, for  $0 < \nu < 1$ , we have*

$$\mathbb{P}\left(\sum_{i=1}^n X_i \geq (1 + \nu)\mathbb{E}\left[\sum_{i=1}^n X_i\right]\right) \leq \exp\left(-\frac{\nu^2\mathbb{E}[\sum_{i=1}^n X_i]}{3}\right).$$

**Lemma A.2** (Corollary 1.7 in Rigollet & Hütter (2015)). *Let  $X_1, \dots, X_n$  be independent random variables with  $\sigma$ -sub-Gaussian distributions. Then, for any  $a = (a_1, \dots, a_n)^\top \in \mathbb{R}^n$  and  $t \geq 0$ , we have*

$$\mathbb{P}\left(\sum_{i=1}^n a_i X_i > t\right) \leq \exp\left(-\frac{t^2}{2\sigma^2\|a\|_2^2}\right) \text{ and } \mathbb{P}\left(\sum_{i=1}^n a_i X_i < -t\right) \leq \exp\left(-\frac{t^2}{2\sigma^2\|a\|_2^2}\right).$$

### A.2. Proof of Theorem 3.1

For showing the lower bound, we will classify arms to be either bad or good depending on the value of their initial mean rewards, in a way that will be defined shortly. We will analyze the number of bad arm pulls which is the main contributor to the regret. For the analysis, we count the number of bad arm pulls until the first or  $m$ -th good arm is pulled over a fixed time horizon in order to preserve i.i.d. property of mean rewards in the proof. To distinguish good from bad arms, we utilize two thresholds in the proof as follows. Let  $0 < \delta < c < 1$  be such that  $1 > 1 - \delta > 1 - c > 0$  for  $\delta$ , which will be set small, and some constant  $c$ . Then,  $1 - \delta$  and  $1 - c$  represent threshold values for distinguishing good arms and bad arms, respectively. In  $\mathcal{A}$ , let  $\bar{a}_1, \bar{a}_2, \dots$ , be a sequence of arms. Given a policy  $\pi$ , without loss of generality let  $\pi$  select arms following the sequence of  $\bar{a}_1, \bar{a}_2, \dots$ .

**Case of small rotting:** When  $\varrho = O(1/T^{3/2})$ , the lower bound of order  $\sqrt{T}$  for the stationary case, from Theorem 3 in Wang et al. (2009), is tight enough for the non-stationary case. This is because we only need to pay the extra regret of at most of order  $\sqrt{T}$  for small  $\varrho$ . From Theorem 3 in Wang et al. (2009), we have

$$\mathbb{E}[R^\pi(T)] = \Omega(\sqrt{T}). \quad (12)$$

We note that even though the mean rewards are rotting in our setting, Theorem 3 in Wang et al. (2009) can be applied to our setting without any proof changes providing a tight regret bound for the near-stationary case. For the sake of completeness, we provide the proof of the theorem. Let  $K_1$  denote the number of “bad” arms  $a$  that satisfy  $\mu_1(a) \leq 1 - c$  before selecting the first “good” arm, which satisfies  $\mu_1(a) > 1 - \delta$ , in the sequence of arms  $\bar{a}_1, \bar{a}_2, \dots$ . Let  $\bar{\mu}$  be the initial mean reward of the best arm among the selected arms by  $\pi$  over time horizon  $T$ . Then for some  $\kappa > 0$ , we have

$$\begin{aligned} R^\pi(T) &= R^\pi(T)\mathbb{1}(\bar{\mu} \leq 1 - \delta) + R^\pi(T)\mathbb{1}(\bar{\mu} > 1 - \delta) \\ &\geq T\delta\mathbb{1}(\bar{\mu} \leq 1 - \delta) + K_1 c\mathbb{1}(\bar{\mu} > 1 - \delta) \\ &\geq T\delta\mathbb{1}(\bar{\mu} \leq 1 - \delta) + \kappa c\mathbb{1}(\bar{\mu} > 1 - \delta, K_1 \geq \kappa). \end{aligned} \quad (13)$$

By taking expectations on the both sides in (13) and setting  $\kappa = T\delta/c$ , we have

$$\mathbb{E}[R^\pi(T)] \geq T\delta\mathbb{P}(\bar{\mu} \leq 1 - \delta) + \kappa c(\mathbb{P}(\bar{\mu} > 1 - \delta) - \mathbb{P}(K_1 < \kappa)) = c\kappa\mathbb{P}(K_1 \geq \kappa).$$

Let  $\delta' = \delta/(1 - c + \delta)$ . Then we observe that  $K_1$  follows a geometric distribution with success probability  $\mathbb{P}(\mu_1(a) > 1 - \delta)/p(\mu_1(a) \notin (1 - c, 1 - \delta)) = \delta'$ , in which the success probability is the probability of selecting a good arm given that the arm is either a good or bad arm. Then by setting  $\delta = 1/\sqrt{T}$  with  $\kappa = \sqrt{T}/c$ , for some constant  $C > 0$  we have

$$\mathbb{E}[R^\pi(T)] \geq c\kappa(1 - \delta')^\kappa = \Omega\left(\sqrt{T}(1 - C/\sqrt{T})^{\sqrt{T}/c}\right) = \Omega(\sqrt{T}),$$

where the last equality is obtained from  $\log x \geq 1 - 1/x$  for all  $x > 0$ .

**Case of large rotting:** When  $\varrho = \omega(1/T^{3/2})$ , however, the lower bound of the stationary case is not tight enough. Here we provide the proof of the lower bound  $T\varrho^{1/3}$  for the case of  $\varrho = \omega(1/T^{3/2})$ . Let  $K_m$  denote the number of “bad” arms  $a$  that satisfy  $\mu_1(a) \leq 1 - c$  before selecting  $m$ -th “good” arm, which satisfies  $\mu_1(a) > 1 - \delta$ , in the sequence of arms  $\bar{a}_1, \bar{a}_2, \dots$ . Let  $N_T$  be the number of selected good arms  $a$  such that  $\mu_1(a) > 1 - \delta$  until  $T$ .

We can decompose  $R^\pi(T)$  into two parts as follows:

$$R^\pi(T) = R^\pi(T)\mathbb{1}(N_T < m) + R^\pi(T)\mathbb{1}(N_T \geq m). \quad (14)$$

We set  $m = \lceil (1/2)T\varrho^{2/3} \rceil$  and  $\delta = \varrho^{1/3}$  with  $\varrho = o(1)$ . For getting a lower bound for the first term in (14),  $R^\pi(T)\mathbb{1}(N_T < m)$ , we use the fact that the minimal regret is obtained from the situation where there are  $m - 1$  arms whose mean rewards are 1. Then we can think of the optimal policy that selects the best  $m - 1$  arms until their mean rewards become below the threshold  $1 - \delta$  (step 1) and then selects the best arm at each time for the remaining time steps (step 2). The required number of pulling each arm for the best  $m - 1$  arms until the mean reward becomes below  $1 - \delta$  is upper bounded by  $\delta/\varrho + 1$ . Therefore, the regret from step 2 is  $R = \Omega((T - m\delta/\varrho)\delta) = \Omega(T\varrho^{1/3})$  in which the optimal policy pulls arms which mean rewards are below  $1 - \delta$  for the remaining time after step 1. Therefore, we have

$$R^\pi(T)\mathbb{1}(N_T < m) = \Omega(R\mathbb{1}(N_T < m)) = \Omega(T\varrho^{1/3}\mathbb{1}(N_T < m)). \quad (15)$$

For getting a lower bound of the second term in (14),  $R^\pi(T)\mathbb{1}(N_T \geq m)$ , we use the minimum number of selected arms  $a$  that satisfy  $\mu_1(a) \leq 1 - c$ . When  $N_T \geq m$  and  $K_m \geq \kappa$ , the policy selects at least  $\kappa$  number of distinct arms  $a$  satisfying  $\mu_1(a) \leq 1 - c$  until  $T$ . Therefore, we have

$$R^\pi(T)\mathbb{1}(N_T \geq m) \geq c\kappa\mathbb{1}(N_T \geq m, K_m \geq \kappa). \quad (16)$$

Let  $\delta' = \delta/(1 - c + \delta)$ . By setting  $\kappa = m/\delta' - m - \sqrt{m}/\delta'$ , with  $\varrho = o(1)$ , we have

$$\kappa = \Theta(T\varrho^{1/3}). \quad (17)$$

Then from (15), (16), and (17), we have

$$\mathbb{E}[R^\pi(T)] = \Omega(T\varrho^{1/3}\mathbb{P}(N_T < m) + T\varrho^{1/3}\mathbb{P}(N_T \geq m, K_m \geq \kappa)) \geq \Omega(T\varrho^{1/3}\mathbb{P}(K_m \geq \kappa)). \quad (18)$$

Next we provide a lower bound for  $\mathbb{P}(K_m \geq \kappa)$ . Observe that  $K_m$  follows a negative binomial distribution with  $m$  successes and the success probability  $\mathbb{P}(\mu_1(a) > 1 - \delta)/\mathbb{P}(\mu_1(a) \notin (1 - c, 1 - \delta)) = \delta/(1 - c + \delta) = \delta'$ , in which the success probability is the probability of selecting a good arm given that the arm is either a good or bad arm. In the following lemma, we provide a concentration inequality for  $K_m$ .

**Lemma A.3.** *For any  $1/2 + \delta'/m < \alpha < 1$ ,*

$$\mathbb{P}(K_m \geq \alpha m(1/\delta') - m) \geq 1 - \exp(-(1/3)(1 - 1/\alpha)^2(\alpha m - \delta')). \quad (19)$$

*Proof.* Let  $X_i$  for  $i > 0$  be i.i.d. Bernoulli random variables with success probability  $\delta'$ . From Section 2 in [Brown \(2011\)](#), we have

$$\mathbb{P}\left(K_m \leq \left\lfloor \alpha m \frac{1}{\delta'} \right\rfloor - m\right) = \mathbb{P}\left(\sum_{i=1}^{\lfloor \alpha m \frac{1}{\delta'} \rfloor} X_i \geq m\right). \quad (20)$$



From (20) and Lemma A.1, for any  $1/2 + \delta'/m < \alpha < 1$  we have

$$\begin{aligned} \mathbb{P}\left(K_m \leq \alpha m \frac{1}{\delta'} - m\right) &= \mathbb{P}\left(K_m \leq \left\lfloor \alpha m \frac{1}{\delta'} \right\rfloor - m\right) \\ &= \mathbb{P}\left(\sum_{i=1}^{\lfloor \alpha m \frac{1}{\delta'} \rfloor} X_i \geq m\right) \\ &\leq \exp\left(-\frac{(1-1/\alpha)^2}{3} \left\lfloor \alpha m \frac{1}{\delta'} \right\rfloor \delta'\right) \\ &\leq \exp\left(-\frac{(1-1/\alpha)^2}{3} (\alpha m - \delta')\right), \end{aligned}$$

in which the first inequality comes from Lemma A.1, which concludes the proof.  $\square$

From Lemma A.3 with  $\alpha = 1 - 1/\sqrt{m}$  and large enough  $T$ , we have

$$\begin{aligned} \mathbb{P}(K_m \geq \kappa) &\geq 1 - \exp\left(-\frac{1}{3}(m - \sqrt{m} - \delta') \left(\frac{1}{\sqrt{m} - 1}\right)^2\right) \\ &\geq 1 - \exp\left(-\frac{1}{6}(m - \sqrt{m}) \left(\frac{1}{\sqrt{m} - 1}\right)^2\right) \\ &= 1 - \exp\left(-\frac{1}{6} \frac{\sqrt{m}}{\sqrt{m} - 1}\right) \\ &\geq 1 - \exp(-1/6). \end{aligned} \tag{21}$$

Therefore, from (18) and (21), we have

$$\mathbb{E}[R^\pi(T)] = \Omega(T \varrho^{1/3}). \tag{22}$$

Finally, from (12) and (22) we conclude that for any policy  $\pi$ , we have

$$\mathbb{E}[R^\pi(T)] = \Omega(\max\{T \varrho^{1/3}, \sqrt{T}\}).$$

### A.3. Proof of Theorem 4.1

Observe that initial mean rewards of selected arms are i.i.d. with a uniform distribution which should simplify analysis of the expected regret. However, by fixing the number of selected arms by a policy over the time horizon  $T$ , the mean rewards of arms become dependent. To deal with this dependence, we analyze the regret by controlling the number of distinct selected arms instead of fixing the time horizon. We explain this in more details in the following proofs.

We set  $\delta = \max\{\varrho^{1/3}, 1/\sqrt{T}\}$ . Then we define an arm  $a$  to be a *good* arm if  $\Delta(a) \leq \delta/2$ , and, otherwise,  $a$  is a *bad* arm. In  $\mathcal{A}$ , let  $\bar{a}_1, \bar{a}_2, \dots$ , be a sequence of arms, which have i.i.d. mean rewards with uniform distribution on  $[0, 1]$ . Given a policy selecting arms in the sequence order, let  $m^{\mathcal{G}}$  be the number of selections of distinct good arms and  $m_i^{\mathcal{B}}$  be the number of consecutive selections of distinct bad arms between the  $i-1$ -st and  $i$ -th selection of a good arm among  $m^{\mathcal{G}}$  good arms. We refer to the period starting from selecting the  $i-1$ -st good arm before selecting the  $i$ -th good arm as the  $i$ -th *episode*. Observe that  $m_1^{\mathcal{B}}, \dots, m_{m^{\mathcal{G}}}^{\mathcal{B}}$  are i.i.d. random variables with geometric distribution with parameter  $\delta/2$ , given a fixed value of  $m^{\mathcal{G}}$ . Therefore, for non-negative integer  $k$  we have  $\mathbb{P}(m_i^{\mathcal{B}} = k) = (1 - \delta/2)^k \delta/2$ , for  $i = 1, \dots, m^{\mathcal{G}}$ . Define  $\tilde{m}$  to be the number of episodes from the policy  $\pi$  over the horizon  $T$ ,  $\tilde{m}^{\mathcal{G}}$  to be the total number of selections of a good arm by the policy  $\pi$  over the horizon  $T$  such that  $\tilde{m}^{\mathcal{G}} = \tilde{m}$  or  $\tilde{m}^{\mathcal{G}} = \tilde{m} - 1$ , and  $\tilde{m}_i^{\mathcal{B}}$  to be the number of selections of a bad arm in the  $i$ -th episode by the policy  $\pi$  over the horizon  $T$ . Without loss of generality, we assume that the policy selects arms in the sequence of  $\bar{a}_1, \bar{a}_2, \dots$ . Let  $\mathcal{A}_T$  be the set of selected arms over the horizon of  $T$  time steps, which satisfies  $|\mathcal{A}_T| \leq T$ . Let

$$\hat{\mu}_t(a) = \frac{\sum_{s=1}^{t-1} r_s \mathbb{1}(a_s = a)}{n_t(a)} \text{ and } \bar{\mu}_t(a) = \frac{\sum_{s=1}^{t-1} \mu_s(a) \mathbb{1}(a_s = a)}{n_t(a)}.$$

We define the event  $E_1 = \{|\hat{\mu}_t(a) - \bar{\mu}_t(a)| \leq \sqrt{2 \log(T^4)/n_t(a)} \text{ for all } t \in [T], a \in \mathcal{A}_T\}$  to guarantee that the estimators of initial mean reward are well estimated. From Lemma A.2, by following similar steps of the proof of Lemma 5 in Auer et al. (2019), we have

$$\mathbb{P}(E_1^c) \leq 2/T^2.$$

Recall that  $R^\pi(T) = \sum_{t=1}^T 1 - \mu_t(a_t^\pi)$ . It is true that  $R^\pi(T) = o(T^2)$  because the maximum mean reward gap from rotting is bounded by  $1 + T\varrho = o(T)$ . Given that  $E_1$  does not hold, the regret is  $\mathbb{E}[R^\pi(T)|E_1^c]\mathbb{P}(E_1^c) = o(1)$ , which is negligible comparing with the regret when  $E_1$  holds true which we show later. Therefore, in the rest of the proof we assume that  $E_1$  holds true.

Under a policy  $\pi$ , let  $R_i^{\mathcal{G}}$  be the regret (summation of mean reward gaps) contributed by pulling the good arm in the  $i$ -th episode and  $R_{i,j}^{\mathcal{B}}$  be the regret contributed by pulling the  $j$ -th bad arm in the  $i$ -th episode. Then let  $R_{m^{\mathcal{G}}}^\pi = \sum_{i=1}^{m^{\mathcal{G}}} (R_i^{\mathcal{G}} + \sum_{j \in [m_i^{\mathcal{B}}]} R_{i,j}^{\mathcal{B}})$ ,<sup>2</sup> which is the regret over the period of  $m^{\mathcal{G}}$  episodes.

For obtaining a regret bound, we first focus on finding a required number of episodes,  $m^{\mathcal{G}}$ , such that  $R^\pi(T) \leq R_{m^{\mathcal{G}}}^\pi$ . Then we provide regret bounds for each bad arm and good arm in an episode. Lastly, we obtain a regret bound for  $\mathbb{E}[R^\pi(T)]$  using the episodic regret bound.

For  $i \in [\tilde{m}^{\mathcal{G}}]$ ,  $j \in [\tilde{m}_i^{\mathcal{B}}]$ , let  $\tilde{n}_i^{\mathcal{G}}$  be the number of pulls of the good arm in the  $i$ -th episode and  $\tilde{n}_{i,j}^{\mathcal{B}}$  be the number of pulls of the  $j$ -th bad arm in the  $i$ -th episode by the policy  $\pi$  over the horizon  $T$ . Let  $\tilde{a}$  be the last selected arm over time horizon  $T$  by  $\pi$ . We denote by  $\hat{m}^{\mathcal{G}}$  and  $\hat{m}_i^{\mathcal{B}}$  for  $i \in [\tilde{m}]$  the number of arms excluding  $\tilde{a}$  in the selected  $\tilde{m}^{\mathcal{G}}$  number of good arms and  $\tilde{m}_i^{\mathcal{B}}$  number of bad arms for  $i \in [\tilde{m}]$  as follows:

$$\hat{m}^{\mathcal{G}} = \begin{cases} \tilde{m}^{\mathcal{G}} - 1 & \text{if } \tilde{a} \text{ is a good arm} \\ \tilde{m}^{\mathcal{G}} & \text{otherwise} \end{cases},$$

$$\hat{m}_i^{\mathcal{B}} = \tilde{m}_i^{\mathcal{B}} \text{ for } i \in [\tilde{m} - 1], \text{ and } \hat{m}_{\tilde{m}}^{\mathcal{B}} = \begin{cases} \tilde{m}_{\tilde{m}}^{\mathcal{B}} & \text{if } \tilde{a} \text{ is a good arm} \\ \tilde{m}_{\tilde{m}}^{\mathcal{B}} - 1 & \text{otherwise.} \end{cases}$$

Those notations are defined only if they exist.<sup>3</sup> Excluding the last arm  $\tilde{a}$  which the policy  $\pi$  may stop to pull suddenly by reaching the horizon  $T$ , we provide lower bounds of the number of pulling a good arm,  $\tilde{n}_i^{\mathcal{G}}$  for  $i \in [\hat{m}^{\mathcal{G}}]$  in the following lemma if they exist.

**Lemma A.4.** *Under  $E_1$ , given  $\hat{m}^{\mathcal{G}}$ , for any  $i \in [\hat{m}^{\mathcal{G}}]$  we have*

$$\tilde{n}_i^{\mathcal{G}} \geq \delta/(2\varrho).$$

*Proof.* Let

$$\hat{\mu}_t^{\mathcal{O}}(a) = \frac{\sum_{s=1}^{t-1} (r_s + \varrho_s n_s(a)) \mathbb{1}(a_s = a)}{n_t(a)},$$

which satisfies  $\tilde{\mu}_t^{\mathcal{O}}(a) \geq \hat{\mu}_t^{\mathcal{O}}(a)$  from  $\varrho_s \leq \varrho$  for all  $s$ . Under  $E_1$ , we can easily show that for all  $t \in [T]$  and  $a \in \mathcal{A}_T$ ,

$$|\hat{\mu}_t^{\mathcal{O}}(a) - \mu_1(a)| \leq \sqrt{8 \log(T)/n_t(a)}.$$

Let  $a(i)$  be a selected good arm in the  $i$ -th episode. Suppose that  $n_t(a(i)) = \tilde{n}_i^{\mathcal{G}} = \lfloor \delta/(2\varrho) \rfloor$  for some  $t > 0$ , then we have

$$\tilde{\mu}_t^{\mathcal{O}}(a(i)) - \varrho \tilde{n}_i^{\mathcal{G}} + \sqrt{8 \log(T)/\tilde{n}_i^{\mathcal{G}}} \geq \hat{\mu}_t^{\mathcal{O}}(a(i)) - \varrho \tilde{n}_i^{\mathcal{G}} + \sqrt{8 \log(T)/\tilde{n}_i^{\mathcal{G}}} \geq \mu_1(a(i)) - \delta/2 \geq 1 - \delta,$$

where the second inequality is obtained from  $E_1$  and  $\tilde{n}_i^{\mathcal{G}} \leq \delta/(2\varrho)$ , and the third inequality is from  $\mu_1(a(i)) \geq 1 - \delta/2$ . Therefore, policy  $\pi$  must pull arm  $a$  more times than  $\lfloor \delta/(2\varrho) \rfloor$ , which implies  $\tilde{n}_i^{\mathcal{G}} \geq \delta/(2\varrho)$ .  $\square$

<sup>2</sup>Note that  $R_{m^{\mathcal{B}}}^\pi$  does not contain undefined  $R_{i,j}^{\mathcal{B}}$  such that  $R_{i,j}^{\mathcal{B}}$  when  $m_i^{\mathcal{B}} = 0$ .

<sup>3</sup> $\tilde{n}_i^{\mathcal{G}}$ ,  $\tilde{n}_{i,j}^{\mathcal{B}}$ , and  $\hat{m}_i^{\mathcal{B}}$  are not defined for  $i \in [0]$  or  $j \in [0]$ .

We first consider the case when  $\varrho = \omega(1/T^{3/2})$ . We have  $\delta = \varrho^{1/3}$ . For getting  $R_{m^G}^\pi$ , here we define the policy  $\pi$  after time  $T$  such that it pulls  $\lceil \delta/(2\varrho) \rceil$  amount for a good arm and 0 for a bad arm. We note that defining how  $\pi$  works after  $T$  is only for the proof to get a regret bound over time horizon  $T$ . For the last arm  $\tilde{a}$  over the horizon  $T$ , it pulls the arm up to  $\lceil \delta/(2\varrho) \rceil$  amounts if  $\tilde{a}$  is a good arm and  $\tilde{n}_{m^G}^G < \delta/(2\varrho)$ . For  $i \in [m^G]$ ,  $j \in [m_i^B]$  let  $n_i^G$  and  $n_{i,j}^B$  be the number of pulling the good arm in  $i$ -th episode and  $j$ -th bad arm in  $i$ -th episode from the policy, respectively. Here we define  $n_i^G$ 's and  $n_{i,j}^B$ 's as follows:

If  $\tilde{a}$  is a good arm and  $\tilde{n}_{m^G}^G < \delta/(2\varrho)$ ,

$$n_i^G = \begin{cases} \tilde{n}_i^G & \text{for } i \in [\tilde{m}^G - 1] \\ \lceil \delta/(2\varrho) \rceil & \text{for } i \in [m^G]/[\tilde{m}^G - 1] \end{cases}, n_{i,j}^B = \begin{cases} \tilde{n}_{i,j}^B & \text{for } i \in [\tilde{m}^G], j \in [\tilde{m}_i^B] \\ 0 & \text{for } i \in [m^G]/[\tilde{m}^G], j \in [m_i^B]/[\tilde{m}_i^B]. \end{cases}$$

Otherwise,

$$n_i^G = \begin{cases} \tilde{n}_i^G & \text{for } i \in [\tilde{m}^G] \\ \lceil \delta/(2\varrho) \rceil & \text{for } i \in [m^G]/[\tilde{m}^G] \end{cases}, n_{i,j}^B = \begin{cases} \tilde{n}_{i,j}^B & \text{for } i \in [\tilde{m}^G], j \in [\tilde{m}_i^B] \\ 0 & \text{for } i \in [m^G]/[\tilde{m}^G - 1], j \in [m_i^B]/[\tilde{m}_i^B]. \end{cases}$$

We note that  $n_i^G$ 's and  $n_{i,j}^B$ 's are defined only if they exist.<sup>4</sup> Then we provide  $m^G$  such that  $R^\pi(T) \leq R_{m^G}^\pi$  in the following lemma.

**Lemma A.5.** *Under  $E_1$ , when  $m^G = \lceil 2T\varrho^{2/3} \rceil$  we have*

$$R^\pi(T) \leq R_{m^G}^\pi.$$

*Proof.* From Lemma A.4, with  $\delta = \varrho^{1/3}$  we have

$$\sum_{i \in [m^G]} \left( n_i^G + \sum_{j \in [m_i^B]} n_{i,j}^B \right) \geq m^G \frac{\delta}{2\varrho} \geq T,$$

which implies that  $R^\pi(T) \leq R_{m^G}^\pi$ . □

From the result of Lemma A.5, we set  $m^G = \lceil 2T\varrho^{2/3} \rceil$ . For getting a bound for  $\mathbb{E}[R_{m^G}^\pi]$ , we provide bounds for  $\mathbb{E}[R_i^G]$  and  $\mathbb{E}[R_{i,j}^B]$  in the following lemma.

**Lemma A.6.** *Under  $E_1$  and policy  $\pi$ , for any  $i \in [m^G]$ ,  $j \in [m_i^B]$ , we have*

$$\mathbb{E}[R_i^G] = \tilde{O} \left( \frac{\delta}{\varrho^{2/3}} + \frac{\delta^2}{\varrho} + \frac{1}{\varrho^{1/3}} \right),$$

and

$$\mathbb{E}[R_{i,j}^B] = \tilde{O} \left( 1 + \frac{\delta}{\varrho^{1/3}} + \frac{\delta^2}{\varrho^{2/3}} \right).$$

*Proof.* First we provide a bound for  $R_i^G$  using an upper bound of  $n_i^G$ . Recall that  $a(i)$  is the selected arm for the  $i$ -th good arm. We have

$$\begin{aligned} \tilde{\mu}_t^G(a(i)) - \varrho n_t(a(i)) &= \hat{\mu}_t(a(i)) + \varrho \frac{\sum_{s=1}^{t-1} n_s(a(i)) \mathbb{1}(a_s = a(i))}{n_t(a(i))} - \varrho n_t(a(i)) \\ &\leq \hat{\mu}_t(a(i)) + \varrho \frac{n_t(a(i)) + 1}{2} - \varrho n_t(a(i)) \\ &= \hat{\mu}_t(a(i)) - \varrho \frac{n_t(a(i)) - 1}{2}. \end{aligned}$$

<sup>4</sup>  $n_i^G$  and  $n_{i,j}^B$  are not defined for  $i \in [0]$  or  $j \in [0]$ .

Then since, under  $E_1$ ,

$$\begin{aligned}\hat{\mu}_t(a(i)) - (\varrho/2)(n_t(a(i)) - 1) &\leq \bar{\mu}_t(a(i)) + \sqrt{8 \log(T)/n_t(a(i))} - (\varrho/2)(n_t(a(i)) - 1) \\ &\leq 1 + \sqrt{8 \log(T)/n_t(a(i))} - (\varrho/2)(n_t(a(i)) - 1),\end{aligned}$$

for  $i \in [\tilde{m}^{\mathcal{G}}]$ , from the policy  $\pi$ , we need to get  $n$  such that

$$1 - \frac{\varrho}{2}(n - 1) + 2\sqrt{8 \log(T)/n} < 1 - \delta, \quad (23)$$

in which  $n + 1$  is an upper bound for  $n_i^{\mathcal{G}}$ . Let  $n_1 = 2(\delta + \varrho^{1/3})/\varrho + 1$  and  $n_2 = C \log(T)/\varrho^{2/3}$  with some large enough constant  $C > 0$ . Then  $n = n_1 + n_2$  satisfies (23) because  $1 - \varrho n_1 + 2\sqrt{8 \log(T)/n_2} < 1 - \delta$ . Therefore, for all  $i \in [\tilde{m}^{\mathcal{G}}]$  we have  $n_i^{\mathcal{G}} = \tilde{O}((\delta + \varrho^{1/3})/\varrho)$ . Then with the fact that  $n_i^{\mathcal{G}} = \lceil \delta/(2\varrho) \rceil$  for  $i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}}]$  if they exist, for any  $i \in [m^{\mathcal{G}}]$  we have

$$n_i^{\mathcal{G}} = \tilde{O}((\delta + \varrho^{1/3})/\varrho).$$

Then for any  $i \in [m^{\mathcal{G}}]$  we have

$$\begin{aligned}\mathbb{E}[R_i^{\mathcal{G}}] &\leq \mathbb{E} \left[ \Delta(a(i))n_i^{\mathcal{G}} + \frac{n_i^{\mathcal{G}}(n_i^{\mathcal{G}} - 1)}{2}\varrho \right] \\ &= \tilde{O} \left( \frac{2}{\delta} \int_0^{\delta/2} \frac{\delta + \varrho^{1/3}}{\varrho} x + \left( \frac{\delta + \varrho^{1/3}}{\varrho} \right)^2 \varrho dx \right) \\ &= \tilde{O} \left( \frac{\delta}{\varrho^{2/3}} + \frac{\delta^2}{\varrho} + \frac{1}{\varrho^{1/3}} \right),\end{aligned}$$

where the first equality is obtained from the fact that  $\Delta(a(i))$  are i.i.d. random variables with uniform distribution on  $[0, \delta/2]$  and  $n_i^{\mathcal{G}} = \tilde{O}((\delta + \varrho^{1/3})/\varrho)$ .

Now we provide an upper bound of  $n_{i,j}^{\mathcal{B}}$  to get a bound of  $R_{i,j}^{\mathcal{B}}$  for  $i \in [m^{\mathcal{G}}], j \in [m_i^{\mathcal{B}}]$ . Let  $a(i, j)$  be a selected arm for  $j$ -th bad arm in the  $i$ -th episode. When  $\delta/2 < \Delta(a(i, j)) \leq \delta + \varrho^{1/3}$ , as in the case of the good arm,  $n = 2(\delta + \varrho^{1/3})/\varrho + 1 + C_1 \log(T)/\varrho^{2/3}$  for some large enough constant  $C_1 > 0$ , satisfies (23) so that  $n_{i,j}^{\mathcal{B}} = \tilde{O}((\delta + \varrho^{1/3})/\varrho)$  for all  $i \in [\tilde{m}], j \in [\tilde{m}_i^{\mathcal{B}}]$ . Since under  $E_1$

$$\begin{aligned}\hat{\mu}_t(a(i, j)) - (\varrho/2)(n_t(a(i, j)) - 1) &\leq \bar{\mu}_t(a(i, j)) + \sqrt{8 \log(T)/n_t(a(i, j))} - (\varrho/2)(n_t(a(i, j)) - 1) \\ &\leq \mu_1(a(i, j)) + \sqrt{8 \log(T)/n_t(a(i, j))} - (\varrho/2)(n_t(a(i, j)) - 1),\end{aligned}$$

when  $\delta + \varrho^{1/3} < \Delta(a(i, j)) \leq 1$ , from the policy  $\pi$  under  $E_1$ , we need to get  $n \geq 1$  such that

$$\mu_1(a(i, j)) - \frac{\varrho}{2}(n - 1) + 2\sqrt{8 \log(T)/n} < 1 - \delta, \quad (24)$$

in which  $n + 1$  is an upper bound of  $n_{i,j}^{\mathcal{B}}$  for  $i \in [\tilde{m}], j \in [\tilde{m}_i^{\mathcal{B}}]$ . From a sufficient condition for (24) to hold such that

$$\mu_1(a(i, j)) + 2\sqrt{8 \log(T)/n} < 1 - \delta,$$

we can find that  $n = C_2 \log(T)/(\Delta(a(i, j)) - \delta)^2$  for some large constant  $C_2 > 0$  satisfies (24). Therefore, when  $\delta + \varrho^{1/3} < \Delta(a(i, j)) \leq 1$ , for all  $i \in [\tilde{m}], j \in [\tilde{m}_i^{\mathcal{B}}]$  we have  $n_{i,j}^{\mathcal{B}} = \tilde{O}(1/(\Delta(a(i, j)) - \delta)^2)$ . Then with the fact that  $n_{i,j}^{\mathcal{B}} = 0$  for  $i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}}], j \in [m_i^{\mathcal{B}}]/[\tilde{m}_i^{\mathcal{B}}]$  if they exist, for any  $i \in [m^{\mathcal{G}}]$  and  $j \in [m_i^{\mathcal{B}}]$ , we have

$$n_{i,j}^{\mathcal{B}} = \begin{cases} \tilde{O}((\delta + \varrho^{1/3})/\varrho) & \text{if } \delta/2 < \Delta(a(i, j)) \leq \delta + \varrho^{1/3} \\ \tilde{O}(1/(\Delta(a(i, j)) - \delta)^2) & \text{if } \delta + \varrho^{1/3} < \Delta(a(i, j)) \leq 1 \end{cases}$$



Then for any  $i \in [m^{\mathcal{G}}]$ ,  $j \in [m_i^{\mathcal{B}}]$ , we have

$$\begin{aligned} \mathbb{E}[R_{i,j}^{\mathcal{B}}] &\leq \mathbb{E} \left[ \Delta(a(i,j))n_{i,j}^{\mathcal{B}} + \frac{n_{i,j}^{\mathcal{B}}(n_{i,j}^{\mathcal{B}} - 1)}{2} \varrho \right] \\ &= \tilde{O} \left( \frac{1}{1 - \delta/2} \left( \int_{\delta/2}^{\delta + \varrho^{1/3}} \frac{\delta + \varrho^{1/3}}{\varrho} x + \left( \frac{\delta + \varrho^{1/3}}{\varrho} \right)^2 \varrho dx + \int_{\delta + \varrho^{1/3}}^1 \frac{1}{(x - \delta)^2} x + \frac{1}{(x - \delta)^4} \varrho dx \right) \right) \\ &= \tilde{O} \left( 1 + \frac{\delta}{\varrho^{1/3}} + \frac{\delta^2}{\varrho^{2/3}} \right). \end{aligned}$$

□

Recall that  $R_{m^{\mathcal{G}}}^{\pi} = \sum_{i=1}^{m^{\mathcal{G}}} (R_i^{\mathcal{G}} + \sum_{j \in [m_i^{\mathcal{B}}]} R_{i,j}^{\mathcal{B}})$ . With  $\delta = \varrho^{1/3}$  and  $m^{\mathcal{G}} = \lceil 2T\varrho^{2/3} \log(1/\delta') \rceil$ , from Lemmas A.5 and A.6, and the fact that  $m_i^{\mathcal{B}}$ 's are i.i.d. random variables with geometric distribution with  $\mathbb{E}[m_i^{\mathcal{B}}] = 2/\delta - 1$ , we have

$$\begin{aligned} \mathbb{E}[R^{\pi}(T)] &= O(\mathbb{E}[R_{m^{\mathcal{G}}}^{\pi}]) \\ &= O \left( \mathbb{E} \left[ \sum_{i=1}^{m^{\mathcal{G}}} \left( R_i^{\mathcal{G}} + \sum_{j \in [m_i^{\mathcal{B}}]} R_{i,j}^{\mathcal{B}} \right) \right] \right) \\ &= \tilde{O} \left( T\varrho^{2/3} \left( \left( \frac{\delta}{\varrho^{2/3}} + \frac{\delta^2}{\varrho} + \frac{1}{\varrho^{1/3}} \right) + \frac{1}{\delta} \left( 1 + \frac{\delta}{\varrho^{1/3}} \right) \right) \right) \\ &= \tilde{O} \left( T\varrho^{2/3} \left( \frac{\delta}{\varrho^{2/3}} + \frac{\delta^2}{\varrho} + \frac{1}{\delta} + \frac{1}{\varrho^{1/3}} \right) \right) \\ &= \tilde{O} \left( T\varrho^{1/3} \right). \end{aligned} \tag{25}$$

Now we consider the case when  $\varrho = O(1/T^{3/2})$ . We have  $\delta = \Theta(1/\sqrt{T})$ . With a slight abuse of notation, we use  $\pi$  for a modified strategy after  $T$ . For getting  $R_{m^{\mathcal{G}}}^{\pi}$ , here we define the policy  $\pi$  after time  $T$  such that it pulls  $T$  amounts for a good arm and once for a bad arm. For the last arm  $\tilde{a}$  over the horizon  $T$ , it pulls the arm up to  $T$  amounts if  $\tilde{a}$  is a good arm and  $\tilde{n}_{\tilde{m}^{\mathcal{G}}}^{\mathcal{G}} < T$ . With slight abuse of notation, for  $i \in [m^{\mathcal{G}}]$ ,  $j \in [m_i^{\mathcal{B}}]$  let  $n_i^{\mathcal{G}}$  and  $n_{i,j}^{\mathcal{B}}$  be the number of pulling the good arm in  $i$ -th episode and  $j$ -th bad arm in  $i$ -th episode from the policy, respectively. Here we define  $n_i^{\mathcal{G}}$ 's and  $n_{i,j}^{\mathcal{B}}$ 's as follows:

If  $\tilde{a}$  is a good arm and  $\tilde{n}_{\tilde{m}^{\mathcal{G}}}^{\mathcal{G}} < T$ ,

$$n_i^{\mathcal{G}} = \begin{cases} \tilde{n}_i^{\mathcal{G}} & \text{for } i \in [\tilde{m}^{\mathcal{G}} - 1] \\ T & \text{for } i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}} - 1] \end{cases}, n_{i,j}^{\mathcal{B}} = \begin{cases} \tilde{n}_{i,j}^{\mathcal{B}} & \text{for } i \in [\tilde{m}^{\mathcal{G}}], j \in [\tilde{m}_i^{\mathcal{B}}] \\ 0 & \text{for } i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}}], j \in [m_i^{\mathcal{B}}]/[\tilde{m}_i^{\mathcal{B}}]. \end{cases}$$

Otherwise,

$$n_i^{\mathcal{G}} = \begin{cases} \tilde{n}_i^{\mathcal{G}} & \text{for } i \in [\tilde{m}^{\mathcal{G}}] \\ T & \text{for } i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}}] \end{cases}, n_{i,j}^{\mathcal{B}} = \begin{cases} \tilde{n}_{i,j}^{\mathcal{B}} & \text{for } i \in [\tilde{m}^{\mathcal{G}}], j \in [\tilde{m}_i^{\mathcal{B}}] \\ 0 & \text{for } i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}} - 1], j \in [m_i^{\mathcal{B}}]/[\tilde{m}_i^{\mathcal{B}}]. \end{cases}$$

From Lemma A.4, under  $E_1$  we can find that  $n_i^{\mathcal{G}} \geq \min\{\delta/(2\varrho), T\}$  for  $i \in [m^{\mathcal{G}}]$ . Then if  $m^{\mathcal{G}} = C_3$  with some large enough constant  $C_3 > 0$ , then with  $\delta = \Theta(1/\sqrt{T})$  and  $\varrho = O(1/T^{3/2})$ , we have

$$\sum_{i \in [m^{\mathcal{G}}]} n_i^{\mathcal{G}} \geq C_3 \min\{\delta/(2\varrho), T\} > T,$$

which implies  $R^{\pi}(T) \leq R_{m^{\mathcal{G}}}^{\pi}$ . Therefore, we set  $m^{\mathcal{G}} = C_3$ . For getting a bound for  $\mathbb{E}[R_{m^{\mathcal{G}}}^{\pi}]$ , we provide bounds for  $\mathbb{E}[R_i^{\mathcal{G}}]$  and  $\mathbb{E}[R_{i,j}^{\mathcal{B}}]$  in the following lemma.

**Lemma A.7.** Under  $E_1$  and policy  $\pi$ , for any  $i \in [m^{\mathcal{G}}]$  and  $j \in [m_i^{\mathcal{B}}]$ , we have

$$\mathbb{E}[R_i^{\mathcal{G}}] = O(\delta T + T^2 \varrho),$$

and

$$\mathbb{E}[R_{i,j}^{\mathcal{B}}] = \tilde{O}\left(T\delta^2 + \sqrt{T}\delta + 1\right).$$

*Proof.* First we provide a bound for  $R_i^{\mathcal{G}}$  using an upper bound of  $n_i^{\mathcal{G}}$ . With the definition of  $n_i^{\mathcal{G}} = T$  for  $i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}}]$ , for any  $i \in [m^{\mathcal{G}}]$  we have

$$n_i^{\mathcal{G}} \leq T.$$

Recall that  $a(i)$  is the selected arm for the  $i$ -th good arm. Then for any  $i \in [m^{\mathcal{G}}]$  we have

$$\begin{aligned} \mathbb{E}[R_i^{\mathcal{G}}] &\leq \mathbb{E}\left[\Delta(a(i))n_i^{\mathcal{G}} + \frac{n_i^{\mathcal{G}}(n_i^{\mathcal{G}} - 1)}{2}\varrho\right] \\ &= O(\delta T + T^2 \varrho), \end{aligned}$$

where the first equality is obtained from the fact that  $\Delta(a(i))$ 's are i.i.d. random variables with uniform distribution on  $[0, \delta/2]$  and  $n_i^{\mathcal{G}} \leq T$ .

Now we provide an upper bound of  $n_{i,j}^{\mathcal{B}}$  to get a bound of  $R_{i,j}^{\mathcal{B}}$  for  $i \in [m^{\mathcal{G}}], j \in [m_i^{\mathcal{B}}]$ . Let  $a(i, j)$  be a selected arm for  $j$ -th bad arm in the  $i$ -th episode. When  $\delta/2 < \Delta(a(i, j)) \leq \delta + 1/\sqrt{T}$ , as in the case of the good arm,  $n_{i,j}^{\mathcal{B}} \leq T$  for all  $i \in [\tilde{m}], j \in [\tilde{m}_i^{\mathcal{B}}]$ . When  $\delta + 1/\sqrt{T} < \Delta(a(i, j)) \leq 1$ , since under  $E_1$

$$\begin{aligned} \tilde{\mu}_i^{\varrho}(a(i, j)) - \varrho n_t(a(i, j)) &= \hat{\mu}_t(a(i, j)) + \varrho \frac{\sum_{s=1}^{t-1} n_s(a(i, j)) \mathbb{1}(a_s = a(i, j))}{n_t(a(i, j))} - \varrho n_t(a(i, j)) \\ &\leq \hat{\mu}_t(a(i, j)) + \varrho \frac{n_t(a(i, j)) + 1}{2} - \varrho n_t(a(i, j)) \\ &= \hat{\mu}_t(a(i, j)) - \varrho \frac{n_t(a(i, j)) - 1}{2} \\ &\leq \bar{\mu}_t(a(i, j)) + \sqrt{8 \log(T)/n_t(a(i, j))} - (\varrho/2)(n_t(a(i, j)) - 1) \\ &\leq \mu_1(a(i, j)) + \sqrt{8 \log(T)/n_t(a(i, j))} - (\varrho/2)(n_t(a(i, j)) - 1), \end{aligned}$$

from the policy  $\pi$  we need to get  $n$  such that

$$\mu_1(a(i, j)) - (\varrho/2)(n - 1) + 2\sqrt{8 \log(T)/n} < 1 - \delta, \quad (26)$$

in which  $n + 1$  is an upper bound of  $n_{i,j}^{\mathcal{B}}$  for  $i \in [\tilde{m}], j \in [\tilde{m}_i^{\mathcal{B}}]$ . From a sufficient condition for (26) to hold such that

$$\mu_1(a(i, j)) + 2\sqrt{8 \log(T)/n} < 1 - \delta,$$

we can find that  $n = C_4 \log(T)/(\Delta(a(i, j)) - \delta)^2$  for some large constant  $C_4 > 0$  satisfies (26). Therefore, when  $\delta + 1/\sqrt{T} < \Delta(a(i, j)) \leq 1$ , for all  $i \in [\tilde{m}], j \in [\tilde{m}_i^{\mathcal{B}}]$  we have  $n_{i,j}^{\mathcal{B}} = \tilde{O}(1/(\Delta(a(i, j)) - \delta)^2)$ . Then with the fact that  $n_{i,j}^{\mathcal{B}} = 0$  for  $i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}}], j \in [m_i^{\mathcal{B}}]/[\tilde{m}_i^{\mathcal{B}}]$ , for any  $i \in [m^{\mathcal{G}}], j \in [m_i^{\mathcal{B}}]$ , if  $\delta/2 < \Delta(a(i, j)) \leq \delta + 1/\sqrt{T}$ , we have

$$n_{i,j}^{\mathcal{B}} \leq T,$$

and if  $\delta + 1/\sqrt{T} < \Delta(a(i, j)) \leq 1$ , we have

$$n_{i,j}^{\mathcal{B}} = \tilde{O}(1/(\Delta(a(i, j)) - \delta)^2).$$

Then for any  $i \in [m^{\mathcal{G}}]$ ,  $j \in [m_i^{\mathcal{B}}]$ , we have

$$\begin{aligned} \mathbb{E}[R_{i,j}^{\mathcal{B}}] &\leq \mathbb{E} \left[ \Delta(a(i,j)) n_{i,j}^{\mathcal{B}} + \frac{n_{i,j}^{\mathcal{B}}(n_{i,j}^{\mathcal{B}} - 1)}{2} \varrho \right] \\ &= \tilde{O} \left( \frac{1}{1 - \delta/2} \left( \int_{\delta/2}^{\delta+1/\sqrt{T}} Tx + T^2 \varrho dx + \int_{\delta+1/\sqrt{T}}^1 \frac{1}{(x - \delta)^2} x + \frac{1}{(x - \delta)^4} \varrho dx \right) \right) \\ &= \tilde{O} \left( T\delta^2 + \sqrt{T}\delta + 1 \right). \end{aligned}$$

□

Then with  $\delta = \Theta(1/\sqrt{T})$  and  $m^{\mathcal{G}} = C_3$ , we have

$$\begin{aligned} \mathbb{E}[R^{\pi}(T)] &= O(\mathbb{E}[R_{m^{\mathcal{G}}}^{\pi}]) \\ &= O \left( \mathbb{E} \left[ \sum_{i \in [m^{\mathcal{G}}]} (R_i^{\mathcal{G}} + \sum_{j \in [m_i^{\mathcal{B}}]} R_{i,j}^{\mathcal{B}}) \right] \right) \\ &= \tilde{O} \left( (\delta T + T^2 \varrho) + \frac{1}{\delta} (T\delta^2 + \sqrt{T}\delta + 1) \right) \\ &= \tilde{O}(\sqrt{T}), \end{aligned} \tag{27}$$

where the third equality is obtained from Lemma A.7 and  $\mathbb{E}[m_i^{\mathcal{B}}] = 2/\delta - 1$ .

**Finally, we can conclude the proof:** From (25) and (27), for  $\varrho = o(1)$ , with  $\delta = \max\{\varrho^{1/3}, 1/\sqrt{T}\}$  we have

$$\mathbb{E}[R^{\pi}(T)] = \tilde{O}(\max\{T\varrho^{1/3}, \sqrt{T}\}).$$

#### A.4. Proof of Theorem 4.2

Let  $\pi_i(\beta)$  for  $\beta \in \mathcal{B}$  denote the base policy for time steps between  $(i-1)H + 1$  and  $i \cdot H \wedge T$  in Algorithm 2 using  $\text{UCB}_{i,t}(a, \beta)$  as a UCB index and  $1 - \beta^{1/3}$  as a threshold. Denote by  $a_t^{\pi_i(\beta)}$  the pulled arm at time step  $t$  by policy  $\pi_i(\beta)$ . Then, for  $\beta^\dagger \in \mathcal{B}$ , which is set later for a near-optimal policy, we have

$$\mathbb{E}[R^{\pi}(T)] = \mathbb{E} \left[ \sum_{t=1}^T 1 - \sum_{i=1}^{\lceil T/H \rceil} \sum_{t=(i-1)H+1}^{i \cdot H \wedge T} \mu_t(a_t^{\pi}) \right] = \mathbb{E}[R_1^{\pi}(T)] + \mathbb{E}[R_2^{\pi}(T)]. \tag{28}$$

where

$$R_1^{\pi}(T) = \sum_{t=1}^T 1 - \sum_{i=1}^{\lceil T/H \rceil} \sum_{t=(i-1)H+1}^{i \cdot H \wedge T} \mu_t(a_t^{\pi_i(\beta^\dagger)})$$

and

$$R_2^{\pi}(T) = \sum_{i=1}^{\lceil T/H \rceil} \sum_{t=(i-1)H+1}^{i \cdot H \wedge T} \mu_t(a_t^{\pi_i(\beta^\dagger)}) - \sum_{i=1}^{\lceil T/H \rceil} \sum_{t=(i-1)H+1}^{i \cdot H \wedge T} \mu_t(a_t^{\pi}).$$

Note that  $R_1^{\pi}(T)$  accounts for the regret caused by the near-optimal base algorithm  $\pi_i(\beta^\dagger)$ 's against the optimal mean reward and  $R_2^{\pi}(T)$  accounts for the regret caused by the master algorithm by selecting a base with  $\beta \in \mathcal{B}$  at every block against the base with  $\beta^\dagger$ . In what follows, we provide upper bounds for each regret component. We first provide an upper bound for  $\mathbb{E}[R_1^{\pi}(T)]$  by following the proof steps in Theorem 4.1. Then we provide an upper bound for  $\mathbb{E}[R_2^{\pi}(T)]$ . We set  $H = \sqrt{T}$  and  $\beta^\dagger$  to be a smallest value in  $\mathcal{B}$  which is larger than  $\max\{\varrho, 1/H^{3/2}\}$ .

**Upper bounding  $\mathbb{E}[R_1^{\pi}(T)]$ .** We refer to the period starting from time step  $(i-1)H + 1$  to time step  $i \cdot H \wedge T$  as the  $i$ -th block. For any  $i \in \lceil T/H \rceil$ , policy  $\pi_i(\beta^\dagger)$  runs over  $H$  time steps independent to other blocks so that each block has the

same expected regret and the last block has a smaller or equal expected regret than other blocks. Therefore, we focus on finding a bound on the regret from the first block equal to  $\sum_{t=1}^H 1 - \mu_t(a_t^{\pi_1(\beta^\dagger)})$ . Denote by  $\mathcal{A}(i)$  the set of selected arms in the  $i$ -th block, which satisfies  $|\mathcal{A}(i)| \leq H$ . For notation simplicity, we use  $n_t(a)$  instead of  $n_{1,t}(a)$  and  $\tilde{\mu}_t^o(a)$  instead of  $\tilde{\mu}_{1,t}^o(a)$ . Let

$$\hat{\mu}_t(a) = \frac{\sum_{s=1}^{t-1} r_s \mathbb{1}(a_s = a)}{n_t(a)} \text{ and } \bar{\mu}_t(a) = \frac{\sum_{s=1}^{t-1} \mu_s(a) \mathbb{1}(a_s = a)}{n_t(a)}.$$

Define the event  $E_1 = \{|\hat{\mu}_t(a) - \bar{\mu}_t(a)| \leq \sqrt{8 \log(H)/n_t(a)}, \text{ for all } t \in [H], a \in \mathcal{A}(1)\}$ . From Lemma A.2, by following similar steps of the proof of Lemma 5 in Auer et al. (2019), we have

$$\mathbb{P}(E_1) \geq 1 - 2/H^2.$$

We assume that  $E_1$  holds true in what follows. Otherwise, the regret for the first block is negligible from  $R^{\pi_1(\beta^\dagger)}(H) = o(H^2)$ . The proof follows similar steps as in the proof of Theorem 4.1.

From  $\pi_1(\beta^\dagger)$  we have  $\delta = \beta^{\dagger 1/3}$ . We define an arm  $a$  to be a *good* arm if  $\Delta(a) \leq \delta/2$ , and, otherwise,  $a$  is a *bad* arm. In  $\mathcal{A}$ , let  $\bar{a}_1, \bar{a}_2, \dots$  be a sequence of arms, which have i.i.d. mean rewards with uniform distribution on  $[0, 1]$ . Given a policy selecting arms in the sequence order, let  $m^G$  be the number of selections of distinct good arms and  $m_i^B$  be the number of consecutive selections of distinct bad arms between the  $i-1$ -st and  $i$ -th selection of a good arm among  $m^G$  good arms. We refer to the period starting from selecting the  $i-1$ -st good arm before selecting the  $i$ -th good arm as the  $i$ -th *episode*. Observe that  $m_1^B, \dots, m_{m^G}^B$ 's are i.i.d. random variables with geometric distribution with parameter  $\delta/2$ , conditional on the value of  $m^G$ . Therefore,  $\mathbb{P}(m_i^B = k) = (1 - \delta/2)^k \delta/2$ , for  $i = 1, \dots, m^G$ . Define  $\tilde{m}$  to be the number of episodes by following policy  $\pi$  over the horizon of  $T$  time steps,  $\tilde{m}^G$  to be the total number of selections of a good arm such that  $\tilde{m}^G = \tilde{m}$  or  $\tilde{m}^G = \tilde{m} - 1$ , and  $\tilde{m}_i^B$  to be the number of selections of a bad arm in the  $i$ -th episode by the policy  $\pi_1(\beta^\dagger)$  over the horizon  $H$ . Without loss of generality, we assume that the policy selects arms in the order of the sequence  $\bar{a}_1, \bar{a}_2, \dots$ .

Under policy  $\pi_1(\beta^\dagger)$ , let  $R_i^G$  be the regret (summation of mean reward gaps) contributed by pulling the good arm in the  $i$ -th episode and  $R_{i,j}^B$  be the regret contributed by pulling the  $j$ -th bad arm in the  $i$ -th episode. Then let  $R_{m^G}^{\pi_1(\beta^\dagger)} = \sum_{i=1}^{m^G} (R_i^G + \sum_{j \in [m_i^B]} R_{i,j}^B)$ <sup>5</sup>, which is the regret over the period of  $m^G$  episodes.

For  $i \in [\tilde{m}^G]$ ,  $j \in [\tilde{m}_i^B]$ , let  $\tilde{n}_i^G$  be the number of pulls of the good arm in the  $i$ -th episode and  $\tilde{n}_{i,j}^B$  be the number of pulls of the  $j$ -th bad arm in the  $i$ -th episode by the policy  $\pi_1(\beta^\dagger)$  over the horizon  $H$ . Let  $\tilde{a}$  be the last selected arm over time horizon  $H$  by  $\pi_1(\beta^\dagger)$ . We denote by  $\hat{m}^G$  and  $\hat{m}_i^B$  for  $i \in [\tilde{m}]$  the number of arms excluding  $\tilde{a}$  in the selected  $\tilde{m}^G$  number of good arms and  $\tilde{m}_i^B$  number of bad arms for  $i \in [\tilde{m}]$  as follows:

$$\hat{m}^G = \begin{cases} \tilde{m}^G - 1 & \text{if } \tilde{a} \text{ is a good arm} \\ \tilde{m}^G & \text{otherwise} \end{cases},$$

$$\hat{m}_i^B = \tilde{m}_i^B \text{ for } i \in [\tilde{m} - 1], \text{ and } \hat{m}_{\tilde{m}}^B = \begin{cases} \tilde{m}_{\tilde{m}}^B & \text{if } \tilde{a} \text{ is a good arm} \\ \tilde{m}_{\tilde{m}}^B - 1 & \text{otherwise.} \end{cases}$$

These notations are defined only if they exist.<sup>6</sup> Excluding the last arm  $\tilde{a}$  which the policy  $\pi_1(\beta^\dagger)$  may stop to pull suddenly by reaching the horizon  $H$ , we provide lower bounds for the number of pulls for each arm,  $\tilde{n}_i^G$  for  $i \in [\hat{m}^G]$  in the following lemma if they exist.

**Lemma A.8.** *Under  $E_1$ , given  $\hat{m}^G$ , for any  $i \in [\hat{m}^G]$  we have*

$$\tilde{n}_i^G \geq \delta/(2\beta^\dagger).$$

<sup>5</sup>Note that  $R_{m^G}^{\pi_1(\beta^\dagger)}$  does not contain undefined  $R_{i,j}^B$ , such that  $R_{i,j}^B$  when  $m_i^B = 0$ .

<sup>6</sup> $\tilde{n}_i^G$ ,  $\tilde{n}_{i,j}^B$ , and  $\tilde{m}_i^B$  are not defined for  $i \in [0]$  or  $j \in [0]$ .



*Proof.* Let

$$\hat{\mu}_t^\rho(a) = \frac{\sum_{s=1}^{t-1} (r_s + \rho_s n_s(a)) \mathbb{1}(a_s = a)}{n_t(a)},$$

which satisfies  $\tilde{\mu}_t^\rho(a, \beta^\dagger) \geq \hat{\mu}_t^\rho(a)$  from  $\beta^\dagger \geq \rho$ . Under  $E_1$ , it is true that for all  $t \in [H]$  and  $a \in \mathcal{A}(1)$ ,

$$|\hat{\mu}_t^\rho(a) - \mu_1(a)| \leq \sqrt{8 \log(H)/n_t(a)}.$$

Let  $a(i)$  be a selected good arm in the  $i$ -th episode. Suppose that  $n_t(a(i)) = \tilde{n}_i^{\mathcal{G}} = \lfloor \delta/(2\beta^\dagger) \rfloor$  for some  $t > 0$ , then we have

$$\tilde{\mu}_t^\rho(a(i), \beta^\dagger) - \beta^\dagger \tilde{n}_i^{\mathcal{G}} + \sqrt{8 \log(H)/\tilde{n}_i^{\mathcal{G}}} \geq \hat{\mu}_t^\rho(a(i)) - \beta^\dagger \tilde{n}_i^{\mathcal{G}} + \sqrt{8 \log(H)/\tilde{n}_i^{\mathcal{G}}} \geq \mu_1(a(i)) - \delta/2 \geq 1 - \delta,$$

where the second inequality is obtained from  $E_1$  and  $\tilde{n}_i^{\mathcal{G}} \leq \delta/(2\beta^\dagger)$ , and the third inequality is from  $\mu_1(a(i)) \geq 1 - \delta/2$ . Therefore, policy  $\pi_1(\beta^\dagger)$  must pull arm  $a$  more times than  $\lfloor \delta/(2\beta^\dagger) \rfloor$ , which implies  $\tilde{n}_i^{\mathcal{G}} \geq \delta/(2\beta^\dagger)$ .  $\square$

**We first consider the case when  $\rho = \omega(1/H^{3/2})$ .** Then we have that  $\beta^\dagger$  is the smallest value in  $\mathcal{B}$  which exceeds  $\rho$  such that  $\rho \leq \beta^\dagger \leq 2\rho$ . For getting  $R_{m^{\mathcal{G}}}^{\pi_1(\beta^\dagger)}$ , here we define how the policy  $\pi_1(\beta^\dagger)$  works after time  $H$  such that it pulls  $\lceil \delta/(2\beta^\dagger) \rceil$  times a good arm and 0 time a bad arm. We note that defining how  $\pi_1(\beta^\dagger)$  works after  $H$  is only for the proof to get a regret bound over time horizon  $H$ . For the last arm  $\tilde{a}$  over the horizon  $H$ , it pulls the arm up to  $\lceil \delta/(2\beta^\dagger) \rceil$  times if  $\tilde{a}$  is a good arm and  $\tilde{n}_{\tilde{m}^{\mathcal{G}}}^{\mathcal{G}} < \delta/(2\beta^\dagger)$ . For  $i \in [m^{\mathcal{G}}]$  and  $j \in [m_i^{\mathcal{B}}]$ , let  $n_i^{\mathcal{G}}$  and  $n_{i,j}^{\mathcal{B}}$  be the number of pulls of the good arm in the  $i$ -th episode and the  $j$ -th bad arm in the  $i$ -th episode by the policy, respectively. Here we define  $n_i^{\mathcal{G}}$ 's and  $n_{i,j}^{\mathcal{B}}$ 's as follows:

If  $\tilde{a}$  is a good arm and  $\tilde{n}_{\tilde{m}^{\mathcal{G}}}^{\mathcal{G}} < \delta/(2\beta^\dagger)$ , then

$$n_i^{\mathcal{G}} = \begin{cases} \tilde{n}_i^{\mathcal{G}} & \text{for } i \in [\tilde{m}^{\mathcal{G}} - 1] \\ \lceil \delta/(2\beta^\dagger) \rceil & \text{for } i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}} - 1] \end{cases}, n_{i,j}^{\mathcal{B}} = \begin{cases} \tilde{n}_{i,j}^{\mathcal{B}} & \text{for } i \in [\tilde{m}^{\mathcal{G}}], j \in [\tilde{m}_i^{\mathcal{B}}] \\ 0 & \text{for } i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}}], j \in [m_i^{\mathcal{B}}]/[\tilde{m}_i^{\mathcal{B}}]. \end{cases}$$

Otherwise,

$$n_i^{\mathcal{G}} = \begin{cases} \tilde{n}_i^{\mathcal{G}} & \text{for } i \in [\tilde{m}^{\mathcal{G}}] \\ \lceil \delta/(2\beta^\dagger) \rceil & \text{for } i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}}] \end{cases}, n_{i,j}^{\mathcal{B}} = \begin{cases} \tilde{n}_{i,j}^{\mathcal{B}} & \text{for } i \in [\tilde{m}^{\mathcal{G}}], j \in [\tilde{m}_i^{\mathcal{B}}] \\ 0 & \text{for } i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}} - 1], j \in [m_i^{\mathcal{B}}]/[\tilde{m}_i^{\mathcal{B}}]. \end{cases}$$

We note that  $n_i^{\mathcal{G}}$ 's and  $n_{i,j}^{\mathcal{B}}$ 's are defined only if they exist.<sup>7</sup>

We provide  $m^{\mathcal{G}}$  such that  $R^{\pi_1(\beta^\dagger)}(H) \leq R_{m^{\mathcal{G}}}^{\pi_1(\beta^\dagger)}$  in the following lemma.

**Lemma A.9.** *Under  $E_1$ , when  $m^{\mathcal{G}} = \lceil 2H\beta^{\dagger 2/3} \rceil$ , we have*

$$R^{\pi_1(\beta^\dagger)}(H) \leq R_{m^{\mathcal{G}}}^{\pi_1(\beta^\dagger)}.$$

*Proof.* From Lemma A.8, with  $\delta = \beta^{\dagger 1/3}$  we have

$$\sum_{i \in [m^{\mathcal{G}}]} \left( n_i^{\mathcal{G}} + \sum_{j \in [m_i^{\mathcal{B}}]} n_{i,j}^{\mathcal{B}} \right) \geq m^{\mathcal{G}} \frac{\delta}{2\beta^\dagger} \geq H,$$

which implies that  $R^{\pi_1(\beta^\dagger)}(H) \leq R_{m^{\mathcal{G}}}^{\pi_1(\beta^\dagger)}$ .  $\square$

From the result of Lemma A.9, we set  $m^{\mathcal{G}} = \lceil 2H\beta^{\dagger 2/3} \rceil$ . For getting a bound for  $\mathbb{E}[R_{m^{\mathcal{G}}}^{\pi_1(\beta^\dagger)}]$ , we provide bounds for  $\mathbb{E}[R_i^{\mathcal{G}}]$  and  $\mathbb{E}[R_{i,j}^{\mathcal{B}}]$  in the following lemma.

<sup>7</sup>  $n_i^{\mathcal{G}}$  and  $n_{i,j}^{\mathcal{B}}$  are not defined for  $i \in [0]$  or  $j \in [0]$ .

**Lemma A.10.** Under  $E_1$  and policy  $\pi_1(\beta^\dagger)$ , for any  $i \in [m^{\mathcal{G}}]$  and  $j \in [m_i^{\mathcal{B}}]$ , we have

$$\mathbb{E}[R_i^{\mathcal{G}}] = \tilde{O}\left(\frac{\delta}{\varrho^{2/3}} + \frac{\delta^2}{\varrho} + \frac{1}{\varrho^{1/3}}\right),$$

and

$$\mathbb{E}[R_{i,j}^{\mathcal{B}}] = \tilde{O}\left(1 + \frac{\delta}{\varrho^{1/3}} + \frac{\delta^2}{\varrho^{2/3}}\right).$$

*Proof.* First we provide a bound for  $R_i^{\mathcal{G}}$  using an upper bound of  $n_i^{\mathcal{G}}$ . Recall that  $a(i)$  is the selected arm for the  $i$ -th good arm. We have

$$\begin{aligned} \tilde{\mu}_t^\varrho(a(i), \beta^\dagger) - \beta^\dagger n_t(a(i)) &= \hat{\mu}_t(a(i)) + \beta^\dagger \frac{\sum_{s=1}^{t-1} n_s(a(i)) \mathbb{1}(a_s = a(i))}{n_t(a(i))} - \beta^\dagger n_t(a(i)) \\ &\leq \hat{\mu}_t(a(i)) + \beta^\dagger \frac{n_t(a(i)) + 1}{2} - \beta^\dagger n_t(a(i)) \\ &= \hat{\mu}_t(a(i)) - \beta^\dagger \frac{n_t(a(i)) - 1}{2}. \end{aligned}$$

Since, under  $E_1$ ,

$$\begin{aligned} \hat{\mu}_t(a(i)) - (\beta^\dagger/2)(n_t(a(i)) - 1) &\leq \bar{\mu}_t(a(i)) + \sqrt{8 \log(H)/n_t(a(i))} - (\beta^\dagger/2)(n_t(a(i)) - 1) \\ &\leq 1 + \sqrt{8 \log(H)/n_t(a(i))} - (\beta^\dagger/2)(n_t(a(i)) - 1), \end{aligned}$$

for  $i \in [\tilde{m}^{\mathcal{G}}]$ , from the policy  $\pi$ , we need to get  $n$  such that

$$1 - \frac{\beta^\dagger}{2}(n - 1) + 2\sqrt{8 \log(H)/n} < 1 - \delta, \quad (29)$$

in which  $n + 1$  is an upper bound for  $n_i^{\mathcal{G}}$ . Let  $n_1 = 2(\delta + \beta^{\dagger 1/3})/\beta^\dagger + 1$  and  $n_2 = C \log(H)/\beta^{\dagger 2/3}$  with some large enough constant  $C > 0$ . Then  $n = n_1 + n_2$  satisfies (29) because  $1 - \beta^\dagger n_1 + 2\sqrt{8 \log(H)/n_2} < 1 - \delta$ . Therefore, for all  $i \in [\tilde{m}^{\mathcal{G}}]$  we have  $n_i^{\mathcal{G}} = \tilde{O}((\delta + \beta^{\dagger 1/3})/\beta^\dagger)$ . Then with the fact that  $n_i^{\mathcal{G}} = \lceil \delta/(2\beta^\dagger) \rceil$  for  $i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}}]$  if they exist and  $\beta^\dagger = \Theta(\varrho)$ , for any  $i \in [m^{\mathcal{G}}]$  we have

$$n_i^{\mathcal{G}} = \tilde{O}((\delta + \beta^{\dagger 1/3})/\beta^\dagger) = \tilde{O}((\delta + \varrho^{1/3})/\varrho).$$

Then for any  $i \in [m^{\mathcal{G}}]$  we have

$$\begin{aligned} \mathbb{E}[R_i^{\mathcal{G}}] &\leq \mathbb{E}\left[\Delta(a(i))n_i^{\mathcal{G}} + \frac{n_i^{\mathcal{G}}(n_i^{\mathcal{G}} - 1)}{2}\varrho\right] \\ &= \tilde{O}\left(\frac{2}{\delta} \int_0^{\delta/2} \frac{\delta + \varrho^{1/3}}{\varrho} x + \left(\frac{\delta + \varrho^{1/3}}{\varrho}\right)^2 \varrho dx\right) \\ &= \tilde{O}\left(\frac{\delta}{\varrho^{2/3}} + \frac{\delta^2}{\varrho} + \frac{1}{\varrho^{1/3}}\right), \end{aligned}$$

where the first equality is obtained from the fact that  $\Delta(a(i))$ 's are i.i.d. random variables with uniform distribution on  $[0, \delta/2]$  and  $n_i^{\mathcal{G}} = \tilde{O}((\delta + \varrho^{1/3})/\varrho)$ .

Now we provide an upper bound of  $n_{i,j}^{\mathcal{B}}$  to get a bound of  $R_{i,j}^{\mathcal{B}}$  for  $i \in [m^{\mathcal{G}}], j \in [m_i^{\mathcal{B}}]$ . Let  $a(i, j)$  be a selected arm for  $j$ -th bad arm in the  $i$ -th episode. When  $\delta/2 < \Delta(a(i, j)) \leq \delta + \varrho^{1/3}$ , as in the case of the good arm,  $n =$

$2(\delta + \beta^\dagger)^{1/3}/\beta^\dagger + 1 + C_1 \log(T)/\beta^\dagger^{2/3}$  for some large enough constant  $C_1 > 0$ , satisfies (29) so that  $n_{i,j}^{\mathcal{B}} = \tilde{O}((\delta + \varrho^{1/3})/\varrho)$  for all  $i \in [\tilde{m}], j \in [\tilde{m}_i^{\mathcal{B}}]$ . When  $\delta + \varrho^{1/3} < \Delta(a(i, j)) \leq 1$ , since under  $E_1$

$$\begin{aligned} \tilde{\mu}_t^o(a(i, j)) - \beta^\dagger n_t(a(i, j)) &= \hat{\mu}_t(a(i, j)) + \beta^\dagger \frac{\sum_{s=1}^{t-1} n_s(a(i, j)) \mathbb{1}(a_s = a(i, j))}{n_t(a(i, j))} - \beta^\dagger n_t(a(i, j)) \\ &\leq \hat{\mu}_t(a(i, j)) + \beta^\dagger \frac{n_t(a(i, j)) + 1}{2} - \beta^\dagger n_t(a(i, j)) \\ &= \hat{\mu}_t(a(i, j)) - \beta^\dagger \frac{n_t(a(i, j)) - 1}{2} \\ &\leq \bar{\mu}_t(a(i, j)) + \sqrt{8 \log(T)/n_t(a(i, j))} - (\beta^\dagger/2)(n_t(a(i, j)) - 1) \\ &\leq \mu_1(a(i, j)) + \sqrt{8 \log(T)/n_t(a(i, j))} - (\beta^\dagger/2)(n_t(a(i, j)) - 1), \end{aligned}$$

from the policy  $\pi$  we need to get  $n \geq 1$  such that

$$\mu_1(a(i, j)) - \frac{\beta^\dagger}{2}(n - 1) + 2\sqrt{8 \log(T)/n} < 1 - \delta, \quad (30)$$

in which  $n + 1$  is an upper bound of  $n_{i,j}^{\mathcal{B}}$  for  $i \in [\tilde{m}], j \in [\tilde{m}_i^{\mathcal{B}}]$ . From a sufficient condition for (30) to hold such that

$$\mu_1(a(i, j)) + 2\sqrt{8 \log(T)/n} < 1 - \delta,$$

we can find that  $n = C_2 \log(T)/(\Delta(a(i, j)) - \delta)^2$  for some large constant  $C_2 > 0$  satisfies (30). Therefore, when  $\delta + \varrho^{1/3} < \Delta(a(i, j)) \leq 1$ , for all  $i \in [\tilde{m}], j \in [\tilde{m}_i^{\mathcal{B}}]$  we have  $n_{i,j}^{\mathcal{B}} = \tilde{O}(1/(\Delta(a(i, j)) - \delta)^2)$ . Then with the fact that  $n_{i,j}^{\mathcal{B}} = 0$  for  $i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}}], j \in [m_i^{\mathcal{B}}]/[\tilde{m}_i^{\mathcal{B}}]$  if they exist, for any  $i \in [m^{\mathcal{G}}]$  and  $j \in [m_i^{\mathcal{B}}]$ , we have

$$n_{i,j}^{\mathcal{B}} = \begin{cases} \tilde{O}((\delta + \varrho^{1/3})/\varrho) & \text{if } \delta/2 < \Delta(a(i, j)) \leq \delta + \varrho^{1/3} \\ \tilde{O}(1/(\Delta(a(i, j)) - \delta)^2) & \text{if } \delta + \varrho^{1/3} < \Delta(a(i, j)) \leq 1 \end{cases}$$

Then for any  $i \in [m^{\mathcal{G}}], j \in [m_i^{\mathcal{B}}]$ , we have

$$\begin{aligned} \mathbb{E}[R_{i,j}^{\mathcal{B}}] &\leq \mathbb{E} \left[ \Delta(a(i, j)) n_{i,j}^{\mathcal{B}} + \frac{n_{i,j}^{\mathcal{B}}(n_{i,j}^{\mathcal{B}} - 1)}{2} \varrho \right] \\ &= \tilde{O} \left( \frac{1}{1 - \delta/2} \left( \int_{\delta/2}^{\delta + \varrho^{1/3}} \frac{\delta + \varrho^{1/3}}{\varrho} x + \left( \frac{\delta + \varrho^{1/3}}{\varrho} \right)^2 \varrho dx + \int_{\delta + \varrho^{1/3}}^1 \frac{1}{(x - \delta)^2} x + \frac{1}{(x - \delta)^4} \varrho dx \right) \right) \\ &= \tilde{O} \left( 1 + \frac{\delta}{\varrho^{1/3}} + \frac{\delta^2}{\varrho^{2/3}} \right). \end{aligned}$$

□

Recall that  $R_{m^{\mathcal{G}}}^{\pi_1(\beta^\dagger)} = \sum_{i=1}^{m^{\mathcal{G}}} (R_i^{\mathcal{G}} + \sum_{j \in [m_i^{\mathcal{B}}]} R_{i,j}^{\mathcal{B}})$ . With  $\beta^\dagger = \Theta(\varrho)$ ,  $\delta = \Theta(\varrho^{1/3})$ , and  $m^{\mathcal{G}} = \lceil 2H\beta^\dagger^{2/3} \rceil$ , from Lemmas A.9 and A.10, and the fact that  $m_i^{\mathcal{B}}$ 's are i.i.d. random variables with geometric distribution with  $\mathbb{E}[m_i^{\mathcal{B}}] = 2/\delta - 1$ , we have

$$\begin{aligned} \mathbb{E}[R^{\pi_1(\beta^\dagger)}(H)] &= O(\mathbb{E}[R_{m^{\mathcal{G}}}^{\pi_1(\beta^\dagger)}]) \\ &= O \left( \mathbb{E} \left[ \sum_{i=1}^{m^{\mathcal{G}}} \left( R_i^{\mathcal{G}} + \sum_{j \in [m_i^{\mathcal{B}}]} R_{i,j}^{\mathcal{B}} \right) \right] \right) \\ &= \tilde{O} \left( H \varrho^{2/3} \left( \left( \frac{\delta}{\varrho^{2/3}} + \frac{\delta^2}{\varrho} + \frac{1}{\varrho^{1/3}} \right) + \frac{1}{\delta} \left( 1 + \frac{\delta}{\varrho^{1/3}} \right) \right) \right) \\ &= \tilde{O} \left( H \varrho^{2/3} \left( \frac{\delta}{\varrho^{2/3}} + \frac{\delta^2}{\varrho} + \frac{1}{\delta} + \frac{1}{\varrho^{1/3}} \right) \right) \\ &= \tilde{O} \left( H \varrho^{1/3} \right). \end{aligned} \quad (31)$$

Now we consider the case when  $\varrho = O(1/H^{3/2})$ . We have set  $\beta^\dagger$  as the smallest element in  $\mathcal{B}$  that exceeds  $\max\{\varrho, 1/H^{3/2}\}$ ; hence we have  $\beta^\dagger = \Theta(1/H^{3/2})$ . For getting  $R_{m^\mathcal{G}}^{\pi_1(\beta^\dagger)}$ , here we define how the policy  $\pi_1(\beta^\dagger)$  works after  $H$  time steps such that it pulls  $H$  times a good arm and once a bad arm. For the last arm  $\tilde{a}$  over the horizon  $H$ , it pulls the arm up to  $H$  times if  $\tilde{a}$  is a good arm and  $\tilde{n}_{\tilde{m}^\mathcal{G}}^\mathcal{G} < H$ . With slight abuse of notation, for  $i \in [m^\mathcal{G}]$  and  $j \in [m_i^\mathcal{B}]$ , let  $n_i^\mathcal{G}$  and  $n_{i,j}^\mathcal{B}$  be the number of pulls of the good arm in the  $i$ -th episode and the  $j$ -th bad arm in the  $i$ -th episode by the policy, respectively. Here we define  $n_i^\mathcal{G}$ 's and  $n_{i,j}^\mathcal{B}$ 's as follows:

If  $\tilde{a}$  is a good arm and  $\tilde{n}_{\tilde{m}^\mathcal{G}}^\mathcal{G} < H$ , then

$$n_i^\mathcal{G} = \begin{cases} \tilde{n}_i^\mathcal{G} & \text{for } i \in [\tilde{m}^\mathcal{G} - 1] \\ H & \text{for } i \in [m^\mathcal{G}]/[\tilde{m}^\mathcal{G} - 1] \end{cases}, n_{i,j}^\mathcal{B} = \begin{cases} \tilde{n}_{i,j}^\mathcal{B} & \text{for } i \in [\tilde{m}^\mathcal{G}], j \in [\tilde{m}_i^\mathcal{B}] \\ 0 & \text{for } i \in [m^\mathcal{G}]/[\tilde{m}^\mathcal{G}], j \in [m_i^\mathcal{B}]/[\tilde{m}_i^\mathcal{B}]. \end{cases}$$

Otherwise,

$$n_i^\mathcal{G} = \begin{cases} \tilde{n}_i^\mathcal{G} & \text{for } i \in [\tilde{m}^\mathcal{G}] \\ H & \text{for } i \in [m^\mathcal{G}]/[\tilde{m}^\mathcal{G}] \end{cases}, n_{i,j}^\mathcal{B} = \begin{cases} \tilde{n}_{i,j}^\mathcal{B} & \text{for } i \in [\tilde{m}^\mathcal{G}], j \in [\tilde{m}_i^\mathcal{B}] \\ 0 & \text{for } i \in [m^\mathcal{G}]/[\tilde{m}^\mathcal{G} - 1], j \in [m_i^\mathcal{B}]/[\tilde{m}_i^\mathcal{B}]. \end{cases}$$

From Lemma A.8, we observe that  $n_i^\mathcal{G} \geq \min\{\delta/(2\beta^\dagger), H\}$  for  $i \in [m^\mathcal{G}]$ . Then under  $E_1$ , if  $m^\mathcal{G} = C_3$  for some large enough constant  $C_3 > 0$ , then with  $\delta = \beta^{\dagger 1/3}$  and  $\beta^\dagger = \Theta(1/H^{3/2})$ , we have

$$\sum_{i \in [m^\mathcal{G}]} n_i^\mathcal{G} \geq H,$$

which implies  $R^{\pi_1(\beta^\dagger)}(H) \leq R_{m^\mathcal{G}}^{\pi_1(\beta^\dagger)}$ . Therefore, we set  $m^\mathcal{G} = C_3$ . For getting a bound for  $\mathbb{E}[R_{m^\mathcal{G}}^{\pi_1(\beta^\dagger)}]$ , we provide bounds for  $\mathbb{E}[R_i^\mathcal{G}]$  and  $\mathbb{E}[R_{i,j}^\mathcal{B}]$  in the following lemma.

**Lemma A.11.** *Under  $E_1$  and policy  $\pi_1(\beta^\dagger)$ , for any  $i \in [m^\mathcal{G}]$  and  $j \in [m_i^\mathcal{B}]$ , we have*

$$\mathbb{E}[R_i^\mathcal{G}] = O(\delta H + H^2 \varrho),$$

and

$$\mathbb{E}[R_{i,j}^\mathcal{B}] = \tilde{O}\left(H\delta^2 + \sqrt{H}\delta + 1\right).$$

*Proof.* First we provide a bound for  $R_i^\mathcal{G}$  using an upper bound of  $n_i^\mathcal{G}$ . With the definition of  $n_i^\mathcal{G} = H$  for  $i \in [m^\mathcal{G}]/[\tilde{m}^\mathcal{G}]$ , for any  $i \in [m^\mathcal{G}]$  we have

$$n_i^\mathcal{G} \leq H.$$

Recall that  $a(i)$  is the selected arm for the  $i$ -th good arm. Then, for any  $i \in [m^\mathcal{G}]$ , we have

$$\begin{aligned} \mathbb{E}[R_i^\mathcal{G}] &\leq \mathbb{E}\left[\Delta(a(i))n_i^\mathcal{G} + \frac{n_i^\mathcal{G}(n_i^\mathcal{G} - 1)}{2}\varrho\right] \\ &= O(\delta H + H^2 \varrho), \end{aligned}$$

where the first equality is obtained from the fact that  $\Delta(a(i))$ 's are i.i.d. random variables with uniform distribution on  $[0, \delta/2]$  and  $n_i^\mathcal{G} \leq H$ .

Now we provide an upper bound of  $n_{i,j}^\mathcal{B}$  to get a bound of  $R_{i,j}^\mathcal{B}$  for  $i \in [m^\mathcal{G}], j \in [m_i^\mathcal{B}]$ . Let  $a(i, j)$  be a selected arm for  $j$ -th bad arm in the  $i$ -th episode. When  $\delta/2 < \Delta(a(i, j)) \leq \delta + 1/\sqrt{H}$ , as in the case of the good arm,  $n_{i,j}^\mathcal{B} \leq H$  for all  $i \in [\tilde{m}], j \in [\tilde{m}_i^\mathcal{B}]$ . When  $\delta + 1/\sqrt{H} < \Delta(a(i, j)) \leq 1$ , from the policy  $\pi_1(\beta^\dagger)$  under  $E_1$ , we need to get  $n \geq 1$  such that

$$\mu_1(a(i, j)) - \frac{\beta^\dagger}{2}(n-1) + 2\sqrt{8\log(T)/n} < 1 - \delta, \quad (32)$$



in which  $n + 1$  is an upper bound of  $n_{i,j}^{\mathcal{B}}$  for  $i \in [\tilde{m}], j \in [\tilde{m}_i^{\mathcal{B}}]$ . From a sufficient condition for (32) to hold such that

$$\mu_1(a(i, j)) + 2\sqrt{8 \log(H)/n} < 1 - \delta,$$

we can find that  $n = C_4 \log(H)/(\Delta(a(i, j)) - \delta)^2$  for some large constant  $C_4 > 0$  satisfies (32). Therefore, when  $\delta + 1/\sqrt{H} < \Delta(a(i, j)) \leq 1$ , for all  $i \in [\tilde{m}], j \in [\tilde{m}_i^{\mathcal{B}}]$  we have  $n_{i,j}^{\mathcal{B}} = \tilde{O}(1/(\Delta(a(i, j)) - \delta)^2)$ . Then with the fact that  $n_{i,j}^{\mathcal{B}} = 0$  for  $i \in [m^{\mathcal{G}}]/[\tilde{m}^{\mathcal{G}}], j \in [m_i^{\mathcal{B}}]/[\tilde{m}_i^{\mathcal{B}}]$ , for any  $i \in [m^{\mathcal{G}}], j \in [m_i^{\mathcal{B}}]$ , if  $\delta/2 < \Delta(a(i, j)) \leq \delta + 1/\sqrt{H}$ , we have

$$n_{i,j}^{\mathcal{B}} \leq H,$$

and if  $\delta + 1/\sqrt{H} < \Delta(a(i, j)) \leq 1$ , we have

$$n_{i,j}^{\mathcal{B}} = \tilde{O}(1/(\Delta(a(i, j)) - \delta)^2).$$

For any  $i \in [m^{\mathcal{G}}], j \in [m_i^{\mathcal{B}}]$ , we obtain

$$\begin{aligned} \mathbb{E}[R_{i,j}^{\mathcal{B}}] &\leq \mathbb{E} \left[ \Delta(a(i, j)) n_{i,j}^{\mathcal{B}} + \frac{n_{i,j}^{\mathcal{B}}(n_{i,j}^{\mathcal{B}} - 1)}{2} \varrho \right] \\ &= \tilde{O} \left( \frac{1}{1 - \delta/2} \left( \int_{\delta/2}^{\delta+1/\sqrt{H}} Hx + H^2 \varrho dx + \int_{\delta+1/\sqrt{H}}^1 \frac{1}{(x - \delta)^2} x + \frac{1}{(x - \delta)^4} \varrho dx \right) \right) \\ &= \tilde{O} \left( H\delta^2 + \sqrt{H}\delta + 1 \right). \end{aligned}$$

□

It follows that, with  $\delta = \Theta(1/\sqrt{H})$  and  $m^{\mathcal{G}} = C_3$ , we have

$$\begin{aligned} \mathbb{E}[R^{\pi_1(\beta^\dagger)}(H)] &= O(\mathbb{E}[R_{m^{\mathcal{G}}}^{\pi_1(\beta^\dagger)}]) \\ &= O \left( \mathbb{E} \left[ \sum_{i \in [m^{\mathcal{G}}]} \left( R_i^{\mathcal{G}} + \sum_{j \in [m_i^{\mathcal{B}}]} R_{i,j}^{\mathcal{B}} \right) \right] \right) \\ &= \tilde{O} \left( (\delta H + H^2 \varrho) + \frac{1}{\delta} (H\delta^2 + \sqrt{H}\delta + 1) \right) \\ &= \tilde{O}(\sqrt{H}), \end{aligned} \tag{33}$$

where the third equality is obtained from Lemma A.11 and  $\mathbb{E}[m_i^{\mathcal{B}}] = 2/\delta - 1$ .

Finally, we can conclude the proof by noting that from (31) and (33), for  $\varrho = o(1)$ , we have

$$\mathbb{E}[R^{\pi_1(\beta^\dagger)}(H)] = \tilde{O}(\max\{H\varrho^{1/3}, \sqrt{H}\}).$$

Therefore, by summing regrets from  $\lceil T/H \rceil$  number of blocks, we have shown that

$$\mathbb{E}[R_1^\pi(T)] = \tilde{O}((T/H) \max\{H\varrho^{1/3}, \sqrt{H}\}) = \tilde{O}(\max\{T\varrho^{1/3}, T/\sqrt{H}\}). \tag{34}$$

**Upper bounding  $\mathbb{E}[R_2^\pi(T)]$ .** We observe that the EXP3 is run for  $\lceil T/H \rceil$  decision rounds and the number of policies (i.e.  $\pi_i(\beta)$  for  $\beta \in \mathcal{B}$ ) is  $B$ . Denote the maximum absolute sum of rewards of any block with length  $H$  by a random variable  $Q'$ . We first provide a bound for  $Q'$  using concentration inequalities. For any block  $i$ , we have

$$\left| \sum_{t=(i-1)H+1}^{i \cdot H \wedge T} \mu_t(a_t^\pi) + \eta_t \right| \leq \left| \sum_{t=(i-1)H+1}^{i \cdot H \wedge T} \mu_t(a_t^\pi) \right| + \left| \sum_{t=(i-1)H+1}^{i \cdot H \wedge T} \eta_t \right|. \tag{35}$$

Denote by  $\mathcal{T}_i$  the set of time steps in the  $i$ -th block. We define the event  $E_2(i) = \{|\hat{\mu}_t(a) - \bar{\mu}_t(a)| \leq \sqrt{7 \log(T)/n_t(a)}, \text{ for all } t \in \mathcal{T}_i, a \in \mathcal{A}(i)\}$  and  $E_2 = \bigcap_{i \in \lceil [T/H] \rceil} E_2(i)$ . From Lemma A.2, by following similar steps of the proof of Lemma 5 in Auer et al. (2019), with  $H = \sqrt{T}$  we have

$$\mathbb{P}(E_2^c) \leq \sum_{i \in \lceil [T/H] \rceil} \frac{2H^2}{T^{7/2}} \leq \frac{2}{T^2}.$$

By assuming that  $E_2$  holds true, we can get a lower bound for  $\mu_t(a_t^\pi)$ , which may be a negative value from rotting, for getting an upper bound for  $|\sum_{t=(i-1)H+1}^{iH \wedge T} \mu_t(a_t^\pi)|$ . Let  $\beta_{\max}$  denote the maximum value in  $\mathcal{B}$ . From the policy  $\pi$  with  $H = \sqrt{T}$ , when  $\mu_1(a) = 0$  for some arm  $a$ , since

$$\begin{aligned} \tilde{\mu}_t^o(a(i, j)) - \varrho n_t(a(i, j)) + \sqrt{\frac{8 \log(H)}{n_t(a(i, j))}} &\leq \hat{\mu}_t(a(i, j)) - (\varrho/2)(n_t(a(i, j)) - 1) + \sqrt{\frac{4 \log(T)}{n_t(a(i, j))}} \\ &= \bar{\mu}_t(a(i, j)) - (\varrho/2)(n_t(a(i, j)) - 1) + \sqrt{\frac{4 \log(T)}{n_t(a(i, j))}} + \sqrt{\frac{7 \log(T)}{n_t(a(i, j))}} \\ &\leq \mu_1(a) - (\varrho/2)(n_t(a(i, j)) - 1) + \sqrt{\frac{4 \log(T)}{n_t(a(i, j))}} + \sqrt{\frac{7 \log(T)}{n_t(a(i, j))}} \\ &\leq \sqrt{\frac{4 \log(T)}{n_t(a(i, j))}} + \sqrt{\frac{7 \log(T)}{n_t(a(i, j))}}, \end{aligned}$$

we need to find an positive integer  $n$  such that

$$\sqrt{4 \log(T)/n} + \sqrt{7 \log(T)/n} \leq 1 - \beta_{\max}^{1/3},$$

in which  $n$  is an upper bound for the number of pulls of arm  $a$ . From  $\beta_{\max}^{1/3} = 1/2$ , we can observe that the condition holds with  $n = \lceil 92 \log(T) \rceil$ . From this fact, we can find that for any selected arm  $a$  from  $\pi$ , we have  $\mu_t(a) \geq -(92 \log(T) + 1)\varrho$ . Then under  $E_2$ , with  $\mu_t(a) \leq 1$ , for any  $i \in \lceil [T/H] \rceil$  we have  $|\sum_{t=(i-1)H+1}^{iH \wedge T} \mu_t(a_t^\pi)| \leq \max\{(92 \log(T) + 1)\varrho H, H\}$ .

Next we provide a bound for  $|\sum_{t=(i-1)H+1}^{iH \wedge T} \eta_t|$ . We define the event  $E_3(i) = \{|\sum_{t=(i-1)H+1}^{iH \wedge T} \eta_t| \leq 2\sqrt{H \log(T)}\}$  and  $E_3 = \bigcap_{i \in \lceil [T/H] \rceil} E_3(i)$ . From Lemma A.2, for any  $i \in \lceil [T/H] \rceil$ , we have

$$\mathbb{P}(E_3(i)^c) \leq \frac{2}{T^2}.$$

Then, under  $E_2 \cap E_3$ , with (35), we have

$$Q' \leq \max\{93H \log(T)\varrho, H\} + 2\sqrt{H \log(T)} \leq 93H \log(T) + 2\sqrt{H \log(T)},$$

which implies  $1/2 + \sum_{t=(i-1)H+1}^{iH \wedge T} r_t / (186H \log T + 4\sqrt{H \log T}) \in [0, 1]$ . With the rescaling and translation of rewards in Algorithm 2, from Corollary 3.2. in Auer et al. (2002b), we have

$$\mathbb{E}[R_2^\pi(T) | E_2 \cap E_3] = \tilde{O}\left((93H \log T + 2\sqrt{H \log T})\sqrt{BT/H}\right) = \tilde{O}\left(\sqrt{HBT}\right). \quad (36)$$

Note that the expected regret from EXP3 is trivially bounded by  $o(H^2(T/H)) = o(TH)$  and  $B = O(\log(T))$ . Then, with (36), we have

$$\begin{aligned} \mathbb{E}[R_2^\pi(T)] &= \mathbb{E}[R_2^\pi(T) | E_2 \cap E_3] \mathbb{P}(E_2 \cap E_3) + \mathbb{E}[R_2^\pi(T) | E_2^c \cup E_3^c] \mathbb{P}(E_2^c \cup E_3^c) \\ &= \tilde{O}\left(\sqrt{HBT}\right) + o(TH) (4/T^2) \\ &= \tilde{O}\left(\sqrt{HBT}\right). \end{aligned} \quad (37)$$

Finally, from (28), (34), and (37), with  $H = \sqrt{T}$ , we have

$$\mathbb{E}[R^\pi(T)] = \tilde{O}\left(\max\{T^{3/4}, \varrho^{1/3}T\}\right),$$

which concludes the proof.