
A Simple and Optimal Policy Design for Online Learning with Safety against Heavy-tailed Risk

David Simchi-Levi

Institute for Data, Systems, and Society
Massachusetts Institute of Technology
Cambridge, MA 02139
dslevi@mit.edu

Zeyu Zheng

Industrial Engineering & Operations Research
University of California, Berkeley
Berkeley, CA 94720
zyzheng@berkeley.edu

Feng Zhu

Institute for Data, Systems, and Society
Massachusetts Institute of Technology
Cambridge, MA 02139
fengzhu@mit.edu

Abstract

We consider the classical multi-armed bandit problem and design simple-to-implement new policies that simultaneously enjoy two properties: worst-case optimality for the expected regret, and safety against heavy-tailed risk for the regret distribution. Recently, [10] showed that information-theoretic optimized bandit policies as well as standard UCB policies suffer from some serious heavy-tailed risk; that is, the probability of incurring a linear regret slowly decays at a polynomial rate of $1/T$, as T (the time horizon) increases. Inspired by their result, we further show that any policy that incurs an instance-dependent $O(\ln T)$ regret must incur a linear regret with probability $\Omega(\text{poly}(1/T))$ and that the heavy-tailed risk actually exists for all “instance-dependent consistent” policies. Next, for the two-armed bandit setting, we provide a simple policy design that (i) has the worst-case optimality for the expected regret at order $\tilde{O}(\sqrt{T})$ and (ii) has the worst-case tail probability of incurring a linear regret decay at an exponential rate $\exp(-\Omega(\sqrt{T}))$. We further prove that this exponential decaying rate of the tail probability is optimal across all policies that have worst-case optimality for the expected regret. Finally, we generalize the policy design and analysis to the general setting with an arbitrary K number of arms. We provide detailed characterization of the tail probability bound for any regret threshold under our policy design. Numerical experiments are conducted to illustrate the theoretical findings. Our results reveal insights on the incompatibility between consistency and light-tailed risk, whereas indicate that worst-case optimality on expected regret and light-tailed risk are compatible.

1 Introduction

The stochastic multi-armed bandit (MAB) problem is a widely studied problem in the domain of sequential decision-making under uncertainty, with many applications such as online advertising, recommendation systems, clinical trials, financial portfolio design, etc. It also has valuable theoretical insights exhibiting the exploration-exploitation trade-off. For policy design and analysis, much of the MAB literature uses the metric of maximizing the *expected* cumulative reward, or equivalently minimizing the *expected regret* (where *regret* is defined as the difference between the cumulative reward obtained by always pulling the best arm and by executing a policy that does not a priori

know the reward distributions). The optimality of a policy is often characterized through its expected regret's rate (order of dependence) on the experiment horizon T .

In this paper, however, we show that renowned policies (such as the standard Upper Confidence Bound (UCB) policy, the Successive Elimination (SE) policy and the Thompson Sampling (TS) policy) that are designed to enjoy optimality in terms of expected regret can incur a “heavy-tailed risk”. That is, the distribution of the regret has a heavy tail — the probability of incurring a linear regret slowly decays at a polynomial rate $\Omega(\text{poly}(1/T))$ as T tends to infinity. In contrast, a “light-tailed” risk means that the probability of a policy incurring a linear regret decays at an exponential rate $\exp(-\Omega(T^\beta))$ for some $\beta > 0$. The heavy-tailed risk can be undesired when an MAB policy is used in applications that are sensitive to tail risks (e.g., finance, healthcare, supply chain). Our work is primarily motivated by trying to answer the following question. *Can we design a simple policy that on one hand enjoys the optimality under the expected regret notion, whereas on the other hand has light-tailed risk?* The results of our work are summarized as follows.

1. Instant-dependent consistency and light-tailed risk are incompatible. Inspired by the results in [8], we adapt their change of measure argument to show that any policy that incurs an instance-dependent $O(\ln T)$ regret must incur a linear regret with probability $\Omega(\text{poly}(1/T))$. We additionally show that any instant-dependent consistent policy cannot be light-tailed: if an instant-dependent consistent policy has the probability of incurring a linear regret decay as $\exp(-f(T))$, then $f(T)$ must be $o(T^\beta)$ as $T \rightarrow +\infty$ for any $\beta > 0$.

2. Worst-case optimality and light-tailed risk can co-exist. Starting from the two-armed bandit setting, we provide a simple policy design and prove that it enjoys both the worst-case optimality $\tilde{O}(\sqrt{T})$ for the expected regret and the light-tailed risk $\exp(-\Omega(\sqrt{T}))$ for the regret distribution. We also show that such exponential decaying rate of the tail probability is optimal within the class of worst-case optimal policies.

3. Extensions and Experiments. We extend the result from the two-armed bandit to the general K -armed bandit and characterize the tail probability bound for any regret threshold in an explicit form. We prove that under our policy design, the worst-case probability of incurring a regret larger than x is bounded by $\exp(-\Omega(x/K\sqrt{T}))$. We find that the associated proof techniques are new and may benefit broader analysis on regret distribution and tail risk. Our result also partially solves an open problem raised in [14]. Experiments are also conducted to illustrate our theoretical findings.

1.1 Related Work

Our work builds upon the vast literature of designing and analyzing policies for the multi-armed bandit (MAB) problem and its various extensions. A comprehensive review can be found in [6, 17, 14]. A standard paradigm for obtaining a near-optimal regret is to first *fix* some confidence parameter $\delta > 0$. Then a “good event” is defined such that good properties are retained conditioned on the event (for example, in the MAB problem, the good event is such that the mean of each arm always lies in the confidence bound). Then one can obtain both high-probability and worst-case expected regret bounds through careful analysis on the good event. It is known that the MAB problem has the following regret bound: for any fixed $\delta \in (0, 1)$, the regret bound of UCB is bounded by $O(\sqrt{KT \ln(T/\delta)})$ with probability at least $1 - \delta$. Or equivalently speaking, the probability of incurring a $\Omega(\sqrt{KT \ln(T/\delta)})$ regret is bounded by δ . However, the parameter δ must be an input parameter for the policy. We will discuss this issue in more details in Section 3. In Section 17.1 of [14], an open question is asked: is it possible to design a single policy such that the worst-case probability of incurring a $\Omega(\sqrt{KT \ln(1/\delta)})$ regret is bounded by δ for *any* $\delta > 0$? We partially solve this question by designing a policy such that for *any* $\delta > 0$, the probability of incurring a $\Omega(K\sqrt{T}(\sqrt{\ln T} + \ln(1/\delta)))$ regret is bounded by δ .

There has been not much work on understanding the tail risk of bandit algorithms. Our work is inspired by [10]. They showed that optimized UCB bandit designs are fragile to mis-specifications and they modified UCB algorithms to ensure a desired polynomial rate of tail risk, which makes the algorithms more robust to mis-specifications. Different from their work, we propose a simple and new policy design that leads to both light-tailed risk (tail bound exponentially decaying with \sqrt{T}) and worst-case optimality (expected regret bounded by $\tilde{O}(\sqrt{T})$). Besides [10], two earlier works are [5, 15] and they study the concentration properties of the regret around the instance-dependent mean $O(\ln T)$. They show that in general the regret of the policies concentrate only at a polynomial

rate. That is, the probability of incurring a regret of $c(\ln T)^p$ (with $c > 0$ and $p > 1$ fixed) is approximately polynomially decaying with T . Different from our work, the concentration in their work is under an instance-dependent environment, and so such polynomial rate might be different across different instances. Nevertheless, their results indicate that standard bandit algorithms generally have undesirable concentration properties. Recently, [4] show that an online learning policy with the goal of obtaining logarithmic regret can be fragile, in the sense that a mis-specified risk parameter (e.g., the parameter for subgaussian noises) in the policy can incur an instance-dependent expected regret polynomially dependent on T . They then focus on robust algorithm development to circumvent the issue. Note that their goal is to handle mis-specification related with risk, but still the task is to minimize the expected regret.

Very recently, there are some works trying to understand the behaviour of UCB and TS policies by considering the diffusion approximations (see, e.g., [3, 19, 9, 12]). Although some distributional characterizations of regret are obtained in these works, the asymptotic regime is typically set such that the gaps between arm means shrink with the total time horizon. This makes their results incomparable to ours because we study how the tail probability decays with T under unchanged environments. Another line of works closely related with ours involve solving risk-averse formulations of the MAB problem (see, e.g., [16, 11, 20, 18, 8, 13]). Compared to standard MAB problems, the main difference in their works is that arm optimality is defined using formulations other than the expected value, such as mean-variance criteria and (conditional) value-at-risk. These formulations consider some single metric that is different compared to the expected regret. From the formulation perspective, our work is different in the sense that we develop policies that on one hand maintain the classical worst-case optimal expected regret, whereas simultaneously achieve light-tailed risk bound. The policy design and analysis in our work are therefore also different from the literature.

2 Setup and Notation

The setting is described as follows. Fix a time horizon of T and the number of arms as K . In each time $t \in [T]$, the decision maker (DM) needs to decide which arm $A_t \in [K]$ should be pulled. To be more precise, let $H_t = \{A_1, X_1, \dots, A_{t-1}, X_{t-1}\}$ be the history prior to time t . When $t = 1$, $H_1 = \emptyset$. In time t , the DM adopt a policy $\pi_t : H_t \mapsto A_t$ that maps the history H_t to an action A_t , where A_t follows a discrete probability distribution on $[K]$ determined by H_t . The environment then independently samples a reward $r_{t,A_t} = \theta_{A_t} + \epsilon_{t,A_t}$ and reveals it to the DM. Here, θ_{A_t} is the mean reward of arm A_t , and ϵ_{t,A_t} is an independent zero-mean noise term. We assume that ϵ_{t,A_t} is σ -subgaussian. That is, there exists a $\sigma > 0$ such that for any time t and arm k ,

$$\max \{ \mathbb{P}(\epsilon_{t,k} \geq x), \mathbb{P}(\epsilon_{t,k} \leq -x) \} \leq \exp(-x^2/(2\sigma^2)).$$

Let $\theta = (\theta_1, \dots, \theta_K)$ be the mean vector. Let $\theta_* = \max\{\theta_1, \dots, \theta_K\}$ be the optimal mean reward among the K arms. Note that DM does not know θ at the beginning, except that $\theta \in [0, 1]^K$. The empirical regret of the policy $\pi = (\pi_1, \dots, \pi_T)$ under the mean vector θ over a time horizon of T is defined as

$$\hat{R}_\theta^\pi(T) = \theta_* \cdot T - \sum_{t=1}^T (\theta_{A_t} + \epsilon_{t,A_t}).$$

Let $\Delta_k = \theta_* - \theta_k$ be the gap between the optimal arm and the k th arm. Let $n_{t,k}$ be the number of times arm k has been pulled up to time t . That is, $n_{t,k} = \sum_{s=1}^t \mathbb{1}\{A_s = k\}$. For simplicity, we will also use $n_k = n_{T,k}$ to denote the total number of times arm k is pulled throughout the whole time horizon. We define $t_k(n)$ as the time period that arm k is pulled for the n th time. Define the pseudo regret as

$$R_\theta^\pi(T) = \sum_{k=1}^K n_k \Delta_k$$

and the genuine noise as

$$N^\pi(T) = \sum_{t=1}^T \epsilon_{t,A_t} = \sum_{k=1}^K \sum_{m=1}^{n_k} \epsilon_{t_k(m),k}.$$

Then the empirical regret can also be written as $\hat{R}_\theta^\pi(T) = R_\theta^\pi(T) - N^\pi(T)$. The following simple lemma gives the mean and the tail probability of the genuine noise $N^\pi(T)$. Intuitively, it shows

when bounding the mean or the tail probability of the empirical regret, one only need to consider the pseudo regret term. We will make it more precise when we discuss the proof of main theorems.

Lemma 1 *We have $\mathbb{E}[N^\pi(T)] = 0$ and*

$$\max \{ \mathbb{P}(N^\pi(T) \geq x), \mathbb{P}(N^\pi(T) \leq -x) \} \leq \exp\left(\frac{-x^2}{2\sigma^2 T}\right).$$

Before proceeding, we introduce some other notations. Throughout the paper, we use $O(\cdot)$ ($\tilde{O}(\cdot)$) and $\Omega(\cdot)$ ($\tilde{\Omega}(\cdot)$) to present upper and lower bounds on the growth rate up to constant (logarithmic) factors, and $\Theta(\cdot)$ ($\tilde{\Theta}(\cdot)$) to characterize the rate when the upper and lower bounds match up to constant (logarithmic) factors. We use $o(\cdot)$ to present strictly dominating upper bounds. In addition, for any $a, b \in \mathbb{R}$, $a \wedge b = \min\{a, b\}$ and $a \vee b = \max\{a, b\}$. For any $a \in \mathbb{R}$, $a_+ = \max\{a, 0\}$.

2.1 Light-tailed Risk, Instance-dependent Consistency, Worst-case Optimality

Now we discuss several core concepts that we are interested in.

1. Light-tailed risk. A policy is called light-tailed, if for any constant $c > 0$, there exists some $\beta > 0$ and constant $C > 0$ such that

$$\limsup_{T \rightarrow +\infty} \frac{\ln \left\{ \sup_{\theta} \mathbb{P} \left(\hat{R}_{\theta}^{\pi}(T) > cT \right) \right\}}{T^{\beta}} \leq -C.$$

Note that here, we allow β and C to be dependent on c . In brief, a policy has light-tailed risk if the probability of incurring a linear regret can be bounded by an exponential term of polynomial T :

$$\sup_{\theta} \mathbb{P}(\hat{R}_{\theta}^{\pi}(T) \geq cT) = \exp(-\Omega(T^{\beta}))$$

for some $\beta > 0$. If a policy is not light-tailed, we say it is *heavy-tailed*.

2. Instance-dependent consistency. A policy is called consistent, if for any underlying true mean vector θ , the policy has that

$$\limsup_{T \rightarrow +\infty} \frac{\mathbb{E} \left[\hat{R}_{\theta}^{\pi}(T) \right]}{T^{\beta}} = 0$$

holds for any $\beta > 0$. We emphasize that here “instant-dependent” means the formula above holds for any *fixed* environment parameter θ . In brief, a policy is consistent if the expected regret can never be polynomially growing in T for any fixed instance.

3. Worst-case optimality. A policy is said to be worst-case optimal, if for any $\beta > 0$, we have

$$\limsup_{T \rightarrow +\infty} \frac{\sup_{\theta} \mathbb{E} \left[\hat{R}_{\theta}^{\pi}(T) \right]}{T^{1/2+\beta}} = 0.$$

In brief, a policy is worst-case optimal if the worst-case expected regret can never be growing in a polynomial rate faster than $T^{1/2}$. Note that here we adopt a relaxed definition of optimality, in the sense that we do not clarify how the regret scale with the number of arms K compared to that in literature. The notion of worst-case optimality in this work focuses on the dependence on T . For example, a policy with worst-case regret $O(\text{poly}(K)\sqrt{T} \cdot \text{poly}(\ln T))$ is also optimal by our definition.

It is well known that for the multi-armed bandit problem, one can design algorithms to achieve both instance-dependent consistency and worst-case optimality. Among them, two types of policies are of prominent interest: Successive Elimination (SE) and Upper Confidence Bound (UCB). We list the algorithm paradigms in Algorithm 1 and 2. The bonus term $\text{rad}(n)$ is typically set as

$$\text{rad}(n) = \sigma \sqrt{\frac{\eta \ln T}{n}} \quad (1)$$

with $\eta > 0$ being some tuning parameter.

Algorithm 1 Successive Elimination

```

1:  $\mathcal{A} = [K]$ .  $t \leftarrow 0$ .
2: while  $t < T$  do
3:   Pull each arm in  $\mathcal{A}$  once.  $t \leftarrow t + K$ .
4:   Eliminate any  $k \in \mathcal{A}$  from  $\mathcal{A}$  if
        $\exists k' : \hat{\mu}_{t,k'} - \text{rad}(n_{t,k'}) > \hat{\mu}_{t,k} + \text{rad}(n_{t,k})$ 
5: end while

```

Algorithm 2 Upper Confidence Bound

```

1:  $\mathcal{A} = [K]$ .  $t \leftarrow 0$ .
2: while  $t < T$  do
3:    $t \leftarrow t + 1$ .
4:   Pull the arm with the highest UCB:
        $\text{UCB}_{t-1,k} = \hat{\mu}_{t-1,k} + \text{rad}(n_{t-1,k})$ .
5: end while

```

3 Main Results: Two-armed Bandit

We start from the simple two-armed bandit setting. The general multi-armed setting is deferred to the next section. We first show that standard policies considered widely in the literature are heavy-tailed. The result reveals a fundamental incompatibility between consistency and light-tailed risk. Then we show how to add a simple twist to standard confidence bound based policies to obtain light-tailed policies. Moreover, we show that our design leads to an optimal tail decaying rate.

3.1 Instance-dependent Consistency Causes Heavy-tailed Risk

Theorem 1 *If a policy π is instance-dependent consistent, then it can never be light-tailed. Moreover, if π satisfies*

$$\limsup_{T \rightarrow +\infty} \frac{\mathbb{E} \left[\hat{R}_\theta^\pi(T) \right]}{\ln T} < +\infty, \quad (2)$$

then we have the following stronger argument. For any $c > 0$, there exists $C_\pi > 0$ such that

$$\liminf_{T \rightarrow +\infty} \frac{\ln \left\{ \sup_\theta \mathbb{P} \left(\hat{R}_\theta^\pi(T) > cT \right) \right\}}{\ln T} \geq -C_\pi. \quad (3)$$

Theorem 1 suggests that a consistent policy must have a risk tail heavier than an exponential one. The proof of Theorem 1 relies on a change of measure argument appeared in [10]. We consider the environment where the noise ϵ is Gaussian with standard deviation σ . Intuitively speaking, if we want a policy to be adaptive enough to handle different instances, then the policy will be sensitive to risk. Moreover, if the policy achieves $O(\ln T)$ regret for any fixed instance θ (the constant is typically dependent on θ), then the probability of incurring a linear regret becomes $\exp(-O(\ln T)) = \Omega(\text{poly}(1/T))$. One special case is the family of confidence bound related policies (SE and UCB) where from Theorem 1 one can see that the standard bonus term (1) will always lead to a tail polynomially dependent on T . Another example is the Thompson Sampling (TS) policy. It has been established that $\pi = \text{TS}$ with Beta or Gaussian priors has the property (2) (see, e.g., Theorem 1 and 2 in [1], proof of Theorem 1.3 in [2]). Theorem 1 then suggests that (3) also holds for $\pi = \text{TS}$.

We need to remark on the difference between Theorem 1 and high-probability bounds in previous literature. It has been well-established that UCB or SE with

$$\text{rad}(n) = \sigma \sqrt{\frac{\eta \ln(1/\delta)}{n}}$$

achieves $\tilde{O}(\sqrt{T \cdot \text{polylog}(T/\delta)})$ regret with probability at least $1 - \delta$ (see, e.g., Section 1.3 in [17], Section 7.1 in [14]). Such design also leads to a consistent policy. However, the bound holds only for fixed δ . In fact, one can see that the bonus design is dependent on the confidence parameter δ . If $\delta = \exp(-\Omega(T^\beta))$ with $\beta > 0$, then the scaling speed of the expected regret with respect to T can only be greater than $1/2$, which is sub-optimal. As a comparison, in our problem, ideally we seek to find a single policy such that it achieves $\tilde{O}(\sqrt{T \cdot \text{polylog}(T/\delta)})$ regret for any $\delta > 0$.

Up till now, we make two core observations. First, from standard MAB results, consistency and optimality can hold simultaneously. Second, from Theorem 1, consistency and light-tailed risk

are always incompatible. Then a natural question arises: Can we design a policy that enjoys both optimality and light-tailed risk? If we can, then can we make the tail risk decay with T in an optimal rate? We answer these two questions with an affirmative “yes” in the next section.

3.2 A New Policy Design Achieving Light-tailed Risk and Worst-case Optimality

In this section, we propose a new policy design that achieves both light-tailed risk and worst-case optimality. The design is simple. We still use the idea of confidence bounds, but instead of setting the bonus as (1), we set

$$\text{rad}(n) = \sigma \frac{\sqrt{\eta T \ln T}}{n} \quad (4)$$

with $\eta > 0$ being a tuning parameter. Theorem 2 gives performance guarantees for the mean and the tail probability of the empirical regret when $\pi = \text{SE}$.

Theorem 2 *For the two-armed bandit problem, the SE policy with $\eta \geq 4$ and the bonus term being (4) satisfies the following two properties. 1. $\sup_{\theta} \mathbb{E}[\hat{R}_{\theta}^{\pi}] = O(\sqrt{T \ln T})$. 2. For any $c > 0$ and any $\alpha \in (1/2, 1]$, we have*

$$\sup_{\theta} \mathbb{P}(\hat{R}_{\theta}^{\pi}(T) \geq cT^{\alpha}) = \exp(-\Omega(T^{\alpha-1/2})).$$

The first item in Theorem 2 means that with the modified bonus term, the worst case regret is still bounded by $O(\sqrt{T \ln T})$, which is the same as the regret bounds for SE and UCB with the standard bonus term (1). The second item shows that the tail probability of incurring a $\Omega(T^{\alpha})$ regret ($\alpha > 1/2$) is exponentially decaying in $\Omega(T^{\alpha-1/2})$, and thus the policy is light-tailed. The detailed proof of Theorem 2 is provided in the supplementary material. An illustrative road-map of the proof is delegated to Section 4, where we provide the proof idea for a theorem that is a strict generalization of Theorem 2. Here, we give some intuition on the new bonus design. Our new bonus term inflates the standard one by a factor of $\sqrt{T/n}$. This means our policy is more conservative than the traditional confidence bound methods, especially at the beginning. In fact, one can observe that for the first $\Theta(\sqrt{T})$ time periods, our policy consistently explores between arm 1 and 2, regardless of the environment. A natural corollary is that our policy is never “consistent”, and this is reasonable following Theorem 1. However, the bonus term (4) decays at a faster rate on the number of pulling times n compared to (1). This means as the experiment goes on, the policy leans towards exploitation. We need to stress that this is not the same as the explore-then commit policy, which is well-known to achieve a sub-optimal $\Theta(T^{2/3})$ regret (see, e.g., Section 1.2 in [17], Section 6.1 in [14]).

The following theorem shows that the risk tail in Theorem 2 is not improvable. That is, if the policy π is worst-case optimal, then for fixed $\alpha \in (1/2, 1]$, the exponent of $\alpha - 1/2$ is tight.

Theorem 3 *Let $c \in (0, 1/2)$. Consider the 2-armed bandit problem with Gaussian noise. Let π be a worst-case optimal policy. That is, for any $\alpha > 1/2$,*

$$\limsup_T \frac{\sup_{\theta} \mathbb{E}[\hat{R}_{\theta}^{\pi}(T)]}{T^{\alpha}} = 0.$$

Then for any $\alpha \in (1/2, 1]$,

$$\liminf_T \frac{\ln \left\{ \sup_{\theta} \mathbb{P}(\hat{R}_{\theta}^{\pi}(T) \geq cT^{\alpha}) \right\}}{T^{\beta}} = 0$$

holds for any $\beta > \alpha - 1/2$.

Theorem 3 also relies on the change of measure argument appeared in Theorem 1. However, there are two important differences: we only have worst-case optimality rather than consistency, and the regret threshold cT^{α} is in general not linear in T . Therefore, we need to take care of constructing the specific “hard” instance when doing the change of measure. The detailed proof is delegated to the supplementary material.

4 Extensions: Multi-armed Bandit

In this section, we provide an extension of our previous tail probability bound in Theorem 2 from the following three aspects: (a) a general K -armed bandit setting; (b) an analysis for UCB aside from SE; (c) a detailed characterization of the tail bound for any fixed regret threshold.

Theorem 4 *For the K -armed bandit problem, $\pi = \text{SE}$ or $\pi = \text{UCB}$ with*

$$\text{rad}(n) = \sigma \frac{\sqrt{\eta T \ln T}}{n}$$

satisfy the following two properties. 1. If $\eta \geq 4$, then $\sup_{\theta} \mathbb{E}[R_{\theta}^{\pi}] \leq 4K + 4K\sigma\sqrt{\eta T \ln T}$. 2. If $\eta > 0$, then for any $x > 0$, we have

$$\begin{aligned} & \sup_{\theta} \mathbb{P}(R_{\theta}^{\pi}(T) \geq x) \\ & \leq \exp\left(-\frac{x^2}{2K\sigma^2 T}\right) + 2K \exp\left(-\frac{(x - 2K - 4K\sigma\sqrt{\eta T \ln T})_+^2}{32\sigma^2 K^2 T}\right) + K^2 T \exp\left(-\frac{x\sqrt{\eta \ln T}}{8\sigma K \sqrt{T}}\right). \end{aligned}$$

Proof Idea. We provide a road-map of proving Theorem 4. The expected regret bound is proved using standard techniques. That is, we define the good event to be such that the mean of each arm always lies in the confidence bounds throughout the whole time horizon. Conditioned on the good event, the regret of each arm is bounded by $O(\sqrt{T \ln T})$, and thus the total expected regret is $O(K\sqrt{T \ln T})$.

The proof of the tail bound requires more effort. Without loss of generality, we assume arm 1 is optimal. We first illustrate the proof for $\pi = \text{SE}$.

Step 1. We use

$$\sup_{\theta} \mathbb{P}(\hat{R}_{\theta}^{\pi}(T) \geq x) \leq \mathbb{P}\left(N^{\pi}(T) \leq -x/\sqrt{K}\right) + \sup_{\theta} \mathbb{P}\left(R_{\theta}^{\pi}(T) \geq x(1 - 1/\sqrt{K})\right)$$

The term with the genuine noise can be easily bounded using Lemma 1. We are left to bound the tail of the pseudo regret. By a union bound, we observe that

$$\mathbb{P}\left(R_{\theta}^{\pi}(T) \geq x(1 - 1/\sqrt{K})\right) \leq \sum_{k \neq 1} \mathbb{P}\left(n_k \Delta_k \geq x/(K + \sqrt{K})\right) \leq \sum_{k \neq 1} \mathbb{P}(n_k \Delta_k \geq x/(2K))$$

Thus, we reduce bounding the sum of the regret incurred by different arms to bounding that by a single sub-optimal arm.

Step 2. For any $k \neq 1$, we define

$$S_k = \{\text{Arm 1 is not eliminated before arm } k\}.$$

With a slight abuse of notation, we let $n_0 = \lceil x/(2K\Delta_k) \rceil - 1$. If both $n_k \Delta_k \geq x/(2K)$ and S_k happen, then arm 1 and k are not eliminated after each of them being pulled n_0 times. This indicates

$$\hat{\mu}_{t_1(n_0),1} - \frac{\sigma\sqrt{\eta T \ln T}}{n_0} \leq \hat{\mu}_{t_k(n_0),k} + \frac{\sigma\sqrt{\eta T \ln T}}{n_0}$$

The probability of this event can be bounded using concentration of subgaussian variables, which yields the second term in the tail bound in Theorem 4. The choice of n_0 is important. We note that at this step, even if we replace our new bonus term by the standard one, the bound still holds.

Step 3. Now consider the situation when \bar{S}_k happens. This means after some phase n , the optimal arm 1 is eliminated by some arm k' , while arm k is not eliminated. Note that $k = k'$ does not necessarily hold when $K > 2$. As a consequence, we have the following two events hold simultaneously:

$$\hat{\mu}_{t_{k'}(n),k'} - \frac{\sigma\sqrt{\eta T \ln T}}{n} \geq \hat{\mu}_{t_1(n),1} + \frac{\sigma\sqrt{\eta T \ln T}}{n}, \quad \hat{\mu}_{t_k(n),k} + \frac{\sigma\sqrt{\eta T \ln T}}{n} \geq \hat{\mu}_{t_1(n),1} + \frac{\sigma\sqrt{\eta T \ln T}}{n}.$$

The first event leads to

$$\text{Mean of some noise terms} \geq \frac{2\sigma\sqrt{\eta T \ln T}}{n} + \Delta_{k'} \geq \frac{2\sigma\sqrt{\eta T \ln T}}{n}.$$

The second event leads to

$$\text{Mean of some noise terms} \geq \Delta_k \geq \frac{x}{2KT}.$$

Now comes the trick to deal with an arbitrary n . We bound the probabilities of the two events separately and take the minimum of the two probabilities ($\mathbb{P}(A \cap B) \leq \min\{\mathbb{P}(A), \mathbb{P}(B)\}$ ($\forall A, B$)). Then such minimum can be further bounded using the basic inequality $(a+b)/2 \geq \sqrt{ab}$ ($\forall a, b \geq 0$). This yields the last term in Theorem 4. We note that at this step, the standard bonus term (1) does not suffice to get an exponential bound. The \sqrt{T}/n design in our new bonus term plays a crucial role.

We next illustrate the proof for $\pi = \text{UCB}$. The proof is in fact simpler. We use the first step in the proof for $\pi = \text{SE}$. For fixed k , we also take the same $n_0 = \lfloor x/(2K\Delta_k) \rfloor - 1$. The difference here is that we do not need to define the event S_k . When arm k is pulled for the $(n_0 + 1)$ th time, by the design of the UCB policy, there exists some n such that

$$\mu_1 + \frac{\sum_{m=1}^n \epsilon_{t_1(m),1} + \sigma\sqrt{\eta T \ln T}}{n} \leq \mu_k + \frac{\sum_{m=1}^{n_0} \epsilon_{t_k(m),k} + \sigma\sqrt{\eta T \ln T}}{n_0}.$$

Now comes the trick. The event is included by a union of two events described as follows:

$$\frac{\sum_{m=1}^{n_0} \epsilon_{t_k(m),k} + \sigma\sqrt{\eta T \ln T}}{n_0} \geq \frac{\Delta_k}{2} \quad \text{and} \quad \frac{\sum_{m=1}^n \epsilon_{t_1(m),1} + \sigma\sqrt{\eta T \ln T}}{n} \leq -\frac{\Delta_k}{2}.$$

The probability of each of the two events can be bounded using similar techniques adopted when $\pi = \text{SE}$. In fact, it is implicitly shown in our proof that UCB can yield better constants than SE. We still need to emphasize that when bounding the second event, similar to the argument for $\pi = \text{SE}$, the choice of our new bonus term is crucial.

Remarks. For the regret bound, we note that compared to the optimal $\tilde{\Theta}(\sqrt{KT})$ bound, we have an additional \sqrt{K} term. We should point out that the additional \sqrt{K} term is not surprising under the bonus term (4). An intuitive explanation is as follows. Compared to the bonus term (1), we widen the bonus term by a factor of \sqrt{T}/n . Among the $K - 1$ arms, there must exist an arm such that it is pulled no more than T/K times throughout the whole time horizon. That is, the bonus term of this arm is always inflated by a factor of at least \sqrt{K} . The standard regret bound analysis will, as a result, lead to an additional \sqrt{K} factor compared to the optimal regret bound $\tilde{\Theta}(\sqrt{KT})$.

For the tail bound, from our proof road-map, one can see that the tail bound in Theorem 4 is also valid for the pseudo-regret $\sup_{\theta} \mathbb{P}(R_{\theta}^{\pi}(T) \geq x)$. In fact, with some simple manipulations on the inequality, if we let

$$y = \frac{(x - 2K - 16\sigma K \sqrt{(\eta \vee 1/\eta)T \ln T})_+}{8\sigma K \sqrt{T}},$$

then for any $x \geq 0$, we can get a neat form (we illustrate this point in the supplementary material)

$$\sup_{\theta} \mathbb{P}(R_{\theta}^{\pi}(T) \geq x) \leq 4K \exp\left(-y^2 \wedge y\sqrt{\eta \ln T}\right).$$

There are two observations:

- (a) For any $\eta > 0$, our policy always yields a $O(\sqrt{T})$ expected regret (although with a constant larger than that in the first result in Theorem 4). In fact, notice that for any $x > 0$

$$\mathbb{E}[\hat{R}_{\theta}^{\pi}(T)] = \mathbb{E}[R_{\theta}^{\pi}(T)] \leq x + \mathbb{P}(R_{\theta}^{\pi}(T) \geq x) \cdot T.$$

If we let $x = 2K + C\sigma K \sqrt{(\eta \vee 1/\eta)T \ln T}$ with $C > 0$ being moderately large, then $\mathbb{P}(R_{\theta}^{\pi}(T) \geq x) \cdot T = O(1)$. As a result, the worst-case regret becomes

$$O\left(K \sqrt{(\eta \vee 1/\eta)T \ln T}\right).$$

This observation shows that the regret order of our policy design is not related to the parameter η , as opposed to the standard UCB policy with (1), where a very small η may possibly make the UCB policy no longer enjoy a $\tilde{O}(\sqrt{T})$ worst-case regret¹.

¹We believe there is work in the literature that has precisely documented that a very small η may possibly make the UCB policy no longer enjoy a $\tilde{O}(\sqrt{T})$ worst-case regret, but we have not been able to identify one. For completeness, we summarize this point as a lemma with a proof in the supplementary material.

(b) If we set $\eta = 1$ and

$$\delta = 4K \exp\left(-(y - \sqrt{\ln T})_+ \sqrt{\ln T}\right) \geq 4K \exp\left(-y^2 \wedge y\sqrt{\ln T}\right),$$

then one can see that for *any* $\delta > 0$, with probability at least $1 - \delta$, the regret of our policy is bounded by

$$O\left(\sigma K \sqrt{T} \left(\sqrt{\ln T} + \frac{\ln(4K/\delta)}{\sqrt{\ln T}}\right)\right) = O\left(\sigma K \sqrt{T} \left(\sqrt{\ln T} + \frac{\ln(1/\delta)}{\sqrt{\ln T}}\right)\right).$$

This partially solve the open question in Section 17.1 of [14] (we've mentioned it in the introduction) up to a logarithmic factor.

5 Numerical Experiments

In this section, we provide a brief two-armed numerical experiment result and defer the experiment details and other multi-armed experiment results to the supplementary material. We consider a two armed-bandit problem with $\theta = (0.2, 0.8)$, $\sigma = 1$, $T = 500$ and Gaussian noise. We test four policies: (i) SE and UCB with the classical bonus design described in (1), and SE_new and UCB_new with the proposed new bonus design in (4). We let $\kappa \triangleq \sigma\sqrt{\eta}$. The tuning parameter κ has 4 choices: $\kappa \in \{0.1, 0.2, 0.4, 0.8\}$. We also consider the TS policy assuming the mean reward of each arm i following the prior $\mathcal{N}(0, 1)$ and the sample from each arm i following $\mathcal{N}(\theta_i, \kappa^2)$. That said, we evaluate TS under mis-specified risk parameters. For each policy and κ , we run 5000 simulation paths and for each path we collect the cumulative reward. We plot the empirical distribution (histogram) for a policy's cumulative reward in Figure 1. Indeed, one can observe that for SE and UCB with (1) and TS, there is a significant part of distribution around 100, indicating a significant risk of incurring a linear regret, especially when κ is small. In contrast, with the new design (4), the reward is highly concentrated for every $\kappa > 0$ with almost no tail risk of getting a low total reward. Particularly, when $\kappa = 0.1$, UCB_new achieves both high empirical mean and light-tailed distribution.




Figure 1: Empirical distribution for the cumulative reward; bottom two are new proposed policies

6 Conclusion

In this work, we consider the MAB problem with a joint goal of minimizing the worst-case expected regret and obtaining light-tailed probability bound of the regret distribution. We characterize the interplay among three concepts: light-tailed risk, instance-dependent consistency, and worst-case optimality. We demonstrate that light-tailed risk and instance-dependent consistency are incompatible, and show that light-tailed risk and worst-case optimality can co-exist through a simple new policy design. We discuss insights and generalizations of our results, and study the empirical performance.

An open question is whether we can improve the regret bound in Theorem 4 to $\Theta(\sqrt{KT \ln T})$ and get a probability bound of

$$\ln \left\{ \sup_{\theta} \mathbb{P}(R_{\theta}^{\pi}(T) \geq x) \right\} = -\Omega \left(\frac{x}{\sqrt{KT}} \right). \quad (5)$$

Another interesting direction is to generalize our policy design to the setting where noises are not sub-Gaussian, but come from general distributions. In fact, there is a separate line of literature studying bandit problems with non sub-Gaussian noise (see, e.g., [7]). We hope our framework can be integrated with this line of literature to handle certain types of non sub-Gaussian noise, provided with controls on the tail behavior of the random noises.

References

- [1] S. Agrawal and N. Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on learning theory*, pages 39–1. JMLR Workshop and Conference Proceedings, 2012.
- [2] S. Agrawal and N. Goyal. Near-optimal regret bounds for thompson sampling. *Journal of the ACM (JACM)*, 64(5):1–24, 2017.
- [3] V. F. Araman and R. Caldentey. Diffusion approximations for a class of sequential testing problems. *arXiv preprint arXiv:2102.07030*, 2021.
- [4] K. Ashutosh, J. Nair, A. Kagracha, and K. Jagannathan. Bandit algorithms: Letting go of logarithmic regret for statistical robustness. In *International Conference on Artificial Intelligence and Statistics*, pages 622–630. PMLR, 2021.
- [5] J.-Y. Audibert, R. Munos, and C. Szepesvári. Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):1876–1902, 2009.
- [6] S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *arXiv preprint arXiv:1204.5721*, 2012.
- [7] S. Bubeck, N. Cesa-Bianchi, and G. Lugosi. Bandits with heavy tail. *IEEE Transactions on Information Theory*, 59(11):7711–7717, 2013.
- [8] A. Cassel, S. Mannor, and A. Zeevi. A general approach to multi-armed bandits under risk criteria. In *Conference On Learning Theory*, pages 1295–1306. PMLR, 2018.
- [9] L. Fan and P. W. Glynn. Diffusion approximations for thompson sampling. *arXiv preprint arXiv:2105.09232*, 2021.
- [10] L. Fan and P. W. Glynn. The fragility of optimized bandit algorithms. *arXiv preprint arXiv:2109.13595*, 2021.
- [11] N. Galichet, M. Sebag, and O. Teytaud. Exploration vs exploitation vs safety: Risk-aware multi-armed bandits. In *Asian Conference on Machine Learning*, pages 245–260. PMLR, 2013.
- [12] A. Kalvit and A. Zeevi. A closer look at the worst-case behavior of multi-armed bandit algorithms. *Advances in Neural Information Processing Systems*, 34, 2021.
- [13] N. Khajonchotpanya, Y. Xue, and N. Rujeerapaiboon. A revised approach for risk-averse multi-armed bandits under cvar criterion. *Operations Research Letters*, 49(4):465–472, 2021.
- [14] T. Lattimore and C. Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- [15] A. Salomon and J.-Y. Audibert. Deviations of stochastic bandit regret. In *International Conference on Algorithmic Learning Theory*, pages 159–173. Springer, 2011.
- [16] A. Sani, A. Lazaric, and R. Munos. Risk-aversion in multi-armed bandits. *Advances in Neural Information Processing Systems*, 25, 2012.
- [17] A. Slivkins. Introduction to multi-armed bandits. *arXiv preprint arXiv:1904.07272*, 2019.

- [18] S. Vakili and Q. Zhao. Risk-averse multi-armed bandit problems under mean-variance measure. *IEEE Journal of Selected Topics in Signal Processing*, 10(6):1093–1111, 2016.
- [19] S. Wager and K. Xu. Diffusion asymptotics for sequential experiments. *arXiv preprint arXiv:2101.09855*, 2021.
- [20] A. Zimin, R. Ibsen-Jensen, and K. Chatterjee. Generalized risk-aversion in stochastic multi-armed bandits. *arXiv preprint arXiv:1405.0833*, 2014.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#)
 - (c) Did you discuss any potential negative societal impacts of your work? [\[N/A\]](#)
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[Yes\]](#)
 - (b) Did you include complete proofs of all theoretical results? [\[Yes\]](#)
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[Yes\]](#)
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[N/A\]](#)
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[N/A\]](#) We actually collect rewards and draw the empirical distribution.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[N/A\]](#)
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [\[N/A\]](#)
 - (b) Did you mention the license of the assets? [\[N/A\]](#)
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[N/A\]](#)
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [\[N/A\]](#)
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[N/A\]](#)
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [\[N/A\]](#)
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [\[N/A\]](#)
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [\[N/A\]](#)

A Proofs for Section 3

To prove Theorem 1, we need the following lemma.

Lemma 2 *Consider the two-armed bandit problem with σ -Gaussian noise. Let π be a policy such that for any true mean vector θ ,*

$$\limsup_{T \rightarrow +\infty} \frac{\mathbb{E} [\hat{R}_\theta^\pi(T)]}{T} = 0.$$

That is, the expected regret under π is always sub-linear in T . Then for any $\tilde{\theta} = (\tilde{\theta}_1, \tilde{\theta}_2)$ where $\tilde{\theta}_1 > \tilde{\theta}_2$, and any $\delta > 0$, we have

$$\limsup_{T \rightarrow +\infty} \mathbb{P}_\theta^\pi(|\hat{\mu}_{T,2} - \tilde{\theta}_2| > \delta) = 0.$$

Proof of Lemma 2. Define

$$E_T = \left\{ |\hat{\mu}_{T,2} - \tilde{\theta}_2| \leq \delta \right\}.$$

Fix any positive integer N , we have

$$\begin{aligned} \mathbb{P}_\theta^\pi(\bar{E}_T) &= \mathbb{P}_\theta^\pi(\bar{E}_T; n_{T,2} < N) + \mathbb{P}_\theta^\pi(\bar{E}_T; n_{T,2} \geq N) \\ &\leq \mathbb{P}_\theta^\pi(n_{T,2} < N) + \sum_{n=N}^{+\infty} \mathbb{P}_\theta^\pi(\bar{E}_T; n_{T,2} = n) \\ &\leq \mathbb{P}_\theta^\pi(n_{T,2} < N) + \sum_{n=N}^{+\infty} 2 \exp\left(-\frac{n\delta^2}{2\sigma^2}\right). \end{aligned}$$

Thus,

$$\limsup_T \mathbb{P}_\theta^\pi(\bar{E}_T) \leq \limsup_T \mathbb{P}_\theta^\pi(n_{T,2} < N) + \sum_{n=N}^{+\infty} 2 \exp\left(-\frac{n\delta^2}{2\sigma^2}\right)$$

holds for any N . Note that the last term converges to 0 as $N \rightarrow +\infty$. It suffices to show $\mathbb{P}_\theta^\pi(n_{T,2} < N) \rightarrow 0$ as $T \rightarrow +\infty$ for any fixed N . Suppose this does not hold, then we can find $p > 0$ and a sequence $\{T(m)\}_{m=1}^{+\infty}$ such that

$$\mathbb{P}_\theta^\pi(n_{T(m),2} < N) > p.$$

Let M be some large number such that $q \triangleq p - N \exp(-\frac{M^2}{2\sigma^2}) > 0$. Consider an alternative environment $\theta = (\theta_1, \theta_2)$ where $\theta_2 > \theta_1 = \tilde{\theta}_1$. Using the change of measure argument, we have

$$\begin{aligned} &\mathbb{P}_\theta^\pi(n_{T(m),2} < N) \\ &= \mathbb{E}_\theta^\pi[1\{n_{T(m),2} < N\}] \\ &= \mathbb{E}_{\tilde{\theta}}^\pi \left[\exp \left(\sum_{n=1}^{n_{T(m),2}} \frac{(X_{t_2(n),2} - \tilde{\theta}_2)^2 - (X_{t_2(n),2} - \theta_2)^2}{2\sigma^2} \right) 1\{n_{T(m),2} < N\} \right] \\ &= \mathbb{E}_{\tilde{\theta}}^\pi \left[\exp \left(n_{T(m),2} \left(\frac{\tilde{\theta}_2^2 - \theta_2^2}{2\sigma^2} + \frac{(\theta_2 - \tilde{\theta}_2)\hat{\theta}_{T(m),2}}{\sigma^2} \right) \right) 1\{n_{T(m),2} < N\} \right] \\ &\geq \mathbb{E}_{\tilde{\theta}}^\pi \left[\exp \left(n_{T(m),2} \left(\frac{\tilde{\theta}_2^2 - \theta_2^2}{2\sigma^2} + \frac{(\theta_2 - \tilde{\theta}_2)\hat{\theta}_{T(m),2}}{\sigma^2} \right) \right) 1\{\hat{\theta}_{T(m),2} > \tilde{\theta}_2 - M, n_{T(m),2} < N\} \right] \\ &\geq \mathbb{E}_{\tilde{\theta}}^\pi \left[\exp \left(N \left(-\frac{(\tilde{\theta}_2 - \theta_2)^2}{2\sigma^2} - \frac{M(\theta_2 - \tilde{\theta}_2)}{\sigma^2} \right) \right) 1\{\hat{\theta}_{T(m),2} > \tilde{\theta}_2 - M, n_{T(m),2} < N\} \right] \\ &= \exp \left(N \left(-\frac{(\tilde{\theta}_2 - \theta_2)^2}{2\sigma^2} - \frac{M(\theta_2 - \tilde{\theta}_2)}{\sigma^2} \right) \right) \mathbb{P}_{\tilde{\theta}}^\pi(\hat{\theta}_{T(m),2} > \tilde{\theta}_2 - M, n_{T(m),2} < N). \end{aligned}$$

Note that

$$\begin{aligned}
& \mathbb{P}_{\tilde{\theta}}^{\pi}(\hat{\theta}_{T(m),2} > \tilde{\theta}_2 - M, n_{T(m),2} < N) \\
&= p - \sum_{n=1}^{N-1} \mathbb{P}_{\tilde{\theta}}^{\pi}(\hat{\theta}_{T(m),2} \leq \tilde{\theta}_2 - M, n_{T(m),2} = n) \\
&\geq p - \sum_{n=1}^{N-1} \exp(-\frac{nM^2}{2\sigma^2}) \geq p - N \exp(-\frac{M^2}{2\sigma^2}) = q > 0.
\end{aligned}$$

Therefore, there exists a constant positive probability such that π pulls arm 2 no more than N times under θ . As a result, π incurs a linear expected regret under θ , leading to a contradiction. \square

Proof of Theorem 1.

We consider the environment where the noise ϵ is gaussian with standard deviation σ . Let $\theta_1 = 1/2$. Let $\theta = (\theta_1, \theta_2)$ and $\tilde{\theta} = (\theta_1, \tilde{\theta}_2)$, where $\theta_2 = \theta_1 + \frac{1}{2}$ and $\tilde{\theta}_2 = \theta_1 - \frac{1}{2}$. Let $c' \in (c, 1/2)$. Define

$$E_T = \left\{ |\hat{\mu}_{T,2} - \tilde{\theta}_2| \leq \delta \right\}$$

where $\delta > 0$ is a small number, and

$$F_T = \{n_2 \leq f(T)\}.$$

Here, $f(T) > 0$ is a strictly increasing function such that

$$\limsup_T \frac{f(T)}{T} = 0.$$

We will detail how $f(T)$ should be chosen under different conditions in the last step of the proof. Then there exists T_0 such that $f(T) < (1 - 2c')T$ for any $T > T_0$. Now we fix any $T > T_0$. Under the environment $\tilde{\theta}$, we have

$$\mathbb{P}_{\tilde{\theta}}^{\pi}(\bar{F}_T) = \mathbb{P}_{\tilde{\theta}}^{\pi}(n_2 > f(T)) \leq \frac{\mathbb{E}_{\tilde{\theta}}^{\pi}[n_2]}{f(T)} \leq \frac{2\mathbb{E}[R_{\tilde{\theta}}^{\pi}(T)]}{f(T)} = \frac{2\mathbb{E}[\hat{R}_{\tilde{\theta}}^{\pi}(T)]}{f(T)}.$$

Combined with Lemma 2, we have

$$\liminf_T \mathbb{P}_{\tilde{\theta}}^{\pi}(E_T, F_T) \geq 1 - \limsup_T \frac{2\mathbb{E}[\hat{R}_{\tilde{\theta}}^{\pi}(T)]}{f(T)}. \quad (6)$$

Notice that

$$\begin{aligned}
& \mathbb{P}\left(\hat{R}_{\tilde{\theta}}^{\pi}(T) \geq cT\right) \\
&\geq \mathbb{P}\left(R_{\tilde{\theta}}^{\pi}(T) \geq c'T, -N^{\pi}(T) \geq -(c' - c)T\right) \\
&= \mathbb{P}\left(R_{\tilde{\theta}}^{\pi}(T) \geq c'T\right) - \mathbb{P}\left(R_{\tilde{\theta}}^{\pi}(T) \geq c'T, N^{\pi}(T) > (c' - c)T\right) \\
&\geq \mathbb{P}\left(R_{\tilde{\theta}}^{\pi}(T) \geq c'T\right) - \mathbb{P}\left(N^{\pi}(T) > (c' - c)T\right) \\
&\geq \mathbb{P}\left(R_{\tilde{\theta}}^{\pi}(T) \geq c'T\right) - \exp\left(-\frac{(c' - c)^2 T}{2\sigma^2}\right)
\end{aligned}$$

The last inequality holds from Lemma 1. Now

$$\begin{aligned}
& \mathbb{P}\left(R_{\tilde{\theta}}^{\pi}(T) \geq c'T\right) \\
&\geq \mathbb{P}_{\tilde{\theta}}^{\pi}(n_1 \geq 2c'T) \\
&\geq \mathbb{P}_{\tilde{\theta}}^{\pi}(n_2 \leq (1 - 2c')T) \\
&\geq \mathbb{P}_{\tilde{\theta}}^{\pi}(n_2 \leq f(T)) \\
&\geq \mathbb{P}_{\tilde{\theta}}^{\pi}(E_T, F_T) \\
&= \mathbb{E}_{\tilde{\theta}}^{\pi}[\mathbf{1}\{E_T F_T\}]
\end{aligned}$$

$$\begin{aligned}
&= \mathbb{E}_{\tilde{\theta}}^{\pi} \left[\exp \left(\sum_{n=1}^{n_2} \frac{(X_{t_2(n),2} - \tilde{\theta}_2)^2 - (X_{t_2(n),2} - \theta_2)^2}{2\sigma^2} \right) \mathbb{1}\{E_T F_T\} \right] \\
&= \mathbb{E}_{\tilde{\theta}}^{\pi} \left[\exp \left(n_2 \left(\frac{\tilde{\theta}_2^2 - \theta_2^2}{2\sigma^2} + \frac{(\theta_2 - \tilde{\theta}_2)\hat{\theta}_{T,2}}{\sigma^2} \right) \right) \mathbb{1}\{E_T F_T\} \right] \\
&\geq \mathbb{E}_{\tilde{\theta}}^{\pi} \left[\exp \left(n_2 \left(\frac{\tilde{\theta}_2^2 - \theta_2^2}{2\sigma^2} + \frac{(\theta_2 - \tilde{\theta}_2)(\tilde{\theta}_2 - \delta)}{\sigma^2} \right) \right) \mathbb{1}\{E_T F_T\} \right] \\
&= \mathbb{E}_{\tilde{\theta}}^{\pi} \left[\exp \left(n_2 \left(-\frac{(\tilde{\theta}_2 - \theta_2)^2}{2\sigma^2} - \frac{\delta(\theta_2 - \tilde{\theta}_2)}{\sigma^2} \right) \right) \mathbb{1}\{E_T F_T\} \right] \\
&\geq \mathbb{E}_{\tilde{\theta}}^{\pi} \left[\exp \left(f(T) \left(-\frac{(\tilde{\theta}_2 - \theta_2)^2}{2\sigma^2} - \frac{\delta(\theta_2 - \tilde{\theta}_2)}{\sigma^2} \right) \right) \mathbb{1}\{E_T F_T\} \right] \\
&= \exp(-f(T)(1/2\sigma^2 + \delta/\sigma^2)) \mathbb{P}_{\tilde{\theta}}^{\pi}(E_T, F_T).
\end{aligned}$$

Therefore,

$$\begin{aligned}
&\liminf_T \frac{\ln \left\{ \sup_{\theta'} \mathbb{P} \left(\hat{R}_{\theta'}^{\pi}(T) \geq cT \right) \right\}}{f(T)} \\
&\geq \liminf_T \frac{\ln \left\{ \exp(-f(T)(1/2\sigma^2 + \delta/\sigma^2)) \mathbb{P}_{\tilde{\theta}}^{\pi}(E_T, F_T) - \exp \left(-\frac{(c'-c)^2 T}{2\sigma^2} \right) \right\}}{f(T)}. \tag{7}
\end{aligned}$$

Now assume that π is consistent. Then we set $f(T) = T^{\beta}$ with $\beta \in (0, 1)$. From (6), we have

$$\liminf_T \mathbb{P}_{\tilde{\theta}}^{\pi}(E_T, F_T) \geq 1 - \limsup_T \frac{2\mathbb{E}[\hat{R}_{\tilde{\theta}}^{\pi}(T)]}{T^{\beta}} = 1.$$

Then from (7), we have

$$\liminf_T \frac{\ln \left\{ \sup_{\theta'} \mathbb{P}(\hat{R}_{\theta'}^{\pi}(T) \geq cT) \right\}}{T^{\beta}} \geq -(1/2\sigma^2 + \delta/\sigma^2).$$

Since $\delta > 0$ is arbitrary, we have

$$\liminf_T \frac{\ln \left\{ \sup_{\theta'} \mathbb{P}(\hat{R}_{\theta'}^{\pi}(T) \geq cT) \right\}}{T^{\beta}} \geq -1/2\sigma^2.$$

Note again that $\beta > 0$ is arbitrary. Now let $0 < \beta' < \beta$, we have

$$\begin{aligned}
\liminf_T \frac{\ln \left\{ \sup_{\theta'} \mathbb{P}(\hat{R}_{\theta'}^{\pi}(T) \geq cT) \right\}}{T^{\beta}} &= \liminf_T \frac{\ln \left\{ \sup_{\theta'} \mathbb{P}(\hat{R}_{\theta'}^{\pi}(T) \geq cT) \right\}}{T^{\beta'}} \cdot \liminf_T T^{\beta' - \beta} \\
&\geq -1/2\sigma^2 \cdot 0 = 0.
\end{aligned}$$

Now assume that π satisfies

$$\limsup_T \frac{\mathbb{E}[\hat{R}_{\tilde{\theta}}^{\pi}(T)]}{\ln T} = c_{\pi}\sigma^2 < +\infty.$$

for any θ . Note that when $\pi = \text{SE}$ and $\pi = \text{UCB}$ with the bonus term (1), the property above always holds. Let $f(T) = 4c_{\pi}\sigma^2 \ln T$. From (6), we have

$$\liminf_T \mathbb{P}_{\tilde{\theta}}^{\pi}(E_T, F_T) \geq 1 - \limsup_T \frac{2\mathbb{E}[\hat{R}_{\tilde{\theta}}^{\pi}(T)]}{4c_{\pi}\sigma^2 \ln T} = 1/2.$$

Then from (7), we have

$$\liminf_T \frac{\ln \left\{ \sup_{\theta'} \mathbb{P}(\hat{R}_{\theta'}^{\pi}(T) \geq cT) \right\}}{4c_{\pi}\sigma^2 \ln T} \geq -(1/2\sigma^2 + \delta/\sigma^2).$$

Since $\delta > 0$ is arbitrary, we have

$$\liminf_T \frac{\ln \left\{ \sup_{\theta'} \mathbb{P}(\hat{R}_{\theta'}^\pi(T) \geq cT) \right\}}{4c_\pi \sigma^2 \ln T} \geq -1/2\sigma^2.$$

Let $C_\pi = 2c_\pi$, we have

$$\liminf_T \frac{\ln \left\{ \sup_{\theta'} \mathbb{P}(\hat{R}_{\theta'}^\pi(T) \geq cT) \right\}}{\ln T} \geq -C_\pi.$$

□

Proof of Theorem 2. Without loss of generality, we assume $\theta_1 > \theta_2$. We prove the results one by one. Since the environment θ is fixed, we will write $\mathbb{P}(\mathbb{E})$ instead of $\mathbb{P}_\theta^\pi(\mathbb{E}_\theta^\pi)$.

1. From Lemma 1,

$$\mathbb{E}[\hat{R}_\theta^\pi(T)] = \mathbb{E}[R_\theta^\pi(T)] = \mathbb{E}[n_2] \cdot \Delta_2.$$

Let G be the event such that

$$G = \{\mu_k \in \text{CI}_{t,k}, \forall (t, k)\}.$$

Then

$$\mathbb{P}(\bar{G}) \leq \sum_{(t,k)} \mathbb{P}(\mu_k \notin \text{CI}_{t,k}) \leq 2 \sum_{n=1}^T 2 \exp(-2 \frac{\eta T \ln T}{n}) \leq 4T^{1-2\eta}.$$

Thus,

$$\begin{aligned} \mathbb{E}[n_2] &= \mathbb{E}[n_2|G]\mathbb{P}(G) + \mathbb{E}[n_2|\bar{G}]\mathbb{P}(\bar{G}) \\ &\leq \mathbb{E}[n_2|G] + T \cdot \mathbb{P}(\bar{G}) \\ &\leq \mathbb{E}[n_2|G] + 4T^{2-2\eta} \leq \mathbb{E}[n_2|G] + 4. \end{aligned}$$

With a slight abuse of notation, we let t be the largest time period such that arm 2 is pulled but subsequently not eliminated from \mathcal{A} . Then under G , we have

$$\begin{aligned} \mu_1 - 2\sigma \frac{\sqrt{\eta T \ln T}}{n_{t,2} - 1} &\leq \mu_1 - 2\sigma \frac{\sqrt{\eta T \ln T}}{n_{t,1}} \leq \hat{\mu}_{t,1} - \sigma \frac{\sqrt{\eta T \ln T}}{n_{t,1}} \\ &\leq \hat{\mu}_{t,2} + \sigma \frac{\sqrt{\eta T \ln T}}{n_{t,2} - 1} \leq \mu_2 + 2\sigma \frac{\sqrt{\eta T \ln T}}{n_{t,2} - 1}. \end{aligned}$$

Therefore,

$$n_{t,2} \leq 1 + 4\sigma \frac{\sqrt{\eta T \ln T}}{\Delta_2}$$

and thus,

$$n_2 \leq 2 + 4\sigma \frac{\sqrt{\eta T \ln T}}{\Delta_2}.$$

As a result,

$$\mathbb{E}[R_\theta^\pi(T)] \leq 2\Delta_2 + 4\sigma \sqrt{\eta T \ln T} + 4 = O(\sqrt{T \ln T}).$$

2. We have

$$\mathbb{P}(R_\theta^\pi(T) \geq cT^\alpha) \leq \mathbb{P}(\hat{R}_\theta^\pi(T) \geq cT^\alpha/2) + \mathbb{P}(N_\theta^\pi(T) \leq -cT^\alpha/2)$$

From Lemma 1, the second term can be bounded as

$$\mathbb{P}(N_\theta^\pi(T) \leq -cT^\alpha/2) \leq \exp\left(-\frac{c^2 T^{2\alpha}}{2\sigma^2 T}\right) = \exp\left(-\frac{c^2 T^{2\alpha-1}}{2\sigma^2}\right). \quad (8)$$

We are left to bound $\mathbb{P}(\hat{R}_\theta^\pi(T) \geq cT^\alpha/2)$. Let S be the event defined as

$$S = \{\text{Arm 1 is never eliminated throughout the whole time horizon}\}.$$

Then

$$\bar{S} = \{\exists t \text{ such that arm 1 is eliminated at time } t\}.$$

So

$$\mathbb{P}(\hat{R}_\theta^\pi(T) \geq cT^\alpha/2) = \mathbb{P}(\hat{R}_\theta^\pi(T) \geq cT^\alpha/2, S) + \mathbb{P}(\hat{R}_\theta^\pi(T) \geq cT^\alpha/2, \bar{S}).$$

Let T be such that

$$cT^\alpha \geq \max\{4, 16\sigma\sqrt{\eta T \ln T}\}.$$

Let $n_0 = \lceil cT^\alpha/2\Delta_2 \rceil - 1$, then

$$n_0 \geq cT^\alpha/4\Delta_2.$$

Also, if $\hat{R}_\theta^\pi(T) \geq cT^\alpha/2$, we must have

$$T \geq cT^\alpha/2\Delta_2,$$

which means $\Delta_2 \geq cT^{\alpha-1}/2$. We have

$$\begin{aligned} & \mathbb{P}(\hat{R}_\theta^\pi(T) \geq cT^\alpha/2, S) \\ &= \mathbb{P}(n_2 \geq cT^\alpha/2\Delta_2, S) \\ &\leq \mathbb{P}(n_2 \geq n_0 + 1, \text{ arm 1 and 2 are pulled in turn for } (n_0 + 1) \text{ times}) \\ &\leq \mathbb{P}(\text{arm 1 and 2 are pulled in turn for } n_0 \text{ times and arm 1 and 2 are both not eliminated}) \\ &\leq \mathbb{P}\left(\hat{\mu}_{t_1(n_0),1} - \frac{\sigma\sqrt{\eta T \ln T}}{n_0} \leq \hat{\mu}_{t_2(n_0),2} + \frac{\sigma\sqrt{\eta T \ln T}}{n_0}\right) \\ &= \mathbb{P}\left(\mu_1 - \frac{\sum_{m=1}^{n_0} \epsilon_{t_1(m),1} + \sigma\sqrt{\eta T \ln T}}{n_0} \leq \mu_2 + \frac{\sum_{m=1}^{n_0} \epsilon_{t_2(m),2} + \sigma\sqrt{\eta T \ln T}}{n_0}\right) \\ &= \mathbb{P}\left(\frac{\sum_{m=1}^{n_0} (\epsilon_{t_1(m),1} - \epsilon_{t_2(m),2})}{n_0} \geq \Delta_2 - \frac{2\sigma\sqrt{\eta T \ln T}}{n_0}\right) \\ &\leq \mathbb{P}\left(\frac{\sum_{m=1}^{n_0} \epsilon_{t_1(m),1}}{n_0} \geq \frac{\Delta_2}{2} - \frac{\sigma\sqrt{\eta T \ln T}}{n_0}\right) + \mathbb{P}\left(\frac{\sum_{m=1}^{n_0} \epsilon_{t_2(m),2}}{n_0} \geq \frac{\Delta_2}{2} - \frac{\sigma\sqrt{\eta T \ln T}}{n_0}\right) \\ &\leq 2 \exp\left(-n_0 \left(\frac{\Delta_2}{2} - \frac{\sigma\sqrt{\eta T \ln T}}{n_0}\right)^2 / 2\sigma^2\right) \\ &= 2 \exp\left(-n_0 \Delta_2^2 \left(1 - \frac{8\sigma\sqrt{\eta T \ln T}}{cT^\alpha}\right)^2 / 2\sigma^2\right) \\ &\leq 2 \exp\left(-\frac{n_0 \Delta_2^2}{8\sigma^2}\right) \\ &\leq 2 \exp\left(-\frac{cT^\alpha \cdot cT^{\alpha-1}}{128\sigma^2}\right) = 2 \exp\left(-\frac{c^2 T^{2\alpha-1}}{128\sigma^2}\right). \end{aligned} \tag{9}$$

Meanwhile,

$$\begin{aligned} & \mathbb{P}(\hat{R}_\theta^\pi(T) \geq cT^\alpha/2, \bar{S}) \\ &\leq \mathbb{P}(\bar{S}) \\ &= \mathbb{P}\left(\exists n \leq T/2 : \hat{\mu}_{t_1(n),1} + \frac{\sigma\sqrt{\eta T \ln T}}{n} < \hat{\mu}_{t_2(n),2} - \frac{\sigma\sqrt{\eta T \ln T}}{n}\right) \\ &= \mathbb{P}\left(\exists n \leq T/2 : \mu_1 + \frac{\sum_{m=1}^n \epsilon_{t_1(m),1} + \sigma\sqrt{\eta T \ln T}}{n} < \mu_2 + \frac{\sum_{m=1}^n \epsilon_{t_2(m),2} - \sigma\sqrt{\eta T \ln T}}{n}\right) \\ &\leq \sum_{n=1}^{\lfloor T/2 \rfloor} \mathbb{P}\left(\frac{\sum_{m=1}^n (\epsilon_{t_2(m),2} - \epsilon_{t_1(m),1})}{n} > \Delta_2 + \frac{2\sigma\sqrt{\eta T \ln T}}{n}\right) \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{n=1}^{\lfloor T/2 \rfloor} \mathbb{P} \left(\frac{\sum_{m=1}^n \epsilon_{t_2(m),2}}{n} > \frac{\Delta_2}{2} + \frac{\sigma \sqrt{\eta T \ln T}}{n} \right) \\
&\quad + \sum_{n=1}^{\lfloor T/2 \rfloor} \mathbb{P} \left(\frac{\sum_{m=1}^n -\epsilon_{t_1(m),1}}{n} > \frac{\Delta_2}{2} + \frac{\sigma \sqrt{\eta T \ln T}}{n} \right) \\
&\leq 2 \sum_{n=1}^{\lfloor T/2 \rfloor} \exp \left(-n \left(\frac{\Delta_2}{2} + \frac{\sigma \sqrt{\eta T \ln T}}{n} \right)^2 / 2\sigma^2 \right) \\
&\quad \text{(The new bonus is important. Otherwise the following steps fail.)} \\
&\leq T \exp \left(-2n\Delta_2 \frac{\sigma \sqrt{\eta T \ln T}}{n} / 2\sigma^2 \right) \\
&\leq T \exp(-\sigma \cdot cT^{\alpha-1} \cdot \sqrt{\eta T \ln T} / 2\sigma^2) \\
&= \exp \left(-\frac{cT^{\alpha-1/2} \sqrt{\eta \ln T} - \sigma \ln T}{2\sigma} \right) \\
&\leq \exp \left(-\frac{cT^{\alpha-1/2} \sqrt{\eta \ln T}}{4\sigma} \right) \leq \exp \left(-\frac{cT^{\alpha-1/2}}{16\sigma} \right). \tag{10}
\end{aligned}$$

Note that the equations above hold for any instance θ . Combining (8), (9), (10) yields

$$\sup_{\theta} \mathbb{P}(R_{\theta}^{\pi}(T) \geq cT^{\alpha}) \leq 4 \exp \left(-\frac{cT^{\alpha-1/2}}{16\sigma} \right).$$

□

Proof of Theorem 3.

We consider the environment where the noise ϵ is gaussian with standard deviation σ . Fix any $\alpha > 1/2$. Let $\theta_1 = 1/2$. Let $\theta(T) = (\theta_1, \theta_2(T))$ and $\tilde{\theta}(T) = (\theta_1, \tilde{\theta}_2(T))$, where $\theta_2(T) = \theta_1 + \frac{1}{2T^{1-\alpha}}$ and $\tilde{\theta}_2(T) = \theta_1 - \frac{1}{2T^{1-\alpha}}$. Let

$$\gamma \in (3/2 - \alpha, \min\{1, \beta + 2 - 2\alpha\}).$$

Such γ always exists because $\beta + 2 - 2\alpha > \alpha - 1/2 + 2 - 2\alpha = 3/2 - \alpha$ and $3/2 - \alpha < 3/2 - 1/2 = 1$. For notation simplicity, we will write $\theta(\tilde{\theta})$ instead of $\theta(T)$ ($\tilde{\theta}(T)$), but we must keep in mind that $\theta(\tilde{\theta})$ is dependent on T . Define

$$E_T = \left\{ |\hat{\mu}_{T,2} - \tilde{\theta}_2| \leq \delta \right\}$$

where $\delta > 0$ is a small number, and

$$F_T = \{n_2 \leq T^{\gamma}\}.$$

Then under the environment $\tilde{\theta}$, we have

$$\mathbb{P}_{\tilde{\theta}}^{\pi}(\bar{F}_T) = \mathbb{P}_{\tilde{\theta}}^{\pi}(n_2 > T^{\gamma}) \leq \frac{\mathbb{E}_{\tilde{\theta}}^{\pi}[n_2]}{T^{\gamma}} \leq \frac{\mathbb{E}[R_{\tilde{\theta}}^{\pi}(T)]}{T^{\gamma+\alpha-1}} \leq \frac{\sup_{\theta'} \mathbb{E}[R_{\theta'}^{\pi}(T)]}{T^{\gamma+\alpha-1}} \rightarrow 0$$

as $T \rightarrow +\infty$. Combined with Lemma 2, we have

$$\liminf_T \mathbb{P}_{\tilde{\theta}}^{\pi}(E_T, F_T) = 1.$$

Let $c' \in (c, 1/2)$. There exists T_0 such that $(1 - 2c')T > T^{\gamma}$ for any $T > T_0$. Fix $T > T_0$. Notice that

$$\begin{aligned}
&\mathbb{P} \left(\hat{R}_{\theta}^{\pi}(T) \geq cT^{\alpha} \right) \\
&\geq \mathbb{P} \left(R_{\theta}^{\pi}(T) \geq c'T^{\alpha}, -N^{\pi}(T) \geq -(c' - c)T^{\alpha} \right) \\
&= \mathbb{P} \left(R_{\theta}^{\pi}(T) \geq c'T - \mathbb{P} \left(R_{\theta}^{\pi}(T) \geq c'T^{\alpha}, N^{\pi}(T) > (c' - c)T^{\alpha} \right) \right)
\end{aligned}$$

$$\begin{aligned}
&\geq \mathbb{P}(R_\theta^\pi(T) \geq c'T^\alpha) - \mathbb{P}(N^\pi(T) > (c' - c)T^\alpha) \\
&\geq \mathbb{P}(R_\theta^\pi(T) \geq c'T^\alpha) - \exp\left(-\frac{(c' - c)^2 T^{2\alpha-1}}{2\sigma^2}\right)
\end{aligned}$$

The last inequality holds from Lemma 1. Now

$$\begin{aligned}
&\mathbb{P}(R_\theta^\pi(T) \geq c'T^\alpha) \\
&\geq \mathbb{P}_\theta^\pi(n_1 \geq 2c'T) \\
&\geq \mathbb{P}_\theta^\pi(n_2 \leq (1 - 2c')T) \\
&\geq \mathbb{P}_\theta^\pi(n_2 \leq T^\gamma) \\
&\geq \mathbb{P}_\theta^\pi(E_T, F_T) \\
&= \mathbb{E}_\theta^\pi[\mathbb{1}\{E_T F_T\}] \\
&= \mathbb{E}_\theta^\pi\left[\exp\left(\sum_{n=1}^{n_2} \frac{(X_{t_2(n),2} - \tilde{\theta}_2)^2 - (X_{t_2(n),2} - \theta_2)^2}{2\sigma^2}\right) \mathbb{1}\{E_T F_T\}\right] \\
&= \mathbb{E}_\theta^\pi\left[\exp\left(n_2 \left(\frac{\tilde{\theta}_2^2 - \theta_2^2}{2\sigma^2} + \frac{(\theta_2 - \tilde{\theta}_2)\hat{\theta}_{T,2}}{\sigma^2}\right)\right) \mathbb{1}\{E_T F_T\}\right] \\
&\geq \mathbb{E}_{\tilde{\theta}}^\pi\left[\exp\left(n_2 \left(\frac{\tilde{\theta}_2^2 - \theta_2^2}{2\sigma^2} + \frac{(\theta_2 - \tilde{\theta}_2)(\tilde{\theta}_2 - \delta)}{\sigma^2}\right)\right) \mathbb{1}\{E_T F_T\}\right] \\
&= \mathbb{E}_{\tilde{\theta}}^\pi\left[\exp\left(n_2 \left(-\frac{(\tilde{\theta}_2 - \theta_2)^2}{2\sigma^2} - \frac{\delta(\theta_2 - \tilde{\theta}_2)}{\sigma^2}\right)\right) \mathbb{1}\{E_T F_T\}\right] \\
&\geq \mathbb{E}_{\tilde{\theta}}^\pi\left[\exp\left(T^\gamma \left(-\frac{(\tilde{\theta}_2 - \theta_2)^2}{2\sigma^2} - \frac{\delta(\theta_2 - \tilde{\theta}_2)}{\sigma^2}\right)\right) \mathbb{1}\{E_T F_T\}\right] \\
&= \exp(-T^{\gamma+2\alpha-2}/2\sigma^2 - \delta T^{\gamma+\alpha-1}/\sigma^2) \mathbb{P}_{\tilde{\theta}}^\pi(E_T, F_T).
\end{aligned}$$

Notice that

$$\gamma + 2\alpha - 2 < 2\alpha - 1, \quad \gamma + 2\alpha - 2 \leq \gamma + \alpha - 1,$$

and $\delta > 0$ can be arbitrary. Therefore,

$$\begin{aligned}
&\liminf_T \frac{\ln \left\{ \sup_{\theta'} \mathbb{P}(\hat{R}_{\theta'}^\pi(T) \geq cT^\alpha) \right\}}{T^\beta} \\
&\geq \liminf_T \frac{-T^{\gamma+2\alpha-2}/2\sigma^2}{T^\beta} \\
&= 0.
\end{aligned}$$

Since $\ln \{\sup_\theta \mathbb{P}(R_\theta^\pi(T) \geq cT^\alpha)\} \leq 0$ always holds, we obtain the result. □

B Proofs for Section 4

Proof of Theorem 4. Without loss of generality, we assume $\theta_1 = \theta_*$. We prove the results one by one.

1. From Lemma 1,

$$\mathbb{E}[R_\theta^\pi] = \mathbb{E}[\hat{R}_\theta^\pi] = \sum_{k=2}^K \mathbb{E}[n_k] \cdot \Delta_k.$$

Let G be the event such that

$$G = \{\mu_k \in \text{CI}_{t,k}, \forall (t, k)\}.$$

Then

$$\mathbb{P}(\bar{G}) \leq \sum_{(t,k)} \mathbb{P}(\mu_k \notin \text{CI}_{t,k}) \leq K \sum_{n=1}^T 2 \exp\left(-\frac{\eta T \ln T}{2n}\right) \leq 2KT^{1-\eta/2}.$$

Thus,

$$\begin{aligned} \mathbb{E}[R_\theta^\pi] &= \sum_{k=2}^K (\mathbb{E}[n_k|G]\mathbb{P}(G) + \mathbb{E}[n_k|\bar{G}]\mathbb{P}(\bar{G})) \Delta_k \\ &\leq \sum_{k=2}^K \mathbb{E}[n_k|G] + T \cdot \mathbb{P}(\bar{G}) \\ &\leq \sum_{k=2}^K \mathbb{E}[n_k|G] + 2KT^{2-\eta/2} \leq \sum_{k=2}^K \mathbb{E}[n_k|G] + 2K. \end{aligned}$$

- (a) Let $\pi = \text{SE}$. Fix any arm $k \neq 1$. We let t'_k be the largest time period such that we have traversed all the arms in \mathcal{A} , and meanwhile arm k is not eliminated from \mathcal{A} . Then $n_k = n_{t'_k,k} + 1$. When doing the elimination after t_k , arm 1 and k are both pulled $n_{t'_k,k}$ times. Under G , we have

$$\begin{aligned} \mu_1 - 2\sigma \frac{\sqrt{\eta T \ln T}}{n_{t'_k,k}} &\leq \mu_1 - 2\sigma \frac{\sqrt{\eta T \ln T}}{n_{t'_k,1}} \leq \hat{\mu}_{t'_k,1} - \sigma \frac{\sqrt{\eta T \ln T}}{n_{t'_k,1}} \\ &\leq \hat{\mu}_{t'_k,k} + \sigma \frac{\sqrt{\eta T \ln T}}{n_{t'_k,k}} \leq \mu_k + 2\sigma \frac{\sqrt{\eta T \ln T}}{n_{t'_k,k}}. \end{aligned}$$

Therefore,

$$n_{t'_k,k} \leq 1 + 4\sigma \frac{\sqrt{\eta T \ln T}}{\Delta_k}$$

and thus,

$$n_k \leq 2 + 4\sigma \frac{\sqrt{\eta T \ln T}}{\Delta_k}.$$

As a result,

$$\mathbb{E}[\hat{R}_\theta^\pi(T)] \leq 2 \sum_{k=2}^K \Delta_k + 4 \sum_{k=2}^K \sigma \sqrt{\eta T \ln T} + 2K \leq 4K + 4K\sigma \sqrt{\eta T \ln T}.$$

- (b) Let $\pi = \text{UCB}$. Fix any arm $k \neq 1$. We let t_k be the largest time period such that arm k is pulled. Then $n_k = n_{t_k,k} = n_{t_k-1,k} + 1$. Under G , we have

$$\mu_1 \leq \hat{\mu}_{t_k-1,1} + \sigma \frac{\sqrt{\eta T \ln T}}{n_{t_k-1,1}} \leq \hat{\mu}_{t_k-1,k} + \sigma \frac{\sqrt{\eta T \ln T}}{n_{t_k-1,k}} \leq \mu_k + 2\sigma \frac{\sqrt{\eta T \ln T}}{n_{t_k-1,k}}$$

Therefore,

$$n_{t_k-1,k} \leq 2\sigma \frac{\sqrt{\eta T \ln T}}{\Delta_k}$$

and thus,

$$n_k \leq 1 + 2\sigma \frac{\sqrt{\eta T \ln T}}{\Delta_k}.$$

As a result,

$$\mathbb{E}[\hat{R}_\theta^\pi(T)] \leq \sum_{k=2}^K \Delta_k + 2 \sum_{k=2}^K \sigma \sqrt{\eta T \ln T} + 2K \leq 3K + 2K\sigma \sqrt{\eta T \ln T}.$$

2. We have

$$\mathbb{P}(\hat{R}_\theta^\pi(T) \geq x) \leq \mathbb{P}\left(R_\theta^\pi(T) \geq x(1 - 1/\sqrt{K})\right) + \mathbb{P}\left(N_\theta^\pi(T) \leq -x/\sqrt{K}\right)$$

From Lemma 1, the second term can be bounded as

$$\mathbb{P}(N_\theta^\pi(T) \leq -x/K) \leq \exp\left(-\frac{x^2}{2K\sigma^2 T}\right). \quad (11)$$

We are left to bound $\mathbb{P}\left(R_\theta^\pi(T) \geq x(1 - 1/\sqrt{K})\right)$.

(a) Let $\pi = \text{SE}$. For any $k \neq 1$, let S_k be the event defined as

$$S_k = \{\text{Arm 1 is not eliminated before arm } k\}.$$

Then

$$\bar{S}_k = \{\text{Arm 1 is eliminated before arm } k\}.$$

So

$$\begin{aligned} \mathbb{P}\left(R_\theta^\pi(T) \geq x(1 - 1/\sqrt{K})\right) &\leq \sum_{k=2}^K \mathbb{P}\left(n_k \Delta_k \geq x/(K + \sqrt{K})\right) \\ &= \sum_{k=2}^K \mathbb{P}(n_k \Delta_k \geq x/2K, S_k) + \mathbb{P}(n_k \Delta_k \geq x/2K, \bar{S}_k). \end{aligned}$$

Let $x > 0$. Fix any $k \neq 1$. With a slight abuse of notation, we let $n_0 = \lceil x/2K\Delta_k \rceil - 1$, then

$$n_0 \geq x/2K\Delta_k - 1 \geq (x - 2K)/2K\Delta_k.$$

Also, if $n_k \Delta_k \geq x/2K$, we must have

$$T \geq x/2K\Delta_k,$$

which means $\Delta_k \geq x/2KT$. By the definition of n_0 , arm k is not eliminated after being pulled n_0 times. So under S_k , after arm 1 being pulled n_0 times, it is still in the active set. We have

$$\begin{aligned} &\mathbb{P}(n_k \Delta_k \geq x/2K, S_k) \\ &= \mathbb{P}(n_k \geq x/2K\Delta_k, S_k) \\ &\leq \mathbb{P}(\text{arm 1 and } k \text{ are both not eliminated after each of them being pulled } n_0 \text{ times}) \\ &\leq \mathbb{P}\left(\hat{\mu}_{t_1(n_0),1} - \frac{\sigma\sqrt{\eta T \ln T}}{n_0} \leq \hat{\mu}_{t_k(n_0),k} + \frac{\sigma\sqrt{\eta T \ln T}}{n_0}\right) \\ &= \mathbb{P}\left(\mu_1 - \frac{\sum_{m=1}^{n_0} \epsilon_{t_1(m),1} + \sigma\sqrt{\eta T \ln T}}{n_0} \leq \mu_k + \frac{\sum_{m=1}^{n_0} \epsilon_{t_k(m),k} + \sigma\sqrt{\eta T \ln T}}{n_0}\right) \\ &= \mathbb{P}\left(\frac{\sum_{m=1}^{n_0} (\epsilon_{t_1(m),1} - \epsilon_{t_k(m),k})}{n_0} \geq \Delta_k - \frac{2\sigma\sqrt{\eta T \ln T}}{n_0}\right) \\ &\leq \mathbb{P}\left(\frac{\sum_{m=1}^{n_0} \epsilon_{t_1(m),1}}{n_0} \geq \frac{\Delta_k}{2} - \frac{\sigma\sqrt{\eta T \ln T}}{n_0}\right) \\ &\quad + \mathbb{P}\left(\frac{\sum_{m=1}^{n_0} \epsilon_{t_k(m),k}}{n_0} \geq \frac{\Delta_k}{2} - \frac{\sigma\sqrt{\eta T \ln T}}{n_0}\right) \\ &\leq 2 \exp\left(-n_0 \left(\frac{\Delta_k}{2} - \frac{\sigma\sqrt{\eta T \ln T}}{n_0}\right)_+^2 / 2\sigma^2\right) \\ &= 2 \exp\left(-n_0 \Delta_k^2 \left(1 - \frac{2\sigma\sqrt{\eta T \ln T}}{n_0 \Delta_k}\right)_+^2 / 8\sigma^2\right) \end{aligned}$$

$$\begin{aligned}
&\leq 2 \exp \left(-\frac{x(x-2K)_+}{4KT} \left(1 - \frac{4K^2\sigma\sqrt{\eta T \ln T}}{x-2K} \right)_+^2 / 8\sigma^2 \right) \\
&\leq 2 \exp \left(-\frac{(x-2K-4K\sigma\sqrt{\eta T \ln T})_+^2}{32\sigma^2 K^2 T} \right). \tag{12}
\end{aligned}$$

In the following, we bound $\mathbb{P}(n_k \Delta_k \geq x/2K, \bar{S}_k)$. Suppose that after n phases, arm 1 is eliminated by arm k' (k' is not necessarily k). By the definition of \bar{S}_k , arm k is not eliminated. Therefore, we have

$$\begin{aligned}
\hat{\mu}_{t_{k'}(n),k'} - \frac{\sigma\sqrt{\eta T \ln T}}{n} &\geq \hat{\mu}_{t_1(n),1} + \frac{\sigma\sqrt{\eta T \ln T}}{n}, \\
\hat{\mu}_{t_k(n),k} + \frac{\sigma\sqrt{\eta T \ln T}}{n} &\geq \hat{\mu}_{t_1(n),1} + \frac{\sigma\sqrt{\eta T \ln T}}{n} \tag{13}
\end{aligned}$$

holds simultaneously. The first inequality holds because arm 1 is eliminated. The second inequality holds because arm k is not eliminated. Now for fixed n ,

$$\begin{aligned}
&\mathbb{P} \left((13) \text{ happens}; \Delta_k \geq \frac{x}{2KT} \right) \\
&\leq \mathbb{P} \left(\exists k' : \hat{\mu}_{t_{k'}(n),k'} - \frac{\sigma\sqrt{\eta T \ln T}}{n} \geq \hat{\mu}_{t_1(n),1} + \frac{\sigma\sqrt{\eta T \ln T}}{n} \right) \\
&\quad \wedge \mathbb{P} \left(\hat{\mu}_{t_k(n),k} + \frac{\sigma\sqrt{\eta T \ln T}}{n} \geq \hat{\mu}_{t_1(n),1} + \frac{\sigma\sqrt{\eta T \ln T}}{n}; \Delta_k \geq x/2KT \right) \\
&\leq \mathbb{P} \left(\exists k' : \frac{\sum_{m=1}^n (\epsilon_{t_{k'}(m),k'} - \epsilon_{t_1(m),1})}{n} \geq \frac{2\sigma\sqrt{\eta T \ln T}}{n} \right) \\
&\quad \wedge \mathbb{P} \left(\frac{\sum_{m=1}^n (\epsilon_{t_k(m),k} - \epsilon_{t_1(m),1})}{n} > \Delta_k; \Delta_k \geq x/2KT \right) \\
&\leq \sum_{k' \neq 1} \mathbb{P} \left(\frac{\sum_{m=1}^n \epsilon_{t_{k'}(m),k'}}{n} \geq \frac{\sigma\sqrt{\eta T \ln T}}{n} \right) + \mathbb{P} \left(\frac{\sum_{m=1}^n \epsilon_{t_1(m),1}}{n} \leq -\frac{\sigma\sqrt{\eta T \ln T}}{n} \right) \\
&\quad \wedge \left(\mathbb{P} \left(\frac{\sum_{m=1}^n \epsilon_{t_k(m),k}}{n} \geq \frac{x}{4KT} \right) + \mathbb{P} \left(\frac{\sum_{m=1}^n \epsilon_{t_1(m),1}}{n} \leq -\frac{x}{4KT} \right) \right) \\
&\leq 2K \exp \left(-\frac{\eta T \ln T}{2n} \right) \wedge 2K \exp \left(-\frac{nx^2}{32\sigma^2 K^2 T^2} \right) \\
&= 2K \exp \left(-\left(\frac{\eta T \ln T}{2n} \vee \frac{nx^2}{32\sigma^2 K^2 T^2} \right) \right) \\
&\leq 2K \exp \left(-\frac{x\sqrt{\eta \ln T}}{8\sigma K \sqrt{T}} \right)
\end{aligned}$$

Therefore,

$$\begin{aligned}
&\mathbb{P}(n_k \Delta_k \geq x/2K, \bar{S}) \\
&= \mathbb{P}(\exists n \leq T/2 : (13) \text{ happens}; n_k \Delta_k \geq x/2K) \\
&\leq \sum_{n=1}^{\lfloor T/2 \rfloor} \mathbb{P} \left((13) \text{ happens}; \Delta_k \geq \frac{x}{2KT} \right) \\
&\leq KT \exp \left(-\frac{x\sqrt{\eta \ln T}}{8\sigma K \sqrt{T}} \right). \tag{14}
\end{aligned}$$

Note that the equations above hold for any instance θ . Combining (11), (12), (14) yields

$$\begin{aligned}
&\sup_{\theta} \mathbb{P}(R_{\theta}^{\pi}(T) \geq x) \\
&\leq \exp \left(-\frac{x^2}{2K\sigma^2 T} \right) + 2K \exp \left(-\frac{(x-2K-4K\sigma\sqrt{\eta T \ln T})_+^2}{32\sigma^2 K^2 T} \right)
\end{aligned}$$

$$+ K^2 T \exp \left(-\frac{x\sqrt{\eta \ln T}}{8\sigma K \sqrt{T}} \right)$$

(b) Let $\pi = \text{UCB}$. From (a), we know that

$$\begin{aligned} \mathbb{P} \left(R_\theta^\pi(T) \geq x(1 - 1/\sqrt{K}) \right) &\leq \sum_{k=2}^K \mathbb{P} \left(n_k \Delta_k \geq x/(K + \sqrt{K}) \right) \\ &\leq \sum_{k=2}^K \mathbb{P}(n_k \Delta_k \geq x/2K). \end{aligned}$$

Let $x > 0$. Fix $k \neq 1$. With a slight abuse of notation, we let $n_0 = \lceil x/2K\Delta_k \rceil - 1$. Remember that $t_k(n_0 + 1)$ is the time period that arm k is pulled for the $(n_0 + 1)$ th time. We emphasize again that $\Delta_k \geq x/(2KT)$. Then

$$\begin{aligned} &\mathbb{P}(n_k \Delta_k \geq x/2K) \\ &= \mathbb{P}(n_k \geq x/2K\Delta_k) \\ &\leq \mathbb{P} \left(\hat{\mu}_{t_k(n_0+1)-1,1} + \frac{\sigma\sqrt{\eta T \ln T}}{n_{t_k(n_0+1)-1,1}} \leq \hat{\mu}_{t_k(n_0+1)-1,k} + \frac{\sigma\sqrt{\eta T \ln T}}{n_0} \right) \\ &= \mathbb{P} \left(\mu_1 + \frac{\sum_{m=1}^{t_k(n_0+1)-1,1} \epsilon_{t_1(m),1} + \sigma\sqrt{\eta T \ln T}}{n_{t_k(n_0+1)-1,1}} \leq \mu_k + \frac{\sum_{m=1}^{n_0} \epsilon_{t_k(m),k} + \sigma\sqrt{\eta T \ln T}}{n_0} \right) \\ &\leq \mathbb{P} \left(\exists n \in [T] : \frac{\sum_{m=1}^{n_0} \epsilon_{t_k(m),k} + \sigma\sqrt{\eta T \ln T}}{n_0} - \frac{\sum_{m=1}^n \epsilon_{t_1(m),1} + \sigma\sqrt{\eta T \ln T}}{n} \geq \Delta_k \right) \\ &\leq \mathbb{P} \left(\frac{\sum_{m=1}^{n_0} \epsilon_{t_k(m),k} + \sigma\sqrt{\eta T \ln T}}{n_0} \geq \frac{\Delta_k}{2} \right) \\ &\quad + \mathbb{P} \left(\exists n \in [T] : \frac{\sum_{m=1}^n \epsilon_{t_1(m),1} + \sigma\sqrt{\eta T \ln T}}{n} \leq -\frac{\Delta_k}{2} \right) \\ &\leq \mathbb{P} \left(\frac{\sum_{m=1}^{n_0} \epsilon_{t_k(m),k} + \sigma\sqrt{\eta T \ln T}}{n_0} \geq \frac{\Delta_k}{2} \right) \\ &\quad + \sum_{n=1}^T \mathbb{P} \left(\frac{\sum_{m=1}^n -\epsilon_{t_1(m),1}}{n} \geq \frac{\Delta_k}{2} + \frac{\sigma\sqrt{\eta T \ln T}}{n} \right) \\ &\leq \mathbb{P} \left(\frac{\sum_{m=1}^{n_0} \epsilon_{t_k(m),k} + \sigma\sqrt{\eta T \ln T}}{n_0} \geq \frac{\Delta_k}{2} \right) + \sum_{n=1}^T \mathbb{P} \left(\frac{\sum_{m=1}^n -\epsilon_{t_1(m),1}}{n} \geq \sqrt{\frac{x\sigma\sqrt{\eta T \ln T}}{nKT}} \right) \\ &\leq \exp \left(-\frac{(x - 2K - 4K\sigma\sqrt{\eta T \ln T})_+^2}{32\sigma^2 KT} \right) + T \exp \left(-\frac{x\sqrt{\eta T \ln T}}{2\sigma KT} \right). \end{aligned} \tag{15}$$

The last inequality holds from (12) and concentration of subgaussian variables. Note that the equations above hold for any instance θ . Combining (11), (15) yields

$$\begin{aligned} &\sup_{\theta} \mathbb{P}(R_\theta^\pi(T) \geq x) \\ &\leq \exp \left(-\frac{x^2}{2K\sigma^2 T} \right) + K \exp \left(-\frac{(x - 2K - 4K\sigma\sqrt{\eta T \ln T})_+^2}{32\sigma^2 KT} \right) \\ &\quad + K^2 T \exp \left(-\frac{x\sqrt{\eta T \ln T}}{8\sigma KT} \right). \end{aligned}$$

□

Remarks for Theorem 4. To get a neat form of the tail bound, one can notice that the last term in the bound can be written as

$$K \exp \left(-\frac{x\sqrt{\eta \ln T} - 8\sigma K \sqrt{T \ln(KT)}}{8\sigma K \sqrt{T}} \right) \leq K \exp \left(-\frac{(x - 16K\sigma\sqrt{1/\eta \cdot T \ln T})\sqrt{\eta \ln T}}{8\sigma K \sqrt{T}} \right)$$

Since the tail probability has a trivial upper bound of 1, the last term can be replaced by

$$K \exp \left(-\frac{(x - 16K\sigma\sqrt{1/\eta \cdot T \ln T})_+ \sqrt{\eta \ln T}}{8\sigma K \sqrt{T}} \right).$$

Then if we let

$$y = \frac{\left(x - 2K - 16\sigma K \sqrt{(\eta \vee 1/\eta)T \ln T}\right)_+}{8\sigma K \sqrt{T}},$$

then the bound can be written as

$$\begin{aligned} & \sup_{\theta} \mathbb{P}(R_{\theta}^{\pi}(T) \geq x) \\ & \leq \exp(-y^2) + K \exp(-y^2) + K \exp\left(-y\sqrt{\eta \ln T}\right) \leq 4K \exp\left(-(y^2 \wedge y\sqrt{\eta \ln T})\right). \end{aligned}$$

(a) **Lemma: SE and UCB may get sub-optimal regret with the standard bonus term (1) when η is too small.**

Let $\theta = (1, 0)$ and $\sigma = 1$ with independent Gaussian noise. We first consider $\pi = \text{SE}$. The probability that arm 1 is eliminated after the first phase is

$$\begin{aligned} & \mathbb{P}\left(\hat{\mu}_{1,1} + \sqrt{\eta \ln T} < \hat{\mu}_{2,2} - \sqrt{\eta \ln T}\right) \\ & = \mathbb{P}\left(\epsilon_{1,1} - \epsilon_{2,2} < -1 - 2\sqrt{\eta \ln T}\right) \\ & \geq \frac{1}{\sqrt{2\pi}} \frac{1/\sqrt{2} + \sqrt{2\eta \ln T}}{1 + (1/\sqrt{2} + \sqrt{2\eta \ln T})^2} \exp(-(1/\sqrt{2} + \sqrt{2\eta \ln T})^2/2) \\ & = \Theta\left(\frac{T^{-\eta}}{\sqrt{\ln T}}\right) \end{aligned}$$

The inequality holds because for a standard normal variable X , it is established that

$$\mathbb{P}(X > t) \geq \frac{1}{\sqrt{2\pi}} \frac{t}{1 + t^2} \exp(-t^2/2).$$

Therefore, the expected regret is at least

$$\Theta\left(\frac{T^{-\eta}}{\sqrt{\ln T}}\right) \cdot (T - 2) = \Theta\left(\frac{T^{1-\eta}}{\sqrt{\ln T}}\right).$$

If η is very small, then apparently the regret is sub-optimal.

Now we consider $\pi = \text{UCB}$. The probability that arm 1 is pulled only once is

$$\begin{aligned} & \mathbb{P}\left(\forall 2 \leq t \leq T : \hat{\mu}_{1,1} + \sqrt{\eta \ln T} < \hat{\mu}_{t,2} + \sqrt{\frac{\eta \ln T}{t-1}}\right) \\ & \geq \mathbb{P}\left(\forall 2 \leq t \leq T : \hat{\mu}_{1,1} + \sqrt{\eta \ln T} < \hat{\mu}_{t,2}\right) \\ & \geq \mathbb{P}\left(\left\{\hat{\mu}_{1,1} < -1 - 2\sqrt{\eta \ln T}\right\} \cap \left\{\forall 1 \leq t < T : \sum_{s=2}^{t+1} \epsilon_{s,2} > -1 - t\sqrt{\eta \ln T}\right\}\right) \\ & = \mathbb{P}\left(\left\{\epsilon_{1,1} < -2 - 2\sqrt{\eta \ln T}\right\}\right) \cdot \mathbb{P}\left(\left\{\forall 1 \leq t < T : \sum_{s=2}^{t+1} \epsilon_{s,2} > -1 - t\sqrt{\eta \ln T}\right\}\right) \end{aligned}$$

We have

$$\begin{aligned} & \mathbb{P}\left(\left\{\epsilon_{1,1} < -2 - 2\sqrt{\eta \ln T}\right\}\right) \\ & \geq \frac{1}{\sqrt{2\pi}} \frac{2 + 2\sqrt{\eta \ln T}}{1 + (2 + 2\sqrt{\eta \ln T})^2} \exp(-(2 + 2\sqrt{\eta \ln T})^2/2) \\ & = \Theta\left(\frac{T^{-2\eta}}{\sqrt{\ln T}}\right) \end{aligned}$$

We use a martingale argument and the optional sampling theorem to bound the second probability. Define $Z_t = \sum_{s=2}^{t+1} \epsilon_{s,2}$. Define the stopping time

$$\tau = \inf_t \{Z_t \leq -1 - t\sqrt{\eta \ln T}\}$$

Then

$$\mathbb{P} \left(\left\{ \forall 1 \leq t < T : \sum_{s=2}^{t+1} \epsilon_{s,2} > -1 - t\sqrt{\eta \ln T} \right\} \right) = \mathbb{P}(\tau \geq T)$$

For fixed T , $\tau \wedge (T-1)$ is finite. Notice that

$$\exp(-2\sqrt{\eta \ln T} Z_t - 2\eta \ln T \cdot t)$$

is a martingale with mean 1. By the optional sampling theorem, we have

$$\begin{aligned} 1 &= \mathbb{E} \left[\exp(-2\sqrt{\eta \ln T} Z_{\tau \wedge (T-1)} - 2\eta \ln T \cdot (\tau \wedge (T-1))) \right] \\ &\geq \mathbb{E} \left[\exp(-2\sqrt{\eta \ln T} Z_{\tau} - 2\eta \ln T \cdot \tau) \mathbb{1}_{\{\tau < T\}} \right] \\ &\geq \exp(2\sqrt{\eta \ln T}) \mathbb{P}(\tau < T) \end{aligned}$$

Therefore, the second probability is bounded by

$$1 - \exp(-2\sqrt{\eta \ln T}).$$

The expected regret is at least

$$\Theta \left(\frac{T^{-2\eta}}{\sqrt{\ln T}} \right) \cdot \left(1 - \exp(-2\sqrt{\eta \ln T}) \right) \cdot (T-2) = \Omega \left(\frac{T^{1-2\eta}}{\sqrt{\ln T}} \right).$$

C Additional Experiments for Section 5

For the first experiment presented in the main body of the paper, we provide the empirical mean for the cumulative reward in Table 1. We note that there is no direct implication by comparing the four different algorithms at the same value of κ , because the algorithms use the parameter κ in different ways and in different format of the bonus term. For example, for some value of κ , SE has a higher empirical mean for the cumulative reward compared to SE_new, whereas for some other value of κ , SE has a smaller empirical mean compared to SE_new. There is no direct implication by fixing a value of κ and comparing different algorithms. Nevertheless, we do observe that TS performs similarly to UCB, from both Figure 1 and Table 1, and so we put our discussion on confidence bound policies.

Here is what Table 1 implies. Table 1 shows that, SE_new (or UCB_new) achieves empirical mean for the cumulative reward as high as that SE (or UCB) can achieve. The highest empirical mean for the cumulative reward that can be achieved by SE_new with various choices of κ is comparable to the highest empirical mean that can be achieved by SE. However, Figure 1 in the main body has shown that, compared to SE, SE_new has much lower probability of incurring a low cumulative reward. The implication is that (i) in terms of the empirical mean of cumulative reward, SE_new is as good as SE; (ii) in terms of the risk of incurring a low cumulative reward, SE_new is much better (i.e., lower risk) than SE. The same implication holds analogously for the comparison between UCB_new and UCB.

Policy \ κ	0.1	0.2	0.4	0.8
SE	311.60	336.46	375.53	374.69
UCB	349.68	359.68	377.17	390.23
TS	351.00	360.71	377.94	390.32
SE_new	388.16	376.69	354.25	309.58
UCB_new	393.27	387.48	377.72	360.69

Table 1: Empirical mean for the cumulative reward

Next, we consider a 4-armed bandit problem with $\theta = (0.2, 0.4, 0.6, 0.8)$, $\sigma = 1$, $T = 500$ and Gaussian noise. We test four policies: SE and UCB with the classical bonus design described in (1), and SE_new and UCB_new with the proposed new bonus design in (4). The tuning parameter has 4 choices: $\kappa \in \{0.1, 0.2, 0.4, 0.8\}$. For each policy and κ , we run 5000 simulation paths and for each

path we collect the cumulative reward. We plot the empirical distribution (histogram) for a policy's cumulative reward in Figure 2. We also report the empirical mean in Table 2. Indeed, one can observe that for TS and both SE and UCB with (1), there is a significant part of distribution around 200 and 300, which means that with a non-negligible probability the two policies always pull arm 2 or 3, incurring a linear regret. Such phenomenon becomes more significant when κ is small. In contrast, with the new design (4), the reward is highly concentrated for every $\kappa > 0$. Particularly, when $\kappa = 0.1$, either SE_new or UCB_new achieves both high empirical mean and light-tailed distribution. When κ is relatively large, e.g., $\kappa = 0.8$, the empirical mean is not very satisfactory. But this is consistent with our analysis in Theorem 4, which indicates an additional \sqrt{K} factor compared to the optimal $\tilde{O}(\sqrt{KT})$ expected regret.




Figure 2: Empirical distribution for the cumulative reward; bottom two are new proposed policies

κ	0.1	0.2	0.4	0.8
Policy				
SE	292.12	307.45	349.35	316.83
UCB	339.41	348.52	360.26	369.25
TS	341.05	349.86	359.82	365.26
SE_new	361.93	334.18	283.69	251.52
UCB_new	370.72	360.02	339.29	308.96

Table 2: Empirical mean for the cumulative reward

D Additional Discussions on the Setup

In our current model, if arm k is pulled in the time period t , we write that the observed reward is $r_{t,k} = \theta_k + \epsilon_{t,k}$, where θ_k is the true expected reward and $\epsilon_{t,k}$ represents a mean zero random noise. We have assumed this mean-zero random noise to be σ sub-Gaussian, which appears to be a standard assumption in the literature and includes distributions such as Gaussian and Bernoulli.

Instead of modeling the random reward using additive noise, one may alternatively model the random reward using multiplicative noise. That is, if arm k is pulled in the time period t , the observed reward $r_{t,k} = \theta_k \cdot \epsilon'_{t,k}$ where θ_k is the true expected reward and $\epsilon'_{t,k}$ represents a mean-one random noise. Some assumptions are needed to describe the distribution of the multiplicative noise $\epsilon'_{t,k}$. If the multiplicative noise $\epsilon'_{t,k}$ is σ sub-Gaussian, then if we rewrite the observed reward as $r_{t,k} = \theta_k \cdot \epsilon'_{t,k} = \theta_k + \theta_k(\epsilon'_{t,k} - 1)$, the multiplicative noise can be transformed into an additive noise with the additive part given by $\theta_k(\epsilon'_{t,k} - 1)$. If $\epsilon'_{t,k}$ is σ sub-Gaussian, then the random term $\theta_k(\epsilon'_{t,k} - 1)$ is sub-Gaussian, and therefore the multiplicative case may be transformed into an additive case with moderate modifications. If otherwise, $\epsilon'_{t,k}$ is not σ sub-Gaussian, then it means that the

random term $\theta_k(\epsilon'_{t,k} - 1)$ is likely also not sub-Gaussian. This would project the question back into how we should deal with policy designs if the additive noise is not sub-Gaussian. That is, noises come from general distributions. In fact, there is a separate line of literature studying bandit problems with non sub-Gaussian noise (see, e.g., [7]). We hope our framework can be integrated with this line of literature to handle certain types of non sub-Gaussian noise, provided with controls on the tail behavior of the random noises.

As is discussed in Section 6, this interesting question is left for future work.