
Bounded Regret for Finite-Armed Structured Bandits

Tor Lattimore

Department of Computing Science
University of Alberta, Canada
tlattimo@ualberta.ca

Rémi Munos

INRIA
Lille, France¹
remi.munos@inria.fr

Abstract

We study a new type of K -armed bandit problem where the expected return of one arm may depend on the returns of other arms. We present a new algorithm for this general class of problems and show that under certain circumstances it is possible to achieve finite expected cumulative regret. We also give problem-dependent lower bounds on the cumulative regret showing that at least in special cases the new algorithm is nearly optimal.

1 Introduction

The multi-armed bandit problem is a reinforcement learning problem with K actions. At each time-step a learner must choose an action i after which it receives a reward distributed with mean μ_i . The goal is to maximise the cumulative reward. This is perhaps the simplest setting in which the well-known exploration/exploitation dilemma becomes apparent, with a learner being forced to choose between exploring arms about which she has little information, and exploiting by choosing the arm that currently appears optimal.

We consider a general class of K -armed bandit problems where the expected return of each arm may be dependent on other arms. This model has already been considered when the dependencies are linear [18] and also in the general setting studied here [12, 1]. Let $\Theta \ni \theta^*$ be an arbitrary

parameter space and define the expected return of arm i by $\mu_i(\theta^*) \in \mathbb{R}$. The learner is permitted to know the functions $\mu_1 \cdots \mu_K$, but not the true parameter θ^* . The unknown parameter θ^* determines the mean reward for each arm. The performance of a learner is measured by the (expected) cumulative regret, which is the difference between the expected return of the optimal policy and the (expected) return of the learner's policy. $R_n := n \max_{i \in 1 \dots K} \mu_i(\theta^*) - \sum_{t=1}^n \mu_{I_t}(\theta^*)$ where I_t is the arm chosen at time-step t .

A motivating example is as follows. Suppose a long-running company must decide each week whether or not to purchase some new form of advertising with unknown expected returns. The problem may be formulated using the new setting by letting $K = 2$ and $\Theta = [-\infty, \infty]$. We assume the base-line performance without purchasing the advertising is known and so define $\mu_1(\theta) = 0$ for all θ . The expected return of choosing to advertise is $\mu_2(\theta) = \theta$ (see Figure (b) above).

Our main contribution is a new algorithm based on UCB [6] for the structured bandit problem with strong problem-dependent guarantees on the regret. The key improvement over UCB is that the new algorithm enjoys finite regret in many cases while UCB suffers logarithmic regret unless all arms have the same return. For example, in (a) and (c) above we show that finite regret is possible for all

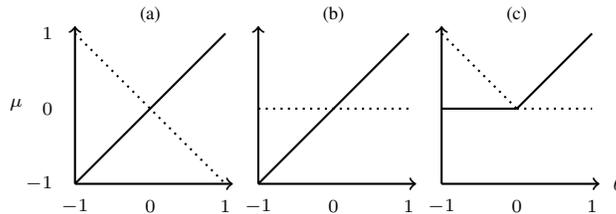


Figure 1: Examples

¹Current affiliation: Google DeepMind.

θ^* , while in the advertising problem finite regret is attainable if $\theta^* \geq 0$. The improved algorithm exploits the known structure and so avoids the famous negative results by Lai and Robbins [17]. One insight from this work is that knowing the return of the optimal arm and a bound on the minimum gap is not the only information that leads to the possibility of finite regret. In the examples given above neither quantity is known, but the assumed structure is nevertheless sufficient for finite regret.

Despite the enormous literature on bandits, as far as we are aware this is the first time this setting has been considered with the aim of achieving finite regret. There has been substantial work on exploiting various kinds of structure to reduce an otherwise impossible problem to one where sub-linear (or even logarithmic) regret is possible [19, 4, 10, and references therein], but the focus is usually on efficiently dealing with large action spaces rather than sub-logarithmic/finite regret. The most comparable previous work studies the case where both the return of the best arm and a bound on the minimum gap between the best arm and some sub-optimal arm is known [11, 9], which extended the permutation bandits studied by Lai and Robbins [16] and more general results by the same authors [15]. Also relevant is the paper by Agrawal et. al. [1], which studied a similar setting, but where Θ was finite. Graves and Lai [12] extended the aforementioned contribution to continuous parameter spaces (and also to MDPs). Their work differs from ours in a number of ways. Most notably, their objective is to compute exactly the asymptotically optimal regret in the case where finite regret is *not* possible. In the case where finite regret is possible they prove only that the optimal regret is sub-logarithmic, and do not present any explicit bounds on the actual regret. Aside from this the results depend on the parameter space being a metric space and they assume that the optimal policy is locally constant about the true parameter.

2 Notation

General. Most of our notation is common with [8]. The indicator function is denoted by $\mathbb{1}\{expr\}$ and is 1 if *expr* is true and 0 otherwise. We use \log for the natural logarithm. Logical and/or are denoted by \wedge and \vee respectively. Define function $\omega(x) = \min\{y \in \mathbb{N} : z \geq x \log z, \forall z \geq y\}$, which satisfies $\log \omega(x) \in O(\log x)$. In fact, $\lim_{x \rightarrow \infty} \log(\omega(x))/\log(x) = 1$.

Bandits. Let Θ be a set. A K -armed structured bandit is characterised by a set of functions $\mu_k : \Theta \rightarrow \mathbb{R}$ where $\mu_k(\theta)$ is the expected return of arm $k \in A := \{1, \dots, K\}$ given unknown parameter θ . We define the mean of the optimal arm by the function $\mu^* : \Theta \rightarrow \mathbb{R}$ with $\mu^*(\theta) := \max_i \mu_i(\theta)$. The true unknown parameter that determines the means is $\theta^* \in \Theta$. The best arm is $i^* := \arg \max_i \mu_i(\theta^*)$. The arm chosen at time-step t is denoted by I_t while $X_{i,s}$ is the s th reward obtained when sampling from arm i . We denote the number of times arm i has been chosen at time-step t by $T_i(t)$. The empiric estimate of the mean of arm i based on the first s samples is $\hat{\mu}_{i,s}$. We define the gap between the means of the best arm and arm i by $\Delta_i := \mu^*(\theta^*) - \mu_i(\theta^*)$. The set of sub-optimal arms is $A' := \{i \in A : \Delta_i > 0\}$. The minimum gap is $\Delta_{\min} := \min_{i \in A'} \Delta_i$ while the maximum gap is $\Delta_{\max} := \max_{i \in A} \Delta_i$. The cumulative regret is defined

$$R_n := \sum_{t=1}^n \mu^*(\theta^*) - \sum_{t=1}^n \mu_{I_t} = \sum_{t=1}^n \Delta_{I_t}$$

Note quantities like Δ_i and i^* depend on θ^* , which is omitted from the notation. As is rather common we assume that the returns are sub-gaussian, which means that if X is the return sampled from some arm, then $\ln \mathbb{E} \exp(\lambda(X - \mathbb{E}X)) \leq \lambda^2 \sigma^2 / 2$. As usual we assume that σ^2 is known and does not depend on the arm. If $X_1 \dots X_n$ are sampled independently from some arm with mean μ and $S_n = \sum_{t=1}^n X_t$, then the following maximal concentration inequality is well-known.

$$\mathbb{P} \left\{ \max_{1 \leq t \leq n} |S_t - t\mu| \geq \varepsilon \right\} \leq 2 \exp \left(-\frac{\varepsilon^2}{2n\sigma^2} \right).$$

A straight-forward corollary is that $\mathbb{P} \{ |\hat{\mu}_{i,n} - \mu_i| \geq \varepsilon \} \leq 2 \exp \left(-\frac{\varepsilon^2 n}{2\sigma^2} \right)$.

It is an important point that Θ is completely arbitrary. The classic multi-armed bandit can be obtained by setting $\Theta = \mathbb{R}^K$ and $\mu_k(\theta) = \theta_k$, which removes all dependencies between the arms. The setting where the optimal expected return is known to be zero and a bound on $\Delta_i \geq \varepsilon$ is known can be regained by choosing $\Theta = (-\infty, -\varepsilon]^K \times \{1, \dots, K\}$ and $\mu_k(\theta_1, \dots, \theta_K, i) = \theta_k \mathbb{1}\{k \neq i\}$. We do not demand that $\mu_k : \Theta \rightarrow \mathbb{R}$ be continuous, or even that Θ be endowed with a topology.

3 Structured UCB

We propose a new algorithm called UCB-S that is a straight-forward modification of UCB [6], but where the known structure of the problem is exploited. At each time-step it constructs a confidence interval about the mean of each arm. From this a subspace $\tilde{\Theta}_t \subseteq \Theta$ is constructed, which contains the true parameter θ with high probability. The algorithm takes the optimistic action over all $\theta \in \tilde{\Theta}_t$.

Algorithm 1 UCB-S

- 1: **Input:** functions $\mu_1, \dots, \mu_k : \Theta \rightarrow [0, 1]$
 - 2: **for** $t \in 1, \dots, \infty$ **do**
 - 3: Define confidence set $\tilde{\Theta}_t \leftarrow \left\{ \tilde{\theta} : \forall i, \left| \mu_i(\tilde{\theta}) - \hat{\mu}_{i, T_i(t-1)} \right| < \sqrt{\frac{\alpha \sigma^2 \log t}{T_i(t-1)}} \right\}$
 - 4: **if** $\tilde{\Theta}_t = \emptyset$ **then**
 - 5: Choose arm arbitrarily
 - 6: **else**
 - 7: Optimistic arm is $i \leftarrow \arg \max_i \sup_{\tilde{\theta} \in \tilde{\Theta}_t} \mu_i(\tilde{\theta})$
 - 8: Choose arm i
-

Remark 1. The choice of arm when $\tilde{\Theta}_t = \emptyset$ does not affect the regret bounds in this paper. In practice, it is possible to simply increase t without taking an action, but this complicates the analysis. In many cases the true parameter θ^* is never identified in the sense that we do not expect that $\tilde{\Theta}_t \rightarrow \{\theta^*\}$. The computational complexity of UCB-S depends on the difficulty of computing $\tilde{\Theta}_t$ and computing the optimistic arm within this set. This is efficient in simple cases, like when μ_k is piecewise linear, but may be intractable for complex functions.

4 Theorems

We present two main theorems bounding the regret of the UCB-S algorithm. The first is for arbitrary θ^* , which leads to a logarithmic bound on the regret comparable to that obtained for UCB by [6]. The analysis is slightly different because UCB-S maintains upper and lower confidence bounds and selects its actions optimistically from the model class, rather than by maximising the upper confidence bound as UCB does.

Theorem 2. *If $\alpha > 2$ and $\theta \in \Theta$, then the algorithm UCB-S suffers an expected regret of at most*

$$\mathbb{E}R_n \leq \frac{2\Delta_{\max}K(\alpha-1)}{\alpha-2} + \sum_{i \in A'} \frac{8\alpha\sigma^2 \log n}{\Delta_i} + \sum_i \Delta_i$$

If the samples from the optimal arm are sufficient to learn the optimal action, then finite regret is possible. In Section 6 we give something of a converse by showing that if knowing the mean of the optimal arm is insufficient to act optimally, then logarithmic regret is unavoidable.

Theorem 3. *Let $\alpha = 4$ and assume there exists an $\varepsilon > 0$ such that*

$$(\forall \theta \in \Theta) \quad |\mu_{i^*}(\theta^*) - \mu_{i^*}(\theta)| < \varepsilon \implies \forall i \neq i^*, \mu_{i^*}(\theta) > \mu_i(\theta). \quad (1)$$

$$\text{Then } \mathbb{E}R_n \leq \sum_{i \in A'} \left(\frac{32\sigma^2 \log \omega^*}{\Delta_i} + \Delta_i \right) + 3\Delta_{\max}K + \frac{\Delta_{\max}K^3}{\omega^*},$$

$$\text{with } \omega^* := \max \left\{ \omega \left(\frac{8\sigma^2 \alpha K}{\varepsilon^2} \right), \omega \left(\frac{8\sigma^2 \alpha K}{\Delta_{\min}^2} \right) \right\}.$$

Remark 4. For small ε and large n the expected regret looks like $\mathbb{E}R_n \in O \left(\sum_{i=1}^K \frac{\log(\frac{1}{\varepsilon})}{\Delta_i} \right)$ (for small n the regret is, of course, even smaller).

The explanation of the bound is as follows. If at some time-step t it holds that all confidence intervals contain the truth and the width of the confidence interval about i^* drops below ε , then by the condition in Equation (1) it holds that i^* is the optimistic arm within $\tilde{\Theta}_t$. In this case UCB-S

suffers no regret at this time-step. Since the number of samples of each sub-optimal arm grows at most logarithmically by the proof of Theorem 2, the number of samples of the best arm must grow linearly. Therefore the number of time-steps before best arm has been pulled $O(\varepsilon^{-2})$ times is also $O(\varepsilon^{-2})$. After this point the algorithm suffers only a constant cumulative penalty for the possibility that the confidence intervals do not contain the truth, which is finite for suitably chosen values of α . Note that Agrawal et. al. [1] had essentially the same condition to achieve finite regret as (1), but specified to the case where Θ is finite.

An interesting question is raised by comparing the bound in Theorem 3 to those given by Bubeck et. al. [11] where if the expected return of the best arm is known and ε is a known bound on the minimum gap, then a regret bound of

$$O\left(\sum_{i \in A'} \left(\frac{\log\left(\frac{2\Delta_i}{\varepsilon}\right)}{\Delta_i} \left(1 + \log \log \frac{1}{\varepsilon}\right)\right)\right) \quad (2)$$

is achieved. If ε is close to Δ_i , then this bound is an improvement over the bound given by Theorem 3, although our theorem is more general. The improved UCB algorithm [7] enjoys a bound on the expected regret of $O(\sum_{i \in A'} \frac{1}{\Delta_i} \log n \Delta_i^2)$. If we follow the same reasoning as above we obtain a bound comparable to (2). Unfortunately though, the extension of the improved UCB algorithm to the structured setting is rather challenging with the main obstruction being the extreme growth of the phases used by improved UCB. Refining the phases leads to super-logarithmic regret, a problem we ultimately failed to resolve. Nevertheless we feel that there is some hope of obtaining a bound like (2) in this setting.

Before the proofs of Theorems 2 and 3 we give some example structured bandits and indicate the regions where the conditions for Theorem 3 are (not) met. Areas where Theorem 3 can be applied to obtain finite regret are unshaded while those with logarithmic regret are shaded.

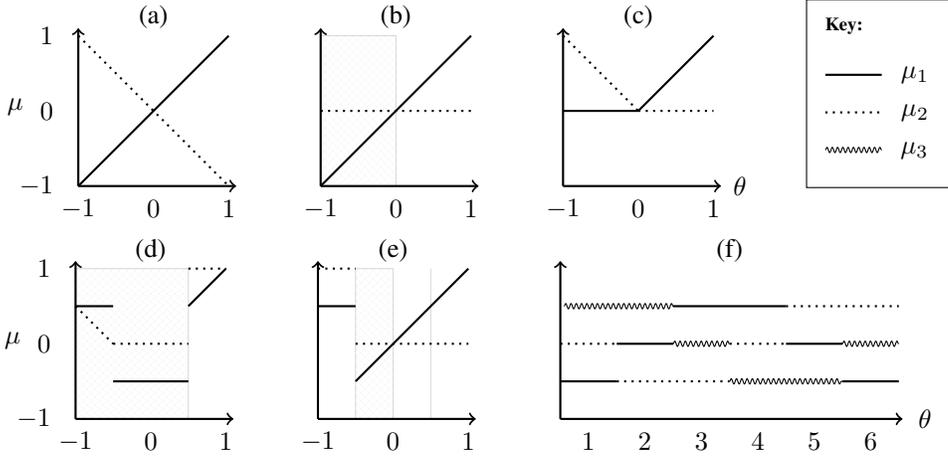


Figure 2: Examples

- (a) The conditions for Theorem 3 are met for all $\theta \neq 0$, but for $\theta = 0$ the regret strictly vanishes for all policies, which means that the regret is bounded by $\mathbb{E}R_n \in O(\mathbb{1}\{\theta^* \neq 0\} \frac{1}{|\theta^*|} \log \frac{1}{|\theta^*|})$.
- (b) Action 2 is uninformative and not globally optimal so Theorem 3 does not apply for $\theta < 1/2$ where this action is optimal. For $\theta > 0$ the optimal action is 1, when the conditions are met and finite regret is again achieved.

$$\mathbb{E}R_n \in O\left(\mathbb{1}\{\theta^* < 0\} \frac{\log n}{|\theta^*|} + \mathbb{1}\{\theta^* > 0\} \frac{\log \frac{1}{\theta^*}}{\theta^*}\right).$$

- (c) The conditions for Theorem 3 are again met for all non-zero θ^* , which leads as in (a) to a regret of $\mathbb{E}R_n \in O(\mathbb{1}\{\theta^* \neq 0\} \frac{1}{|\theta^*|} \log \frac{1}{|\theta^*|})$.

Examples (d) and (e) illustrate the potential complexity of the regions in which finite regret is possible. Note especially that in (e) the regret for $\theta^* = \frac{1}{2}$ is logarithmic in the horizon, but finite for θ^* arbitrarily close. Example (f) is a permutation bandit with 3 arms where it can be clearly seen that the conditions of Theorem 3 are satisfied.

5 Proof of Theorems 2 and 3

We start by bounding the probability that some mean does not lie inside the confidence set.

Lemma 5. $\mathbb{P}\{F_t = 1\} \leq 2Kt \exp(-\alpha \log(t))$ where

$$F_t = \mathbb{1} \left\{ \exists i : |\hat{\mu}_{i, T_i(t-1)} - \mu_i| \geq \sqrt{\frac{2\alpha\sigma^2 \log t}{T_i(t-1)}} \right\}.$$

Proof. We use the concentration guarantees:

$$\begin{aligned} \mathbb{P}\{F_t = 1\} &\stackrel{(a)}{=} \mathbb{P} \left\{ \exists i : |\mu_i(\theta^*) - \hat{\mu}_{i, T_i(t-1)}| \geq \sqrt{\frac{2\alpha\sigma^2 \log t}{T_i(t-1)}} \right\} \\ &\stackrel{(b)}{\leq} \sum_{i=1}^K \mathbb{P} \left\{ |\mu_i(\theta^*) - \hat{\mu}_{i, T_i(t-1)}| \geq \sqrt{\frac{2\alpha\sigma^2 \log t}{T_i(t-1)}} \right\} \\ &\stackrel{(c)}{\leq} \sum_{i=1}^K \sum_{s=1}^t \mathbb{P} \left\{ |\mu_i(\theta^*) - \hat{\mu}_{i, s}| \geq \sqrt{\frac{2\alpha\sigma^2 \log t}{s}} \right\} \stackrel{(d)}{\leq} \sum_{i=1}^K \sum_{s=1}^t 2 \exp(-\alpha \log t) \stackrel{(e)}{=} 2Kt^{1-\alpha} \end{aligned}$$

where (a) follows from the definition of F_t . (b) by the union bound. (c) also follows from the union bound and is the standard trick to deal with the random variable $T_i(t-1)$. (d) follows from the concentration inequalities for sub-gaussian random variables. (e) is trivial. \square

Proof of Theorem 2. Let i be an arm with $\Delta_i > 0$ and suppose that $I_t = i$. Then either F_t is true or

$$T_i(t-1) < \left\lceil \frac{8\sigma^2 \alpha \log n}{\Delta_i^2} \right\rceil =: u_i(n) \quad (3)$$

Note that if F_t does not hold then the true parameter lies within the confidence set, $\theta^* \in \tilde{\Theta}_t$. Suppose on the contrary that F_t and (3) are both false.

$$\begin{aligned} \max_{\tilde{\theta} \in \tilde{\Theta}_t} \mu_{i^*}(\tilde{\theta}) &\stackrel{(a)}{\geq} \mu^*(\theta^*) \stackrel{(b)}{=} \mu_i(\theta^*) + \Delta_i \stackrel{(c)}{>} \Delta_i + \hat{\mu}_{i, T_i(t-1)} - \sqrt{\frac{2\sigma^2 \alpha \log t}{T_i(t-1)}} \\ &\stackrel{(d)}{\geq} \hat{\mu}_{i, T_i(t-1)} + \sqrt{\frac{2\alpha\sigma^2 \log t}{T_i(t-1)}} \stackrel{(e)}{\geq} \max_{\tilde{\theta} \in \tilde{\Theta}_t} \mu_i(\tilde{\theta}), \end{aligned}$$

where (a) follows since $\theta^* \in \tilde{\Theta}_t$. (b) is the definition of the gap. (c) since F_t is false. (d) is true because (3) is false. Therefore arm i is not taken. We now bound the expected number of times that arm i is played within the first n time-steps by

$$\begin{aligned} \mathbb{E}T_i(n) &\stackrel{(a)}{=} \mathbb{E} \sum_{t=1}^n \mathbb{1}\{I_t = i\} \stackrel{(b)}{\leq} u_i(n) + \mathbb{E} \sum_{t=u_i+1}^n \mathbb{1}\{I_t = i \wedge (3) \text{ is false}\} \\ &\stackrel{(c)}{\leq} u_i(n) + \mathbb{E} \sum_{t=u_i+1}^n \mathbb{1}\{F_t = 1 \wedge I_t = i\} \end{aligned}$$

where (a) follows from the linearity of expectation and definition of $T_i(n)$. (b) by Equation (3) and the definition of $u_i(n)$ and expectation. (c) is true by recalling that playing arm i at time-step t implies that either F_t or (3) must be true. Therefore

$$\mathbb{E}R_n \leq \sum_{i \in A'} \Delta_i \left(u_i(n) + \mathbb{E} \sum_{t=u_i+1}^n \mathbb{1}\{F_t = 1 \wedge I_t = i\} \right) \leq \sum_{i \in A'} \Delta_i u_i(n) + \Delta_{\max} \mathbb{E} \sum_{t=1}^n \mathbb{1}\{F_t = 1\} \quad (4)$$

Bounding the second summation

$$\mathbb{E} \sum_{t=1}^n \mathbb{1}\{F_t = 1\} \stackrel{(a)}{=} \sum_{t=1}^n \mathbb{P}\{F_t = 1\} \stackrel{(b)}{\leq} \sum_{t=1}^n 2Kt^{1-\alpha} \stackrel{(c)}{\leq} \frac{2K(\alpha-1)}{\alpha-2}$$

where (a) follows by exchanging the expectation and sum and because the expectation of an indicator function can be written as the probability of the event. (b) by Lemma 5 and (c) is trivial. Substituting into (4) leads to

$$\mathbb{E}R_n \leq \frac{2\Delta_{\max}K(\alpha-1)}{\alpha-2} + \sum_{i \in A'} \frac{8\alpha\sigma^2 \log n}{\Delta_i} + \sum_i \Delta_i. \quad \square$$

Before the proof of Theorem 3 we need a high-probability bound on the number of times arm i is pulled, which is proven along the lines of similar results by [5].

Lemma 6. *Let $i \in A'$ be some sub-optimal arm. If $z > u_i(n)$, then $\mathbb{P}\{T_i(n) > z\} \leq \frac{2Kz^{2-\alpha}}{\alpha-2}$.*

Proof. As in the proof of Theorem 2, if $t \leq n$ and F_t is false and $T_i(t-1) > u_i(n) \geq u_i(t)$, then arm i is not chosen. Therefore

$$\mathbb{P}\{T_i(n) > z\} \leq \sum_{t=z+1}^n \mathbb{P}\{F_t = 1\} \stackrel{(a)}{\leq} \sum_{t=z+1}^n 2Kt^{1-\alpha} \stackrel{(b)}{\leq} 2K \int_z^n t^{1-\alpha} dt \stackrel{(c)}{\leq} \frac{2Kz^{2-\alpha}}{\alpha-2}$$

where (a) follows from Lemma 5 and (b) and (c) are trivial. \square

Lemma 7. *Assume the conditions of Theorem 3 and additionally that $T_{i^*}(t-1) \geq \left\lceil \frac{8\alpha\sigma^2 \log t}{\varepsilon^2} \right\rceil$ and F_t is false. Then $I_t = i^*$.*

Proof. Since F_t is false, for $\tilde{\theta} \in \tilde{\Theta}_t$ we have:

$$|\mu_{i^*}(\tilde{\theta}) - \mu_{i^*}(\theta^*)| \stackrel{(a)}{\leq} |\mu_{i^*}(\tilde{\theta}) - \hat{\mu}_{i^*, T_i(t-1)}| + |\hat{\mu}_{i^*, T_i(t-1)} - \mu_{i^*}(\theta^*)| \stackrel{(b)}{<} 2\sqrt{\frac{2\sigma^2\alpha \log t}{T_{i^*}(t-1)}} \stackrel{(c)}{\leq} \varepsilon$$

where (a) is the triangle inequality. (b) follows by the definition of the confidence interval and because F_t is false. (c) by the assumed lower bound on $T_{i^*}(t-1)$. Therefore by (1), for all $\tilde{\theta} \in \tilde{\Theta}_t$ it holds that the best arm is i^* . Finally, since F_t is false, $\theta^* \in \tilde{\Theta}_t$, which means that $\tilde{\Theta}_t \neq \emptyset$. Therefore $I_t = i^*$ as required. \square

Proof of Theorem 3. Let ω^* be some constant to be chosen later. Then the regret may be written as

$$\mathbb{E}R_n \leq \mathbb{E} \sum_{t=1}^{\omega^*} \sum_{i=1}^K \Delta_i \mathbb{1}\{I_t = i\} + \Delta_{\max} \mathbb{E} \sum_{t=\omega^*+1}^n \mathbb{1}\{I_t \neq i^*\}. \quad (5)$$

The first summation is bounded as in the proof of Theorem 2 by

$$\mathbb{E} \sum_{t=1}^{\omega^*} \sum_{i \in A} \Delta_i \mathbb{1}\{I_t = i\} \leq \sum_{i \in A'} \left(\Delta_i + \frac{8\alpha\sigma^2 \log \omega^*}{\Delta_i} \right) + \sum_{t=1}^{\omega^*} \mathbb{P}\{F_t = 1\}. \quad (6)$$

We now bound the second sum in (5) and choose ω^* . By Lemma 6, if $\frac{n}{K} > u_i(n)$, then

$$\mathbb{P}\left\{T_i(n) > \frac{n}{K}\right\} \leq \frac{2K}{\alpha-2} \left(\frac{K}{n}\right)^{\alpha-2}. \quad (7)$$

Suppose $t \geq \omega^* := \max\left\{\omega\left(\frac{8\sigma^2\alpha K}{\varepsilon^2}\right), \omega\left(\frac{8\sigma^2\alpha K}{\Delta_{\min}^2}\right)\right\}$. Then $\frac{t}{K} > u_i(t)$ for all $i \neq i^*$ and $\frac{t}{K} \geq \frac{8\sigma^2\alpha \log t}{\varepsilon^2}$. By the union bound

$$\mathbb{P}\left\{T_{i^*}(t) < \frac{8\sigma^2\alpha \log t}{\varepsilon^2}\right\} \stackrel{(a)}{\leq} \mathbb{P}\left\{T_{i^*}(t) < \frac{t}{K}\right\} \stackrel{(b)}{\leq} \mathbb{P}\left\{\exists i : T_i(t) > \frac{t}{K}\right\} \stackrel{(c)}{<} \frac{2K^2}{\alpha-2} \left(\frac{K}{t}\right)^{\alpha-2} \quad (8)$$

where (a) is true since $\frac{t}{K} \geq \frac{8\sigma^2\alpha \log t}{\varepsilon^2}$. (b) since $\sum_{i=1}^K T_i(t) = t$. (c) by the union bound and (7). Now if $T_i(t) \geq \frac{8\sigma^2\alpha \log t}{\varepsilon^2}$ and F_t is false, then the chosen arm is i^* . Therefore

$$\begin{aligned} \mathbb{E} \sum_{t=\omega^*+1}^n \mathbb{1}\{I_t \neq i^*\} &\leq \sum_{t=\omega^*+1}^n \mathbb{P}\{F_t = 1\} + \sum_{t=\omega^*+1}^n \mathbb{P}\left\{T_i(t-1) < \frac{8\sigma^2\alpha \log t}{\varepsilon^2}\right\} \\ &\stackrel{(a)}{\leq} \sum_{t=\omega^*+1}^n \mathbb{P}\{F_t = 1\} + \frac{2K^2}{\alpha-2} \sum_{t=\omega^*+1}^n \left(\frac{K}{t}\right)^{\alpha-2} \\ &\stackrel{(b)}{\leq} \sum_{t=\omega^*+1}^n \mathbb{P}\{F_t = 1\} + \frac{2K^2}{(\alpha-2)(\alpha-3)} \left(\frac{K}{\omega^*}\right)^{\alpha-3} \end{aligned} \quad (9)$$

where (a) follows from (8) and (b) by straight-forward calculus. Therefore by combining (5), (6) and (9) we obtain

$$\begin{aligned} \mathbb{E}R_n &\leq \sum_{i:\Delta_i>0} \Delta_i \left[\frac{8\sigma^2\alpha \log \omega^*}{\Delta_i^2} \right] + \frac{2\Delta_{\max}K^2}{(\alpha-2)(\alpha-3)} \left(\frac{K}{\omega^*}\right)^{\alpha-3} + \Delta_{\max} \sum_{t=1}^n \mathbb{P}\{F_t = 1\} \\ &\leq \sum_{i:\Delta_i>0} \Delta_i \left[\frac{8\sigma^2\alpha \log \omega^*}{\Delta_i^2} \right] + \frac{2\Delta_{\max}K^2}{(\alpha-2)(\alpha-3)} \left(\frac{K}{\omega^*}\right)^{\alpha-3} + \frac{2\Delta_{\max}K(\alpha-1)}{\alpha-2} \end{aligned}$$

Setting $\alpha = 4$ leads to $\mathbb{E}R_n \leq \sum_{i=1}^K \left(\frac{32\sigma^2 \log \omega^*}{\Delta_i} + \Delta_i \right) + 3\Delta_{\max}K + \frac{\Delta_{\max}K^3}{\omega^*}$. \square

6 Lower Bounds and Ambiguous Examples

We prove lower bounds for two illustrative examples of structured bandits. Some previous work is also relevant. The famous paper by Lai and Robbins [17] shows that the bound of Theorem 2 cannot in general be greatly improved. Many of the techniques here are borrowed from Bubeck et. al. [11]. Given a fixed algorithm and varying θ we denote the regret and expectation by $R_n(\theta)$ and \mathbb{E}_θ respectively. Returns are assumed to be sampled from a normal distribution with unit variance, so that $\sigma^2 = 1$.

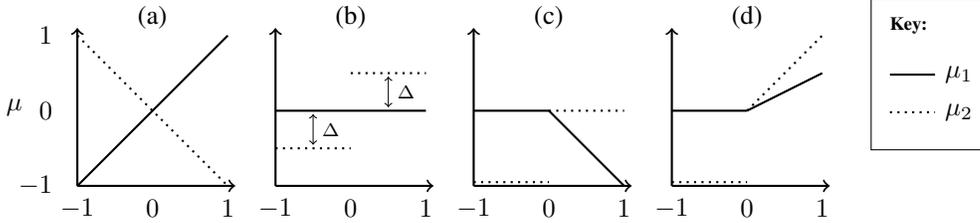


Figure 3: Counter-examples

Theorem 8. Given the structured bandit depicted in Figure 3.(a) or Figure 2.(c), then for all $\theta > 0$ and all algorithms the regret satisfies $\max\{\mathbb{E}_{-\theta}R_n(-\theta), \mathbb{E}_\theta R_n(\theta)\} \geq \frac{1}{8\theta}$ for sufficiently large n .

Proof. The proof uses the same technique as the proof of Theorem 5 in the paper by [11]. Fix an algorithm and let $\mathbb{P}_{\theta,t}$ be the probability measure on the space of outcomes up to time-step t under the bandit determined by parameter θ .

$$\begin{aligned} \mathbb{E}_{-\theta}R_n(-\theta) + \mathbb{E}_\theta R_n(\theta) &\stackrel{(a)}{=} 2\theta \left(\mathbb{E}_{-\theta} \sum_{t=1}^n \mathbb{1}\{I_t = 1\} + \mathbb{E}_\theta \sum_{t=1}^n \mathbb{1}\{I_t = 2\} \right) \\ &\stackrel{(b)}{=} 2\theta \sum_{t=1}^n (\mathbb{P}_{-\theta,t}\{I_t = 1\} + \mathbb{P}_{\theta,t}\{I_t = 2\}) \stackrel{(c)}{\geq} \theta \sum_{t=1}^n \exp(-\text{KL}(\mathbb{P}_{-\theta,t}, \mathbb{P}_{\theta,t})) \\ &\stackrel{(d)}{=} \theta \sum_{t=1}^n \exp(-4t\theta^2) \stackrel{(e)}{\geq} \frac{1}{8\theta} \end{aligned}$$

where (a) follows since $2|\theta|$ is the gap between the expected returns of the two arms given parameter θ and by the definition of the regret. (b) by replacing the expectations with probabilities. (c) follows from Lemma 4 by [11] where $\text{KL}(\mathbb{P}_{-\theta,t}, \mathbb{P}_{\theta,t})$ is the relative entropy between measures $\mathbb{P}_{-\theta,t}$ and $\mathbb{P}_{\theta,t}$. (d) is true by computing the relative entropy between two normals with unit variance and means separated by 2θ , which is $4\theta^2$. (e) holds for sufficiently large n . \square

Theorem 9. *Let $\Theta, \{\mu_1, \mu_2\}$ be a structured bandit where returns are sampled from a normal distribution with unit variance. Assume there exists a pair $\theta_1, \theta_2 \in \Theta$ and constant $\Delta > 0$ such that $\mu_1(\theta_1) = \mu_1(\theta_2)$ and $\mu_1(\theta_1) \geq \mu_2(\theta_1) + \Delta$ and $\mu_2(\theta_2) \geq \mu_1(\theta_2) + \Delta$. Then the following hold:*

- (1) $\mathbb{E}_{\theta_1} R_n(\theta_1) \geq \frac{1 + \log 2n\Delta^2}{8\Delta} - \frac{1}{2} \mathbb{E}_{\theta_2} R_n(\theta_2)$
- (2) $\mathbb{E}_{\theta_2} R_n(\theta_2) \geq \frac{n\Delta}{2} \exp(-4\mathbb{E}_{\theta_1} R_n(\theta_1)\Delta) - \mathbb{E}_{\theta_1} R_n(\theta_1)$

A natural example where the conditions are satisfied is depicted in Figure 3.(b) and by choosing $\theta_1 = -1, \theta_2 = 1$. We know from Theorem 3 that UCB-S enjoys finite regret of $\mathbb{E}_{\theta_2} R_n(\theta_2) \in O(\frac{1}{\Delta} \log \frac{1}{\Delta})$ and logarithmic regret $\mathbb{E}_{\theta_1} R_n(\theta_1) \in O(\frac{1}{\Delta} \log n)$. Part 1 of Theorem 9 shows that if we demand finite regret $\mathbb{E}_{\theta_2} R_n(\theta_2) \in O(1)$, then the regret $\mathbb{E}_{\theta_1} R_n(\theta_1)$ is necessarily logarithmic. On the other hand, part 2 shows that if we demand $\mathbb{E}_{\theta_1} R_n(\theta_1) \in o(\log(n))$, then the regret $\mathbb{E}_{\theta_2} R_n(\theta_2) \in \Omega(n)$. Therefore the trade-off made by UCB-S essentially cannot be improved.

Proof of Theorem 9. Again, we make use of the techniques of [11].

$$\begin{aligned} \mathbb{E}_{\theta_1} R_n(\theta_1) + \mathbb{E}_{\theta_2} R_n(\theta_2) &\stackrel{(a)}{\geq} \Delta (\mathbb{E}_{\theta_1} T_2(n) + \mathbb{E}_{\theta_2} T_1(n)) \stackrel{(b)}{\geq} \Delta \sum_{t=1}^n (\mathbb{P}_{\theta_1} \{I_t = 2\} + \mathbb{P}_{\theta_2} \{I_t = 1\}) \\ &\stackrel{(c)}{\geq} \frac{\Delta}{2} \sum_{t=1}^n \exp(-\text{KL}(\mathbb{P}_{\theta_1,t}, \mathbb{P}_{\theta_2,t})) \stackrel{(d)}{\geq} \frac{n\Delta}{2} \exp(-\text{KL}(\mathbb{P}_{\theta_1,n}, \mathbb{P}_{\theta_2,n})) \\ &\stackrel{(e)}{\geq} \frac{n\Delta}{2} \exp(-4\Delta^2 \mathbb{E}_{\theta_1} T_2(n)) \stackrel{(f)}{\geq} \frac{n\Delta}{2} \exp(-4\Delta \mathbb{E}_{\theta_1} R_n(\theta_1)) \end{aligned} \quad (10)$$

where (a) follows from the definition of the regret and the bandits used. (b) by the definition of $T_k(n)$. (c) by Lemma 4 of [11]. (d) since the relative entropy $\text{KL}(\mathbb{P}_{\theta_1,t}, \mathbb{P}_{\theta_2,t})$ is increasing with t . (e) By checking that $\text{KL}(\mathbb{P}_{\theta_1,n}, \mathbb{P}_{\theta_2,n}) = 4\Delta^2 \mathbb{E}_{\theta_1} T_2(n)$. (f) by substituting the definition of the regret. Now part 2 is completed by rearranging (10). For part 1 we also rearrange (10) to obtain

$$\mathbb{E}_{\theta_1} R_n(\theta_1) \geq \frac{n\Delta}{2} \exp(-4\Delta \mathbb{E}_{\theta_1} R_n(\theta_1)) - \mathbb{E}_{\theta_2} R_n(\theta_2)$$

Letting $x = \mathbb{E}_{\theta_1} R_n(\theta_1)$ and using the constraint above we obtain:

$$x \geq \frac{x}{2} + \frac{1}{2} \left(\frac{n\Delta}{2} \exp(-4\Delta x) - \mathbb{E}_{\theta_2} R_n(\theta_2) \right).$$

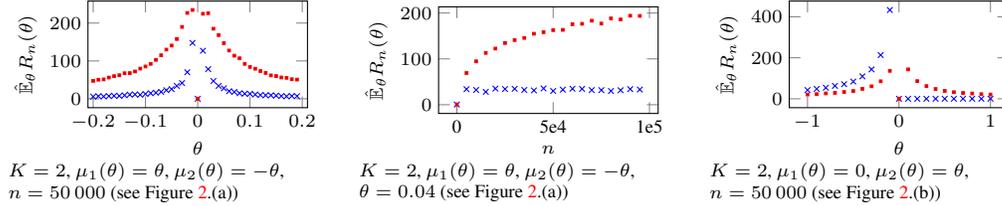
But by simple calculus the function on the right hand side is minimised for $x = \frac{1}{4\Delta} \log(2n\Delta^2)$, which leads to

$$\mathbb{E}_{\theta_1} R_n(\theta_1) \geq \frac{\log(2n\Delta^2)}{8\Delta} + \frac{1}{8\Delta} - \frac{1}{2} \mathbb{E}_{\theta_2} R_n(\theta_2). \quad \square$$

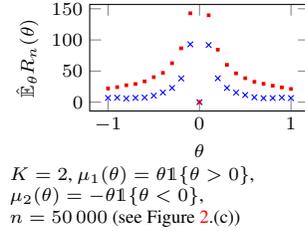
Discussion of Figure 3.(c/d). In both examples there is an ambiguous region for which the lower bound (Theorem 9) does not show that logarithmic regret is unavoidable, but where Theorem 3 cannot be applied to show that UCB-S achieves finite regret. We managed to show that finite regret is possible in both cases by using a different algorithm. For (c) we could construct a carefully tuned algorithm for which the regret was at most $O(1)$ if $\theta \leq 0$ and $O(\frac{1}{\theta} \log \log \frac{1}{\theta})$ otherwise. This result contradicts a claim by Bubeck et. al. [11, Thm. 8]. Additional discussion of the ambiguous case in general, as well as this specific example, may be found in the supplementary material. One observation is that unbridled optimism is the cause of the failure of UCB-S in these cases. This is illustrated by Figure 3.(d) with $\theta \leq 0$. No matter how narrow the confidence interval about μ_1 , if the second action has not been taken sufficiently often, then there will still be some belief that $\theta > 0$ is possible where the second action is optimistic, which leads to logarithmic regret. Adapting the algorithm to be slightly risk averse solves this problem.

7 Experiments

We tested Algorithm 1 on a selection of structured bandits depicted in Figure 2 and compared to UCB [6, 8]. Rewards were sampled from normal distributions with unit variances. For UCB we chose $\alpha = 2$, while we used the theoretically justified $\alpha = 4$ for Algorithm 1. All code is available in the supplementary material. Each data-point is the average of 500 independent samples with the blue crosses and red squares indicating the regret of UCB-S and UCB respectively.



The results show that Algorithm 1 typically out-performs regular UCB. The exception is the top right experiment where UCB performs slightly better for $\theta < 0$. This is not surprising, since in this case the structured version of UCB cannot exploit the additional structure and suffers due to worse constant factors. On the other hand, if $\theta > 0$, then UCB endures logarithmic regret and performs significantly worse than its structured counterpart. The superiority of Algorithm 1 would be accentuated in the top left and bottom right experiments by increasing the horizon.



8 Conclusion

The limitation of the new approach is that the proof techniques and algorithm are most suited to the case where the number of actions is relatively small. Generalising the techniques to large action spaces is therefore an important open problem. There is still a small gap between the upper and lower bounds, and the lower bounds have only been proven for special examples. Proving a general problem-dependent lower bound is an interesting question, but probably extremely challenging given the flexibility of the setting. We are also curious to know if there exist problems for which the optimal regret is somewhere between finite and logarithmic. Another question is that of how to define Thompson sampling for structured bandits. Thompson sampling has recently attracted a great deal of attention [13, 2, 14, 3, 9], but so far we are unable even to define an algorithm resembling Thompson sampling for the general structured bandit problem. Not only because we have not endowed Θ with a topology, but also because choosing a reasonable prior seems rather problem-dependent. An advantage of our approach is that we do not rely on knowing the distribution of the rewards while with one notable exception [9] this is required for Thompson sampling.

Acknowledgements. Tor Lattimore was supported by the Google Australia Fellowship for Machine Learning and the Alberta Innovates Technology Futures, NSERC. The majority of this work was completed while Rémi Munos was visiting Microsoft Research, New England. This research was partially supported by the European Community’s Seventh Framework Programme under grant agreements no. 270327 (project CompLACS).

References

- [1] Rajeev Agrawal, Demosthenis Teneketzis, and Venkatachalam Anantharam. Asymptotically efficient adaptive allocation schemes for controlled markov chains: Finite parameter space. *Automatic Control, IEEE Transactions on*, 34(12):1249–1259, 1989.
- [2] Shipra Agrawal and Navin Goyal. Analysis of Thompson sampling for the multi-armed bandit problem. In *In Proceedings of the 25th Annual Conference on Learning Theory*, 2012.
- [3] Shipra Agrawal and Navin Goyal. Further optimal regret bounds for thompson sampling. In *In Proceedings of the 16th International Conference on Artificial Intelligence and Statistics*, volume 31, pages 99–107, 2013.

- [4] Kareem Amin, Michael Kearns, and Umar Syed. Bandits, query learning, and the haystack dimension. *Journal of Machine Learning Research-Proceedings Track*, 19:87–106, 2011.
- [5] Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Variance estimates and exploration function in multi-armed bandit. Technical report, research report 07-31, Certis-Ecole des Ponts, 2007.
- [6] Peter Auer, Nicoló Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002.
- [7] Peter Auer and Ronald Ortner. UCB revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61(1-2):55–65, 2010.
- [8] Sébastien Bubeck and Nicolò Cesa-Bianchi. *Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems*. Foundations and Trends in Machine Learning. Now Publishers Incorporated, 2012.
- [9] Sébastien Bubeck and Che-Yu Liu. Prior-free and prior-dependent regret bounds for thompson sampling. In *Advances in Neural Information Processing Systems*, pages 638–646, 2013.
- [10] Sébastien Bubeck, Rémi Munos, Gilles Stoltz, and Csaba Szepesvári. Online optimization in X-armed bandits. In *NIPS*, pages 201–208, 2008.
- [11] Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet. Bounded regret in stochastic multi-armed bandits. In *In Proceedings of the 26th Annual Conference on Learning Theory*, 2013.
- [12] Todd L Graves and Tze Leung Lai. Asymptotically efficient adaptive choice of control laws in controlled Markov chains. *SIAM journal on control and optimization*, 35(3):715–743, 1997.
- [13] Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In *Algorithmic Learning Theory*, pages 199–213. Springer, 2012.
- [14] Nathaniel Korda, Emilie Kaufmann, and Rémi Munos. Thompson sampling for 1-dimensional exponential family bandits. In *Advances in Neural Information Processing Systems*, pages 1448–1456, 2013.
- [15] Tze Leung Lai and Herbert Robbins. Asymptotically optimal allocation of treatments in sequential experiments. In T. J. Santner and A. C. Tamhane, editors, *Design of Experiments: Ranking and Selection*, pages 127–142. 1984.
- [16] Tze Leung Lai and Herbert Robbins. Optimal sequential sampling from two populations. *Proceedings of the National Academy of Sciences*, 81(4):1284–1286, 1984.
- [17] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- [18] Adam J Mersereau, Paat Rusmevichientong, and John N Tsitsiklis. A structured multiarmed bandit problem and the greedy policy. *Automatic Control, IEEE Transactions on*, 54(12):2787–2802, 2009.
- [19] Dan Russo and Benjamin Van Roy. Eluder dimension and the sample complexity of optimistic exploration. In *Advances in Neural Information Processing Systems*, pages 2256–2264, 2013.

A Ambiguous Case

We assume for convenience that $K = 2$ and $\mu_1(\theta) \neq \mu_2(\theta)$ for all $\theta \in \Theta$. The second assumption is non-restrictive, since an algorithm cannot perform badly on the θ for which $\mu_1(\theta) = \mu_2(\theta)$, so we can simply remove these points from the parameter space. Now Θ can be partitioned into three sets according to whether or not finite regret is expected by Theorem 3, or impossible by Theorem 9.

$$\begin{aligned} \Theta_{\text{easy}} &:= \{\theta \in \Theta : \exists \varepsilon > 0 \text{ such that } |\mu_{i^*(\theta)}(\theta') - \mu_{i^*(\theta)}(\theta)| < \varepsilon \implies i^*(\theta') = i^*(\theta)\} \\ \Theta_{\text{hard}} &:= \{\theta \in \Theta : \exists \theta' \in \Theta \text{ such that } \mu_{i^*(\theta)}(\theta) = \mu_{i^*(\theta)}(\theta') \text{ and } i^*(\theta') \neq i^*(\theta)\} \\ \Theta_{\text{amb}} &:= \Theta - \Theta_{\text{easy}} - \Theta_{\text{hard}} \end{aligned}$$

The topic of this section is to study whether or not finite regret is possible on Θ_{amb} , and what sacrifices need to be made in order to achieve this. Some examples are given in Figure 4. Note that (a) was considered by Bubeck et. al. [11, Thm. 8] and will receive special attention here.

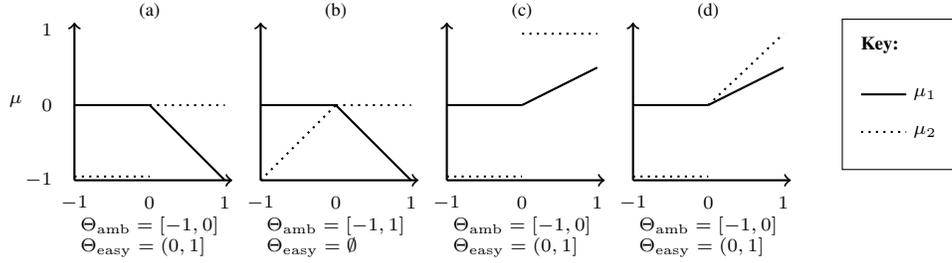


Figure 4: Ambiguous examples

The main theorem in this section shows that finite regret is indeed possible for many $\theta \in \Theta_{\text{amb}}$ without incurring significant additional regret for $\theta \in \Theta_{\text{easy}}$ and retaining logarithmic regret for $\theta \in \Theta_{\text{hard}}$. The following algorithm is similar to Algorithm 1, but favours actions which may be optimal for some plausible ambiguous θ . Theorems will be given subsequently, but proofs are omitted.

Algorithm 2

- 1: **Input:** functions $\mu_1, \dots, \mu_k : \Theta \rightarrow [0, 1]$, $\{\beta_t\}_{t=1}^\infty$
 - 2: $\kappa_1 = 0$
 - 3: **for** $t \in 1, \dots, \infty$ **do**
 - 4: Define confidence set: $\tilde{\Theta}_t \leftarrow \left\{ \tilde{\theta} : \forall i, \left| \mu_i(\tilde{\theta}) - \hat{\mu}_{i, T_i(t-1)} \right| < \sqrt{\frac{\alpha \sigma^2 \log t}{T_i(t-1)}} \right\}$
 - 5: **if** $\kappa_t = 0 \wedge \exists \tilde{\theta}_t \cap \Theta_{\text{amb}} \neq \emptyset$ **then**
 - 6: Choose $\theta \in \tilde{\Theta}_t \cap \Theta_{\text{amb}}$ arbitrarily and set $\kappa_t = i^*(\theta)$
 - 7: Choose $I_t = \arg \max_k \sup_{\theta \in \tilde{\Theta}_t} \mu_k(\theta) + \mathbb{1}\{k = \kappa_t\} \sqrt{\frac{\beta_t \log t}{T_k(t-1)}}$
 - 8: $\kappa_{t+1} \leftarrow \mathbb{1}\{I_t = \kappa_t\}$
-

Theorem 10. Suppose $K = 2$, $\theta \in \Theta$ and $i^* = 1$ and $\beta_t = \log \log t$ and $\Delta := \mu_1(\theta) - \mu_2(\theta)$. Then Algorithm 2 satisfies:

1. $\limsup_{n \rightarrow \infty} \mathbb{E}_\theta R_n(\theta) / \log n < \infty$.
2. If $\theta \in \Theta_{\text{easy}}$, then $\mathbb{E}_\theta R_n(\theta) \in O\left(\frac{1}{\Delta} \left(\log \frac{1}{\Delta}\right) \left(\log \log \frac{1}{\Delta}\right)\right)$.
3. If θ is such that

$$\lim_{\delta \rightarrow 0} \sup_{\theta' : |\mu_1(\theta) - \mu_1(\theta')| < \delta} \frac{\mu_2(\theta') - \mu_1(\theta')}{|\mu_1(\theta) - \mu_1(\theta')|} < \infty, \quad (11)$$

then $\lim_{n \rightarrow \infty} \mathbb{E}_\theta R_n(\theta) < \infty$.

Remark 11. The condition (11) not satisfied for $\theta \in \Theta_{\text{hard}}$, since in this case there exists some θ' with $\mu_1(\theta) = \mu_1(\theta')$ but where $\mu_2(\theta') - \mu_1(\theta') > 0$. The condition may not be satisfied even for $\theta \in \Theta_{\text{amb}}$. See, for example, Figure 4.(c). The condition is satisfied for all other ambiguous θ for the problems shown in Figure 4.(a,b,d) where the risk $\mu_2(\theta') - \mu_1(\theta')$ decreases linearly as $\mu_1(\theta) - \mu_1(\theta')$ converges to zero.

The following theorem shows that you cannot get finite regret for the ambiguous case where (11) is not satisfied without making sacrifices in the easy case.

Theorem 12. Suppose $\theta \in \Theta_{\text{amb}}$ with $i^*(\theta) = 1$ and $\mathbb{E}_\theta R_n(\theta) \in O(1)$. Then there exists a constant $c > 0$ such that for each θ' with $i^*(\theta') = 2$ we have

$$\mathbb{E}_{\theta'} R_n(\theta') \geq c \frac{\mu_2(\theta') - \mu_1(\theta')}{(\mu_1(\theta) - \mu_1(\theta'))^2}.$$

Therefore if the condition (3) in the statement of Theorem 10 is not satisfied for some ambiguous θ , then we can construct a sequence $\{\theta'_i\}_{i=1}^\infty$ such that $\lim_{i \rightarrow \infty} \mu_1(\theta'_i) = \mu_1(\theta)$ and where

$$\lim_{i \rightarrow \infty} (\mu_2(\theta'_i) - \mu_1(\theta')) \mathbb{E}_{\theta'_i} R_n(\theta'_i) = \infty,$$

which means that the regret must grow faster than the inverse of the gap. The situation becomes worse the faster the quantity below diverges to infinity.

$$\sup_{\theta': |\mu_1(\theta) - \mu_1(\theta')| < \delta} \frac{\mu_2(\theta') - \mu_1(\theta)}{|\mu_1(\theta) - \mu_1(\theta')|}.$$

In summary, finite regret is often possible in the ambiguous case, but may lead to worse regret guarantees in the easy case. Ultimately we are not sure how to optimise these trade-offs and there are still many interesting unanswered questions.

Analysis of Figure 4.(a)

We now consider a case of special interest that was previously studied by Bubeck et. al. [11] and is depicted in Figure 4.(a). The structured bandit falls into the ambiguous case when $\theta \leq 0$, since no interval about $\mu_1(\theta) = 0$ is sufficient to rule out the possibility that the second action is in fact optimal. Nevertheless, using a carefully crafted algorithm we show that the optimal regret is smaller than one might expect. The new algorithm operates in phases, choosing each action a certain number of times. If all evidence points to the first action being best, then this is taken until its optimality is proven to be implausible, while otherwise the second action is taken. The algorithm is heavily biased towards choosing the first action where estimation is more challenging, and where the cost of an error tends to be smaller.

Algorithm 3

```

1:  $\alpha \leftarrow 5$ 
2: for  $\ell \in 2, \dots, \infty$  do // Iterate over phases
3:    $n_{1,\ell} = 2^\ell$  and  $n_{2,\ell} = \ell^2$ 
4:   Choose each arm  $k \in \{1, 2\}$  exactly  $n_{k,\ell}$  times and let  $\hat{\mu}_{k,\ell,n_{k,\ell}}$  be the average return
5:    $s \leftarrow 0$ 
6:   if  $\hat{\mu}_{1,\ell,n_{1,\ell}} \geq -\sqrt{\frac{\alpha}{n_{1,\ell}} \log \log n_{1,\ell}}$  and  $\hat{\mu}_{2,\ell,n_{2,\ell}} < -1/2$  then
7:     while  $\hat{\mu}_{1,\ell,n_{1,\ell}+s} \geq -\sqrt{\frac{\alpha \log \log(n_{1,\ell}+s)}{n_{1,\ell}+s}}$  do
8:       Choose action 1 and  $s \leftarrow s + 1$  and  $\hat{\mu}_{1,\ell,n_{1,\ell}+s}$  is average return of arm 1 this phase
9:   else
10:    while  $\hat{\mu}_{2,\ell,n_{2,\ell}+s} \geq -\frac{1}{2}$  do
11:      Choose action 2 and  $s \leftarrow s + 1$  and  $\hat{\mu}_{2,\ell,n_{2,\ell}+s}$  is average return of arm 2 this phase

```

Theorem 13. Let $\Theta = [-1, 1]$ and $\mu_1(0) = -\theta \mathbb{1}\{\theta > 0\}$ and $\mu_2(\theta) = -\mathbb{1}\{\theta \leq 0\}$. Assume returns are normally distributed with unit variance. Then Algorithm 3 suffers regret bounded by

$$\mathbb{E}_\theta R_n(\theta) \in \begin{cases} O\left(\frac{1}{\theta} \log \log \frac{1}{\theta}\right) & \text{if } \theta > 0 \\ O(1) & \text{otherwise.} \end{cases}$$

Remark 14. Theorem 13 contradicts a result by Bubeck et. al. [11, Thm. 8], which states that for any algorithm

$$\max \left\{ \mathbb{E}_0 R_n(0), \sup_{\theta > 0} \theta \cdot \mathbb{E}_\theta R_n(\theta) \right\} \in \Omega(\log n).$$

But by Theorem 13 there exists an algorithm for which

$$\max \left\{ \mathbb{E}_0 R_n(0), \sup_{\theta > 0} \theta \cdot \mathbb{E}_\theta R_n(\theta) \right\} \in O \left(\sup_{\theta > 0} \min \left\{ \theta^2 n, \log \log \frac{1}{\theta} \right\} \right) = O(\log \log n).$$

We are currently unsure whether or not the dependence on $\log \log \frac{1}{\theta}$ can be dropped from the bound given in Theorem 13. Note that Theorem 3 cannot be applied when $\theta = 0$, so Algorithm 1 suffers logarithmic regret in this case. Algorithm 3 is carefully tuned and exploits the asymmetry in the problem. It is possible that the result of Bubeck et. al. can be saved in spirit by using the symmetric structured bandit depicted in Figure 4.(b). This would still only give a worst-case bound and does not imply that finite problem-dependent regret is impossible.

Proof of Theorem 13. It is enough to consider only $\theta \in [0, 1]$, since the returns on the arms is constant for $\theta \in [-1, 0]$. We let L be the number phases (times that the outer loop is executed) and T_ℓ be the number of times the sub-optimal action is taken in the ℓ th phase. Recall that $\hat{\mu}_{k,\ell,t}$ denotes the empirical estimate of μ based on t samples taken in the ℓ th phase.

Step 1: Decomposing the regret

The regret is decomposed:

$$\begin{aligned}
(\theta = 0) : \quad \mathbb{E}_0 R_n(0) &= \mathbb{E}_0 \sum_{\ell=0}^L T_\ell = \sum_{\ell=0}^{\infty} \mathbb{P}_0 \{L \geq \ell\} \mathbb{E}_0 [T_\ell | L \geq \ell] \\
(\theta > 0) : \quad \mathbb{E}_\theta R_n(\theta) &= \theta \mathbb{E}_\theta \sum_{\ell=0}^L T_\ell = \theta \sum_{\ell=0}^{\infty} \mathbb{P}_\theta \{L \geq \ell\} \mathbb{E}_\theta [T_\ell | L \geq \ell]
\end{aligned}$$

Step 2: Bounding $\mathbb{E}_\theta [T_\ell | L \geq \ell]$

We need to consider the cases when $\theta = 0$ and $\theta > 0$ separately. If $s \geq 1$, then

$$\begin{aligned}
\mathbb{P}_0 \{T_\ell \geq n_{2,\ell} + s | L \geq \ell\} &\stackrel{(a)}{\leq} \mathbb{P}_0 \left\{ \hat{\mu}_{2,\ell,n_{2,\ell}+s-1} \geq -\frac{1}{2} \right\} \stackrel{(b)}{=} \mathbb{P}_0 \left\{ \hat{\mu}_{2,\ell,n_{2,\ell}+s-1} - \mu_2(0) \geq \frac{1}{2} \right\} \\
&\stackrel{(c)}{\leq} \exp \left(-\frac{1}{2} (n_{2,\ell} + s - 1) \right) \stackrel{(d)}{\leq} \exp \left(-\frac{s}{2} \right),
\end{aligned}$$

where (a) follows since if the second action is chosen more than $n_{2,\ell}$ times in the ℓ th phase, then that phase ends when $\hat{\mu}_{2,\ell,t} < -\frac{1}{2}$, (b) by noting that $\mu_2(0) = -1$, (c) follows from the standard concentration inequality and the fact that unit variance is assumed, (d) since $n_{2,\ell} \geq 1$. Therefore by Lemma 17 we have that $\mathbb{E}_0 [T_\ell | L \geq \ell] \leq n_{2,\ell} + 2e^{1/2}$. Now assume $\theta > 0$ and define

$$\omega_2(x) = \min \{z : y \geq x \log \log y, \forall y \geq z\},$$

which satisfies $\omega_2(x) \in O(x \log \log x)$. If $n_{1,\ell} + s - 1 \geq \omega_2 \left(\frac{4\alpha}{\theta^2} \right)$, then

$$\begin{aligned}
\mathbb{P}_\theta \{T_\ell \geq n_{1,\ell} + s | L \geq \ell\} &\stackrel{(a)}{\leq} \mathbb{P}_\theta \left\{ \hat{\mu}_{1,\ell,n_{1,\ell}+s-1} \geq -\sqrt{\frac{\alpha}{n_{1,\ell} + s - 1} \log \log (n_{1,\ell} + s - 1)} \right\} \\
&\stackrel{(b)}{=} \mathbb{P}_\theta \left\{ \hat{\mu}_{1,\ell,n_{1,\ell}+s-1} - \mu_1(\theta) \geq \theta - \sqrt{\frac{\alpha}{n_{1,\ell} + s - 1} \log \log (n_{1,\ell} + s - 1)} \right\} \\
&\stackrel{(c)}{\leq} \mathbb{P}_\theta \left\{ \hat{\mu}_{1,\ell,n_{1,\ell}+s-1} - \mu_1(\theta) \geq \theta/2 \right\} \stackrel{(d)}{\leq} \exp \left(-\frac{\theta^2}{8} (n_{1,\ell} + s - 1) \right) \stackrel{(e)}{\leq} \exp \left(-\frac{\theta^2}{8} s \right),
\end{aligned}$$

where (a) follows since if the first arm (which is now sub-optimal) is chosen more than $n_{1,\ell}$ times, then the phase ends if $\hat{\mu}_{1,\ell,t}$ drops below the confidence interval. (b) since $\mu_1(\theta) = -\theta$. (c) since $n_{1,\ell} + s - 1 \geq \omega_2 \left(\frac{4\alpha}{\theta^2} \right)$. (d) by the usual concentration inequality and (e) since $n_{1,\ell} \geq 1$. Another application of Lemma 17 yields

$$\mathbb{E}_\theta [T_\ell | L \geq \ell] \leq \max \left\{ n_{1,\ell}, \omega_2 \left(\frac{4\alpha}{\theta^2} \right) \right\} + \frac{8e^{\theta^2/8}}{\theta^2},$$

where the max appears because we demanded that $n_{1,\ell} + s - 1 \geq \omega_2 \left(\frac{4\alpha}{\theta^2} \right)$ and since at the start of each phase the first action is taken at least $n_{1,\ell}$ times before the phase can end.

Bounding the number of phases

Again we consider the cases when $\theta = 0$ and $\theta \geq 0$ separately.

$$\begin{aligned}
\mathbb{P}_0 \{L > \ell\} &\stackrel{(a)}{\leq} \mathbb{P}_0 \left\{ \hat{\mu}_{2,\ell,n_{2,\ell}} \geq -\frac{1}{2} \vee \exists s : \hat{\mu}_{1,\ell,n_{1,\ell}+s} \leq -\sqrt{\frac{\alpha}{n_{1,\ell} + s} \log \log (n_{1,\ell} + s)} \right\} \\
&\stackrel{(b)}{\leq} \mathbb{P}_0 \left\{ \hat{\mu}_{2,\ell,n_{2,\ell}} \geq -\frac{1}{2} \right\} + \mathbb{P}_0 \left\{ \exists s : \hat{\mu}_{1,\ell,n_{1,\ell}+s} \leq -\sqrt{\frac{\alpha}{n_{1,\ell} + s} \log \log (n_{1,\ell} + s)} \right\} \\
&\stackrel{(c)}{\leq} \exp \left(-\frac{n_{2,\ell}}{8} \right) + \mathbb{P}_0 \left\{ \exists s : \hat{\mu}_{1,\ell,n_{1,\ell}+s} \leq -\sqrt{\frac{\alpha}{n_{1,\ell} + s} \log \log (n_{1,\ell} + s)} \right\}, \quad (12)
\end{aligned}$$

where (a) is true since the ℓ th phase will not end if $\hat{\mu}_{2,\ell,n_{2,\ell}} < -1/2$ and if $\hat{\mu}_{1,\ell,t}$ never drops below the confidence interval. (b) follows from the union bound and (c) by the concentration inequality. The second term is bounded using the maximal inequality and the peeling technique.

$$\begin{aligned}
& \mathbb{P}_0 \left\{ \exists s : \hat{\mu}_{1,\ell,n_{1,\ell}+s} \leq -\sqrt{\frac{\alpha}{n_{1,\ell}+s} \log \log(n_{1,\ell}+s)} \right\} \\
& \stackrel{(a)}{\leq} \sum_{k=0}^{\infty} \mathbb{P}_0 \left\{ \exists t : 2^k n_{1,\ell} \leq t \leq 2^{k+1} n_{1,\ell} \wedge \hat{\mu}_{1,\ell,t} \leq -\sqrt{\frac{\alpha}{t} \log \log t} \right\} \\
& \stackrel{(b)}{\leq} \sum_{k=0}^{\infty} \mathbb{P}_0 \left\{ \exists t \leq 2^{k+1} n_{1,\ell} : \hat{\mu}_{1,\ell,t} \leq -\sqrt{\frac{\alpha}{n_{1,\ell} 2^k} \log \log 2^k n_{1,\ell}} \right\} \\
& \stackrel{(c)}{\leq} \sum_{k=0}^{\infty} \exp(-\alpha \log \log(2^k n_{1,\ell})) \stackrel{(d)}{=} \sum_{k=0}^{\infty} \left(\frac{1}{\log 2^k + \log n_{1,\ell}} \right)^\alpha \stackrel{(e)}{\leq} \frac{2}{\log 2} \left(\frac{1}{\ell \log 2} \right)^{\alpha-1}
\end{aligned}$$

where (a) follows by the union bound, (b) by bounding t in the interval $2^k n_{1,\ell} \leq t \leq 2^{k+1} n_{1,\ell}$. (c) follows from the maximal inequality. (d) is trivial while (e) follows by approximating the sum by an integral. By combining with (12) we obtain

$$\mathbb{P}_0 \{L > \ell\} \leq \exp\left(-\frac{n_{2,\ell}}{8}\right) + \frac{2}{\log 2} \left(\frac{1}{\ell \log 2}\right)^{\alpha-1} = \exp\left(-\frac{\ell}{8}\right) + \frac{2}{\log 2} \left(\frac{1}{\ell \log 2}\right)^{\alpha-1}.$$

More straight-forwardly, if $\theta > 0$, then

$$\mathbb{P}_\theta \{L > \ell\} \leq \mathbb{P}_\theta \left\{ \exists s : \hat{\mu}_{2,n_{2,\ell}+s} < -\frac{1}{2} \right\} \leq 5 \exp\left(-\frac{n_{2,\ell}}{16}\right),$$

where in the last inequality we used Lemma 16 and naive bounding.

Putting it together

We now combine the results of the previous components to obtain the required bound on the regret. Recall that $\alpha = 5$.

$$\begin{aligned}
(\theta = 0) : \quad \mathbb{E}_0 R_n(0) &= \mathbb{E}_\theta \sum_{\ell=2}^{\infty} T_\ell = \sum_{\ell=2}^{\infty} \mathbb{P}_0 \{L \geq \ell\} \mathbb{E}_0 [T_\ell | L \geq \ell] \\
&\leq \sum_{\ell=2}^{\infty} \left(\exp\left(-\frac{n_{2,\ell}}{8}\right) + \frac{2}{\log 2} \left(\frac{1}{\ell \log 2}\right)^{\alpha-1} \right) (n_{2,\ell} + 2e^{1/2}) \\
&= \sum_{\ell=2}^{\infty} \left(\exp\left(-\frac{\ell}{8}\right) + \frac{2}{\log 2} \left(\frac{1}{\ell \log 2}\right)^{\alpha-1} \right) (\ell^2 + 2e^{1/2}) \in O(1) \\
(\theta > 0) : \quad \mathbb{E}_\theta R_n(\theta) &= \theta \mathbb{E}_\theta \sum_{\ell=2}^{\infty} T_\ell = \theta \sum_{\ell=1}^{\infty} \mathbb{P} \{L \geq \ell\} \mathbb{E}[T_\ell | L \geq \ell] \\
&\leq 5\theta \sum_{\ell=2}^{\infty} \exp\left(-\frac{n_{2,\ell}}{16}\right) \left(\max \left\{ n_{1,\ell}, \omega_2 \left(\frac{4\alpha}{\theta^2} \right) \right\} + \frac{8e^{\theta^2/8}}{\theta^2} \right) \\
&\in O\left(\theta \cdot \omega_2 \left(\frac{\alpha}{\theta^2} \right)\right) = O\left(\frac{1}{\theta} \log \log \frac{1}{\theta}\right). \quad \square
\end{aligned}$$

B Technical Lemmas

Lemma 15. Define functions ω and ω_2 by

$$\begin{aligned}
\omega(x) &:= \min \{z > 1 : y \geq x \log y, \forall y \geq z\} \\
\omega_2(x) &:= \min \{z > e : y \geq x \log \log y, \forall y \geq z\}.
\end{aligned}$$

Then $\omega(x) \in O(x \log x)$ and $\omega_2(x) \in O(x \log \log x)$.

Lemma 16. Let $\{X_i\}_{i=1}^\infty$ be sampled from some sub-gaussian distributed arm with mean μ and unit sub-gaussian constant. Define $\hat{\mu}_t = \frac{1}{t} \sum_{s=1}^t X_s$. Then for $s \geq 6/\Delta^2$ we have

$$\mathbb{P}\{\exists t \geq s : \hat{\mu}_t - \mu \geq \Delta\} \leq p + \frac{1}{\log 2} \log \frac{1}{1-p}$$

where $p = \exp\left(-\frac{s\Delta^2}{4}\right)$.

Proof. We assume without loss of generality that $\mu = 0$ and use a peeling argument combined with Azuma's maximal inequality

$$\begin{aligned} \mathbb{P}\{\exists t > s : \hat{\mu}_t \geq \Delta\} &\stackrel{(a)}{=} \mathbb{P}\{\exists k \in \mathbb{N}, 2^k s \leq t < 2^{k+1} s : \hat{\mu}_t \geq \Delta\} \\ &\stackrel{(b)}{=} \mathbb{P}\{\exists k \in \mathbb{N}, 2^k s \leq t < 2^{k+1} s : t\hat{\mu}_t \geq t\Delta\} \\ &\stackrel{(c)}{\leq} \mathbb{P}\{\exists k \in \mathbb{N}, 2^k s \leq t < 2^{k+1} s : t\hat{\mu}_t \geq 2^k s\Delta\} \\ &\stackrel{(d)}{\leq} \sum_{k=0}^{\infty} \mathbb{P}\{\exists 2^k s \leq t < 2^{k+1} s : t\hat{\mu}_t \geq 2^k s\Delta/2\} \\ &\stackrel{(e)}{\leq} \sum_{k=0}^{\infty} \mathbb{P}\{\exists t < 2^{k+1} s : t\hat{\mu}_t \geq 2^k s\Delta/2\} \\ &\stackrel{(f)}{\leq} \sum_{k=0}^{\infty} \exp\left(-\frac{1}{2} \frac{(2^k s\Delta)^2}{2^{k+1} s}\right) \stackrel{(g)}{=} \sum_{k=0}^{\infty} \exp\left(-\frac{2^k s\Delta^2}{4}\right) \\ &\stackrel{(h)}{=} \sum_{k=0}^{\infty} \exp\left(-\frac{s\Delta^2}{4}\right)^{2^k} \stackrel{(i)}{\leq} p + \frac{1}{\log 2} \log \frac{1}{1-p} \end{aligned}$$

where (a) follows by splitting the sum over an exponential grid. (b) by comparing cumulative differences rather than the means. (c) since $t > 2^k s$. (d) by the union bound over all k . (e) follows by increasing the range. (f) by Azuma's maximal inequality. (g) and (h) are true by straight-forward arithmetic while (i) follows from Lemma 18. \square

Lemma 17. Suppose z is a positive random variable and for some $\alpha > 0$ it holds for all natural numbers k that $\mathbb{P}\{z \geq k\} \leq \exp(-k\alpha)$. Then $\mathbb{E}z \leq \frac{e^\alpha}{\alpha}$

Proof. Let $\delta \in (0, 1)$. Then

$$\begin{aligned} \mathbb{P}\left\{z \geq \log \frac{1}{\delta}\right\} &\leq \mathbb{P}\left\{z \geq \left\lfloor \log \frac{1}{\delta} \right\rfloor\right\} \leq \exp\left(-\alpha \left\lfloor \log \frac{1}{\delta} \right\rfloor\right) \\ &\leq e^\alpha \cdot \exp\left(-\alpha \log\left(\frac{1}{\delta}\right)\right) = e^\alpha \cdot \delta^\alpha. \end{aligned}$$

To complete the proof we use a standard identity to bound the expectation

$$\mathbb{E}z \leq \int_0^1 \frac{1}{\delta} \mathbb{P}\left\{z \geq \log \frac{1}{\delta}\right\} d\delta \leq e^\alpha \int_0^1 \delta^{\alpha-1} d\delta = \frac{e^\alpha}{\alpha}. \quad \square$$

Lemma 18. Let $p \in (0, 7/10)$. Then $\sum_{k=0}^{\infty} p^{2^k} \leq p + \frac{1}{\log 2} \log \frac{1}{1-p}$.

Proof. Splitting the sum and comparing to an integral yields:

$$\begin{aligned} \sum_{k=0}^{\infty} p^{2^k} &\stackrel{(a)}{=} p + \sum_{k=1}^{\infty} p^{2^k} \stackrel{(b)}{\leq} p + \int_0^{\infty} p^{2^k} dk \stackrel{(c)}{=} p + \frac{1}{\log 2} \int_1^{\infty} p^u / u du \\ &\stackrel{(d)}{\leq} p + \frac{1}{\log 2} \sum_{u=1}^{\infty} p^u / u \stackrel{(e)}{=} p + \frac{1}{\log 2} \log \frac{1}{1-p} \end{aligned}$$

where (a) follows by splitting the sum. (b) by noting that p^{2^k} is monotone decreasing and comparing to an integral. (c) by substituting $u = 2^k$. (d) by reverting back to a sum. (e) follows from a standard formula. \square

C Table of Notation

K	number of arms
Θ	parameter space
θ^*	unknown parameter $\theta^* \in \Theta$
I_t	arm played at time-step t
$T_i(n)$	number times arm i has been played after time-step n
$X_{i,s}$	sth reward obtained when playing arm i
Δ_i	gap between the means of the best arm and the i th arm
Δ_{\min}	minimum gap, $\Delta_{\min} := \min_{i:\Delta_i>0} \Delta_i$
Δ_{\max}	maximum gap, $\Delta_{\max} := \max_i \Delta_i$
A	set of arms $A := \{1, 2, \dots, K\}$
A'	set of suboptimal arms $A := \{i : \Delta_i > 0\}$
R_n	regret at time-step n given unknown true parameter θ^*
$R_n(\theta)$	regret at time-step n given parameter θ
$\mu_i(\theta)$	mean of arm i given θ
$\hat{\mu}_{i,s}$	empiric estimate of the mean of arm i after s plays
$\mu^*(\theta)$	maximum return at θ . $\mu^*(\theta) := \max_i \mu_i(\theta)$
i^*	optimal arm given θ^*
$i^*(\theta)$	optimal arm given θ
$\omega(x)$	minimum value y such that $z \geq x \log z$ for all $z \geq y$
$\omega_2(x)$	minimum value y such that $z \geq x \log \log z$ for all $z \geq y$
F_t	event that the true value of some mean is outside the confidence interval about the empiric estimate at time-step t
α	parameter controlling how exploring the algorithm UCB-S is
σ^2	known parameter controlling the tails of the distributions governing the return of the arms
$u_i(n)$	critical number of samples for arm i . $u_i(n) := \left\lceil \frac{8\sigma^2 \alpha \log n}{\Delta_i^2} \right\rceil$