
Crush Optimism with Pessimism: Structured Bandits Beyond Asymptotic Optimality

Kwang-Sung Jun
The University of Arizona
kjun@cs.arizona.edu

Chicheng Zhang
The University of Arizona
chichengz@cs.arizona.edu

Abstract

In this paper, we study stochastic structured bandits for minimizing regret. The fact that the popular optimistic algorithms do not achieve the asymptotic instance-dependent regret optimality (asymptotic optimality for short) has recently allured researchers. On the other hand, it is known that one can achieve a bounded regret (i.e., does not grow indefinitely with n) in certain instances. Unfortunately, existing asymptotically optimal algorithms rely on forced sampling that introduces an $\omega(1)$ term w.r.t. the time horizon n in their regret, failing to adapt to the “easiness” of the instance. In this paper, we focus on the finite hypothesis class and ask if one can achieve the asymptotic optimality while enjoying bounded regret whenever possible. We provide a positive answer by introducing a new algorithm called CRush Optimism with Pessimism (CROP) that eliminates optimistic hypotheses by pulling the informative arms indicated by a pessimistic hypothesis. Our finite-time analysis shows that CROP (i) achieves a constant-factor asymptotic optimality and, thanks to the forced-exploration-free design, (ii) adapts to bounded regret, and (iii) its regret bound scales not with the number of arms K but with an effective number of arms K_ψ that we introduce. We also discuss a problem class where CROP can be exponentially better than existing algorithms in *nonasymptotic* regimes. Finally, we observe that even a clairvoyant oracle who plays according to the asymptotically optimal arm pull scheme may suffer a linear worst-case regret, indicating that it may not be the end of optimism.

1 Introduction

We consider the stochastic structured multi-armed bandit problem with a fixed arm set. In this problem, we are given a known structure that encodes how mean rewards of the arms are inter-dependent. Specifically, the learner is given a space of arms \mathcal{A} and a space of hypotheses \mathcal{F} where each $f \in \mathcal{F}$ maps each arm $a \in \mathcal{A}$ to its mean reward $f(a)$. Define $[n] := \{1, \dots, n\}$. At each time step $t \in [n]$, the learner chooses an arm $a_t \in \mathcal{A}$ and observes a (stochastic) noisy version of its mean reward $f^*(a)$ where $f^* \in \mathcal{F}$ is the ground truth hypothesis determined before the game starts and not known to her. After n time steps, the learner’s performance is evaluated by her cumulative expected (pseudo-)regret:

$$\mathbb{E} \text{Reg}_n = \mathbb{E} \left[n \cdot \max_{a \in \mathcal{A}} f^*(a) - \sum_{t=1}^n f^*(a_t) \right]. \quad (1)$$

Minimizing this regret poses a well-known challenge in balancing between exploration and exploitation; we refer to Lattimore and Szepesvári [21] for the backgrounds on bandits. We define our problem precisely in Section 2.

Linear bandits, a special case of structured bandits, have gained popularity over the last decade with exciting applications (e.g., news recommendation) [6, 11, 1, 22, 9]. While these algorithms use the

celebrated optimistic approaches to obtain near-optimal *worst-case* regret bounds (i.e., $\tilde{O}(\sqrt{dn})$ where \tilde{O} hides logarithmic factors and d is the dimensionality of the model), Lattimore and Szepesvári [20] have pointed out that their *instance-dependent* regret is often far from achieving the asymptotic instance-dependent optimality (hereafter, asymptotic optimality). This observation has spurred a flurry of research activities in asymptotically optimal algorithms for structured bandits and beyond, including OSSB [10], OAM [16] for linear bandits, and DEL [23] for reinforcement learning, although structured bandits and their optimality have been studied earlier in more general settings [3, 14].

The asymptotically optimal regret in structured bandits is of order $c(f) \cdot \ln(n)$ for instance $f \in \mathcal{F}$ where $c(f)$ is characterized by the optimization problem in (2). Its solution $\gamma \in [0, \infty)^K$ represents the optimal allocation of the arm pulls over \mathcal{A} , and some arms may receive zero arm pulls; we call those with nonzero arm pulls the *informative arms*. On the other hand, it is well-known that structured bandits can admit bounded regret [8, 19, 5, 25, 15, 27]; i.e., $\limsup_{n \rightarrow \infty} \mathbb{E} \text{Reg}_n < \infty$. This is because the hypothesis space, which encodes the side information or constraints, can contain a hypothesis f whose best arm alone is informative enough so that exploration is not needed, asymptotically.

However, existing asymptotically optimal strategies such as OSSB [10] cannot achieve bounded regret by design. The closest one we know is OAM [16] that can have a sub-logarithmic regret bound. The main culprit is their forced sampling, a widely-used mechanism for asymptotic optimality in structured bandits [10, 16]. Forced sampling, though details vary, ensures that we pull each arm proportional to an increasing but *unbounded* function of the time horizon n , which necessarily forces a *non-finite* regret. Furthermore, they tend to introduce the dependence on the number of arms K in the regret unless a structure-specific sampling is performed, e.g., pulling a barycentric spanner in the linear structure [16].¹ While the dependence on K disappears as $n \rightarrow \infty$, researchers have reported that the lower-order terms do matter in practice [16]. Such a dependence also goes against the well-known merit of exploiting the structure that their regret guarantees can have a mild dependence on the number of arms or may not scale with the number of arms at all (e.g., the worst-case regret of linear bandits mentioned above). We discuss more related work in the appendix (found in our supplementary material) due to space constraints, though important papers are discussed and cited throughout.²

Towards adapting to the easiness of the instance while achieving the asymptotic optimality, we turn to the simple case of the finite hypothesis space (i.e., $|\mathcal{F}| < \infty$) and ask: *can we design an algorithm with a constant-factor asymptotic optimality while adapting to finite regret?* Our main contribution is to answer the question above in the affirmative by designing a new algorithm and analyzing its finite-time regret. Departing from the forced sampling, we take a fundamentally different approach, which we call CROP (CRush Optimism with Pessimism). In a nutshell, at each time step t , CROP maintains a confidence set $\mathcal{F}_t \subseteq \mathcal{F}$ designed to capture the ground truth hypothesis f^* and identifies two hypothesis sets: the optimistic set $\tilde{\mathcal{F}}_t$ and the pessimistic set $\bar{\mathcal{F}}_t$ (defined in Algorithm 1). The key idea is to first pick carefully a $\bar{f}_t \in \bar{\mathcal{F}}_t$ that we call “pessimism”, and then pull the *informative arms* indicated by \bar{f}_t . This, as we show, eliminates either the optimistic set $\tilde{\mathcal{F}}_t$ or the pessimism \bar{f}_t from the confidence set. Our analysis shows that repeating this process achieves the asymptotic optimality within a constant factor. Furthermore, our regret bound reduces to a finite quantity whenever the instance allows it and does not depend on the number of arms K in general; rather it depends on an effective number of arms K_ψ defined in (6). We elaborate more on CROP and the role of pessimism in Section 3. We present the main theoretical result in Section 4 and show a particular problem class where CROP’s regret bound can be exponentially better than that of forced-sampling-based ones. Our regret bound of CROP includes an interesting $\ln \ln(n)$ term. In Section 5, we show a lower bound result indicating that such a $\ln \ln(n)$ term is unavoidable in general.

¹Some algorithms like OSSB [10] parameterize the exploration rate as ϵ , introducing $\epsilon K g_n$ for some $g_n = \omega(1)$ in the regret bound. One may attempt to set $\epsilon = 1/K$ to remove the dependence, but there is another term K/ϵ in the bound (see [10, Appendix 2.3]). Above all, we believe the dependence on K has to appear somewhere in the regret if forced sampling is used.

²Concurrent studies by Degenne et al. [12] and Saber et al. [24] avoid forced sampling but still have an explicit linear dependence on K in the regret.

Finally, we conclude with discussions in Section 6 where we report a surprising finding that UCB can be in fact better than a clairvoyant oracle algorithm (that, of course, achieves the asymptotic optimality) in nonasymptotic regimes. We also show that such an oracle can suffer a linear worst-case regret under some families of problems including linear bandits, which we find to be disturbing, but this leaves numerous open problems.

2 Problem definition and preliminaries

In the structured multi-armed bandit problem, the learner is given a discrete arm space $\mathcal{A} = [K]$, and a finite hypothesis class $\mathcal{F} \subset (\mathcal{A} \rightarrow \mathbb{R})$ where we color definitions in blue, hereafter. There exists an unknown $f^* \in \mathcal{F}$ that is the ground truth mean reward function. Denote by n the time horizon of the problem. For every $f \in \mathcal{F}$, denote by $a^*(f) = \arg \max_{a \in \mathcal{A}} f(a)$ and $\mu^*(f) = \max_{a \in \mathcal{A}} f(a)$ the arm and the mean reward supported by f , respectively. We remark that the focus of our paper is not computational complexity but the achievable regret bounds. For ease of exposition, we make the unique best arm assumption as follows:³

Assumption 1 (Unique best arm). *For every $f \in \mathcal{F}$, there exists a unique best arm $a^*(f)$, i.e., $a^*(f)$ is singleton.*

For an arm a and a hypothesis f , denote by $\Delta_a(f) = \mu^*(f) - f(a)$ the gap between the arm a and the optimal arm, if the true reward function were f . Given a set of hypotheses \mathcal{G} , we denote by $a^*(\mathcal{G}) = \{a^*(f) : f \in \mathcal{G}\}$ and $\mu^*(\mathcal{G}) = \{\mu^*(f) : f \in \mathcal{G}\}$ the set of arms and mean rewards supported by \mathcal{G} respectively.

The learning protocol is as follows: for each round $t \in [n]$, the learner pulls an arm $a_t \in \mathcal{A}$ and then receives a reward $r_t = f^*(a_t) + \xi_t$ where ξ_t is an independent σ^2 -sub-Gaussian random variable. The performance of the learner is measured by its expected cumulative regret over n rounds defined in (1). Given an arm a and time step t , denote by $T_a(t) = \sum_{s=1}^t \mathbb{1}\{a_s = a\}$ the arm pull count of a up to round t . With this notation, $\mathbb{E} \text{Reg}_n = \sum_{a \in \mathcal{A}} \mathbb{E}[T_a(n)] \Delta_a(f^*)$.

Asymptotically optimal regret. Our aim is to achieve an asymptotic instance-dependent regret guarantee. Hereafter we abbreviate ‘asymptotic optimality’ to **AO**. Specifically, we would like to develop *uniformly good* algorithms, in that for any problem instance, the algorithm satisfies $\mathbb{E} \text{Reg}_n = o(n^p)$ for any $p > 0$ where the little-o here is w.r.t. n only. The regret lower bound of structured bandits is based on the *competing* class of functions $\mathcal{C}(f) = \{g : g(a^*(f)) = f(a^*(f)) \wedge a^*(g) \neq a^*(f)\}$. The class $\mathcal{C}(f)$ consists of hypotheses $g \in \mathcal{F}$ such that pulling arm $a^*(f)$ provides no statistical evidence to distinguish g from f . Thus, even if the learner is confident that f is the ground truth, she has to pull arms other than $a^*(f)$ to guard against the case where the true hypothesis is actually g (in which case she suffers a linear regret); see the example in Figure 1(a) where $\mathcal{C}(f_4) = \{f_1\}$. The lower bound precisely captures such a requirement as constraints in the following optimization problem:

$$c(f) := \min_{\gamma \in [0, \infty)^K : \gamma_{a^*(f)} = 0} \sum_a \gamma_a \Delta_a(f) \quad \text{s.t.} \quad \forall g \in \mathcal{C}(f), \quad \sum_a \gamma_a \cdot \text{KL}(f(a), g(a)) \geq 1. \quad (2)$$

where $\text{KL}(f(a), g(a))$ is the KL-divergence between the two reward distributions when the arm a is pulled under f and g respectively. For the discussion of optimality, we focus on Gaussian rewards with variance σ^2 , which means $\text{KL}(f(a), g(a)) = \frac{(f(a) - g(a))^2}{2\sigma^2}$, though our proposed algorithm has a regret guarantee for more generic sub-Gaussian rewards. We denote by $\gamma(f)$ the solution of (2). Then, $c(f) = \sum_{a \in \mathcal{A}} \gamma_a(f) \cdot \Delta_a(f)$. The intuition is that if one could play arms in proportion to $\gamma^* = \{\gamma_a(f^*)\}_{a \in \mathcal{A}}$, then, by the constraints of the optimization problem, she would have enough statistical power to distinguish f^* from all members of $\mathcal{C}(f^*)$; furthermore, γ^* is the most cost-efficient arm allocation due to the objective function. The value of γ^* can be viewed as the allocation that balances optimally between maximizing the *information gap* (i.e., the KL divergence in (2)) and minimizing the *reward gap* (i.e., $\Delta_a(f)$).

It is known from the celebrated works of Agrawal et al. [3] and Graves and Lai [14] that any uniformly good algorithm must have regret at least $(1 - o(1))c(f) \ln(n)$ for large enough n , under

³Our algorithms and theorems can be easily extended to the setting where optimal actions w.r.t. f can be non-unique. This requires us to redefine the equivalence relationship, which we omit for brevity.




Figure 1: (a) An example instance. (b) A diagram of various hypothesis classes w.r.t. the ground truth hypothesis f^* . Best viewed in colors.

environment with ground truth reward function $f^* = f$. In other words, if an algorithm has a regret of $(1 - \Omega(1))c(f) \ln n$ under the ground truth f , then for large enough n , its expected arm pull scheme $\gamma = \left(\frac{\mathbb{E}[T_a(n)]}{\ln n}\right)_{a \in \mathcal{A}}$ must violate the constraint in (2) for some $g \in \mathcal{C}(f)$, implying that the algorithm must not be a uniformly good algorithm (i.e., suffer a polynomial regret under g). They also show the lower bound is tight by developing algorithms with asymptotic regret bound of $(1 + o(1))c(f) \ln n$.

The oracle. The lower bound suggests that one should strive to ensure $\mathbb{E}[T_a(n)] \approx \gamma_a^* \ln(n)$. Indeed a clairvoyant oracle (the oracle, hereafter) who knows f^* would, at round t , pull the arm a such that $T_a(t-1) \leq \gamma_a^* \ln(t)$ if there exists such an arm (i.e., exploration), and otherwise pull the best arm (i.e., exploitation). The oracle will initially pull the informative arms only, but as t increases, exploitation will crowd out exploration. We believe mimicking the oracle is what most algorithms with AO are after. Particularly, the most common strategy is to replace γ^* with the Empirical Risk Minimizer (ERM) $\gamma(\widehat{f}_t)$ where $\widehat{f}_t \in \mathcal{F}$ is the one that best fits the observed rewards. Unlike supervised learning, however, the observed rewards are controlled by the algorithm itself, making the ERM brittle; i.e., the ERM may not converge to f^* . Thus, most studies employ a form of forced sampling to ensure that \widehat{f}_t converges to f^* so that $\gamma(\widehat{f}_t)$ converges to γ^* . As discussed before, this is precisely where the issues begin, and we will see that CROP avoids forced sampling and $\gamma(\widehat{f}_t)$ altogether.

Example: cheating code. We describe an example inspired by Amin et al. [4] when algorithms with AO provide an improvement over the popular optimistic algorithms. Let $K_0 \in \mathbb{N}_+$ and e_i be the i -th indicator vector. The idea is to first consider a hypothesis like $f = (1, 1 - \epsilon, 1 - \epsilon, \dots, 1 - \epsilon)$ and then add those hypotheses that copy f , pick one of its non-best arms, and replace its mean reward with $1 + \epsilon$. This results in total $K_0 - 1$ competing hypotheses. Specifically, let e_i be the i -th indicator vector and define $h(i, j) \in \mathbb{R}^{K_0}$ as follows: $\forall i \in [K_0], h(i, 0) = (1 - \epsilon)\mathbf{1} + \epsilon e_i$ and $\forall j \in [K_0] \setminus \{i\}, [h(i, j)]_k = \begin{cases} 1 + \epsilon & \text{if } k = j \\ h(i, 0) & \text{otherwise} \end{cases}$. Let $\mathcal{F}_0 = \{h(i, j) : i \in [K_0], j \in \{0, 1, \dots, K_0\} \setminus \{i\}\}$, $k = \lceil \log_2(K_0) \rceil$, and $\Lambda \in [0, 1/2]$. Finally, we define the ‘‘cheating code’’ class with $K = K_0 + k$ arms:

$$\mathcal{F}^{\text{code}} = \{(g_{1:K_0}, \Lambda \cdot b_{1:k}) \in \mathbb{R}^{K_0+k} : g \in \mathcal{F}_0, b \in \{0, 1\}^k : \text{binary representation of } a^*(g) - 1\},$$

which appends k ‘‘cheating arms’’ that tells us the index of the best arm. Let us fix $f^* \in \mathcal{F}^{\text{code}}$ such that $\mu^*(f^*) = 1$. Assume $\frac{1}{2\epsilon} > \frac{2}{\Lambda^2}$ so that the informative arms of f^* are the cheating arms (see the appendix for reasoning) where we color in green for emphasis, throughout. Let $\sigma^2 = 1$. For the instance f^* , an algorithm with a constant-factor AO would have regret $O\left(\frac{\log_2 K}{\Lambda^2} \ln(n)\right)$ (elaborated more in the appendix). In contrast, optimistic algorithms such as UCB [7] (i.e., run naively without using the structure) or UCB-S [19], would pull the arm \tilde{a}_t where

$$(\tilde{a}_t, \tilde{f}_t) = \arg \max_{a \in \mathcal{A}, f \in \mathcal{F}_t} f(a) \quad (3)$$

and \mathcal{F}_t is a confidence set designed to trap f^* with high probability. One can show that \tilde{a}_t is always one of the first K_0 arms and that their regret is $O\left(\frac{K}{\epsilon} \ln(n)\right)$, which can be much larger. In fact, the gap between the two bounds can be arbitrarily large as ϵ approaches to 0.

The anatomy of the function classes. There are function classes besides $\mathcal{C}(f)$ that will become useful in our study. We first define an equivalence relationship between hypotheses: we call $f \sim g$

Algorithm 1 CRush Optimism with Pessimism (CROP)

Require: The hypothesis class \mathcal{F} , parameters $z, \hat{z}, \alpha, \hat{\alpha} > 1$

- 1: **for** $t = 1, 2, \dots, n$ **do**
 - 2: Let $\mathcal{F}_t = \{f \in \mathcal{F} : L_{t-1}(f) - \min_{g \in \mathcal{F}} L_{t-1}(g) \leq \beta_t := 2\sigma^2 \ln(zt(\log_2 t)^2)\}$
 - 3: **if** $a^*(\mathcal{F}_t)$ is singleton **then**
 - 4: (**Exploit**) Pull the arm $a_t \in a^*(\mathcal{F}_t)$, observe the reward r_t .
 - 5: Continue to the next iteration.
 - 6: **end if**
 - 7: Let $\mathcal{B}_t = \{(a^*(f), \mu^*(f)) : f \in \mathcal{F}_t\}$ be the best arm candidate set.
 - 8: Find the optimistic arm, mean, and set:

$$(\tilde{a}_t, \tilde{\mu}_t) = \arg \max_{(a, \mu) \in \mathcal{B}_t} \mu, \quad \tilde{\mathcal{F}}_t = \mathcal{F}_t(\tilde{a}_t, \tilde{\mu}_t).$$
 - 9: Find the pessimistic arm, mean, set, and hypothesis:

$$(\bar{a}_t, \bar{\mu}_t) = \arg \min_{(a, \mu) \in \mathcal{B}_t: a = \bar{a}_t} \mu, \quad \bar{\mathcal{F}}_t = \mathcal{F}_t(\bar{a}_t, \bar{\mu}_t), \quad \bar{f}_t = \arg \min_{f \in \bar{\mathcal{F}}_t} L_{t-1}(f).$$
 - 10: Define $\hat{\mathcal{F}}_t = \{f \in \bar{\mathcal{F}}_t : L_{t-1}(f) - L_{t-1}(\bar{f}_t) \leq \hat{\beta}_t := 2\sigma^2 \ln(\hat{z}(\log_2(t))^{\hat{\alpha}})\}$. (let $\hat{\beta}_1 = \infty$)
 - 11: **if** $\exists f, g \in \hat{\mathcal{F}}_t$ s.t. $\gamma(f) \not\leq \gamma(g)$ **then**
 - 12: (**Conflict**) $\pi_t = \phi(\bar{f}_t)$. (see (5))
 - 13: **else if** $\gamma(\bar{f}_t)$ satisfies that $\forall f \in \tilde{\mathcal{F}}_t, \sum_a \gamma_a(\bar{f}_t) \frac{(\bar{f}_t(a) - f(a))^2}{2\sigma^2} \geq 1$, **then**
 - 14: (**Feasible**) $\pi_t = \gamma(\bar{f}_t)$.
 - 15: **else**
 - 16: (**Fallback**) $\pi_t = \psi(\bar{f}_t)$. (see (4))
 - 17: **end if**
 - 18: Pull arm $a_t = \arg \min_a \frac{T_a(t-1)}{\pi_{t,a}}$ (take $\frac{x}{0}$ with $x \geq 0$ as ∞ ; break ties arbitrarily), and then observe the reward r_t .
 - 19: **end for**
-

if $a^*(f) = a^*(g)$ and $\mu^*(f) = \mu^*(g)$; one can verify that it satisfies reflexiveness, symmetry, and transitivity, and induces a partition over \mathcal{F} . Given $f \in \mathcal{F}$, we denote by $\mathcal{E}(f)$ the equivalent class f belongs to and by $\mathcal{D}(f) = \{g : g(a^*(f)) \neq f(a^*(f))\}$ its *docile* class that can be easily distinguished from f as we describe later. One can show that for every $f \in \mathcal{F}$, the class \mathcal{F} is a disjoint union of $\mathcal{E}(f)$, $\mathcal{D}(f)$, and $\mathcal{C}(f)$. We also define $\mathcal{O}(f) = \{g : \mu^*(g) \geq \mu^*(f)\}$ (and $\mathcal{P}(f) = \{g : \mu^*(g) \leq \mu^*(f)\}$) as the set of hypotheses that support mean rewards that are not lower (and not higher) than $\mu^*(f)$ (respectively). We use shorthands $\mathcal{E}^* := \mathcal{E}(f^*)$ and $\mathcal{D}^*, \mathcal{C}^*, \mathcal{O}^*$, and \mathcal{P}^* defined similarly. We draw a Venn diagram of these classes in Figure 1(b) along with the example hypotheses in Figure 1(a); we recommend that the readers verify the example themselves to get familiar with these classes.

Bounded regret. When the ground truth f^* enjoys $\mathcal{C}^* = \emptyset$, then $c(f^*) = 0$ and the algorithm can achieve bounded regret, which is well-known as mentioned in our introduction. This is because, when f^* has no competing hypothesis, pulling the best arm $a^*(f^*)$ alone provides a nonzero statistical evidence that distinguishes f^* from $\mathcal{F} \setminus \mathcal{E}^* = \mathcal{D}^*$. That is, there is no need to explore as exploitation alone provides sufficient exploration.

3 Crush Optimism with Pessimism (CROP)

We now introduce our algorithm CROP. First, some definitions: for any \mathcal{G} , define $\mathcal{G}(a, \mu) = \{f \in \mathcal{G} : a^*(f) = a, \mu^*(f) = \mu\}$. Given a set of observations $\{(a_s, r_s) : s \in [t]\}$ up to time step t , and $f \in \mathcal{F}$, denote by $L_t(f) = \sum_{s=1}^t (f(a_s) - r_s)^2$ the cumulative squared loss of f up to time step t . We use this loss to construct a confidence set that captures the ground truth f^* , which is inspired by Agarwal et al. [2], but we extend theirs to allow sub-Gaussian rewards. The loss $L_t(f)$ gives a measure of goodness of fit of hypothesis class f , in that f^* is the Bayes optimal regressor that minimizes $\mathbb{E} L_t(f)$.

We describe CROP in Algorithm 1, where the parameters $\{\alpha, \hat{\alpha}\}$ are numerical constants and $\{z, \hat{z}\}$ should be set to $|\mathcal{F}|$ (precise defined in Theorem 1). CROP has four main branches: **Exploit**, **Conflict**, **Feasible**, and **Fallback**. Note that **Feasible** is the main insight of the algorithm that we focus first while **Conflict** deals with some difficult cases, which we describe the last.

Exploit. At every round t , CROP maintains a confidence set \mathcal{F}_t , the set of hypotheses f in \mathcal{F} that fits well with the data observed so far w.r.t. $L_{t-1}(f)$. This is designed so that the probability of failing to trap the ground truth hypothesis f^* is $O(\frac{1}{t^\alpha})$. We first check if $a^*(\mathcal{F}_t)$ is a singleton. If true, we pull the arm $a^*(\mathcal{F}_t)$ that is unanimously supported by all f in \mathcal{F}_t . Note that the equivalence relationship \sim induces a partition of \mathcal{F}_t . If we do not enter the exploit case, we select the equivalence class $\tilde{\mathcal{F}}_t$ that maximizes its shared supported mean reward; we call this the *optimistic set*. This is related to the celebrated “optimism in the face of uncertainty” (OFU) principle that pulls arm a_t by (3). In line 9, we deviate from the OFU and define the *pessimistic set* $\bar{\mathcal{F}}_t$, which is the equivalence class in \mathcal{F}_t that *minimizes* its shared supported mean reward $\mu^*(\bar{\mathcal{F}}_t)$ with a constraint that they support an action other than \tilde{a}_t . We then define \bar{f}_t , which we call *the pessimism*, as the Empirical Risk Minimizer (ERM) over $\bar{\mathcal{F}}_t$. Next, we compute $\hat{\mathcal{F}}_t$, a refined confidence set inside the pessimistic set $\bar{\mathcal{F}}_t$, and then test a condition to enter **Conflict**; we will discuss it later as mentioned above. For now, suppose that we did not enter **Conflict** and are ready to test the condition for **Feasible** (line 13).

Feasible. The condition in line 13 first computes $\gamma(\bar{f}_t)$ and then tests whether all the hypotheses in $\tilde{\mathcal{F}}_t$ satisfy the information constraint that takes the same form as those in the optimization problem for $c(\bar{f}_t)$. If this is true, then we set $\pi_t = \gamma(\bar{f}_t)$ and then move onto line 18 to choose which arm to pull. The intention here is to pull the arm that is most far away from the pull scheme of $\gamma(\bar{f}_t)$, which is often referred to as “tracking” [13]. Note that the arm \bar{a}_t is never pulled because $\gamma_{\bar{a}_t}(\bar{f}_t) = 0$.

Why the pessimism? To motivate the design choice of tracking the pessimism, consider the example hypothesis space \mathcal{H} in Figure 2. Suppose that at time t we have $\mathcal{F}_t = \mathcal{H} = \{f_1, f_2, f_3\}$. Which arms should we pull? The OFU tells us to pull the optimistic arm \tilde{a}_t as done in Lattimore and Munos [19], but it does not achieve the instance optimality. Another idea mentioned in Section 2 is find the ERM $\hat{f}_t = \arg \min_{f \in \mathcal{F}} L_{t-1}(f)$ and then pull the arms by tracking $\gamma(\hat{f}_t)$; i.e., $a_t = \arg \min_{a \neq a^*(\hat{f}_t)} T_a(t-1)/\gamma_a(\hat{f}_t)$. This is essentially the main idea of OSSB [10].⁴ OAM [16] also relies on the ERM \hat{f}_t , though they partly use the optimism. However, ERMs are brittle in bandits.

For example, when $f^* = f_3$, in earlier rounds, the ERM \hat{f}_t can be f_2 with nontrivial probability. Pulling the informative arm of f_2 , which is A4, eliminates f_1 but will not eliminate f_3 , and we get stuck at pulling A4 indefinitely. To avoid such a trap, researchers have introduced forced sampling.

What are the robust alternatives to the ERM? For now, suppose that f^* is always in the confidence set \mathcal{F}_t . Among $\{\gamma(f_1), \gamma(f_2), \gamma(f_3)\}$, which one should we track? We claim that we should follow the pessimism, which is f_3 in this case. Specifically, if $f^* = f_3$, we are lucky and following the pessimism will soon remove both f_1 and f_2 from \mathcal{F}_t . We then keep entering **Exploit** and pull the best arm A3 for a while. Note that f_1 or f_2 will come back to \mathcal{F}_t again as pulling A3 provides the same loss to every $f \in \mathcal{H}$ but the threshold β_t of the confidence set \mathcal{F}_t increases over time. In this case, the principle of pessimism will do the right thing, again.

What if f^* was actually f_2 ? Following the pessimism f_3 is not optimal, but it does eliminate f_3 from \mathcal{F}_t because f_2 appears in the constraint of the optimization (2); after the elimination, we have $\mathcal{F}_t = \{f_1, f_2\}$ and $\bar{f}_t = f_2$, so the pessimism is back in charge. In sum, the key observation is that the optimal pull scheme $\gamma(f)$ is designed to differentiate f from its *competing hypotheses* that support arms with higher mean rewards than that of f . Assuming the confidence set works properly, tracking the pessimism either does the right thing or, if \bar{f}_t is not the ground truth, removes \bar{f}_t from the confidence set (also the right thing to do). However, to make it work beyond this example, we need other mechanisms: **Fallback** and **Conflict**.

Arms	A1	A2	A3	A4	A5
f_1	1	.99	.98	0	0
f_2	.98	.99	.98	<u>.25</u>	0
f_3	.97	.97	.98	.25	<u>.25</u>
f_4 (optional)	.98	.99	.98	.25	<u>.50</u>

Figure 2: The “staircase” example. Define $\mathcal{H} = \{f_1, f_2, f_3\}$ and $\mathcal{H}^+ = \{f_1, f_2, f_3, f_4\}$. We boldface the best arm and underline the informative arms of each hypothesis.

⁴OSSB in fact does not find the ERM but rather uses the empirical means of the arms to solve the optimization problem (2), which can work for some problem families. Still, we believe extending OSSB to use the ERM with suitable loss function should achieve (near) asymptotic optimality for the finite \mathcal{F} .

Fallback. When the condition of **Feasible** is not satisfied, we know that the arm pull scheme $\gamma(\bar{f}_t)$ will not be sufficient to remove every $f \in \bar{\mathcal{F}}_t$ – or, it could even be impossible. Thus, we should not track $\gamma(\bar{f}_t)$. Instead, we design a different arm pull scheme $\psi(f)$ defined below so that tracking $\psi(\bar{f}_t)$ can remove all members of $\bar{\mathcal{F}}_t$ in a cost-efficient manner. With the notation $\Delta_{\min}(f) = \min_{a \neq a^*(f)} \Delta_{a^*(f)}(f)$,

$$\begin{aligned} \psi(f) &:= \arg \min_{\gamma \in [0, \infty)^K} \Delta_{\min}(f) \cdot \gamma_{a^*(f)} + \sum_{a \neq a^*(f)} \Delta_a(f) \cdot \gamma_a \\ \text{s.t. } &\forall g \in \mathcal{O}(f) \setminus \mathcal{E}(f) : \sum_a \gamma_a \frac{(f(a) - g(a))^2}{2\sigma^2} \geq 1 \\ &\gamma \geq \phi(f) \vee \gamma(f) \end{aligned} \quad (4)$$

where ϕ is defined in (5) and explained below and $x \geq y$ means $x_a \geq y_a, \forall a$. The constraints above now ensure that $\psi(\bar{f}_t)$ provides a sufficient arm pull scheme to eliminate $\bar{\mathcal{F}}_t$ even if the condition of **Feasible** is not satisfied. Another difference from $\gamma(f)$ is that $\gamma_{a^*(f)}$ can be nonzero, but we use $\Delta_{\min}(f)$ instead of $\Delta_{a^*(f)}(f) = 0$ to avoid $\gamma_{a^*(f)} = \infty$. That said, there are other design choices for $\psi(f)$, especially given that ψ appears only with the finite terms in the regret bound. We discuss more on the motivation and alternative designs of (4) in the appendix.

Conflict. This is an interesting case where the learner faces the challenge not in finding which arm is the best arm, but rather which *informative* arms and their pull scheme one should track. Specifically, consider the other example of \mathcal{H}^+ in Figure 2. Suppose at time t we have $\mathcal{F}_t = \{f_1, f_2, f_4\}$ and the ground truth is f_4 , which means $\mathcal{E}^* = \{f_2, f_4\}$. If CROP does not have the **Conflict** mechanism, it will use \bar{f}_t , the ERM among $\bar{\mathcal{F}}_t$, which can be either f_2 or f_4 . However, as explained before, ERMs are brittle; one can see that it can get stuck at tracking f_2 with nontrivial probability and pull less informative arms. Interestingly, this would not incur a linear regret. Rather, the regret would still be like $\ln(n)$ but with a suboptimal constant of $c(f_2)$ rather than $c(f_4)$; one can adjust our example to make this gap $c(f_2) - c(f_4)$ arbitrarily large, making it arbitrarily far from the AO. On the other hand, a closer look at f_2 and f_4 reveals that A5 is the only arm that can help distinguish f_2 from f_4 . One might attempt to change CROP so that it pulls A5 in such a case, which results in either removing f_4 from the confidence set if $f^* = f_2$ or removing f_2 if $f^* = f_4$. However, if $f^* = f_2$, this would introduce an extra $\ln(n)$ term in the regret bound since A5 is a noninformative arm, which again can lead to a suboptimal regret bound.

CROP resolves this issue by constructing a refined confidence set $\dot{\mathcal{F}}_t$ with a more aggressive failure rate of $1/\ln(t)$ rather than the usual $1/t$, and use this confidence set to weed out conflicting pull schemes. If the refined set $\dot{\mathcal{F}}_t$ still contains hypotheses with conflicting pull schemes, then CROP enters **Conflict** and computes a different allocation scheme:

$$\begin{aligned} \phi(f) &= \arg \min_{\gamma \in [0, \infty)^K : \gamma_{a^*(f)} = 0} \sum_a \Delta_a(f) \cdot \gamma_a \\ \text{s.t. } &\forall g \in \mathcal{E}(f) : \gamma(g) \not\subseteq \gamma(f), \quad \sum_a \gamma_a \frac{(f(a) - g(a))^2}{2\sigma^2} \geq 1 \end{aligned}, \quad (5)$$

where we use the convention $0 \infty = 0$. Consider \mathcal{H}^+ in Figure 2 with $\sigma^2 = 1$. Then, $\phi(f_2) = \phi(f_4) = (0, 0, 0, 0, \frac{2}{(\frac{2}{5})^2} = 8)$. Our regret analysis will show that the quantity $\phi(f)$ appears in the regret bound with $\ln(\ln(n))$ term only instead of $\ln(n)$, allowing us to achieve the AO within constant-factor.

4 Analysis

Before presenting our analysis, we define the effective number of arms K_ψ as the size of the union of the supports of $\psi(f)$ for all $f \in \mathcal{F}$:

$$K_\psi = \left| \left\{ a : \exists f \in \mathcal{F}, \psi_a(f) \neq 0 \right\} \right|. \quad (6)$$

Define $\phi(\mathcal{G}) = \left(\max_{f \in \mathcal{G}} \phi_a(f) \right)_{a \in \mathcal{A}}$ and $\psi(\mathcal{G})$ similarly. Let $\Delta_{\min} = \min_{f \in \mathcal{D}^*} \frac{|f(a^*) - f^*(a^*)|}{\sigma}$ the smallest information gap where $a^* := a^*(f^*)$. We use the shorthand $\Delta_{\max} := \max_a \Delta_a(f^*)$. We present our main theorem on the regret bound of CROP.

Theorem 1. Let $(\alpha, \hat{\alpha}, z, \hat{z}) = (2, 3, |\mathcal{F}|, |\mathcal{F}|)$. Suppose we run CROP with hypothesis class \mathcal{F} with the environment $f^* \in \mathcal{F}$. Then, CROP has the following anytime regret guarantee: $\forall n \geq 2$,

$$\mathbb{E} \text{Reg}_n \leq c_1 \cdot \left(P_1 \ln(n) + P_2 \ln(\ln(n)) + P_3 (\ln(|\mathcal{F}|) + \ln(Q_1)) + K_\psi \Delta_{\max} \right),$$

where c_1 is a numerical constant, and

$$P_1 = \sum_a \Delta_a \gamma_a^*, \quad P_2 = \sum_a \Delta_a \phi_a(\mathcal{E}^*), \quad P_3 = \sum_a \Delta_a \psi_a(\mathcal{F}), \quad \text{and} \quad Q_1 = \Lambda_{\min}^{-2} + K_\psi (1 + \max_i \psi_i(\mathcal{F})).$$

Furthermore, when $\gamma^* = 0$, we have $P_1 = P_2 = 0$, achieving a bounded regret.

Proof. The main proof is deferred to the appendix. One technical challenge is to deal with `Conflict` in CROP via our refined confidence set $\hat{\mathcal{F}}_t$. The failure rate of $\hat{\mathcal{F}}_t$ is set $\text{poly}(1/\ln(t))$ rather than the usual $\text{poly}(1/t)$.⁵ For example, there is an event where $\hat{\mathcal{F}}_t$ fails to capture f^* but f^* is still in \mathcal{F}_t , which would lead to a $\ln(n)$ regret; we manage to prove that this scenario contributes to an $O(1)$ term *in expectation* by showing that it happens with probability like $1/\ln(n)$ times (roughly speaking) using a technique that we call “regret peeling”. To bound other $O(1)$ terms that are attributed to the docile class \mathcal{D}^* , we borrow techniques from Lattimore and Munos [19]. \square

Our main theorem provide a sharp non-asymptotic instance-dependent regret guarantee. The leading term $O(\sum_a \Delta_a \gamma_a^* \log(n))$ implies that we achieve the AO up to a constant factor. The second term is of order $\ln \ln(n)$, which comes from our analysis on `Conflict`. The remaining terms are $O(1)$, which depends on properties of $\psi_a(\mathcal{F})$ and Λ_{\min} . Unlike many strategies that perform forced exploration on all arms [16, 10] to achieve the asymptotic optimality, our bound has no dependency on the number of arms K at all, even in the finite terms, but rather depends on the effective number of arms K_ψ .

Note that K -free regret bounds still happens with optimistic algorithms; e.g., in $\mathcal{F}^{\text{code}}$ (defined in Section 2), UCB depends on K_0 rather than K , and one can add arbitrarily many cheating arms to make $K \gg K_0$. Bounded regrets also have been shown via optimism [19, 15, 27], but they are far from the AO in general. Our novelty is to obtain instance optimality, remove the dependency on K , and achieve bounded regret whenever possible, simultaneously. We make more remarks on $\ln \ln(n)$ term and how one can get rid of $\ln(|\mathcal{F}|)$ and handle infinite hypothesis spaces in the appendix.

Example: Cheating code. Let $\mathcal{F} = \mathcal{F}^{\text{code}}$ and $\sigma^2 = 1$, and fix $f^* \in \mathcal{F}$ such that $\mu^*(f^*) = 1$. Assume $\frac{1}{2\epsilon} > \frac{2}{\Lambda^2}$. Then, one can show that $\gamma^* = \psi(\mathcal{F}) = (0, \dots, 0, \frac{2}{\Lambda^2}, \dots, \frac{2}{\Lambda^2})$, where the first K_0 coordinates are zeros, and $\phi(\mathcal{E}^*) = 0$. We also have $|\mathcal{F}| = K_0^2$, $K_\psi = \lceil \log_2(K_0) \rceil$, and $\Lambda_{\min} = \epsilon$. Then, using $K_0 \leq K$,

$$\mathbb{E} \text{Reg}_n = O\left(\frac{\ln(K)}{\Lambda^2} \ln\left(Kn \cdot \left(\frac{1}{\epsilon} + \frac{\ln(K)}{\Lambda^2}\right)\right) + \ln(K)\right),$$

which is $\approx \frac{\ln^2(K)}{\Lambda^2} \ln(\frac{n}{\epsilon})$ when taking the highest-order factors for each $(n, K, \epsilon, \Lambda)$. We speculate that $\ln(1/\epsilon)$ can be removed with a tighter analysis. We compare CROP to algorithms with AO that use forced sampling (FS in short). Say, during n rounds, FS pulls every arm $\ln \ln(n)$ times each, introducing a term $O(\epsilon K \ln \ln(n))$ in the regret, but let us ignore the $\ln \ln(\cdot)$ factor. For FS, the best regret bound one can hope for is $O(K\epsilon + \frac{\log_2(K)}{\Lambda^2} \ln(n))$. To satisfy the condition $\frac{1}{2\epsilon} > \frac{2}{\Lambda^2}$, set $\Lambda = 1/2$ and $\epsilon = 1/32$. Then, CROP’s regret is $O(\ln^2(K) \ln(n))$ whereas FS’s regret is $O(K + \ln(K) \ln(n))$. When $K \approx n$, FS has a linear regret whereas CROP has $\ln^3(n)$ regret. If $K = 2^d$ for some d , then the gap between the two becomes more dramatic: $O(2^d + d \ln(n))$ of FS vs $O(d^2 \ln(n))$ of CROP, an exponential gap in the nonasymptotic regime.

5 Lower bound: Necessity of the $\Omega(\ln \ln(n))$ term

One may wonder if the $\ln(\ln(n))$ term in our upper bound is necessary to achieve the asymptotic optimality up to constant factors. We show that there exist cases where such a dependence is indeed required. In fact, our lower bound statement is stronger; in a hypothesis class we construct for lower

⁵Similar aggressive definitions of confidence sets have also appeared in recent works for other purposes [e.g. 20].

bound, even polynomial-regret algorithms must also pull a non-informative arm at least $\ln \ln(n)$ times.

Theorem 2 (Informal version). *Assume the Gaussian noise model with $\sigma^2 = 1$. There exists a hypothesis class \mathcal{F} and an absolute constant C that satisfies the following: If an algorithm has $\mathbb{E} \text{Reg}_n \leq O(n^u)$ for some $u \in [0, 1)$ under an instance $f \in \mathcal{F}$, then there exists another instance $f' \in \mathcal{F}$ and an arm i with $\gamma_i(f') = 0$ (i.e., non-informative arm) such that*

$$\mathbb{E}_{f'} T_i(n) = \Omega(\ln(1 + (1 - u) \ln(n)))$$

where $\mathbb{E}_{f'}$ is the expectation under the instance f' .

The constructed instance for the lower bound is a variation of \mathcal{H}^+ in Figure 2. Our theorem shows that, just because an arm is noninformative, it does not mean that we can pull it a finite number of times.

Our result also has an implication for algorithms with forced sampling. To be specific, suppose that an algorithm A performs forced sampling by requiring each arm to be pulled at least τ where τ is fixed at the beginning. Then, to achieve sublinear regret bounds, it is required that A use $\tau = \omega(\ln(\ln(n)))$. We emphasize that even $\tau = \Theta(\ln(\ln(n)))$ can suffer a polynomial regret, let alone being uniformly good. This is because the constant factor matters and is a function of the problem. We provide the precise constants, the full statement of the theorem, and its proof in our appendix.

6 Discussion

There are improvements to be made including more examples and studying properties of alternative designs of ϕ and ψ , which we discuss more in the appendix. Meanwhile, we make a few observations and open problems below.

It may not be the end of the optimism [20]. Let us forget about CROP and consider the oracle described in Section 2. Consider $\mathcal{F}^{\text{code}}$ with $\frac{1}{2\epsilon} > \frac{2}{\Lambda^2}$. Note that UCB in fact has a regret bound of $O(\min\{\frac{K}{\epsilon} \ln(n), \epsilon n\})$; the first argument can be vacuous (i.e., $\geq n$) in which case we know the regret so far is ϵn since UCB by design only pulls arm i with $\Delta_i(f) = \epsilon$. The oracle has regret $\Theta(\min\{\frac{\ln(K)}{\Lambda^2} \ln(n), n\})$ where we have n rather than ϵn because she pulls informative arms. However, this implies that, until $n \lesssim \frac{1}{\Lambda^2}$, the oracle has a linear regret. In fact, all known algorithms with AO would be the same, to our knowledge. This is not just a theoretical observation. In Hao et al. [16, Figure 1], their algorithm with AO performs worse than an optimistic one until $n \approx 2000$. We ask if one can achieve the minimum of the two; i.e., obtain a finite-time regret bound of $O(\min\{\frac{\ln(K)}{\Lambda^2} \ln(n), \epsilon n\})$. This is a reminiscent of the “sub-UCB” criterion by Lattimore [18] (also discussed in Tirinzoni et al. [27]) in the sense that we like to perform no worse than UCB. For $\mathcal{F} = \mathcal{F}^{\text{code}}$, we provide a positive answer in the appendix, but a more generic algorithm that enjoys the AO and performs no worse than UCB for any \mathcal{F} is an open problem.

The worst-case regret. For more on the worst-case regret and how it is different from the instance-dependent regret, see our related work section in the appendix. The example above shows that the oracle suffers a linear worst-case regret over the family of problems $\{\mathcal{F}^{\text{code}, \epsilon, \Lambda} : \epsilon, \Lambda \leq (0, 1/2]\}$. That is, for any given problem complexity n and K , one can always find ϵ and Λ for which the oracle suffers a linear regret. This is in stark contrast to UCB that has $\tilde{O}(\sqrt{Kn})$ worst-case regret over this family. In fact, the oracle suffers a linear regret in linear bandits, too. In Example 4 of Lattimore and Szepesvári [20], it is easy to see that the oracle has regret $\min\{2\alpha^2 \ln(n), n\}$ when ϵ satisfies $2/\epsilon > 2\alpha^2$. Thus, given n , this bandit problem with $\alpha \approx \sqrt{n}$ for some ϵ with $2/\epsilon > 2\alpha^2$ would make the oracle suffer a linear regret for the instance $\theta = (1, 0)$. To our knowledge, all known AO algorithms share the same trait as they do not have any device to purposely avoid pulling the informative arm in the small- n regime.

We believe the issue is not that we study instance-dependent regret but that we tend to focus too much on the leading term w.r.t. n in the asymptotic regime, which we attribute to the fact that it is the one where the optimality can be claimed as of now. Less is known about the optimality on the lower order terms along with other instance-dependent parameters. This is studied a bit more in pure exploration problems [26, 17]. We hope to see more research on precise instance-dependent regret bounds in nonasymptotic regimes and practical structured bandit algorithms.

Broader Impact

Our study is mainly about a novel approach to solve structured bandits algorithms where we try to overcome some shortcomings of existing methods. Algorithmic developments in bandits have a huge impact in many potential applications including dose-finding trials. In this application, a structured bandits that encode a proper inductive bias and can help resolve health issues of many people by significantly reducing time/trials needed to find dosage or the right types of drugs, leading to maximum efficacy with minimal side-effects.

Acknowledgments

We thank the anonymous reviewers, Lalit Jain, Kevin Jamieson, Akshay Krishnamurthy, Tor Lattimore, Robert Nowak, Ardhendu Tripathy, and the organizers and participants of RL Theory Virtual Seminars for providing valuable feedback and helpful discussions. This work is supported in part by the startup fund from the University of Arizona.

References

- [1] Yasin Abbasi-Yadkori, David Pál, and Csaba Szepesvári. Improved Algorithms for Linear Stochastic Bandits. *Advances in Neural Information Processing Systems (NIPS)*, pages 1–19, 2011.
- [2] Alekh Agarwal, Dean P Foster, Daniel J Hsu, Sham M Kakade, and Alexander Rakhlin. Stochastic convex optimization with bandit feedback. In *Advances in Neural Information Processing Systems*, pages 1035–1043, 2011.
- [3] Rajeev Agrawal, Demosthenis Teneketzis, and Venkatachalam Anantharam. Asymptotically Efficient Adaptive Allocation Schemes for Controlled I.I.D. Processes: Finite Parameter Space. *IEEE Transactions on Automatic Control*, 1989.
- [4] Kareem Amin, Michael Kearns, and Umar Syed. Bandits, query learning, and the haystack dimension. volume 19 of *Proceedings of Machine Learning Research*, pages 87–106, Budapest, Hungary, 09–11 Jun 2011. JMLR Workshop and Conference Proceedings.
- [5] Onur Atan, Cem Tekin, and Mihaela Schaar. Global multi-armed bandits with Hölder continuity. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 28–36, 2015.
- [6] Peter Auer. Using Confidence Bounds for Exploitation-Exploration Trade-offs. *Journal of Machine Learning Research*, 3:2002, 2002.
- [7] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning*, 47(2–3):235–256, 2002.
- [8] Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet. Bounded regret in stochastic multi-armed bandits. In *Journal of Machine Learning Research*, 2013.
- [9] Wei Chu, Lihong Li, Lev Reyzin, and Robert E Schapire. Contextual Bandits with Linear Payoff Functions. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 15, pages 208–214, 2011.
- [10] Richard Combes, Stefan Magureanu, and Alexandre Proutiere. Minimal Exploration in Structured Stochastic Bandits. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1763–1771. 2017.
- [11] Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic Linear Optimization under Bandit Feedback. In *Proceedings of the Conference on Learning Theory (COLT)*, pages 355–366, 2008.

- [12] Rémy Degenne, Han Shao, and Wouter M Koolen. Structure Adaptive Algorithms for Stochastic Bandits. *Proceedings of the International Conference on Machine Learning (ICML)*, pages 8125–8134, 2020.
- [13] Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *Proceedings of the Conference on Learning Theory (COLT)*, pages 998–1027, 2016.
- [14] Todd L Graves and Tze Leung Lai. Asymptotically Efficient Adaptive Choice of Control Laws In Controlled Markov Chains. *SIAM J. Control Optim.*, 35(3):715–743, 1997.
- [15] Samarth Gupta, Shreyas Chaudhari, Subhojyoti Mukherjee, Gauri Joshi, and Osman Yağan. A unified approach to translate classical bandit algorithms to the structured bandit setting, 2018.
- [16] Botao Hao, Tor Lattimore, and Csaba Szepesvari. Adaptive exploration in linear contextual bandit. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- [17] Julian Katz-Samuels, Lalit Jain, Zohar Karnin, and Kevin Jamieson. An empirical process approach to the union bound: Practical algorithms for combinatorial and linear bandits. In *Advances in Neural Information Processing Systems*, 2020.
- [18] Tor Lattimore. Refining the Confidence Level for Optimistic Bandit Strategies. *Journal of Machine Learning Research*, 19(20):1–32, 2018.
- [19] Tor Lattimore and Rémi Munos. Bounded regret for finite-armed structured bandits. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.
- [20] Tor Lattimore and Csaba Szepesvári. The end of optimism? An asymptotic analysis of finite-armed linear bandits. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 728–737, 2017.
- [21] Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020.
- [22] Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. A Contextual-Bandit Approach to Personalized News Article Recommendation. *Proceedings of the International Conference on World Wide Web (WWW)*, pages 661–670, 2010.
- [23] Jungseul Ok, Alexandre Proutiere, and Damianos Tranos. Exploration in structured reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 8874–8882, 2018.
- [24] Hassan Saber, Pierre Ménard, and Odalric-Ambrym Maillard. Forced-exploration free Strategies for Unimodal Bandits, 2020.
- [25] Cong Shen, Ruida Zhou, Cem Tekin, and Mihaela van der Schaar. Generalized global bandit and its application in cellular coverage optimization. *IEEE Journal of Selected Topics in Signal Processing*, 12(1):218–232, 2018.
- [26] Max Simchowitz, Kevin Jamieson, and Benjamin Recht. The simulator: Understanding adaptive sampling in the moderate-confidence regime. In *Conference on Learning Theory*, pages 1794–1834, 2017.
- [27] Andrea Tirinzoni, Alessandro Lazaric, and Marcello Restelli. A novel confidence-based algorithm for structured bandits. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.

Appendix

Table of Contents

A Related work	12
B Details of the arguments	13
B.1 More on problem definition and preliminaries	13
B.2 More on CROP	15
B.3 More on analysis	16
B.4 More on discussion	16
C Upper bound proof	17
C.1 Concentration inequalities	17
C.2 Generic lemmas for analyzing bandit algorithms	20
C.3 Lemmas on the execution of CROP	21
C.4 Main proofs	23
C.5 Bounding the regret in each individual case	25
C.6 Miscellaneous lemmas	33
D Lower bound	34

A Related work

Structured bandits, generally defined, consist of bandit problems where there exist pieces of side-information or constraints in (and among) the mean rewards. While, of course, the standard K -armed bandit problem [38] is a special case, researchers usually use the term structured bandits for nontrivial structure (e.g., beyond simple constraints like having mean rewards in $[0, 1]$). We group related work by the kinds of guarantees each study aim to achieve.

Worst-case regret bounds. Suppose we are given a predefined family of problems $\Gamma = \{\mathcal{F}\}$, The worst-case regret bound of an algorithm π over the family Γ is the one that answers the following question: given a set of problem complexity parameters like n and K (and d in linear bandits, for example), what is the largest regret bound that π can suffer over the family Γ ? As an example, for the linear bandit problem with a fixed arm set, the family Γ contains any bandit problem \mathcal{F} with an arm set \mathcal{A} and a feature map $v : \mathcal{A} \rightarrow \mathbb{R}^d$ for which every $f \in \mathcal{F}$ and $a \in \mathcal{A}$ satisfy $f(a) = \langle \theta, v(a) \rangle, \forall a \in \mathcal{A}$ for some $\theta \in \mathbb{R}^d$.

Many works in the structured bandits focus on specific parametric reward models, such as linear and generalized linear models [35, 1, 14, 17], and regret bounds of order \sqrt{n} has been obtained in these settings. Russo and Van Roy [36] propose a general notion called eluder dimension that facilitates analysis for structured bandits for general function classes. For nonlinear reward models, worst-case regret guarantees have been obtained on reward structures such as unimodality [43], convexity [2], and Lipschitzness [25].

Instance-dependent regret bounds. Instance-dependent regret bounds aim to capture finer structures of problem instances beyond the complexity of hypothesis classes. In the asymptotic regime (i.e., fixing a problem instance and letting time horizon n go to infinity), many works [3, 20, 29, 22] derive matching asymptotic regret upper and lower bounds for *uniformly-good* algorithms (defined in Section 2). However, their analysis cannot be easily converted to obtain a finite-sample guarantee.

In the finite-sample regime, instance-optimal algorithms under specific model classes have been developed, such as unimodal reward [12] and Lipschitz rewards [31]. Under general reward function classes, Combes et al. [13] provide an asymptotically optimal algorithm that is amenable to non-asymptotic analyses, which is later extended to reinforcement learning [33]. However, their finite

sample guarantees depend strongly on the size of the action space due to the forced sampling as discussed in Section 4, although such a dependence goes away in the asymptotic regime.

We remark that the forced sampling is not bad when the problem is *unstructured* because the sampling complexity must scale with K anyways. For example, in the pure-exploration version of the bandit problem [7, 8], the algorithm by Garivier and Kaufmann [19] relies on forced sampling, but there is no known evidence that the forced sampling degrades its performance while without it their algorithm can provably fail.

Other instance-dependent regret bounds. Many algorithms do not achieve the asymptotic optimality (not even within a constant factor) but possess instance-dependent guarantees in structured bandits, which is mostly based on the optimism, defined in (3), or often elimination (e.g., see Jamieson [23, Section 6] for linear bandits and Tirinzoni et al. [39] for generic structure). We remark that, unlike the standard K -armed bandits without structure, elimination-based approaches in structured bandits can have a different regret bound from that of the optimism (i.e., one is better than the other and vice versa depending on the problem instance) as described in Tirinzoni et al. [39].

Many linear bandit studies obtain regret guarantees that depend on the gap Δ between the best mean reward and the second-best mean reward [14, 17, 1]. As another example, in Lipschitz bandits, regret guarantees that depend on the zooming dimension or near-optimality dimension have been shown [24, 9]. Other particular structures including univariate linearity [32], global bandits [4, 37], regional bandits [42], and K -armed bandits with side-information [10, 40, 15] show that it is possible to achieve bounded regret.

For the generic structures, many studies provide finite-time instance-dependent guarantees that do not achieve a constant-factor asymptotically optimality [28, 6, 39, 21], except for known special cases where the reward function class has a factorized representation across different arms, i.e., in the settings of Lai and Robbins [26], and Bernetas and Katehakis [11]. These studies, however, have regret bounds that reflect the structure of the instance beyond the gap Δ mentioned above. We conjecture that the suboptimality in their bounds is rooted in their confidence sets that are an intersection of confidence intervals for each arm, let alone their sampling strategies. In contrast, CROP maintains a confidence set that captures the structure of the hypothesis class, which, as we show in the proof, allows us to connect the constraints of $c(f)$ in (2) to the concentration inequality and thus to the confidence set as well.

Although we focus on the finite hypothesis space for simplicity, our ultimate goal of the paper is to find a generic algorithmic principle for any hypothesis space. Our algorithm CROP achieves the asymptotic optimal regret with a constant-factor, enjoys bounded regret whenever possible, and has mild dependence on K , all thanks to our novel forced-sampling-free design. This partially resolves the open question raised by Tirinzoni et al. [39, Section 7] where they ask if one can design confidence-based strategies (as opposed to solving the optimization problem (2) along with forced sampling) that are optimal for general structures with good finite-time performance. We believe that such an advancement is not an artifact of the finite hypothesis space setting but the fact that we rely on the pessimism, the key novelty in our algorithm design.

B Details of the arguments

In this section, we provide more explanations on our argument in the main body in the order they appear, section by section.

B.1 More on problem definition and preliminaries

Let us elaborate more on the cheating code example in Section 2.

The cheating code example. Let us elaborate more on the cheating code example in Section 2. To understand the optimization problem $c(f)$ for the cheating code class, let us consider a simple case where $K_0 = 2$ and $k = 1$. Without loss of generality, let $f^* = (1, 1 - \epsilon, 0)$. In this case, the only

competing hypothesis is $(1, 1 + \epsilon, \Lambda)$. Then, $c(f^*)$ is written as follows:

$$c(f^*) = \min_{\gamma_1=0, \gamma_2 \geq 0, \gamma_3 \geq 0} \epsilon \gamma_2 + 1 \cdot \gamma_3$$

$$\text{s.t. } \gamma_2 \frac{(2\epsilon)^2}{2} + \gamma_3 \frac{\Lambda^2}{2} \geq 1$$

One can see that a solution $\gamma_2 = \alpha \cdot \frac{1}{(2\epsilon)^2/2}$ and $\gamma_3 = (1 - \alpha) \cdot \frac{1}{\Lambda^2/2}$ for some $\alpha \in [0, 1]$ is feasible and satisfies the constraint with equality. Furthermore, one can see that any feasible solution that cannot be expressed by the solution above has a strictly larger objective value due to $\gamma \geq 0$ and $\epsilon > 0$. The objective function is then $\alpha \cdot \frac{1}{2\epsilon} + (1 - \alpha) \cdot \frac{2}{\Lambda^2}$. This means that whenever $\frac{1}{2\epsilon} > \frac{2}{\Lambda^2}$, setting $\alpha = 0$ achieves the minimum.

For generic K_0 , we first provide an example of $f^* = (1, 1 - \epsilon, 1 - \epsilon, 1 - \epsilon, 0, 0)$ for the case of $K_0 = 4$.

$$c(f^*) = \min_{\gamma_1=0, \gamma_2, \gamma_3 \in (0, \infty]^{K-1}} \epsilon \cdot \left(\sum_{a=2}^{K_0} \gamma_a \right) + 1 \cdot \left(\sum_{a=K_0+1}^K \gamma_a \right)$$

$$\text{s.t. } \begin{aligned} \gamma_2 \frac{(2\epsilon)^2}{2} &+ \gamma_6 \frac{\Lambda^2}{2} &\geq 1 \\ &+ \gamma_3 \frac{(2\epsilon)^2}{2} &+ \gamma_5 \frac{\Lambda^2}{2} &\geq 1 \\ &+ \gamma_4 \frac{(2\epsilon)^2}{2} &+ \gamma_5 \frac{\Lambda^2}{2} &+ \gamma_6 \frac{\Lambda^2}{2} &\geq 1 \end{aligned} \quad (7)$$

We use shorthands $\gamma_{i:j}$ for $\gamma_i, \gamma_{i+1}, \dots, \gamma_j$. It is now nontrivial to see how the optimal solution would look like. The following proposition provides a characterization of the optimal solution.

Proposition 1. Consider $\mathcal{F}^{\text{code}}$. Let $K_0 = 2^k$ for some integer k . We claim that if $\frac{1}{2\epsilon} > \frac{2}{\Lambda^2}$, then the solution of the optimization problem (2) is

$$\gamma_a^\dagger = \begin{cases} 0, & a \in [K_0] \\ \frac{2}{\Lambda^2}, & a \in \{K_0 + 1, \dots, K\} \end{cases}. \quad (8)$$

Proof. For clarity, our convention is that the coordinate K is for the least significant bit of the code. For example, when $K_0 = 4$, then the hypothesis f with $a^*(f) = 2$ is $f = (1, 1 + \epsilon, 1 - \epsilon, 1 - \epsilon, 0, \Lambda)$.

The plan is to suppose that u is a feasible solution to the optimization problem $c(f^*)$ for $\mathcal{F}^{\text{code}}$ for which there exists a coordinate $a \in [K_0] \setminus \{a^*\}$ with $u_a > 0$. Then, we show:

- First, we prove that it is possible to construct a feasible solution v that is strictly better than u where v does not have nonzero entries for the first K_0 coordinates; this proves that the optimal solution must be supported only on the cheating arms.
- Second, we show that γ^\dagger is the optimal solution.

As a starter, consider the example of (7). Suppose we have a feasible solution $u = (0, \frac{1}{4\epsilon^2}, 0, 0, \frac{2}{\Lambda^2}, \frac{1}{\Lambda^2})$. Then, the coordinate $q = 2$ is nonzero. Consider a hypothesis g that has arm q as the best arm: $g = (1, 1 + \epsilon, 1 - \epsilon, 1 - \epsilon, 0, \Lambda)$. The arm 6 is the cheating arm whose mean reward is different from that of f^* ; let $j = 2$ so that $6 = K_0 + j$. Then, we can modify u by zeroing out γ_q and adding more mass to γ_{K_0+j} so that the first constraint of (7) is satisfied. This modification leads to $h^* = (0, 0, 0, 0, \frac{2}{\Lambda^2}, \frac{2}{\Lambda^2})$. Note that this operation does not make other constraints violated because the variable γ_q appears in one constraint only, and adding more mass to γ_{K_0+j} never harms. We now generalize this example.

Let $q \in [K_0] \setminus \{a^*\}$ satisfy $u_q > 0$. Consider $g \in \mathcal{C}^*$ with $a^*(g) = q$. Let $j \in [k]$ be the largest index j for which $g(K_0 + j) \neq f^*(K_0 + j)$ whose existence is certified by the definition of $\mathcal{F}^{\text{code}}$. Let e_i be the i -th indicator vector. Let $\delta = \frac{2\epsilon^2 u_q}{\Lambda^2/2}$. We then define

$$h^*(u, q) = u - u_q e_q + \delta \cdot e_{K_0+j}$$

One can show that our choice of δ indeed ensures that $h^*(u, q)$ is a feasible solution. We now show that this modified version $h^* := h^*(u, q)$ has a strictly smaller objective function value:

$$\left(\sum_{a=1}^{K_0} \epsilon u_a + \sum_{a=K_0+1}^K \Delta_a u_a \right) - \left(\sum_{a=1}^{K_0} \epsilon h_a^* + \sum_{a=K_0+1}^K \Delta_a h_a^* \right) = \epsilon u_q - \Delta_j \delta$$

$$\begin{aligned}
&\geq \epsilon u_q - \delta && (\because \Delta_j \leq 1) \\
&= \epsilon u_q - \frac{2\epsilon^2 u_q}{\Lambda^2/2} \\
&= \epsilon u_q \left(1 - \frac{2\epsilon}{\Lambda^2/2}\right) \\
&> 0 && (\because \text{the assumption of the proposition})
\end{aligned}$$

Therefore, one can perform the following one sweep of coordinate descent.

- Input: f^*, u : a feasible solution of $c(f^*)$.
- $v^{(1)} \leftarrow u$
- For $a = 1, \dots, K_0$,
 - $v^{(a+1)} \leftarrow h^*(v^{(a)}, a)$
- Output: $v^* := v^{(K_0+1)}$

We conclude the first part of the proof by the following observations:

- The output v^* above has a strictly smaller objective function than that of the input u whenever $\exists a \in [K_0] : u_a > 0$.
- The output v^* satisfies that $v_{1:K_0}^* = 0$.

We now show the second part of the proof. Let γ be a feasible solution that is supported only on the cheating arms. We claim that $\forall j \in [k]$, the value γ_{K_0+j} must be at least $\frac{2}{\Lambda^2}$.

The reason is that, given $j \in [k]$, we can find $g \in \mathcal{C}^*$ for which $(g(K_0+1), \dots, g(K))$ differs from $(f^*(K_0+1), \dots, f^*(K))$ at the coordinate j only, by the definition of $\mathcal{F}^{\text{code}}$; i.e., the binary representation of $a^*(g) - 1$ differs from that of $a^* - 1$ only at the $(k-j+1)$ -th least significant bit. In the example of (7), for $j = 1$, $g = (1, 1+\epsilon, 1-\epsilon, 1-\epsilon, 0, \Lambda)$ and for $j = 2$, $g = (1, 1-\epsilon, 1+\epsilon, 1-\epsilon, \Lambda, 0)$. Then, the constraint induced by that hypothesis g is:

$$\gamma_{a^*(g)} \frac{(2\epsilon)^2}{2} + \gamma_{K_0+j} \frac{\Lambda^2}{2} \geq 1.$$

Plugging the fact that $\gamma_{a^*(g)} = 0$, we have $\gamma_{K_0+j} \frac{\Lambda^2}{2} \geq 1$, implying that $\gamma_{K_0+j} \geq \frac{2}{\Lambda^2}$. This proves the claim above, establishing a coordinate-wise lower bound on γ .

We observe that the objective function of this lower bound on γ is a lower bound for the optimal solution, which is achieved by γ^\dagger ; this concludes the proof. \square

B.2 More on CROP

We here discuss more on the design choices for CROP. We do not claim our design of ψ and ϕ is the best; fine-tuning of those and the algorithm itself is left as future work. In what follows, we focus on describing our intention on the current design *at the time of development*.

The design of $\psi(f)$ in Fallback. The reason why we now allow $\gamma_{a^*(f)} \neq 0$ is for the following case where $\epsilon > 0$ and we boldface the best arms:

Arms	A1	A2	A3
f_1	1	.25	.25
f_2	.75	.25	.25
f_3	0	.25 - 2 ϵ	.25 - ϵ
f_4	0	.25	0

Suppose $f^* = f_2$ and $\epsilon > 0$ is small enough. At the beginning, $\mathcal{F}_t = \{f_1, f_2, f_3, f_4\}$, $\tilde{\mathcal{F}}_t = \{f_1\}$, $\bar{f}_t = \{f_3\}$. If we do not allow $\gamma_{a^*(f)}$ to be nonzero in (4), then $\psi(\bar{f}_t)$ will assign a very large number of pulls to arm A2 whereas pulling A3 will eliminate f_4 quickly. On the other hand, those troublesome hypotheses f_3 and f_4 belong to the docile class w.r.t. f_2 , and incurring high regret for those (albeit finite terms) seems unreasonable.




Figure 3: A cartoon showing the regret bound of UCB and AO where AO is the asymptotically optimal oracle described in Section 2. The minimum of the two curves is shaded. Can we achieve this regret bound?

We have a constraint $\gamma \geq \phi(f) \vee \gamma(f)$ in (4), which we call the *extension constraint*. The intention is, if f is the ground truth, we will have to make $\gamma_a(f) \ln(n)$ pulls for the informative arms $\{a : \gamma_a(f) > 0\}$ anyways, so we add it to the constraint. This guards against the case where we inadvertently pull noninformative arms $\{a : \gamma_a(f) = 0\}$ too much. Our understanding is that, in general, this does not affect the regret bound too much, at least in the current analysis.

B.3 More on analysis

The proof of Theorem 1 can be found in Appendix C.

Remark 1. Note we have a $\ln(|\mathcal{F}|)$ dependence, which comes from the naive union bounds. One can extend CROP to a larger or even infinite hypothesis space using the covering number argument as done in Foster et al. [18, Lemma 4], although we have focused on the finite hypothesis for simplicity.

B.4 More on discussion

We provide missing details from the discussion section.

Regret no more than the optimism. Consider $\mathcal{F}^{\text{code}}$. Assume $f^* = (1, 1 - \epsilon, \dots, 1 - \epsilon, 0, \dots, 0) \in \mathcal{F}^{\text{code}}$ without loss of generality, and assume $\frac{1}{2\epsilon} > \frac{2}{\Lambda^2}$. The goal is to achieve a regret bound of $\mathbb{E} \text{Reg}_n = O\left(\min\left\{\frac{\ln(K)}{\Lambda^2} \ln(n), \epsilon n\right\}\right)$, which is depicted in Figure 3. Note that, in the fixed budget setting (i.e., the target n is given before running the algorithm), it is trivial to perform no worse than the optimism. Specifically, check the regret bound of the oracle or those that mimic the oracle (simply call them the oracle, hereafter) and see if it is larger than ϵn . If true, then simply pull any of the first K_0 arms uniformly throughout; otherwise, invoke the oracle.

For achieving the goal in the anytime setting, note that the optimism, such as UCB1 [5] run without knowing the structure or UCB-S [27], achieves an anytime regret bound of

$$\mathbb{E} \text{Reg}_n \leq \min\left\{c_1 \frac{K}{\epsilon} \ln(n), \epsilon n\right\}$$

for $n \geq 2$ and some numerical constant c_1 . The oracle achieves an anytime regret bound of

$$\mathbb{E} \text{Reg}_n \leq \min\left\{c_2 \frac{\ln(K)}{\Lambda^2} \ln(n), n\right\}$$

for $n \geq 2$ and some numerical constant c_2 .

Suppose K is large enough so that $c_1 \frac{K}{\epsilon} > c_2 \frac{\ln(K)}{\Lambda^2}$. Let t_0 be the t such that $c_2 \frac{\ln(K)}{\Lambda^2} \ln(t) = \epsilon t$, the time step after which the oracle outperforms UCB. The idea is simple: run UCB up to time step $t \leq t_0$ and then start the oracle as if we are starting from the beginning (in words, throw away all the samples so far). Then,

$$\begin{aligned} \mathbb{E} \text{Reg}_n &= \mathbb{1}\{n \leq t_0\} \epsilon n + \mathbb{1}\{n > t_0\} \mathbb{E}\left[\epsilon t_0 + c_2 \frac{\ln(K)}{\Lambda^2} \ln(n - t_0)\right] \\ &\leq \mathbb{1}\{n \leq t_0\} 2 \cdot \epsilon n + \mathbb{1}\{n > t_0\} 2 \cdot c_2 \frac{\ln(K)}{\Lambda^2} \ln(n) \end{aligned}$$

$$\leq 2 \cdot \min \left\{ c_2 \frac{\ln(K)}{\Lambda^2} \ln(n), \epsilon n \right\},$$

which achieves our goal.

Note, however, that this simple strategy was possible because we knew ϵ for the given class $\mathcal{F}^{\text{code}}$. Consider $\mathcal{F}^{\text{code}2} := [1/2, 1]^{K_0} \times \{0, \Lambda\}^k$ where the last $k = \lceil \log_2(K_0) \rceil$ arms are binary codes encoding the best arm index (break ties with the smallest index). Then, the class contains hypotheses that have ϵ gap for the first K_0 arms for any ϵ . Even if we do know that f^* is one of those cases, we do not know ϵ and thus have to adaptively decide when to start pulling the informative arms via the rewards collected throughout. Achieving the asymptotic optimality while maintaining the standard worst-case regret (e.g., $\sqrt{dn \ln(K)}$ for linear bandits) seems to be an interesting open problem.

Future work. Our study unlocks numerous open problems besides achieving both the worst-case regret and asymptotic optimality. Since many bandit studies have focused on the leading term $\ln(n)$, the optimal sampling strategy that minimize the arm pulls of those noninformative arms is not studied well. We like to study the optimality of the number of arm pulls of noninformative arms w.r.t. not just n but also problem-dependence parameters such as information gaps, which tend to matter when the number of arms is very large. Towards practical algorithms, we believe this is more important than getting the exact asymptotic optimality. Furthermore, we like to investigate if we can extend our pessimism to more popular structures such as linear bandits or Lipschitz bandits so one can achieve the asymptotic optimality without forced sampling (and without the dependence on the number of arms).

C Upper bound proof

In this section, we discuss the proof of Theorems 1 and auxiliary lemmas for it.

Theorem 3. *Let $(\alpha, \hat{\alpha}, z, \hat{z}) = (2, 3, |\mathcal{F}|, |\mathcal{F}|)$. Suppose we run CROP with hypothesis class \mathcal{F} with the environment $f^* \in \mathcal{F}$. Then, CROP has the following anytime regret guarantee: $\forall n \geq 2$,*

$$\mathbb{E} \text{Reg}_n \leq P_1 \ln n + c_1 \left(P_1 \left(\sqrt{\ln n \ln |\mathcal{F}|} + \ln \ln n + \ln |\mathcal{F}| \right) + P_2 \ln(\ln(n)) + P_3 \left(\ln(|\mathcal{F}|) + \ln(Q_1) \right) + K_\psi \Delta_{\max} \right),$$

where c_1 is a numerical constant, and

$$P_1 = \sum_a \Delta_a \gamma_a^*, \quad P_2 = \sum_a \Delta_a \phi_a(\mathcal{E}^*), \quad P_3 = \sum_a \Delta_a \psi_a(\mathcal{F})$$

and $Q_1 = \Lambda_{\min}^{-2} + K_\psi (1 + \max_a \psi_a(\mathcal{F}))$. Furthermore, when $\gamma^* = 0$, we have $P_1 = P_2 = 0$, achieving a bounded regret.

Before going into details we highlight some of the technical aspects of our proof. In our proof, Lemma 8 plays a key role for analyzing the confidence set $\hat{\mathcal{F}}_t$ that has an aggressive confidence level, which we believe is not commonly dealt with in the standard bandit settings. One can see why this is needed in Section C.5.2 and Section C.5.3 where we apply the ‘‘regret peeling.’’ The proofs in Section C.5.4 are a bit lengthy, but the main idea stems from Lattimore and Munos [27]. Other proofs are relatively standard, we believe.

We first present some martingale concentration inequalities for our confidence set.

C.1 Concentration inequalities

This section establishes a few important concentration inequalities on the losses of reward regressors, which are instrumental in our analysis.

Additional notations. Recall that at time step t , the learner pulls arm a_t and receives reward $r_t = f_*(a_t) + \xi_t$, where ξ_t is σ^2 -sub-Gaussian. To avoid double subscripting, we will sometimes use $a(t)$ to denote a_t . Define the instantaneous loss of regressor f as time s as $\ell_s(f) := (f(a_s) - r_s)^2$. At the end of time step t , we define the cumulative loss of f as $L_t(f) := \sum_{s=1}^t \ell_s(f)$. Define the instantaneous regret of f as $M_s(f) := \ell_s(f) - \ell_s(f^*)$. We define the *information gap* between f and f^* on action a as: $\Lambda_a(f) := \frac{f(a(s)) - f^*(a(s))}{\sigma}$. A larger $\Lambda_a(f)$ implies that f is easier to

be distinguished from f^* by pulling arm a . This should not be confused with the *reward gap* $\Delta_a(f) = \mu^*(f) - f_a$.

We define the filtration $\{\Sigma_t\}_{t=0}^\infty$ as follows: $\Sigma_t = \sigma(a_1, r_1, \dots, a_t, r_t, a_{t+1})$. We abbreviate $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot | \Sigma_t]$. Define $\text{IC}(f, g, \pi) = \sum_{a \in \mathcal{A}} \pi(a) \frac{(f(a) - g(a))^2}{2\sigma^2}$ as the **information constraint** between f and g w.r.t. π . Define $\text{IC}^*(f, \pi) = \text{IC}(f, f^*, \pi)$.

Define $T(t) = (T_a(t))_{a \in \mathcal{A}}$ as the vector that encodes the number of arm pulls for each arm up to time step t . For vectors $u = (u_a)_{a \in \mathcal{A}}$, $v = (v_a)_{a \in \mathcal{A}}$, we denote $u \geq v$ if $\forall a \in \mathcal{A}, u_a \geq v_a$. We use shorthands $a^* = a^*(f^*)$ and $\mu^* = \mu^*(f^*)$.

We establish fundamental concentration result on partial sums of $\{M_s(f)\}$.

Lemma 4. (i) For any f and any $\beta > 0$,

$$\mathbb{P}(\exists t \in \mathbb{N} \cdot L_t(f^*) - L_t(f) \geq \beta) \leq \exp\left(-\frac{\beta}{2\sigma^2}\right).$$

(ii) For any f and any $r > 0, \beta > 0$,

$$\mathbb{P}(\exists t \in \mathbb{N} \cdot \text{IC}(f, f^*, T(t)) \geq r \wedge L_t(f) - L_t(f^*) \leq \beta) \leq \exp\left(-\frac{1}{4r} \left(r - \frac{\beta}{2\sigma^2}\right)_+^2\right) \quad (9)$$

$$\leq \exp\left(-\frac{\sigma^2 r - \beta}{4\sigma^2}\right). \quad (10)$$

Proof. Throughout the proof, we will abbreviate $M_s(f)$ as M_s to avoid notation clutter. It can be easily seen that $\sum_{s=1}^t M_s = L_t(f) - L_t(f^*)$. Define $d_s = \sigma \Lambda_{a(s)}(f) = f(a) - f^*(a)$.

The first item follows from Equation (12) in Lemma 5 with $\lambda = \frac{1}{\sigma^2}$ and $\delta = \exp\left(-\frac{\beta}{2\sigma^2}\right)$.

For the second item, we first show Equation (9). The equation is trivially true when $2\sigma^2 r < \beta$, so we focus on the case of $2\sigma^2 r \geq \beta$. With foresight, consider setting $\delta = \exp\left(-\frac{1}{4r} \left(r - \frac{\beta}{2\sigma^2}\right)_+^2\right)$ and

$\lambda = \frac{1}{2\sigma^2} \sqrt{\frac{\ln \frac{1}{\delta}}{r}} = \frac{1}{4\sigma^2 r} \left(r - \frac{\beta}{2\sigma^2}\right)$ in Equation (12) in Lemma 5, we get

$$\mathbb{P}\left(\exists t \in \mathbb{N} \cdot \sum_{s=1}^t M_s \leq \left(1 - \sqrt{\frac{\ln \frac{1}{\delta}}{r}}\right) \left(\sum_{s=1}^t d_s^2\right) - 2\sigma^2 \sqrt{r \ln \frac{1}{\delta}}\right) \leq \delta.$$

As the setting of δ implies that $r \geq \ln \frac{1}{\delta}$, we have $\left(1 - \sqrt{\frac{\ln \frac{1}{\delta}}{r}}\right) \geq 0$. This implies that

$$\mathbb{P}\left(\exists t \in \mathbb{N} \cdot \sum_{s=1}^t d_s^2 \geq 2\sigma^2 r, \sum_{s=1}^t M_s \leq 2\sigma^2 r - 4\sigma^2 \sqrt{r \ln \frac{1}{\delta}}\right) \leq \delta.$$

The setting of δ implies that $2\sigma^2 r - 4\sigma^2 \sqrt{r \ln \frac{1}{\delta}} = \beta$. In addition, observing that $\sum_{s=1}^t d_s^2 = 2\sigma^2 \text{IC}(f, f^*, T(t))$, and $\sum_{s=1}^t M_s = L_t(f) - L_t(f^*)$, the above inequality is precisely (9).

For Equation (10), it suffices to show $\frac{1}{4r} \left(r - \frac{\beta}{2\sigma^2}\right)_+^2 \geq \frac{\sigma^2 r - \beta}{4\sigma^2}$. This is trivially true if $\sigma^2 r < \beta$. Otherwise, $\sigma^2 r \geq \beta$; in this case, the above is a direct consequence of the AM-GM inequality that $\left(r - \frac{\beta}{2\sigma^2}\right)_+^2 \geq r \cdot \left(r - \frac{\beta}{\sigma^2}\right)$. \square

Lemma 5. Suppose f is in \mathcal{F} . Define $d_s(f) = \sigma \Lambda_{a(s)}(f) = f(a) - f^*(a)$. Then, for any $\lambda > 0$ and $\delta > 0$,

$$\mathbb{P}\left(\exists t \in \mathbb{N} \cdot \sum_{s=1}^t M_s(f) \geq (1 + 2\sigma^2 \lambda) \sum_{s=1}^t d_s(f)^2 + \frac{1}{\lambda} \ln \frac{1}{\delta}\right) \leq \delta, \quad (11)$$

$$\mathbb{P}\left(\exists t \in \mathbb{N} \cdot \sum_{s=1}^t M_s(f) \leq (1 - 2\sigma^2\lambda) \sum_{s=1}^t d_s(f)^2 - \frac{1}{\lambda} \ln \frac{1}{\delta}\right) \leq \delta. \quad (12)$$

Proof. Throughout the proof, we will abbreviate $M_s(f)$ as M_s , and abbreviate $d_s(f)$ as d_s , to avoid notation clutter.

First, observe that $\ell_s(f^*) = \xi_s^2$. in addition,

$$M_s = \ell_s(f) - \ell_s(f^*) = (d_s - \xi_s)^2 - \xi_s^2 = d_s^2 - 2d_s\xi_s,$$

which implies that

$$\mathbb{E}_{s-1}[M_s] = d_s^2.$$

For any $\theta \in \mathbb{R}$, define $H_t = \exp(\theta \sum_{s=1}^t M_s - \theta(1 + 2\sigma^2\theta) \sum_{s=1}^t d_s^2)$ with the convention $H_0 = 1$. We now show $\{H_t\}_{t=0}^n$ is a nonnegative supermartingale.

By the calculations on M_s above, $\theta M_s - \theta(1 + 2\sigma^2\theta)d_s^2 = -2\theta d_s \xi_s - 2\sigma^2\theta^2 d_s^2$. Using this, we have

$$\mathbb{E}_{t-1}[H_t] = \mathbb{E}_{t-1}\left[H_{t-1} \cdot \exp(\theta M_t - \theta(1 + 2\sigma^2\theta)d_t^2)\right] = H_{t-1} \cdot \mathbb{E}_{t-1}\left[\exp(-2\theta d_t \xi_t - 2\sigma^2\theta^2 d_t^2)\right].$$

Observe that by the σ^2 -subgaussian property of ξ_t ,

$$\begin{aligned} \mathbb{E}_{t-1}\left[\exp(-2\theta d_t \xi_t - 2\sigma^2\theta^2 d_t^2)\right] &= \mathbb{E}_{t-1}\left[\exp(-2\theta d_t \xi_t)\right] \exp(-2\sigma^2\theta^2 d_t^2) \\ &\leq \exp\left(\frac{(2\theta)^2 d_t^2 \sigma^2}{2}\right) \exp(-2\sigma^2\theta^2 d_t^2) = 1. \end{aligned}$$

This shows that $\mathbb{E}_{t-1}[H_t] \leq H_{t-1}$, proving $\{H_t\}_{t=0}^n$ is a nonnegative supermartingale. Therefore, by Ville's maximal inequality [41] (see also Durrett [16, Exercise 5.7.1]), we have that for any $\delta > 0$,

$$\mathbb{P}(\exists t \in \mathbb{N} \cdot H_t \geq \delta^{-1}) \leq \delta.$$

Choosing $\theta = \lambda$, and noting that $H_t \geq \frac{1}{\delta}$ is equivalent to $\sum_{s=1}^t M_s \geq (1 + 2\sigma^2\lambda) \sum_{s=1}^t d_s^2 + \frac{1}{\lambda} \ln \frac{1}{\delta}$, we get Equation (11). Likewise, choose $\theta = -\lambda$, and noting that $H_t \geq \frac{1}{\delta}$ is equivalent to $\sum_{s=1}^t M_s \leq (1 - 2\sigma^2\lambda) \sum_{s=1}^t d_s^2 - \frac{1}{\lambda} \ln \frac{1}{\delta}$, we get Equation (12). \square

Lemma 6. Let $\beta_t = 2\sigma^2 \ln(zt(\log_2 t)^2)$ with $z \geq |\mathcal{F}|$ and $\alpha \geq 1$. Define $B_t = \{L_{t-1}(f^*) - \min_{f \in \mathcal{F}} L_{t-1}(f) > \beta_t\}$. Let $q \geq 1$ be an integer. Then,

$$\mathbb{P}(\exists t \geq q : B_t) \leq \frac{1}{q(\log_2 q)^2}.$$

Proof.

$$\begin{aligned} \mathbb{P}(\exists t \geq q : B_t) &= \mathbb{P}(\exists t \geq q : L_{t-1}(f^*) - \min_{f \in \mathcal{F}} L_{t-1}(f) > \beta_t) \\ &\leq \mathbb{P}(\exists t \geq q : L_{t-1}(f^*) - \min_{f \in \mathcal{F}} L_{t-1}(f) > \beta_q) \\ &\leq \mathbb{P}(\exists t \in \mathbb{N}, f \in \mathcal{F} : L_{t-1}(f^*) - L_{t-1}(f) > \beta_q) \\ &\leq \sum_{f \in \mathcal{F}} \mathbb{P}(\exists t \in \mathbb{N} : L_{t-1}(f^*) - L_{t-1}(f) > \beta_q) \\ &\leq |\mathcal{F}| \cdot \exp\left(-\frac{\beta_q}{2\sigma^2}\right) \leq \frac{1}{q(\log_2 q)^2} \end{aligned}$$

where the first inequality is from the fact that β_t is monotonically increasing; the second inequality is by relaxing the range of t ; the third inequality is from union bound; the fourth inequality is from item (i) of Lemma 4; the last inequality is by algebra. \square

C.2 Generic lemmas for analyzing bandit algorithms

The following is a standard inequality used in the UCB analysis.

Lemma 7. *Consider any bandit algorithm. Let a_t be the index of the arm pulled at time t and $T_i(t-1)$ be the number of times arm i is pulled up to (and including) time $t-1$. Let τ be an integer and Q_t be an event. Then,*

$$\sum_{t=1}^n \mathbb{1}\{a_t = i, Q_t\} \leq \tau + \sum_{t=\tau+1}^n \mathbb{1}\{a_t = i, Q_t, T_i(t-1) \geq \tau\}$$

Proof.

$$\sum_{t=1}^n \mathbb{1}\{a_t = i, Q_t\} = \sum_{t=1}^n \mathbb{1}\{a_t = i, Q_t, T_i(t-1) < \tau\} + \sum_{t=1}^n \mathbb{1}\{a_t = i, Q_t, T_i(t-1) \geq \tau\}$$

The first summation is bounded by τ , for the following reason: if there are $\tau+1$ time steps $t_1 < \dots < t_{\tau+1}$ in which $a_t = i, Q_t, T_i(t-1) < \tau$ holds, we have that $T_i(t_{\tau+1}-1) \geq \tau$, which contradicts with the fact that $T_i(t_{\tau+1}-1) < \tau$.

Furthermore, for $t \leq \tau$, it must be the case that $T_i(t-1) \leq t-1 \leq \tau-1$, therefore the first τ terms of the second sum must be zero. This implies that the second sum equals $\sum_{t=\tau+1}^n \mathbb{1}\{a_t = i, Q_t, T_i(t-1) \geq \tau\}$. \square

In fact, the indicator terms in the lemma above have dependencies between different time steps that is not being explicitly captured. When we take the expectation and apply concentration inequalities, these dependencies are lost. The following lemma extends Lemma 7 so that such a dependency becomes explicit, which help prove tighter bounds. The basic idea is that whenever we pull an arm a at time t , the count of arm a increases by 1, so by the time we pull arm a , the pull count must be larger; this helps, when taking the expectation, obtain a tighter concentration of measure.

Lemma 8. *Under the same assumptions in Lemma 7,*

$$\sum_{t=1}^n \mathbb{1}\{a_t = i, Q_t\} \leq \tau + \sum_{m=1}^{\infty} \mathbb{1}\{\exists t \geq \tau+1 : a_t = i, T_i(t-1) \geq \tau+m-1, Q_t\}$$

Proof. By Lemma 7,

$$\sum_{t=1}^n \mathbb{1}\{a_t = i, Q_t\} \leq \tau + \sum_{t=\tau+1}^n \mathbb{1}\{a_t = i, Q_t, T_i(t-1) \geq \tau\}$$

Define $t^- = t-1$ Define an event:

$$A_t = \{a_t = i, T_i(t^-) \geq \tau, Q_t\}$$

We aim to bound $\sum_{t=\tau+1}^n \mathbb{1}\{A_t\}$. Define t_m to be the m -th time step after (and including) $t = \tau+1$ that $A_t = 1$ is true; i.e.,

$$t_1 := \min\{t \in [\tau+1, n] : A_t \text{ is true}\} \\ \forall m \geq 2, t_m := \min\{t \in [t_{m-1}+1, n] : A_t \text{ is true}\}$$

where we take $\min \emptyset = \infty$. One can verify that, if $t_m < \infty$, then $T_i(t_m^-) \geq \tau+m-1$. Then,

$$\begin{aligned} \sum_{t=\tau+1}^n \mathbb{1}\{A_t\} &= \sum_{m=1}^{\infty} \mathbb{1}\{t_m < \infty\} \mathbb{1}\{A_{t_m}\} \\ &= \sum_{m=1}^{\infty} \mathbb{1}\{t_m < \infty, a_{t_m} = i, T_i(t_m^-) \geq \tau, Q_{t_m}\} \\ &\stackrel{(a)}{\leq} \sum_{m=1}^{\infty} \mathbb{1}\{t_m < \infty, a_{t_m} = i, T_i(t_m^-) \geq \tau+m-1, Q_{t_m}\} \\ &\leq \sum_{m=1}^{\infty} \mathbb{1}\{t_m < \infty, \exists t \geq \tau+1 : a_t = i, T_i(t^-) \geq \tau+m-1, Q_t\} \end{aligned}$$

$$\leq \sum_{m=1}^{\infty} \mathbb{1} \{ \exists t \geq \tau + 1 : a_t = i, T_i(t^-) \geq \tau + m - 1, Q_t \}$$

where (a) is by our observation above. \square

C.3 Lemmas on the execution of CROP

Depending on the execution trace of CROP, at time t , we define four events that form a disjoint union of the sample space:

1. ‘E’xploit: Ex_t , i.e. $a^*(\mathcal{F}_t)$ is singleton (line 4)
2. ‘C’on‘f’lict: Cf_t (line 12)
3. ‘F’ea‘s’ibility: Fs_t (line 14)
4. ‘F’all‘b’ack: Fb_t (line 16)

The following lemma becomes useful when showing that, even if one utilizes f to eliminate g while neither f nor g is the ground truth, she will successfully eliminate either f or g , up to a constant-factor w.r.t. the arm pulls.

Lemma 9. *For any f, g and $\pi \geq 0$, we have $\text{IC}(f, g, \pi) \leq \max(\text{IC}^*(f, 4\pi), \text{IC}^*(g, 4\pi))$.*

Proof. Given $\pi \geq 0$, it can be easily seen that $d(f, g) := \sqrt{\text{IC}(f, g, \pi)} = \|\vec{f} - \vec{g}\|_{M_\pi}$, where we denote $\vec{h} = (h_a)_{a \in \mathcal{A}}$ and $M_\pi = \text{diag}(\frac{\pi_a}{2\sigma^2} : a \in \mathcal{A})$; therefore $d(f, g)$ is a Mahalanobis distance, hence satisfying triangle inequality. Specifically,

$$d(f, g) \leq d(f, f^*) + d(g, f^*) \leq 2 \max(d(f, f^*), d(g, f^*)).$$

Squaring both sides, we have $d(f, g)^2 \leq 4 \max(d(f, f^*)^2, d(g, f^*)^2)$, which implies that

$$\text{IC}(f, g, \pi) \leq \max(\text{IC}^*(f, 4\pi), \text{IC}^*(g, 4\pi)). \quad \square$$

Lemma 10. *In Algorithm 1, for every t , if we do not enter **Exploit**, then $\exists f \in \mathcal{F}_t : \text{IC}^*(f, 4\pi_t) \geq 1$; furthermore, if Cf_t happens, then $\exists f \in \mathring{\mathcal{F}}_t : \text{IC}^*(f, 4\pi_t) \geq 1$.*

Proof. There are three cases: $\text{Cf}_t, \text{Fs}_t, \text{Fb}_t$.

1. If Cf_t , then by the definition of $\phi(\vec{f}_t)$ and Cf_t , there exists $f \in \mathring{\mathcal{F}}_t$ such that $\text{IC}(f, \vec{f}_t, \phi(\vec{f}_t)) \geq 1$. As \vec{f}_t is also in $\mathring{\mathcal{F}}_t$, by Lemma 9, $\exists f \in \mathring{\mathcal{F}}_t : \text{IC}^*(f, 4\phi(\vec{f}_t)) \geq 1$. Therefore, we can state $\exists f \in \mathcal{F}_t : \text{IC}^*(f, 4\phi(\vec{f}_t)) \geq 1$.
2. If Fs_t , with a similar logic, either $\forall f \in \tilde{\mathcal{F}}_t : \text{IC}^*(f, 4\gamma(\vec{f}_t)) \geq 1$ or $\text{IC}^*(\vec{f}_t, 4\gamma(\vec{f}_t)) \geq 1$. As $\tilde{\mathcal{F}}_t \subseteq \mathcal{F}_t$ and $\vec{f}_t \in \mathcal{F}_t$, we have $\exists f \in \mathcal{F}_t : \text{IC}^*(f, 4\gamma(\vec{f}_t)) \geq 1$ as well.
3. If Fb_t , then with the same reasoning as the Fs_t case, either $\forall f \in \tilde{\mathcal{F}}_t : \text{IC}^*(f, 4\psi(\vec{f}_t)) \geq 1$ or $\text{IC}^*(\vec{f}_t, 4\psi(\vec{f}_t)) \geq 1$, so we have $\exists f \in \mathcal{F}_t : \text{IC}^*(f, 4\psi(\vec{f}_t)) \geq 1$ as well. \square

Recall that we use $a(t) := a_t$ to avoid double subscripts.

Lemma 11. *If Algorithm 1 enters state Cf_t, Fs_t , or Fb_t , we have $\pi_t \neq 0$. In addition, $\pi_{a(t)} \neq 0$ holds with probability 1.*

Proof. We consider three cases:

- If Cf_t , there must exist $g \in \mathcal{E}(\vec{f}_t) \setminus \{\vec{f}_t\}, \gamma(g) \not\subseteq \gamma(f)$. By the definition of ϕ , we must have $\text{IC}(\vec{f}_t, g, \phi(\vec{f}_t)) \geq 1$, implying that $\pi_t = \phi(\vec{f}_t) \neq 0$.
- If Fs_t , then by the entering condition and the fact that $\vec{f}_t \in \tilde{\mathcal{F}}_t$, $\text{IC}(\vec{f}_t, \vec{f}_t, \gamma(\vec{f}_t)) \geq 1$, implying that $\pi_t = \gamma(\vec{f}_t) \neq 0$.
- If Fb_t , then by the definition of ψ and the fact that $\mu^*(\vec{f}_t) \geq \mu^*(\vec{f}_t)$ and $a^*(\vec{f}_t) \neq a^*(\vec{f}_t)$, we must have $\forall \tilde{f} \in \tilde{\mathcal{F}}_t, \text{IC}(\vec{f}_t, \tilde{f}, \psi(\vec{f}_t)) \geq 1$, implying that $\pi_t = \psi(\vec{f}_t) \neq 0$.

For the other claim, if $\pi_{a(t)} = 0$, then we have $\frac{T_{a(t)}(t-1)}{\pi_{t,a(t)}} = \infty$, so the selection of $a(t)$ implies that $\pi_t = 0$, which is impossible by the first claim here. \square

We now present a useful lemma that formalizes the intuition that tracking (line 18) controls the number of arm pulls and thus the statistical power (i.e., information) to distinguish f^* from the rest. We remark that the key variable below is ζ , which appears three times in the LHS below.

Lemma 12. *Let $\zeta \in [0, \infty)^K$. Then, for any hypothesis f ,*

$$\overline{\text{Ex}}_t, \pi_t \propto \zeta, a_t = a, T_a(t-1) \geq \rho\zeta_a, \text{IC}^*(f, \zeta) \geq c \implies \text{IC}^*(f, T(t-1)) \geq \rho c.$$

Proof. By the definition of $a_t = \arg \min_a \frac{T_a(t-1)}{\pi_{t,a}} = \arg \min_a \frac{T_a(t-1)}{\zeta_a}$ and the condition that $T_a(t-1) \geq \rho\zeta_a$, we have, for every $b \in \mathcal{A}$,

$$\frac{T_b(t-1)}{\zeta_b} \geq \frac{T_a(t-1)}{\zeta_a} \geq \rho.$$

This implies that

$$T(t-1) \geq \rho\zeta. \quad (13)$$

As a consequence, $\text{IC}^*(f, T(t-1)) \geq \text{IC}^*(f, \rho\zeta) = \rho\text{IC}^*(f, \zeta) \geq \rho c$. \square

As an application of Lemma 12, we have the following lemma that will be useful for bounding the number of times different branches of CROP (Cf_t, Fs_t , and Fb_t) are entered.

Lemma 13. *The following statements hold:*

(i) *For any $a \in \mathcal{A}$, $t \in \mathbb{N}$ and $\rho > 0$,*

$$\mathbb{P}(\exists s \in [t]. \text{Fs}_s, \pi_s \propto \gamma_a(f^*), a_s = a, T_a(s-1) \geq \rho\gamma_a(f^*), \tilde{\mathcal{F}}_s \subseteq \mathcal{C}^*) \leq \exp\left(-\frac{1}{4\rho} \left(\rho - \frac{\beta_t}{2\sigma^2}\right)_+^2\right)$$

(ii) *For any $a \in \mathcal{A}$, $t \in \mathbb{N}$ and $\rho > 0$,*

$$\mathbb{P}(\exists s \in [t]. \overline{\text{Ex}}_s, a_s = a, T_a(s-1) \geq \rho\pi_{s,a}) \leq |\mathcal{F}| \cdot \exp\left(-\frac{\sigma^2\rho - 4\beta_t}{16\sigma^2}\right).$$

(iii) *For any $a \in \mathcal{A}$, $t \in \mathbb{N}$ and $\rho > 0$,*

$$\mathbb{P}(\exists s \in [t]. \text{Cf}_s, a_s = a, T_a(s-1) \geq \rho\pi_{s,a}) \leq |\mathcal{F}| \cdot \exp\left(-\frac{\sigma^2\rho - 4\beta_t}{16\sigma^2}\right).$$

Proof.

(i) If the event inside $\mathbb{P}(\cdot)$ happens, we have the following: there exists some $s_0 \in [t]$, such that Fs_{s_0} happens, $\pi_{s_0} \propto \gamma_a(f^*)$, $a_{s_0} = a$, $T_a(s_0-1) \geq \rho\gamma_a(f^*)$, $\tilde{\mathcal{F}}_{s_0} \subseteq \mathcal{C}^*$. As $\tilde{\mathcal{F}}_{s_0} \subseteq \mathcal{C}^*$, there must exist some $f_0 \in \mathcal{F}_{s_0}$ such that $\text{IC}(f_0, f^*, \gamma(f^*)) \geq 1$. For this $f_0 \in \mathcal{F}_{s_0}$, $L_{s_0-1}(f_0) - \min_{g \in \mathcal{F}} L_{s_0-1}(g) \leq \beta_{s_0} \implies L_{s_0-1}(f_0) - L_{s_0-1}(f^*) \leq \beta_t$. Taking $\zeta = \gamma(f^*)$ in Lemma 12, we have $\text{IC}(f_0, f^*, T(s_0-1)) \geq \rho$. Therefore,

$$\begin{aligned} & \mathbb{P}(\exists s \in [t], \text{Fs}_s, \pi_s \propto \gamma_a(f^*), a_s = a, T_a(s-1) \geq \rho\gamma_a(f^*), \tilde{\mathcal{F}}_s \subseteq \mathcal{C}^*) \\ & \leq \mathbb{P}(\exists s \in [t], f \in \mathcal{F} \cdot L_{s-1}(f) - L_{s-1}(f^*) \leq \beta_t, \text{IC}(f, f^*, T(s-1)) \geq \rho) \\ & \leq \sum_{f \in \mathcal{F}} \mathbb{P}(\exists s \in [t]. L_{s-1}(f) - L_{s-1}(f^*) \leq \beta_t, \text{IC}(f, f^*, T(s-1)) \geq \rho) \\ & \leq |\mathcal{F}| \exp\left(-\frac{1}{4\rho} \left(\rho - \frac{\beta_t}{2\sigma^2}\right)_+^2\right), \end{aligned}$$

where the second inequality is from union bound, and the last inequality is from Equation (9) in Lemma 4(ii) and algebra.

- (ii) If the event inside $\mathbb{P}(\cdot)$ happens, we have the following: there exists $s_0 \in [t]$, such that $\overline{\text{Ex}}_{s_0}$ happens, $a_{s_0} = a$, and $T_a(s_0 - 1) \geq \rho\pi_{s_0, a}$. From Lemma 10, there exists $f_0 \in \mathcal{F}_{s_0} \subseteq \{f : L_{s_0-1}(f) - L_{s_0-1}(f^*) \leq \beta_{s_0}\}$ such that $\text{IC}(f_0, f^*, 4\pi_{s_0}) \geq 1$, implying that $\text{IC}(f_0, f^*, \pi_{s_0}) \geq 1/4$. Taking $\zeta = \pi_{s_0}$ in Lemma 12, we have that $\text{IC}(f_0, f^*, T(s_0 - 1)) \geq \rho/4$. Therefore,

$$\begin{aligned} & \mathbb{P}(\exists s \in [t] \cdot \overline{\text{Ex}}_s, a_s = a, T_a(s - 1) \geq \rho\pi_{s, a}) \\ & \leq \mathbb{P}(\exists s \in [t], f \in \mathcal{F} \cdot L_{s-1}(f) - L_{s-1}(f^*) \leq \beta_s, \text{IC}(f, f^*, T(s - 1)) \geq \rho/4) \\ & \leq \mathbb{P}(\exists s \in [t], f \in \mathcal{F} \cdot L_{s-1}(f) - L_{s-1}(f^*) \leq \beta_t, \text{IC}(f, f^*, T(s - 1)) \geq \rho/4) \\ & \leq \sum_{f \in \mathcal{F}} \mathbb{P}(\exists s \in [t] \cdot L_{s-1}(f) - L_{s-1}(f^*) \leq \beta_t, \text{IC}(f, f^*, T(s - 1)) \geq \rho/4) \\ & \leq |\mathcal{F}| \exp\left(-\frac{\sigma^2 \rho - 4\beta_t}{16\sigma^2}\right), \end{aligned}$$

where the second inequality uses the fact that $\beta_s \leq \beta_t$ for all $s \in [t]$; the third inequality is by union bound; the last inequality follows from Equation (10) of Lemma 4(ii) and algebra.

- (iii) If the event inside $\mathbb{P}(\cdot)$ happens, we have the following: there exists $s_0 \in [t]$, such that Cf_{s_0} happens, $a_{s_0} = a$, and $T_a(s_0 - 1) \geq \rho\pi_{s_0, a}$. As Cf_{s_0} happens, by Lemma 10, there exists f_0 in $\mathcal{F}_{s_0} \subseteq \{f \in \mathcal{F} : L_{s_0-1}(f) - L_{s_0-1}(\bar{f}_{s_0}) \leq \dot{\beta}_{s_0}\}$ such that $\text{IC}(f_0, f^*, 4\pi_{s_0}) \geq 1$. Taking $\zeta = \pi_{s_0} = \phi(\bar{f}_{s_0})$ in Lemma 12, we have for that f_0 , $\text{IC}(f_0, f^*, T(s_0 - 1)) \geq \rho/4$. Therefore,

$$\begin{aligned} & \mathbb{P}(\exists s \in [t] \cdot \text{Cf}_s, a_s = a, T_a(s - 1) \geq \rho\pi_{s, a}) \\ & \leq \mathbb{P}(\exists s \in [t], f \in \mathcal{F} \cdot L_{s-1}(f) - L_{s-1}(f^*) \leq \dot{\beta}_s, \text{IC}(f, f^*, T(s - 1)) \geq \rho/4) \\ & \leq \mathbb{P}(\exists s \in [t], f \in \mathcal{F} \cdot L_{s-1}(f) - L_{s-1}(f^*) \leq \dot{\beta}_t, \text{IC}(f, f^*, T(s - 1)) \geq \rho/4) \\ & \leq \sum_{f \in \mathcal{F}} \mathbb{P}(\exists s \in [t] \cdot L_{s-1}(f) - L_{s-1}(f^*) \leq \dot{\beta}_t, \text{IC}(f, f^*, T(s - 1)) \geq \rho/4) \\ & \leq |\mathcal{F}| \exp\left(-\frac{\sigma^2 \rho - 4\dot{\beta}_t}{16\sigma^2}\right). \end{aligned}$$

where the second inequality uses the fact that $\dot{\beta}_s \leq \dot{\beta}_t$ for all $s \in [t]$; the third inequality is by union bound; the last inequality follows from Equation (10) of Lemma 4(ii) and algebra. \square

C.4 Main proofs

Recall that f^* is the ground truth mean rewards unless mentioned otherwise. Throughout, we use shorthands for the ground truth: $a^* := a^*(f^*)$, $\mu^* := \mu^*(f^*)$, and $\Delta_a := \Delta_a(f^*)$.

Recall that we have define $\psi(\mathcal{G})$, $\phi(\mathcal{G})$ and K_ψ in the main text. Throughout we frequently use the notation $\mathcal{A}_\zeta = \{a \in \mathcal{A} : \zeta_a > 0\}$ for vector $\zeta \in [0, \infty)^K$.

Unlike observable states such as Ex_t , Cf_t , Fs_t , and Fb_t , there are hidden states that become useful for the purpose of analysis. Based on the relationship between the hypothesis sets constructed by CROP and the hypothesis classes related to the ground truth hypothesis f^* , we define four *hidden states* of CROP:

1. 'B'ad: $B_t = \{f^* \notin \mathcal{F}_t\}$.
2. Strongly steady state: $S_t^+ = \{\tilde{\mathcal{F}}_t \subseteq \mathcal{C}^*, \bar{f}_t \in \mathcal{E}^*, \gamma(\bar{f}_t) \propto \gamma^*\}$.
3. Weakly steady state: $S_t^0 = \{\tilde{\mathcal{F}}_t \subseteq \mathcal{C}^*, \bar{f}_t \in \mathcal{E}^*, \gamma(\bar{f}_t) \not\propto \gamma^*\}$.
4. Non-steady state: $S_t^- = \{\tilde{\mathcal{F}}_t \not\subseteq \mathcal{C}^* \vee \bar{f}_t \notin \mathcal{E}^*\}$.

Note that the last three states forms a partition of the sample space, and can potentially overlap with B_t . In addition, if $\gamma^* = 0$, this would imply that $\mathcal{C}^* = \emptyset$; in this case, states S_t^+ and S_t^- will not ever be entered.

We first show a simple lemma that explains how non-steady states are related to having docile hypotheses in \mathcal{F}_t .

Lemma 14. Suppose \overline{B}_t happens. Then, the event S_t^- implies that $\exists f \in \mathcal{D}^* \cdot f \in \mathcal{F}_t$.

Proof. Recall that $S_t^- = \{\overline{f}_t \in \mathcal{E}^* \vee \widetilde{\mathcal{F}}_t \not\subseteq \mathcal{C}^*\} = \{\overline{f}_t \in \mathcal{D}^*\} \cup \{\overline{f}_t \in \mathcal{C}^*\} \cup \{\overline{f}_t \in \mathcal{E}^*, \widetilde{\mathcal{F}}_t \not\subseteq \mathcal{C}^*\}$. In addition, \overline{B}_t gives that f^* is in \mathcal{F}_t . These imply one of the following:

1. $\overline{f}_t \in \mathcal{D}^*$; in this case, we are done.
2. $\overline{f}_t \in \mathcal{C}^*$. We first note that in this case, by the definition of \mathcal{C}^* along with the unique best arm assumption (defined in Section 1), we have $\forall f \in \mathcal{C}^*, \mu^*(f) > \mu^*(f^*) = \mu^*$; this implies that $\overline{\mu}_t = \mu^*(\overline{f}_t) > \mu^*$. We consider two further subcases:
 - (a) If $\widetilde{a}_t \neq a^*$, then by the definition of $(\overline{a}_t, \overline{\mu}_t) = \arg \min_{(a, \mu) \in \mathcal{B}_t: a \neq \widetilde{a}_t} \mu$, the range of (a, μ) 's in the minimum includes (a^*, μ^*) , we have $\mu^*(\overline{f}_t) = \overline{\mu}_t \leq \mu^*$. This contradicts with our premise that $\overline{\mu}_t > \mu^*$.
 - (b) If $\widetilde{a}_t = a^*$, then consider any hypothesis $f_0 \in \widetilde{\mathcal{F}}_t$. We have

$$f_0(a^*) = \widetilde{\mu}_t \geq \overline{\mu}_t > \mu^* = f^*(a^*),$$

implying that $f_0 \in \mathcal{D}^*$.

3. $\overline{f}_t \in \mathcal{E}^* \wedge \widetilde{\mathcal{F}}_t \not\subseteq \mathcal{C}^*$. In this case, there must exist an element $f_0 \in \widetilde{\mathcal{F}}_t$ such that $f_0 \in \mathcal{D}^*$ or $f_0 \in \mathcal{E}^*$. We claim that the latter cannot happen. To see why, by the definition of $\widetilde{\mathcal{F}}_t$, it must be true that $\overline{\mathcal{F}}_t$ and $\widetilde{\mathcal{F}}_t$ belong to two different equivalence classes induced by relationship \sim . In addition, as $\overline{f}_t \in \mathcal{E}^*$, $\widetilde{\mathcal{F}}_t$ is a subset of the equivalence class \mathcal{E}^* . This implies that $\widetilde{\mathcal{F}}_t \cap \mathcal{E}^* = \emptyset$. Therefore, $f_0 \in \mathcal{D}^*$ must hold.

In summary, in all cases, we have that there exists some f in \mathcal{F}_t such that $f \in \mathcal{D}^*$. \square

We will bound the expected regret of CROP by a case-by-case analysis on the combination of observable and hidden states at each time step:

$$\begin{aligned} \mathbb{E} \text{Reg}_n &= \mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a\} \\ &\leq \mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, B_t\} + \mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \overline{B}_t, \text{Ex}_t\} \\ &\quad + \underbrace{\mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \overline{B}_t, \overline{\text{Ex}}_t, S_t^-\}}_{(Z1)} + \underbrace{\mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \overline{B}_t, \overline{\text{Ex}}_t, \overline{S}_t^-\}}_{(Z2)} \end{aligned} \quad (14)$$

Note that the first two terms are easy to bound. First,

$$\begin{aligned} \mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a\} &\leq \Delta_{\max} \mathbb{E} \sum_{t=1}^n \mathbb{1}\{B_t\} \\ &\leq \Delta_{\max} \sum_{t=1}^n \mathbb{P}\left(L_{t-1}(f^*) - \min_{f \in \mathcal{F}} L_{t-1}(f) > \beta_t\right) \\ &\leq \Delta_{\max} \left(1 + |\mathcal{F}| \sum_{t=2}^n \frac{1}{zt(\log_2 t)^2}\right) \leq 3\Delta_{\max}. \end{aligned}$$

where the second to last inequality is from Lemma 4(i), and the last inequality is by the values of z and α stated in Theorem 3, and the elementary fact that $\sum_{t=2}^{\infty} \frac{1}{t(\log_2 t)^2} \leq \frac{1}{2} + \int_2^{\infty} \frac{1}{t(\log_2 t)^2} dt \leq 2$.

Second, if \overline{B}_t and Ex_t happens, $a_t = a^*(\mathcal{F}_t) = a^*$. Therefore,

$$\mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \overline{B}_t, \text{Ex}_t\} = \mathbb{E} \Delta_{a^*} \sum_{t=1}^n \mathbb{1}\{a_t = a^*, \overline{B}_t, \text{Ex}_t\} = 0.$$

To bound the third term (Z1), we use Lemma 18 in Section C.5.

For (Z2), we decompose (Z2) to a few more sub-terms. We first have the following claim:

Claim 1. If S_t^- does not happen, then Fb_t does not happen, i.e. either Cf_t or Fs_t happens.

Proof. If S_t^- does not happen, then we have $\bar{f}_t \in \mathcal{E}^*$ and $\bar{\mathcal{F}}_t \subseteq \mathcal{C}^*$ both hold. These imply that $\mathcal{C}^* = \mathcal{C}(\bar{f}_t)$, and consequently, $\bar{\mathcal{F}}_t \subseteq \mathcal{C}(\bar{f}_t)$. But this would imply that the condition of line 13 is satisfied and the state Fb_t will not be entered. \square

The above claim indicates that (Z2) can be bounded by:

$$\begin{aligned}
(\text{Z2}) &= \mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t, \bar{S}_t^-\} + \mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Cf}_t, \bar{S}_t^-\} \\
&= \underbrace{\mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t, S_t^+\}}_{(\text{Z2-a})} + \underbrace{\mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Cf}_t, \bar{S}_t^-\}}_{(\text{Z2-b})} \\
&\quad + \underbrace{\mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t, S_t^0\}}_{(\text{Z2-c})}.
\end{aligned}$$

Now, using Lemmas 15, 16, and 17 in Section C.5, we get

$$\begin{aligned}
(\text{Z2}) &\leq P_1 \cdot \left(\ln n + 8\sqrt{\ln n \ln |\mathcal{F}|} + 2 \ln \ln n + 23 \ln |\mathcal{F}| \right) + 2K_\psi \Delta_{\max} + 64P_2 \cdot \left(\ln(\ln(n)) + \ln(|\mathcal{F}|) \right) + 2K_\psi \Delta_{\max} \\
&\quad + 200 \cdot \left(\sum_a \Delta_a \gamma_a(\mathcal{E}^*) \right) \ln |\mathcal{F}| + 5K_\psi \Delta_{\max} \\
&\leq P_1 \ln n + O \left(P_1 \left(\sqrt{\ln n \ln |\mathcal{F}|} + \ln \ln n + \ln |\mathcal{F}| \right) + P_2 \ln(\ln n) + P_3 \cdot \ln |\mathcal{F}| + K_\psi \Delta_{\max} \right)
\end{aligned}$$

where the second inequality is by algebra, and the fact that $P_3 \geq P_1 \vee P_2 \vee (\sum_a \Delta_a \gamma_a(\mathcal{E}^*))$, which in turn is from the constraint in (4) we have $\psi_a(f) \geq \phi_a(f) \vee \gamma_a(f)$ for all $a \in \mathcal{A}$ and $f \in \mathcal{F}$.

Combining the above bound on (Z2) with Lemma 18 and Equation (14), we can bound the regret of CROP as follows:

$$\begin{aligned}
\text{Reg}_n &\leq (\text{Z1}) + (\text{Z2}) + 3\Delta_{\max} \\
&\leq P_1 \ln n + 3\Delta_{\max} + O \left(P_3 (\ln |\mathcal{F}| + \ln(Q_1)) + K_\psi \Delta_{\max} \right) \\
&\quad + O \left(P_1 \left(\sqrt{\ln n \ln |\mathcal{F}|} + \ln \ln n + \ln |\mathcal{F}| \right) \right) \\
&= P_1 \ln n + O \left(P_1 \left(\sqrt{\ln n \ln |\mathcal{F}|} + \ln \ln n + \ln |\mathcal{F}| \right) + P_2 \ln(\ln(n)) + P_3 \left(\ln(|\mathcal{F}|) + \ln(Q_1) \right) + K_\psi \Delta_{\max} \right).
\end{aligned}$$

If $\gamma^* = 0$, we have $P_1 = \sum_{a \in \mathcal{A}} \Delta_a \gamma_a^* = 0$. In addition, we must have $\mathcal{C}^* = \emptyset$, implying that for all $f \in \mathcal{E}^*$, $\gamma^*(f) = 0$. This in turn implies that for all f, g in \mathcal{E}^* , $f \propto g$ is trivially true, and consequently $\phi(f) = 0$ for all $f \in \mathcal{E}^*$. Therefore, $\phi(\mathcal{E}^*) = 0$ and $P_2 = \sum_{a \in \mathcal{A}} \Delta_a \phi_a(\mathcal{E}^*) = 0$. The proof of Theorem 1 is complete. \square

C.5 Bounding the regret in each individual case

C.5.1 Bounding (Z2-a).

Recall that we use the shortcut $\gamma^* := \gamma(f^*)$. In addition, we have defined $P_1 = \sum_a \Delta_a \gamma_a^*$ and $\mathcal{A}_{\gamma^*} = \{a \in \mathcal{A} : \gamma_a^* \neq 0\}$. Note $|\mathcal{A}_{\gamma^*}| \leq K_\psi$.

Lemma 15.

$$\mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t, S_t^+\} \leq P_1 (\ln n + 8\sqrt{\ln n \ln |\mathcal{F}|} + 2 \ln \ln n + 23 \ln |\mathcal{F}|) + 3|\mathcal{A}_{\gamma^*}| \Delta_{\max}.$$

Proof. By linearity of expectation,

$$\mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t, S_t^+\} = \sum_a \Delta_a \mathbb{E} \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t, S_t^+\}. \quad (15)$$

By Lemma 11, $\{a_t = a, \text{Fs}_t, S_t^+\}$ will happen only for those a 's in \mathcal{A}_{γ^*} ; thus only the arms in \mathcal{A}_{γ^*} will contribute to the sum, which we focus on, hereafter.

With foresight, we pick $q_{1,a} = \left\lceil \frac{\beta_n}{2\sigma^2} \cdot \gamma_a^* \right\rceil$; clearly $q_{1,a} \leq \frac{\beta_n}{2\sigma^2} \cdot \gamma_a^* + 1 \leq \gamma_a^* \ln(|\mathcal{F}| n (\ln n)^2) + 1$.

Then,

$$\begin{aligned} & \mathbb{E} \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t, S_t^+\} \\ & \leq q_{1,a} + 1 + \sum_{m=1}^{\infty} \mathbb{P}\{\exists t. a_t = a, \text{Fs}_t, S_t^+, T_a(t-1) \geq q_{1,a} + m\} \\ & = q_{1,a} + 1 + \sum_{m=1}^{\infty} \mathbb{P}\{\exists t. a_t = a, \text{Fs}_t, S_t^+, T_a(t-1) \geq \gamma_a^* \cdot \frac{(q_{1,a} + m)}{\gamma_a^*}\} \\ & \leq q_{1,a} + 1 + \sum_{m=1}^{\infty} \min \left(1, |\mathcal{F}| \exp \left(-\frac{\gamma_a^*}{4(q_{1,a} + m)} \cdot \left(\frac{q_{1,a} + m}{\gamma_a^*} - \frac{\beta_n}{2\sigma^2} \right)^2 \right) \right) \\ & \leq q_{1,a} + 1 + \sum_{m=1}^{\infty} \min \left(1, |\mathcal{F}| \exp \left(-\frac{m^2}{4(q_{1,a} + m)\gamma_a^*} \right) \right), \end{aligned}$$

where the first inequality is from Lemma 8; the first equality is by decomposing $q_{1,a} + m = \gamma_a^* \cdot \frac{(q_{1,a} + m)}{\gamma_a^*}$; the second inequality is from Lemma 13(i); the third inequality is from the fact that $q_{1,a} = \left\lceil \frac{\beta_n}{2\sigma^2} \cdot \gamma_a^* \right\rceil \geq \frac{\beta_n}{2\sigma^2} \cdot \gamma_a^*$.

We now bound the last term of the last expression; observe that $\exp\left(-\frac{m^2}{4(q_{1,a} + m)\gamma_a^*}\right)$ is at most $\max\left(\exp\left(-\frac{m^2}{8q_{1,a}\gamma_a^*}\right), \exp\left(-\frac{m}{8\gamma_a^*}\right)\right)$. Therefore,

$$\begin{aligned} & \sum_{m=1}^{\infty} \min \left(1, |\mathcal{F}| \exp \left(-\frac{m^2}{4(q_{1,a} + m)\gamma_a^*} \right) \right) \\ & \leq \sum_{m=1}^{\infty} \min \left(1, |\mathcal{F}| \max \left(\exp \left(-\frac{m^2}{8q_{1,a}\gamma_a^*} \right), \exp \left(-\frac{m}{8\gamma_a^*} \right) \right) \right) \\ & \leq \sum_{m=1}^{\infty} \min \left(1, |\mathcal{F}| \exp \left(-\frac{m^2}{8q_{1,a}\gamma_a^*} \right) \right) + \sum_{m=1}^{\infty} \min \left(1, |\mathcal{F}| \exp \left(-\frac{m}{8\gamma_a^*} \right) \right) \\ & \leq \int_0^{\infty} \min \left(1, |\mathcal{F}| \exp \left(-\frac{z^2}{8q_{1,a}\gamma_a^*} \right) \right) dz + \int_0^{\infty} \min \left(1, |\mathcal{F}| \exp \left(-\frac{z}{8\gamma_a^*} \right) \right) dz. \end{aligned}$$

For the first integral, consider the boundary of its two segments $z_1 = \sqrt{8q_{1,a}\gamma_a^* \ln|\mathcal{F}|}$; it can then be simplified as

$$\begin{aligned} & \int_0^{\infty} \min \left(1, |\mathcal{F}| \exp \left(-\frac{z^2}{8q_{1,a}\gamma_a^*} \right) \right) dz \\ & = z_1 + \int_{z_1}^{\infty} |\mathcal{F}| \exp \left(-\frac{z^2}{8q_{1,a}\gamma_a^*} \right) dz \\ & \leq z_1 + \int_{z_1}^{\infty} |\mathcal{F}| \exp \left(-\frac{z^2}{8q_{1,a}\gamma_a^*} \right) \frac{z}{4q_{1,a}\gamma_a^*} dz \cdot \frac{4q_{1,a}\gamma_a^*}{z_1} \end{aligned}$$

$$\begin{aligned}
&= z_1 + (-|\mathcal{F}|) \exp\left(-\frac{z^2}{8q_{1,a}\gamma_a^*}\right) \Bigg|_{z_1}^{\infty} \cdot \frac{4q_{1,a}\gamma_a^*}{z_1} \\
&\leq 2z_1.
\end{aligned}$$

For the second integral, consider the boundary between its two segments $z_2 = 8\gamma_a^* \ln|\mathcal{F}|$; it can then be simplified as

$$\begin{aligned}
&\int_0^{\infty} \min\left(1, |\mathcal{F}| \exp\left(-\frac{z}{8\gamma_a^*}\right)\right) dz \\
&= z_2 + \int_{z_2}^{\infty} |\mathcal{F}| \exp\left(-\frac{z}{8\gamma_a^*}\right) dz \\
&= z_2 + (-8\gamma_a^*) |\mathcal{F}| \exp\left(-\frac{z}{8\gamma_a^*}\right) \Bigg|_{z_0}^{\infty} \\
&\leq 2z_2.
\end{aligned}$$

In summary, we have shown that for all a ,

$$\begin{aligned}
\mathbb{E} \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t, S_t^+\} &\leq q_{1,a} + 1 + 2z_1 + 2z_2 \\
&\leq q_{1,a} + 1 + \sqrt{32(q_{1,a} + 1)\gamma_a^* \ln|\mathcal{F}|} + 8\gamma_a^* \ln|\mathcal{F}| \\
&\leq \gamma_a^* \ln(|\mathcal{F}| n (\ln n)^2) + 2 + \sqrt{32(\gamma_a^* \ln(|\mathcal{F}| n (\ln n)^2) + 1)\gamma_a^* \ln|\mathcal{F}|} + 8\gamma_a^* \ln|\mathcal{F}| \\
&\leq \gamma_a^* (\ln n + 8\sqrt{\ln n \ln|\mathcal{F}|} + 2 \ln \ln n + 23 \ln|\mathcal{F}|) + 3.
\end{aligned}$$

Continuing Equation (15), summing over all actions $a \in \mathcal{A}_{\gamma^*}$ with weight Δ_a 's, we have

$$\begin{aligned}
&\sum_a \Delta_a \mathbb{E} \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t \wedge S_t^+\} \\
&\leq \left(\sum_a \Delta_a \gamma_a^* \right) \cdot (\ln n + 8\sqrt{\ln n \ln|\mathcal{F}|} + 2 \ln \ln n + 23 \ln|\mathcal{F}|) + 3|\mathcal{A}_{\gamma^*}| \Delta_{\max}.
\end{aligned}$$

where the first inequality is from the fact that $\Delta_a \leq \Delta_{\max}$, and the second inequality is from the definition of β_n . The lemma follows from the definition of P_1 . \square

C.5.2 Bounding (Z2-b).

Recall that $P_2 = \sum_a \Delta_a \phi_a(\mathcal{E}^*)$ and $\mathcal{A}_{\phi(\mathcal{E}^*)} = \{a \in \mathcal{A} : \phi_a(\mathcal{E}^*) > 0\}$. Note $|\mathcal{A}_{\phi(\mathcal{E}^*)}| \leq K_\psi$.

Lemma 16.

$$\mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Cf}_t, \overline{S}_t^-\} \leq 64P_2 \cdot (\ln(\ln(n)) + \ln(|\mathcal{F}|)) + 2|\mathcal{A}_{\phi(\mathcal{E}^*)}| \Delta_{\max}.$$

Proof. First, we note that by Lemma 11, if a is not in $\mathcal{A}_{\phi(\mathcal{E}^*)}$, it does not contribute to the sum, as Cf_t implies that only actions in $\mathcal{A}_{\phi(\mathcal{E}^*)}$ are taken with nonzero probability.

Next, by the linearity of expectation, we rewrite the expectation as follows:

$$\mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Cf}_t, \overline{S}_t^-\} = \sum_{a \in \mathcal{A}_{\phi(\mathcal{E}^*)}} \Delta_a \mathbb{E} \left[\sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Cf}_t, \overline{S}_t^-\} \right]$$

For any $a \in \mathcal{A}_{\phi(\mathcal{E}^*)}$,

$$\mathbb{E} \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Cf}_t, \overline{S}_t^-\} \leq 1 + \mathbb{E} \sum_{m=1}^{\infty} \mathbb{1}\{\exists t \leq n. a_t = a, T_a(t-1) \geq m, \text{Cf}_t, \overline{S}_t^-\}$$

$$\begin{aligned}
&\leq 1 + \sum_{m=1}^{\infty} \mathbb{P} \left(\exists t \leq n \cdot a_t = a, \text{Cf}_t, T_a(t-1) \geq \frac{m}{\phi_a(\mathcal{E}^*)} \pi_{t,a} \right) \\
&\leq 1 + \sum_{m=1}^{\infty} \min \left(1, \exp \left(-\frac{\frac{\sigma^2 m}{\phi_a(\mathcal{E}^*)} - 4\hat{\beta}_n}{16\sigma^2} \right) \right)
\end{aligned}$$

where the first inequality is from Lemma 8 with $\tau = 1$; the second inequality is from the fact that if $T_a(t-1) \geq m$ and \overline{S}_t^- happens, then $\overline{f}_t \in \mathcal{E}^*$, and therefore $T_a(t-1) \geq \frac{m}{\phi_a(\mathcal{E}^*)} \phi_a(\overline{f}_t) = \frac{m}{\phi_a(\mathcal{E}^*)} \pi_{t,a}$; the third inequality is from Lemma 13(iii) and $\mathbb{P}(A) \leq 1$ for any event A .

We remark that naively applying Lemma 7 instead of Lemma 8 as used in the case (Z2-a) (also used in the proofs of UCB [5] and UCB-S [27]), does not lead to the desired bound because of the aggressive confidence level of $\hat{\mathcal{F}}_t$.

Denote by $N_m = \min \left(1, \exp \left(-\frac{\frac{\sigma^2 m}{\phi_a(\mathcal{E}^*)} - 4\hat{\beta}_n}{16\sigma^2} \right) \right)$ and let $m_0 = \lceil 4\phi_a(\mathcal{E}^*) \frac{\hat{\beta}_n}{\sigma^2} \rceil$.

For $m \leq m_0 - 1$, we use the fact that $N_m \leq 1$. For $m \geq m_0$, $\{N_m\}_{m \geq m_0}$ is a geometric progression with initial value $N_{m_0} \leq 1$ and common ratio $\exp(-\frac{1}{16\phi_a(\mathcal{E}^*)})$. This implies that

$$\begin{aligned}
1 + \sum_{m=1}^{\infty} N_m &\leq m_0 + \sum_{m=m_0}^{\infty} N_m \leq m_0 + \frac{1}{1 - \exp(-\frac{1}{16\phi_a(\mathcal{E}^*)})} \\
&\leq 1 + 4\phi_a(\mathcal{E}^*) \frac{\hat{\beta}_n}{\sigma^2} + (1 + 16\phi_a(\mathcal{E}^*)) \\
&\leq 2 + 64\phi_a(\mathcal{E}^*) \ln(|\mathcal{F}| \ln(n)).
\end{aligned}$$

Chicheng: With the new setting of $\hat{\beta}_t$, the constant here can be improved. But I think it is OK not to do that where the first two inequalities are by algebra, the third inequality is from the definition of m_0 and the elementary fact that $\frac{1}{1 - \exp(-1/x)} = 1 + \frac{1}{\exp(1/x) - 1} \leq 1 + \frac{1}{((1/x)+1) - 1} = 1 + x$ for $x > 0$; the last inequality is from the definition of $\hat{\beta}_n$, $|\mathcal{F}| \geq 2$, $n \geq 2$ and algebra. Consequently,

$$\begin{aligned}
&\sum_{a \in \mathcal{A}_{\phi}(\mathcal{E}^*)} \Delta_a \mathbb{E} \left[\sum_{t=1}^n \mathbb{1} \{a_t = a, \text{Cf}_t, \overline{S}_t^-\} \right] \\
&\leq \sum_{a \in \mathcal{A}_{\phi}(\mathcal{E}^*)} \Delta_a (2 + 64\phi_a(\mathcal{E}^*) \ln(|\mathcal{F}| \ln(n))) \\
&\leq 64 \left(\sum_a \Delta_a \phi_a(\mathcal{E}^*) \right) \cdot (\ln(\ln(n)) + \ln(|\mathcal{F}|)) + 2|\mathcal{A}_{\phi}(\mathcal{E}^*)| \Delta_{\max},
\end{aligned}$$

where the second inequality uses the facts that $\Delta_a \leq \Delta_{\max}$ and algebra. The lemma follows from the definition of P_2 . \square

C.5.3 Bounding (Z2-c).

Recall that $\mathcal{A}_{\gamma}(\mathcal{E}^*) = \{a \in \mathcal{A} : \gamma_a(\mathcal{E}^*) > 0\}$. Note $|\mathcal{A}_{\gamma}(\mathcal{E}^*)| \leq K_{\psi}$.

Lemma 17.

$$\mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1} \{a_t = a, \text{Fs}_t, S_t^0\} \leq 200 \left(\sum_a \Delta_a \gamma_a(\mathcal{E}^*) \right) \ln|\mathcal{F}| + 5|\mathcal{A}_{\gamma}(\mathcal{E}^*)| \Delta_{\max}.$$

Proof. First, we note that if a is not in $\mathcal{A}_{\gamma}(\mathcal{E}^*)$, it does not contribute to the sum, as Fs_t implies that only actions in $\mathcal{A}_{\gamma}(\mathcal{E}^*)$ are taken with nonzero probability.

By linearity of expectation,

$$\mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t, S_t^0\} \leq \sum_a \Delta_a \mathbb{E} \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t, S_t^0\}. \quad (16)$$

For every $a \in \mathcal{A}_{\gamma(\mathcal{E}^*)}$, we will upper bound $\mathbb{E} \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t, S_t^0\}$. To this end, we will upper bound $C_{a,n_0} := \mathbb{E} \sum_{t=n_0+1}^{2n_0} \mathbb{1}\{a_t = a, \text{Fs}_t, S_t^0\}$, for every $n_0 \in \{2^k : k \in \{1, 2, \dots\}\}$.

We first note that if Fs_t and S_t^0 both happen, then by the definition of S_t^0 , $\bar{f}_t \sim f^*$, and $\dot{\mathcal{F}}_t \subseteq \bar{\mathcal{F}}_t \subseteq \mathcal{E}^*$; we also have $\gamma(\bar{f}_t) \not\prec \gamma(f^*)$. In addition, by the definition of Fs_t , for all $f, g \in \dot{\mathcal{F}}_t$, $\gamma(f) \prec \gamma(g)$. Therefore, it must be the case that $f^* \notin \dot{\mathcal{F}}_t$, implying that $\exists f \in \mathcal{E}^* \cdot L_{t-1}(f^*) - L_{t-1}(f) > \dot{\beta}_t$. We use this observation in the subsequent proof that we call ‘‘regret peeling’’.

With foresight, we pick $u_a = \lceil 5\gamma_a(\mathcal{E}^*) \frac{\beta_{2n_0}}{\sigma^2} \rceil$. We can bound C_{a,n_0} as follows:

$$\begin{aligned} C_{a,n_0} &= \mathbb{E} \sum_{t=n_0+1}^{2n_0} \mathbb{1}\{a_t = a, \text{Fs}_t, S_t^0\} \\ &\leq \mathbb{E} \sum_{t=n_0+1}^{2n_0} \mathbb{1}\{a_t = a, \text{Fs}_t, \exists f \in \mathcal{E}^* \cdot L_{t-1}(f^*) - L_{t-1}(f) > \dot{\beta}_t\} \\ &\leq \mathbb{E} \sum_{t=n_0+1}^{2n_0} \mathbb{1}\{a_t = a, \text{Fs}_t, \exists f \in \mathcal{E}^* \cdot L_{t-1}(f^*) - L_{t-1}(f) > \dot{\beta}_{n_0}\} \\ &\leq \mathbb{E} \mathbb{1}\left\{\exists s \in \mathbb{N}, f \in \mathcal{E}^* \cdot L_s(f^*) - L_s(f) > \dot{\beta}_{n_0}\right\} \cdot \sum_{t=n_0+1}^{2n_0} \mathbb{1}\{a_t = a, \text{Fs}_t\} \\ &\leq \mathbb{E} \mathbb{1}\left\{\exists s \in \mathbb{N}, f \in \mathcal{E}^* \cdot L_s(f^*) - L_s(f) > \dot{\beta}_{n_0}\right\} \cdot \left(u_a + \sum_{t=n_0+1}^{2n_0} \mathbb{1}\{a_t = a, \text{Fs}_t, T_a(t-1) \geq u_a\}\right) \\ &\leq \mathbb{P}\left(\exists s \in \mathbb{N}, f \in \mathcal{E}^* \cdot L_s(f^*) - L_s(f) > \dot{\beta}_{n_0}\right) \cdot u_a + \sum_{t=n_0+1}^{2n_0} \mathbb{P}(a_t = a, \text{Fs}_t, T_a(t-1) \geq u_a) \end{aligned}$$

where the second inequality uses the basic fact that $\dot{\beta}_t > \dot{\beta}_{n_0}$ for all $t \geq n_0 + 1$; the third inequality is from the basic fact that $\mathbb{1}\{A, B\} = \mathbb{1}\{A\} \cdot \mathbb{1}\{B\}$ and the fact that for predicate p , $p(t)$ implies $\exists s \cdot p(s)$.

The first term can be bounded by Lemma 4(i) and the union bound as follows:

$$\begin{aligned} &\mathbb{P}\left(\exists s \in \mathbb{N}, f \in \mathcal{E}^* \cdot L_s(f^*) - L_s(f) > \dot{\beta}_{n_0}\right) \cdot u_a \\ &\leq |\mathcal{F}| \exp\left(-\frac{\dot{\beta}_{n_0}}{2\sigma^2}\right) \cdot u_a \\ &\leq \frac{1}{(\log_2 n_0)^3} + 20\gamma_a(\mathcal{E}^*) \cdot \left(\frac{\ln|\mathcal{F}|}{(\log_2 n_0)^3} + \frac{2}{(\log_2 n_0)^2}\right). \end{aligned}$$

where the last inequality uses $n_0 \geq 2$ and $\frac{\ln(2n_0 \log(2n_0)^2)}{(\log_2(n_0))} \leq 3 \cdot \frac{\ln 2 + \ln n_0}{\log_2 n_0} \leq 5$.

Remark 2. We remark that the inequality above is the one that reflects our intuition that, even if we track a wrong $\gamma(f)$ with $f \in \mathcal{E}^*$ and suffer regret like $\sum_a \Delta_a \gamma_a(f) \ln(t)$ up to time t (that can be much larger than $\sum_a \Delta_a \gamma_a^* \ln(t)$), such an event happens with small enough probability like $O(\frac{1}{\ln(t)})$. Therefore, *in expectation*, this event contributes to the regret only as a finite term w.r.t. n . This intuition is manifested in the proof in a bit more complicated way, unfortunately, because the algorithm is designed to enjoy an anytime regret bound rather than the fixed-budget setting. Specifically, the failure rate of the confidence set $\dot{\mathcal{F}}_t$ changes over time, and we use the common technique called ‘‘peeling device’’ from concentration of measure to deal with it. If we knew the time horizon n , then one can set $\dot{\beta}_t = 4\sigma^2 \ln(|\mathcal{F}| \ln n)$ for all t to obtain the same guarantee.

Meanwhile, each subterm in the second term can be bounded using Lemma 13(ii) as follows:

$$\begin{aligned} \mathbb{P}(a_t = a, \text{Fs}_t, T_a(t-1) \geq u_a) &\leq \mathbb{P}\left(a_t = a, \text{Fs}_t, T_a(t-1) \geq \frac{u_a}{\gamma_a(\mathcal{E}^*)} \pi_{t,a}\right) \\ &\leq |\mathcal{F}| \cdot \exp\left(-\frac{\sigma^2 \frac{u_a}{\gamma_a(\mathcal{E}^*)} - \beta_t}{4\sigma^2}\right) \leq \frac{1}{(2n_0)^2}. \end{aligned}$$

where the last inequality is from the definition of u_a (which in turn implies that $\sigma^2 \frac{u_a}{\gamma_a(\mathcal{E}^*)} \geq 5\beta_{2n_0}$) and $\beta_{2n_0} \geq \beta_t$, and $\exp(-\frac{\beta_{2n_0}}{\sigma^2}) \leq \frac{1}{(2n_0)^2}$. This implies that

$$\sum_{t=n_0+1}^{2n_0} \mathbb{P}(a_t = a, \text{Fs}_t, T_a(t-1) \geq u_a) \leq \frac{1}{4n_0}.$$

In summary, we have

$$C_{a,n_0} \leq \left(\frac{1}{(\log_2 n_0)^3} + \frac{1}{4n_0}\right) + 20\gamma_a(\mathcal{E}^*) \cdot \left(\frac{\ln|\mathcal{F}|}{(\log_2 n_0)^3} + \frac{5}{(\log_2 n_0)^2}\right).$$

Now, we can upper bound $\mathbb{E} \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t, S_t^0\}$ as follows:

$$\begin{aligned} \mathbb{E} \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t \wedge S_t^0\} &\leq 2 + \sum_{k=1}^{\infty} C_{a,2^k} \\ &\leq 2 + \sum_{k=1}^{\infty} \left(\frac{1}{k^3} + \frac{1}{2^k}\right) + 16\gamma_a(\mathcal{E}^*) \cdot \left(\sum_{k=1}^{\infty} \frac{\ln|\mathcal{F}|}{k^3} + \frac{2}{k^2}\right) \\ &\leq 5 + 200\gamma_a(\mathcal{E}^*) \cdot \ln|\mathcal{F}|, \end{aligned}$$

where the last inequality is by algebra and $1 \leq 2 \ln|\mathcal{F}|$ due to $|\mathcal{F}| \geq 2$.

Using the bound above, continuing Equation (16), we have

$$\begin{aligned} \sum_{a \in \mathcal{A}_{\gamma(\mathcal{E}^*)}} \Delta_a \mathbb{E} \sum_{t=1}^n \mathbb{1}\{a_t = a, \text{Fs}_t \wedge S_t^0\} &\leq \sum_{a \in \mathcal{A}_{\gamma(\mathcal{E}^*)}} \Delta_a (5 + 200\gamma_a(\mathcal{E}^*) \cdot \ln|\mathcal{F}|) \\ &\leq 200 \left(\sum_a \Delta_a \gamma_a(\mathcal{E}^*)\right) \ln|\mathcal{F}| + 5|\mathcal{A}_{\gamma(\mathcal{E}^*)}| \Delta_{\max}, \end{aligned}$$

where the last inequality uses $\Delta_a \leq \Delta_{\max}$ for all $a \in \mathcal{A}$. \square

C.5.4 Bounding (Z1)

Recall that $P_3 = \sum_a \Delta_a \psi_a(\mathcal{F})$, $\Lambda_{\min} = \min_{f \in \mathcal{D}^*} \frac{|f(a^*) - \mu^*|}{\sigma}$ is the smallest information gap, and $Q_1 = \Lambda_{\min}^{-2} + K_{\psi}(1 + \max_a \psi_a(\mathcal{F}))$.

Lemma 18.

$$\mathbb{E} \sum_a \Delta_a \sum_{t=1}^n \mathbb{1}\{a_t = a, \overline{B}_t, \overline{\text{Ex}}_t, S_t^-\} \leq O\left(P_3(\ln|\mathcal{F}| + \ln(Q_1)) + K_{\psi} \Delta_{\max}\right). \quad (17)$$

Proof. Recall that $\alpha = 2$. With foresight, define

$$\tau = \max \left\{ t \in \mathbb{N}_+ : (t < 8) \vee \left(\frac{t}{2} < 1 + \frac{6}{\Lambda_{\min}^2 \sigma^2} \beta_t\right) \vee \left(\frac{t}{4K_{\psi}} < 1 + 12 \frac{\beta_t}{\sigma^2} (\max_a \psi_a(\mathcal{F}))\right) \right\}. \quad (18)$$

We upper bound the LHS of Equation (17) with three terms:

$$\sum_a \Delta_a \mathbb{E} \sum_{t=1}^n \mathbb{1}\{a_t = a, \overline{B}_t, \overline{\text{Ex}}_t, S_t^-\}$$

$$\begin{aligned}
&\leq \mathbb{E} \sum_a \Delta_a \sum_{t=1}^{\tau} \mathbb{1}\{a_t = a, \bar{B}_t, \bar{\mathbf{E}}x_t, S_t^-\} + \mathbb{E} \sum_{t=\tau+1}^n \sum_a \Delta_a \mathbb{1}\{a_t = a, \bar{B}_t, \bar{\mathbf{E}}x_t, S_t^-, T_{a^*}(t-1) \geq \frac{t}{2} - 1\} \\
&\quad + \mathbb{E} \sum_{t=\tau+1}^n \sum_a \Delta_a \mathbb{1}\{a_t = a, \bar{B}_t, \bar{\mathbf{E}}x_t, S_t^-, T_{a^*}(t-1) < \frac{t}{2} - 1\} \\
&\leq \underbrace{\sum_a \Delta_a \mathbb{E} \sum_{t=1}^{\tau} \mathbb{1}\{a_t = a, \bar{B}_t, \bar{\mathbf{E}}x_t, S_t^-\}}_{(Z1-a)} + \underbrace{\Delta_{\max} \sum_{t=\tau+1}^n \mathbb{P}\left(\bar{B}_t, \bar{\mathbf{E}}x_t, S_t^-, T_{a^*}(t-1) \geq \frac{t}{2} - 1\right)}_{(Z1-b)} \\
&\quad + \underbrace{\Delta_{\max} \sum_{t=\tau+1}^n \mathbb{P}\left(T_{a^*}(t-1) < \frac{t}{2} - 1\right)}_{(Z1-c)}
\end{aligned}$$

where the first inequality is by algebra; the second inequality uses the fact that $\Delta_a \leq \Delta_{\max}$ for all a , and linearity of expectation.

We bound each term respectively.

Bounding (Z1-a). Recall $\mathcal{A}_\psi = \{a \in \mathcal{A} : \psi(\mathcal{F}) > 0\}$. Define $\mathcal{A}'_\psi = \mathcal{A}_\psi \setminus \{a^*\}$. First, we note that if a is not in \mathcal{A}'_ψ , by Lemma 11, it does not contribute to the sum. This is because, if \bar{B}_t happens, the only arms being pulled is either a^* or from \mathcal{A}'_ψ (and $\Delta_{a^*} = 0$).

With the choice of $q_{3,a} = \lceil 12 \frac{\beta_\tau}{\sigma^2} \psi_a(\mathcal{F}) \rceil$, we have

$$\begin{aligned}
&\mathbb{E} \sum_{t=1}^{\tau} \mathbb{1}\{a_t = a, \bar{B}_t, \bar{\mathbf{E}}x_t, S_t^-\} \\
&\leq \mathbb{E} \sum_{t=1}^{\tau} \mathbb{1}\{a_t = a, \bar{\mathbf{E}}x_t, S_t^-\} \\
&\leq q_{3,a} + \sum_{t=1}^{\tau} \mathbb{P}\left(a_t = a, \bar{\mathbf{E}}x_t, T_a(t-1) \geq q_{3,a}\right) \quad (\because \text{Lemma 7}) \\
&\leq q_{3,a} + \sum_{t=1}^{\tau} \mathbb{P}\left(a_t = a, \bar{\mathbf{E}}x_t, T_a(t-1) \geq 12 \frac{\beta_\tau}{\sigma^2} \pi_{t,a}\right) \quad (\because q_{3,a} \geq 12 \frac{\beta_\tau}{\sigma^2} \psi_a(\mathcal{F})) \\
&\leq q_{3,a} + \sum_{t=1}^{\tau} |\mathcal{F}| \exp\left(-\frac{12\beta_\tau - 4\beta_t}{16\sigma^2}\right) \quad (\because \text{Lemma 13(ii)}) \\
&\leq 12 \frac{\beta_\tau}{\sigma^2} \psi_a(\mathcal{F}) + 1 + \sum_{t=1}^{\tau} \frac{1}{\tau (\log_2 \tau)^2} \quad (\because \beta_t \leq \beta_\tau, \text{ the definition of } \beta_\tau) \\
&\leq 24\psi_a(\mathcal{F})(\ln|\mathcal{F}| + \ln \tau) + 3 \quad (\because \sum_{t=1}^{\tau} \frac{1}{\tau (\log_2 \tau)^2} = \frac{1}{(\log_2 \tau)^2} \text{ and } \tau \geq 8) \\
&= O(\psi_a(\mathcal{F}) \cdot (\ln|\mathcal{F}| + \ln Q_1) + 1),
\end{aligned}$$

where the last inequality is from Lemma 20 below where we show $\ln \tau = O(\ln(Q_1) + \ln \ln|\mathcal{F}|)$.

Summing over all $a \in \mathcal{A}'_\psi$, we have

$$\sum_a \Delta_a \mathbb{E} \sum_{t=1}^{\tau} \mathbb{1}\{a_t = a, \bar{B}_t, \bar{\mathbf{E}}x_t, S_t^-\} = O\left(\left(\sum_a \Delta_a \psi_a(\mathcal{F})\right) (\ln|\mathcal{F}| + \ln(Q_1)) + K_\psi \Delta_{\max}\right).$$

Bounding (Z1-b). In subsequent derivations, we denote by $X_t = \frac{6}{\Lambda_{\min}^2} \frac{\beta_t}{\sigma^2}$.

$$\begin{aligned}
&\sum_{t=\tau+1}^n \mathbb{P}\left(\bar{B}_t, \bar{\mathbf{E}}x_t, S_t^-, T_{a^*}(t-1) \geq \frac{t}{2} - 1\right) \\
&\leq \sum_{t=\tau+1}^n \mathbb{P}\left(\bar{B}_t, \bar{\mathbf{E}}x_t, S_t^-, T_{a^*}(t-1) \geq X_t\right) \\
&\leq \sum_{t=\tau+1}^n \mathbb{P}\left(T_{a^*}(t-1) \geq X_t, \exists f \in \mathcal{D}^* \cdot L_{t-1}(f) - L_{t-1}(f^*) \leq \beta_t\right)
\end{aligned}$$

$$\begin{aligned}
&\leq \sum_{t=\tau+1}^n \mathbb{P} \left(\exists f \in \mathcal{D}^* \cdot \text{IC}^*(f, T(t-1)) \geq 3 \frac{\beta_t}{\sigma^2}, L_{t-1}(f) - L_{t-1}(f^*) \leq \beta_t \right) \\
&\leq \sum_{t=\tau+1}^n |\mathcal{F}| \exp \left(-\frac{2\beta_t}{4\sigma^2} \right) \\
&\leq \sum_{t=2}^n \frac{1}{t(\log_2 t)^2} \leq 2,
\end{aligned}$$

where the first inequality is from the definition of τ : for every $t > \tau$, $\frac{t}{2} - 1 \geq X_t$; the second inequality is from Lemma 14; the third inequality is from the observation that $\text{IC}^*(f, T(t-1)) \geq \frac{1}{2} T_{a^*}(t-1) \Lambda_{a^*}(f)^2 \geq \frac{1}{2} T_{a^*}(t-1) \Lambda_{\min}^2 \geq 3 \frac{\beta_t}{\sigma^2}$; the fourth inequality is from Lemma 4(ii); the last two inequalities are by algebra.

Bounding (Z1-c). We first bound $\mathbb{P}(T_{a^*}(t-1) < \frac{t}{2} - 1)$ for each t . First, denote by $I = [\lfloor \frac{t}{4} \rfloor + 1, t-1]$. In this notation, we claim that the following implication holds:

$$\left\{ T_{a^*}(t-1) < \frac{t}{2} - 1 \right\} \cap \left(\bigcap_{s \in I} \overline{B_s} \right) \subseteq \left\{ \exists a \in \mathcal{A}'_\psi, s \in I \cdot \overline{B_s}, a_s = a, T_a(s-1) \geq \frac{t}{4K_\psi} - 1 \right\}. \quad (19)$$

Indeed, if $\bigcap_{s \in I} \overline{B_s}$ holds, then the chosen arm a_s at time step s must come from $\mathcal{A}_\psi = \mathcal{A}'_\psi \cup \{a^*\}$; the reason is as follows:

1. if Ex_s , then $a_s = a^*$ is pulled;
2. otherwise, a_s is drawn from π_s which is supported on \mathcal{A}_ψ .

Throughout time interval I , we note that there are $\geq t-1 - \lfloor \frac{t}{4} \rfloor \geq \frac{3}{4}t - 1$ time steps. Given the premise that $T_{a^*}(t-1) < \frac{t}{2} - 1$, the number of arm pulls of a^* in I must be $< \frac{t}{2} - 1$; this implies that the total number of arm pulls in \mathcal{A}'_ψ in I must be greater than $(\frac{3t}{4} - 1) - (\frac{t}{2} - 1) \geq \frac{t}{4}$. By pigeonhole's principle, there exists an arm $a_0 \in \mathcal{A}'_\psi$ such that the number of arm pulls of a_0 in time span I is at least $\frac{t}{4K_\psi}$. Let s be the last time step in I when a_0 is pulled; therefore, we have $a_s = a_0$, $T_{a_0}(s-1) \geq \frac{t}{4K_\psi} - 1$, and $\overline{B_s}$ holding simultaneously, proving the above implication.

Equation (19) is equivalent to

$$\left\{ T_{a^*}(t-1) < \frac{t}{2} - 1 \right\} \cap \left(\bigcap_{s \in I} \overline{B_s} \right) \subseteq \bigcup_{s \in I} \left\{ \overline{B_s}, \exists a \in \mathcal{A}'_\psi \cdot a_s = a, T_a(s-1) \geq \frac{t}{4K_\psi} - 1 \right\}.$$

Therefore, by the elementary fact that $\mathbb{P}(U) \leq P(V) + P(\overline{V} \cap U)$ and De Morgan's Law, we have

$$\mathbb{P} \left(T_{a^*}(t-1) < \frac{t}{2} - 1 \right) \leq \mathbb{P} \left(\bigcup_{s \in I} B_s \right) + \mathbb{P} \left(\bigcup_{s \in I} \left\{ \overline{B_s}, \exists a \in \mathcal{A}'_\psi \cdot a_s = a, T_a(s-1) \geq \frac{t}{4K_\psi} - 1 \right\} \right).$$

For the first term, by Lemma 6, we have

$$\mathbb{P} \left(\bigcup_{s \in I} B_s \right) \leq \mathbb{P} \left(\bigcup_{s \geq \lfloor \frac{t}{4} \rfloor + 1} B_s \right) \leq \frac{1}{\frac{t}{4} (\log_2 \frac{t}{4})^2} \leq \frac{36}{t(\log_2 t)^2},$$

where the last inequality uses the elementary fact that $\log_2 \frac{t}{4} \geq \frac{1}{3} \log_2 t$ for $t \geq \tau \geq 8$. For the second term, we have:

$$\begin{aligned}
&\mathbb{P} \left(\bigcup_{s \in I} \left\{ \overline{B_s}, \exists a \in \mathcal{A}'_\psi \cdot a_s = a, T_a(s-1) \geq \frac{t}{4K_\psi} - 1 \right\} \right) \\
&\leq \sum_{a \in \mathcal{A}'_\psi} \mathbb{P} \left(\exists s \in I \cdot \overline{B_s}, a_s = a, T_a(s-1) \geq \frac{t}{4K_\psi} - 1 \right)
\end{aligned}$$

$$\begin{aligned}
&\leq \sum_{a \in \mathcal{A}'_\psi} \mathbb{P} \left(\exists s \in I \cdot \overline{\text{Ex}}_s, a_s = a, T_a(s-1) \geq 12 \frac{\beta_t}{\sigma^2} \cdot \psi_a(\mathcal{F}) \right) \\
&\leq \sum_{a \in \mathcal{A}'_\psi} \mathbb{P} \left(\exists s \in I \cdot \overline{\text{Ex}}_s, a_s = a, T_a(s-1) \geq 12 \frac{\beta_t}{\sigma^2} \cdot \pi_{s,a} \right) \\
&\leq K_\psi \cdot |\mathcal{F}| \cdot \frac{1}{zt^\alpha} = \frac{K_\psi}{t(\log_2 t)^2},
\end{aligned}$$

where the first inequality is by union bound; the second inequality is from the definition of τ : for all $t \geq \tau + 1$, $\frac{t}{4K_\psi} - 1 \geq 12 \frac{\beta_t}{\sigma^2} (\max_a \psi_a(\mathcal{F})) \geq 12 \frac{\beta_t}{\sigma^2} \cdot \psi_a(\mathcal{F})$ and the fact that \mathcal{A}'_ψ does not contain a^* ; the third inequality is from the observation that $\psi_a(\mathcal{F}) \geq \psi_a(\bar{f}_s) \geq \pi_{s,a}$; the fourth inequality is from Lemma 13(ii) and the fact that $\beta_t = \exp\left(-\frac{\beta_t}{2\sigma^2}\right) \leq \frac{1}{|\mathcal{F}|t(\log_2 t)^2}$; the last inequality is by algebra.

To summarize,

$$\mathbb{P} \left(T_{a^*}(t-1) < \frac{t}{2} - 1 \right) \leq \frac{K_\psi + 36}{t(\log_2 t)^2}.$$

Summing over all t 's, we get that

$$\Delta_{\max} \sum_{t=\tau+1}^n \mathbb{P} \left(T_{a^*}(t-1) < \frac{t}{2} - 1 \right) \leq \Delta_{\max} \sum_{t=3}^{\infty} \frac{K_\psi + 36}{t(\log_2 t)^2} \leq 2(K_\psi + 36)\Delta_{\max}.$$

Putting all together. Combining the bounds on (Z1-a), (Z1-b), (Z1-c), we have

$$(Z1) \leq O \left(\left(\sum_a \Delta_a \psi_a(\mathcal{F}) \right) \cdot (\ln|\mathcal{F}| + \ln(Q_1)) + K_\psi \Delta_{\max} \right).$$

Applying the definition of P_3 concludes the proof. \square

C.6 Miscellaneous lemmas

Lemma 19. Let $A, B > 0$. Then, $t < A + B \log(t) \implies \ln t < 2 \ln(B + \sqrt{A})$.

Proof. We use $\log(t) \leq \sqrt{t}$:

$$\begin{aligned}
t &< A + B \log(t) \\
&\leq A + B\sqrt{t} \\
\implies \sqrt{t} &< \frac{B + \sqrt{B^2 + 4A}}{2}
\end{aligned}$$

Taking logs on both sides,

$$\frac{1}{2} \log t = \log(\sqrt{t}) \leq \ln\left(\frac{B + \sqrt{B^2 + 4A}}{2}\right) \leq \ln(B + \sqrt{A}).$$

\square

Lemma 20. Suppose τ is defined as in Equation (18). Then,

$$\log(\tau) = O(\ln(Q_1) + \ln(\ln(|\mathcal{F}|))). \quad (20)$$

where Q_1 is defined in Theorem 3.

Proof. Let $R := \max_a \psi_a(\mathcal{F})$. The constraint set that defines τ , according to Equation (18), can be rewritten as:

$$\left\{ \tau \leq 8 \right\} \cup \left\{ (\tau \geq 9) \wedge \left(\frac{\tau}{2} < 1 + \frac{6}{\Lambda_{\min}^2} \cdot 2 \ln(|\mathcal{F}| \tau (\log_2 \tau)^2) \right) \right\} \cup \left\{ (\tau \geq 9) \wedge \left(\frac{\tau}{4K_\psi} < 1 + 12 \cdot 2R \cdot \ln(|\mathcal{F}| \tau (\log_2 \tau)^2) \right) \right\}$$

Observe that when $\tau \geq 9$, $(\log_2 \tau)^2 \leq \tau$. This implies that,

$$\begin{aligned} (\tau \geq 9) \wedge \left(\frac{\tau}{2} < 1 + \frac{6}{\Lambda_{\min}^2} \cdot 2 \ln(|\mathcal{F}| \tau (\log_2 \tau)^2) \right) &\implies \tau < 2 + \frac{64}{\Lambda_{\min}^2} \ln |\mathcal{F}| + \frac{64}{\Lambda_{\min}^2} \ln(\tau), \\ (\tau \geq 9) \wedge \left(\frac{\tau}{4K_\psi} < 1 + 12 \cdot 2R \cdot \ln(|\mathcal{F}| \tau (\log_2 \tau)^2) \right) &\implies \tau < 4K_\psi + 256K_\psi R \ln |\mathcal{F}| + 256K_\psi R \ln(\tau). \end{aligned}$$

By the definition of τ , we have

$$\begin{aligned} \tau &\leq \max \left\{ 8, 2 + \frac{64}{\Lambda_{\min}^2} \ln |\mathcal{F}| + \frac{64}{\Lambda_{\min}^2} \ln(\tau), 4K_\psi + 256K_\psi R \ln |\mathcal{F}| + 256K_\psi R \ln(\tau) \right\} \\ &\leq \max \left\{ 2 + \frac{64}{\Lambda_{\min}^2} \ln |\mathcal{F}| + \frac{64}{\Lambda_{\min}^2} \ln(\tau), 8K_\psi, 512K_\psi R \ln |\mathcal{F}| + 512K_\psi R \ln(\tau) \right\} \end{aligned}$$

where the second inequality is by $a + b \leq \max\{2a, 2b\}$. We can compactly write down $\tau \leq \max\{8K_\psi, A + B \ln(\tau)\}$ with

$$\begin{aligned} A &= \max \left\{ 2 + \frac{64}{\Lambda_{\min}^2} \ln |\mathcal{F}|, 512K_\psi R \ln |\mathcal{F}| \right\} \\ B &= \max \left\{ \frac{64}{\Lambda_{\min}^2}, 512K_\psi R \right\}. \end{aligned}$$

Then, by Lemma 19, we have

$$\ln \tau \leq \max(\ln 8K_\psi, 2 \ln(B + \sqrt{A})) \leq \max(\ln 8K_\psi, 2 \ln(A + B))$$

Let $\xi = \Lambda_{\min}^{-2} + K_\psi R$. Because $A = \Theta(1 + \xi \ln(|\mathcal{F}|))$, $B = \Theta(\xi)$, we have

$$\begin{aligned} \ln \tau &\leq O(\ln K_\psi + \ln(1 + \xi \ln(|\mathcal{F}|))) \\ &= O(\ln(K_\psi + \xi \ln(|\mathcal{F}|))) \\ &= O(\ln(\Lambda_{\min}^{-2} + K_\psi(1 + R)) + \ln(\ln(|\mathcal{F}|))) \end{aligned}$$

where the first equality uses $\ln(a) + \ln(b) = 2 \max(\ln(a), \ln(b)) \leq 2 \ln(\max(a, b))$, and the second inequality is from the definition of ξ and algebra. \square

D Lower bound

For our lower bound, we consider the following instance that resembles \mathcal{H}^+ from Figure 2.

Example 1. Let $\epsilon, \Lambda > 0$ and $r > 1$. Suppose ϵ is small enough to ensure that f_2 has the only informative arm of 3 and f_3 has the only informative arm of 4.

- $f_1 = (1, 1 + \epsilon, 0, 0)$
- $f_2 = (1, 1 - \epsilon, \Lambda, 0)$
- $f_3 = (1, 1 - \epsilon, \Lambda, r\Lambda)$

Let us denote by $\mathbb{E}_i T_j(n)$ the expected number of pulls of arm j under the instance f_i . We state our lower bound result in the following theorem.

Theorem 21. Consider Example 1. Assume the Gaussian noise model with $\sigma^2 = 1$. Suppose a bandit algorithm has $\mathbb{E}_1 T_1(n) = O(n^u)$ for some $u \in [0, 1)$, then, for sufficiently large n ,

$$\mathbb{E}_3 T_2(n) \vee \mathbb{E}_3 T_3(n) \vee \mathbb{E}_2 T_4(n) \geq \frac{6}{5} \frac{1}{r^2 \Lambda^2} \ln \left(1 + \frac{(1-u) \ln(n)}{48} \right) = \Omega(\ln(1 + (1-u) \ln n)).$$

At first sight, intuition on why the statement is true is not obvious since, assuming $\mathbb{E}_3 T_2(n) \vee \mathbb{E}_3 T_3(n)$ is $O(\ln(\ln((1-u)n)))$, somehow the fact that arm 1 is not pulled sufficiently under f_1 implies a lower bound on arm 4 under f_2 . To explain this, let us consider the contraposition: If, $\mathbb{E}_2 T_4(n) < O(\ln(\ln((1-u)n)))$ for large enough n , then $T_1(n) = \Omega(n^u)$. What happens in a nutshell is as

follows. The fact that $\mathbb{E}_2 T_4(n)$ is not sufficient means that the algorithm cannot distinguish between f_2 and f_3 with probability approaching to 1. Thus, roughly speaking, the behavior of the algorithm under f_2 and f_3 must be very similar for most times. Together with the assumption on $\mathbb{E}_3 T_2(n)$ and $\mathbb{E}_3 T_3(n)$, the algorithms does not collect sufficient number of samples from arm 2, 3, and 4 under both f_2 and f_3 . The implication is that (i) the remaining pulls all go to arm 1 and (ii) the algorithm cannot distinguish between f_1 from $\{f_2, f_3\}$ either for most times, so it collects a lot of samples from arm 1 even under f_1 . The specific reason why $\ln(\ln(n))$ appears is quite technical.

This has an implication on forced sampling, as discussed in Section 5. That is, the contraposition of Theorem 21 implies that naïvely pulling $\Theta(\ln(\ln(n)))$ for each arm may lead to regret that is arbitrarily close to being linear, let alone being uniformly good or being around the asymptotic optimality!

Proof. Let $Y_1, \dots, Y_4 > 0$ be some constants such that $Y_1 + Y_2 + Y_3 + Y_4 = n$, to be tuned later. Define $Y_{i:j}$ as Y_i, Y_{i+1}, \dots, Y_j for convenience. Denote by \mathbb{E}_i and \mathbb{P}_i the expectation and probability under f_i , respectively. Using divergence decomposition and Bretagnolle–Huber inequality [30, Lemma 15.1 and Theorem 14.2, respectively], we have

$$\begin{aligned} & \frac{1}{2} \exp\left(-\mathbb{E}_2 \left[\sum_{t=1}^n \text{KL}(f_2(a_t), f_3(a_t)) \right]\right) \\ &= \frac{1}{2} \exp\left(-\frac{r^2 \Lambda^2}{2} \mathbb{E}_2 T_4(n)\right) \\ &\leq \mathbb{P}_2(T_4(n) \geq Y_4) + \mathbb{P}_3(T_4(n) < Y_4) \\ &= \mathbb{P}_2(T_4(n) \geq Y_4) + \mathbb{P}_3(T_4(n) < Y_4, T_3(n) \geq Y_3) \\ &\quad + \mathbb{P}_3(T_4(n) < Y_4, T_3(n) < Y_3, T_2(n) \geq Y_2) \\ &\quad + \mathbb{P}_3(T_4(n) < Y_4, T_3(n) < Y_3, T_2(n) < Y_2, T_1(n) \geq Y_1) \end{aligned}$$

On the other hand,

$$\begin{aligned} & \mathbb{E}_2 T_4(n) + \mathbb{E}_3 T_3(n) + \mathbb{E}_3 T_2(n) \\ &\geq Y_4 \mathbb{P}_2(T_4(n) \geq Y_4) + Y_3 \mathbb{P}_3(T_3(n) \geq Y_3) + Y_2 \mathbb{P}_3(T_2(n) \geq Y_2) \\ &\geq \min\{Y_2, Y_3, Y_4\} \cdot (\mathbb{P}_2(T_4(n) \geq Y_4) + \mathbb{P}_3(T_3(n) \geq Y_3) + \mathbb{P}_3(T_2(n) \geq Y_2)) \end{aligned}$$

Together,

$$\begin{aligned} \frac{1}{2} \exp\left(-\frac{r^2 \Lambda^2}{2} \mathbb{E}_2 T_4(n)\right) &\leq \frac{1}{\min\{Y_{2:4}\}} (\mathbb{E}_2 T_4(n) + \mathbb{E}_3 T_3(n) + \mathbb{E}_3 T_2(n)) \\ &\quad + \underbrace{\mathbb{P}_3(T_4(n) < Y_4, T_3(n) < Y_3, T_2(n) < Y_2, T_1(n) \geq Y_1)}_{=: Q_n} \end{aligned} \quad (21)$$

The main effort is spent on bounding Q_n .

Recall that reward distribution is Gaussian with variance $\sigma^2 = 1$. Let $p_f(r_s | a_s)$ be the pdf of the reward distribution under $f^* = f$ when arm a_s is pulled at time s .

We have the following anytime inequality. Denote by $\mathbb{P}_{f^*}(\cdot)$ be the probability of an event when f^* is the ground truth. The following lemma states that under f^* the empirical KL-divergence is not too far from the KL-divergence (that is controlled by expected arm pulls made by the algorithm) with high probability.

Lemma 22. *For every $\rho > 0$,*

$$\mathbb{P}_{f^*} \left(B(f^*, f) := \left\{ \exists t \geq 1, \sum_{s=1}^t \ln \frac{p_{f^*}(r_s | a_s)}{p_f(r_s | a_s)} \geq (1 + \rho) \sum_{s=1}^t \text{KL}(f^*(a_s), f(a_s)) + \frac{1}{\rho} \ln(\delta^{-1}) \right\} \right) \leq \delta$$

Proof. Using Equation (11) of Lemma 5, we have

$$\mathbb{P}_{f^*} \left(\exists t \geq 1, \sum_s^t M_s(f) - (1 + 2\sigma^2\lambda) \sum_s^t \mathbb{E}_s[M_s(f)] \geq \frac{1}{\lambda} \ln(\delta^{-1}) \right) \leq \delta$$

Notice that $M_s = 2\sigma^2 \ln \frac{p_{f^*}(r_s | a_s)}{p_f(r_s | a_s)}$ and $\mathbb{E}_s[M_s(f)] = 2\sigma^2 \text{KL}(f^*(a_s), f(a_s))$. Then,

$$\mathbb{P}_{f^*} \left(\exists t, \sum_s^t \ln \frac{p_{f^*}(r_s | a_s)}{p_f(r_s | a_s)} - (1 + 2\sigma^2\lambda) \sum_{s=1}^t \text{KL}(f^*(a_s), f(a_s)) \geq \frac{1}{2\sigma^2\lambda} \ln(\delta^{-1}) \right) \leq \delta.$$

A simple change of variable concludes the proof. \square

Let us define $\beta_4 = \frac{2}{r^2\Lambda^2}$, $Y = \frac{w}{3}\beta_4 \ln(n)$, and $A_n = \{T_2(n), T_3(n), T_4(n) \leq Y\}$ for some $w \in (0, 1)$ that we tune later. Assume that $\mathbb{E}_1 T_1(n) = O(n^u)$ for some $u \in [0, 1)$. Recall that we want to upper bound $\mathbb{P}_3(A_n)$ from (21) for which we plan to use the change of measure argument. Specifically, we observe that, for large enough n ,

$$\mathbb{P}_1(A_n) \leq \mathbb{P}_1(T_1(n) \geq n - w\beta_4 \ln(n)) \leq \frac{\mathbb{E}_1 T_1(n)}{n - w\beta_4 \ln(n)} \leq \frac{\mathbb{E}_1 T_1(n)}{n/2} \stackrel{(a)}{\leq} c_1 n^{u-1}$$

for some constant $c_1 > 0$ where (a) is by our assumption on $\mathbb{E}_1 T_1(n)$.

Because we have set ϵ to be small enough, we have $\beta_4 = \frac{2}{r^2\Lambda^2} \leq \frac{2}{\Lambda^2} \wedge \frac{2}{4\epsilon^2}$. Then, under A_n , we have

$$\begin{aligned} \sum_{s=1}^t \text{KL}(f_3(a_s), f_1(a_s)) &= T_2(n) \cdot \frac{4\epsilon^2}{2} + T_3(n) \cdot \frac{\Lambda^2}{2} + T_4(n) \cdot \frac{r^2\Lambda^2}{2} \\ &\leq 3Y \max\left(\frac{4\epsilon^2}{2}, \frac{\Lambda^2}{2}, \frac{r^2\Lambda^2}{2}\right) \leq w \ln(n) \end{aligned} \quad (22)$$

We now lower bound $\mathbb{P}_1(A_n)$. Recall the definition of $B(\cdot, \cdot)$ from Lemma 22.

$$\begin{aligned} \mathbb{P}_1(A_n) &\geq \mathbb{P}_1\left(A_n, \overline{B(f_3, f_1)}\right) \\ &= \mathbb{E}_3 \left[\mathbb{1}\{A_n, \overline{B(f_3, f_1)}\} \prod_{t=1}^n \frac{p_1(r_t | a_t)}{p_3(r_t | a_t)} \right] \\ &= \mathbb{E}_3 \left[\mathbb{1}\{A_n, \overline{B(f_3, f_1)}\} \exp\left(-\sum_{t=1}^n \ln \frac{p_3(r_t | a_t)}{p_1(r_t | a_t)}\right) \right] \\ &\geq \mathbb{E}_3 \left[\mathbb{1}\{A_n, \overline{B(f_3, f_1)}\} \exp\left(-\left((1+\rho) \sum_{t=1}^n \text{KL}(f_3(a_t), f_1(a_t)) + \frac{1}{\rho} \ln(\delta^{-1})\right)\right) \right] \\ &\stackrel{(22)}{\geq} \mathbb{E}_3 \left[\mathbb{1}\{A_n, \overline{B(f_3, f_1)}\} \exp\left(-\left((1+\rho) \cdot w \ln(n) + \frac{1}{\rho} \ln(\delta^{-1})\right)\right) \right] \\ &= \mathbb{P}_3\left(A_n, \overline{B(f_3, f_1)}\right) \cdot (1/n)^{(1+\rho)w} \cdot \delta^{\frac{1}{\rho}} \\ &\geq (\mathbb{P}_3(A_n) - \mathbb{P}_3(B(f_3, f_1))) \cdot (1/n)^{(1+\rho)w} \cdot \delta^{\frac{1}{\rho}} \quad (\because \mathbb{P}(A) \leq \mathbb{P}(A, \overline{B}) + \mathbb{P}(B)) \\ &\geq (\mathbb{P}_3(A_n) - \delta) \cdot (1/n)^{(1+\rho)w} \cdot \delta^{\frac{1}{\rho}} \quad (\because \text{Lemma 22}) \end{aligned}$$

Combining the lower and upper bound on $\mathbb{P}_1(A_n)$ above,

$$\begin{aligned} (\mathbb{P}_3(A_n) - \delta) \cdot (1/n)^{(1+\rho)w} \cdot \delta^{\frac{1}{\rho}} &\leq c_1 n^{u-1} \\ \implies \mathbb{P}_3(A_n) &\leq \delta + c_1 n^{u-1} \cdot (1/n)^{-(1+\rho)w} \cdot \delta^{-\frac{1}{\rho}} \\ &\leq (1/n)^q + c_1 (1/n)^{1-u-(1+\rho)w-\frac{q}{\rho}} \quad (\text{set } \delta = (1/n)^q) \end{aligned}$$

By setting $w = q = \frac{1-u}{4}$ and $\rho = 1$, we have $1 - u - (1 + \rho)w - \frac{q}{\rho} = \frac{1-u}{4}$. Then, With this choice, we have

$$\mathbb{P}_3(A_n) \leq (1/n)^{\frac{1-u}{4}} + c_1 \cdot (1/n)^{\frac{1-u}{4}}$$

Using our choice of $Y_{2:4} = \frac{\beta_4}{6} \ln(n)$, we go back to where we began:

$$\begin{aligned} & \frac{1}{2} \exp\left(-\frac{r^2 \Lambda^2}{2} \mathbb{E}_2 T_4(n)\right) \\ & \leq \frac{1}{\min\{Y_{2:4}\}} \left(\mathbb{E}_2 T_4(n) + \mathbb{E}_3 T_3(n) + \mathbb{E}_3 T_2(n)\right) + Q_n \\ & \leq \frac{1}{(w/3)\beta_4 \ln(n)} \cdot 3 \cdot \left(\mathbb{E}_2 T_4(n) \vee \mathbb{E}_3 T_3(n) \vee \mathbb{E}_3 T_2(n)\right) + (1/n)^{\frac{1-u}{4}} + c_1 \cdot (1/n)^{\frac{1-u}{4}} \end{aligned}$$

Denote by $R = \mathbb{E}_3 T_3(n) \vee \mathbb{E}_3 T_2(n) \vee \mathbb{E}_2 T_4(n)$. One can see that, if R is uniformly bounded w.r.t. n , we get a contradiction because the LHS is bounded below but the RHS gets smaller with n . Therefore, R must grow indefinitely over time. This implies that, for large enough n , we have $C \leq R$ and $(1/n)^{\frac{1-u}{4}} + c_1 \cdot (1/n)^{\frac{1-u}{4}} \leq \frac{R}{(w/3)\beta_4 \ln(n)}$. Then,

$$\frac{1}{2} \exp\left(-\frac{r^2 \Lambda^2}{2} R\right) \leq \frac{4R}{(w/3)\beta_4 \ln(n)} = \frac{(12/w)R}{\frac{2}{r^2 \Lambda^2} \ln(n)}$$

It remains to solve the above for R . We do so by inverting the Lambert function. Let $Z = \frac{r^2 \Lambda^2 R}{2}$. Then,

$$\begin{aligned} \exp(-Z) & \leq \frac{(12/w)Z}{\ln(n)} \\ \frac{w}{12} \ln(n) & \leq Z \exp(Z) =: X \end{aligned}$$

We like to find $Z(X)$ that satisfies $Z(X) \exp(Z(X)) = X$. Using Orabona and Pal [34, Lemma 17], we have $\frac{3}{5} \ln(X+1) \leq Z(X) \leq \ln(X+1)$. Therefore,

$$Z(X) \geq \frac{3}{5} \ln(X+1) \geq \frac{3}{5} \ln\left(\frac{w \ln(n)}{12} + 1\right)$$

Substituting $Z(X)$ and w with their definitions concludes the proof. \square

References

- [1] Yasin Abbasi-Yadkori, David Pál, and Csaba Szepesvári. Improved Algorithms for Linear Stochastic Bandits. *Advances in Neural Information Processing Systems (NIPS)*, pages 1–19, 2011.
- [2] Alekh Agarwal, Dean P Foster, Daniel J Hsu, Sham M Kakade, and Alexander Rakhlin. Stochastic convex optimization with bandit feedback. In *Advances in Neural Information Processing Systems*, pages 1035–1043, 2011.
- [3] Rajeev Agrawal, Demosthenis Teneketzis, and Venkatachalam Anantharam. Asymptotically Efficient Adaptive Allocation Schemes for Controlled I.I.D. Processes: Finite Parameter Space. *IEEE Transactions on Automatic Control*, 1989.
- [4] Onur Atan, Cem Tekin, and Mihaela Schaar. Global multi-armed bandits with Hölder continuity. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 28–36, 2015.
- [5] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning*, 47(2–3):235–256, 2002.

- [6] Mohammad G Azar, Alessandro Lazaric, Emma Brunskill, et al. Sequential transfer in multi-armed bandit with finite set of models. In *Advances in Neural Information Processing Systems*, pages 2220–2228, 2013.
- [7] Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure Exploration in Multi-armed Bandits Problems. In *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, pages 23–37, 2009.
- [8] Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure exploration in finitely-armed and continuous-armed bandits. *Theoretical Computer Science*, 412(19):1832–1852, apr 2011. URL <https://hal-hec.archives-ouvertes.fr/hal-00609550>.
- [9] Sébastien Bubeck, Rémi Munos, Gilles Stoltz, and Csaba Szepesvári. X-armed bandits. *Journal of Machine Learning Research*, 12(May):1655–1695, 2011.
- [10] Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet. Bounded regret in stochastic multi-armed bandits. In *Journal of Machine Learning Research*, 2013.
- [11] Apostolos N Burnetas and Michael N Katehakis. Optimal adaptive policies for sequential allocation problems. *Advances in Applied Mathematics*, 17(2):122–142, 1996.
- [12] Richard Combes and Alexandre Proutiere. Unimodal bandits: Regret lower bounds and optimal algorithms. In *International Conference on Machine Learning*, pages 521–529, 2014.
- [13] Richard Combes, Stefan Magureanu, and Alexandre Proutiere. Minimal Exploration in Structured Stochastic Bandits. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1763–1771. 2017.
- [14] Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic Linear Optimization under Bandit Feedback. In *Proceedings of the Conference on Learning Theory (COLT)*, pages 355–366, 2008.
- [15] Rémy Degenne, Evrard Garcelon, and Vianney Perchet. Bandits with Side Observations: Bounded vs. Logarithmic Regret. *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2018.
- [16] Rick Durrett. *Probability: Theory and Examples*. Cambridge university press, 4th edition, 2010.
- [17] Sarah Filippi, Olivier Cappé, Aurélien Garivier, and Csaba Szepesvári. Parametric Bandits: The Generalized Linear Case. In *Advances in Neural Information Processing Systems (NIPS)*, pages 586–594. 2010.
- [18] Dylan Foster, Alekh Agarwal, Miroslav Dudík, Haipeng Luo, and Robert Schapire. Practical Contextual Bandits with Regression Oracles. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1534–1543, 2018.
- [19] Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *Proceedings of the Conference on Learning Theory (COLT)*, pages 998–1027, 2016.
- [20] Todd L Graves and Tze Leung Lai. Asymptotically Efficient Adaptive Choice of Control Laws In Controlled Markov Chains. *SIAM J. Control Optim.*, 35(3):715–743, 1997.
- [21] Samarth Gupta, Shreyas Chaudhari, Subhojyoti Mukherjee, Gauri Joshi, and Osman Yağan. A unified approach to translate classical bandit algorithms to the structured bandit setting, 2018.
- [22] Botao Hao, Tor Lattimore, and Csaba Szepesvári. Adaptive exploration in linear contextual bandit. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.

- [23] Kevin Jamieson. Some notes on multi-armed bandits. 2020. URL https://courses.cs.washington.edu/courses/cse599i/20wi/resources/bandit_notes.pdf.
- [24] Robert Kleinberg, Aleksandrs Slivkins, and Eli Upfal. Multi-armed bandits in metric spaces. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 681–690, 2008.
- [25] Robert D Kleinberg. Nearly tight bounds for the continuum-armed bandit problem. In *Advances in Neural Information Processing Systems*, pages 697–704, 2005.
- [26] T. L. Lai and Herbert Robbins. Asymptotically Efficient Adaptive Allocation Rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- [27] Tor Lattimore and Rémi Munos. Bounded regret for finite-armed structured bandits. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.
- [28] Tor Lattimore and Rémi Munos. Bounded regret for finite-armed structured bandits. In *Advances in Neural Information Processing Systems*, pages 550–558, 2014.
- [29] Tor Lattimore and Csaba Szepesvári. The end of optimism? An asymptotic analysis of finite-armed linear bandits. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 728–737, 2017.
- [30] Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020.
- [31] Stefan Magureanu, Richard Combes, and Alexandre Proutiere. Lipschitz bandits: Regret lower bound and optimal algorithms. In *Conference on Learning Theory*, pages 975–999, 2014.
- [32] Adam J Mersereau, Paat Rusmevichientong, and John N Tsitsiklis. A structured multiarmed bandit problem and the greedy policy. *IEEE Transactions on Automatic Control*, 54(12):2787–2802, 2009.
- [33] Jungseul Ok, Alexandre Proutiere, and Damianos Tranos. Exploration in structured reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 8874–8882, 2018.
- [34] Francesco Orabona and David Pal. Coin Betting and Parameter-Free Online Learning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 577–585, 2016.
- [35] Paat Rusmevichientong and John N Tsitsiklis. Linearly Parameterized Bandits. *Math. Oper. Res.*, 35(2):395–411, 2010.
- [36] Daniel 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.
- [37] Cong Shen, Ruida Zhou, Cem Tekin, and Mihaela van der Schaar. Generalized global bandit and its application in cellular coverage optimization. *IEEE Journal of Selected Topics in Signal Processing*, 12(1):218–232, 2018.
- [38] William R. Thompson. On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples. *Biometrika*, 25(3/4):285, 1933.
- [39] Andrea Tirinzoni, Alessandro Lazaric, and Marcello Restelli. A novel confidence-based algorithm for structured bandits. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- [40] Sattar Vakili and Qing Zhao. Achieving complete learning in Multi-Armed Bandit problems. In *Asilomar Conference on Signals, Systems and Computers*, pages 1778–1782, 2013.
- [41] Jean Ville. Etude critique de la notion de collectif. *Bull. Amer. Math. Soc*, 45(11):824, 1939.

- [42] Zhiyang Wang, Ruida Zhou, and Cong Shen. Regional multi-armed bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 510–518, 2018.
- [43] Jia Yuan Yu and Shie Mannor. Unimodal bandits. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, pages 41–48, 2011.