
Structure Adaptive Algorithms for Stochastic Bandits

Rémy Degenne¹ Han Shao² Wouter M. Koolen³

Abstract

We study reward maximisation in a wide class of structured stochastic multi-armed bandit problems, where the mean rewards of arms satisfy some given structural constraints, e.g. linear, unimodal, sparse, etc. Our aim is to develop methods that are *flexible* (in that they easily adapt to different structures), *powerful* (in that they perform well empirically and/or provably match instance-dependent lower bounds) and *efficient* in that the per-round computational burden is small. We develop asymptotically optimal algorithms from instance-dependent lower-bounds using iterative saddle-point solvers. Our approach generalises recent iterative methods for pure exploration to reward maximisation, where a major challenge arises from the estimation of the sub-optimality gaps and their reciprocals. Still we manage to achieve all the above desiderata. Notably, our technique avoids the computational cost of the full-blown saddle point oracle employed by previous work, while at the same time enabling finite-time regret bounds. Our experiments reveal that our method successfully leverages the structural assumptions, while its regret is at worst comparable to that of vanilla UCB.

1. Introduction

Stochastic multi-armed bandits are online learning problems in which an algorithm sequentially chooses among a finite set of actions (it “pulls an arm”) and receives a stochastic reward in return. The goal is to obtain the maximal cumulative reward over time. Introduced by Thompson (1933), bandits have been the subject of intense study for many years now, starting with asymptotic results in the 80s and 90s (Lai & Robbins, 1985; Graves & Lai, 1997) and moving to the fi-

¹INRIA - DIENS - PSL Research University, Paris, France ²Toyota Technological Institute at Chicago ³Centrum Wiskunde & Informatica. Correspondence to: Rémy Degenne <remydegenne@gmail.com>.

nite time analysis of algorithms since the beginning of the new millennium, notably with the introduction of strategies based on the Optimism in Face of Uncertainty principle (Agrawal, 1995; Auer et al., 2002). For an overview of the field, we point the reader to (Bubeck & Cesa-Bianchi, 2012; Lattimore & Szepesvári, 2019).

Recently many studies focused on problems inspired from practical applications, in which additional prior information is available. For example the means of the arms might be known to be sparse (Kwon et al., 2017), to form a Lipschitz function (Magureanu et al., 2014) or have a linear structure (Dani et al., 2008; Abbasi-Yadkori et al., 2011), or a combinatorial one (Cesa-Bianchi & Lugosi, 2012; Kveton et al., 2015). All these assumptions have been regrouped under the name “structured bandit” (Combes et al., 2017). As was the case in that last work, we aim at providing a family of algorithms that adapt to the known structure.

Bandit strategies realize a trade-off between exploitation (pulling the apparently best arm) and exploration. In the related problem of bandit pure exploration, in which no reward is gained but the goal is to answer a query, algorithms that adapt to any structure (and query) have been recently developed, first with asymptotic (Garivier & Kaufmann, 2016a), then with non-asymptotic guarantees (Degenne et al., 2019). Our approach here will be to develop these techniques further, and obtain results for reward maximisation.

1.1. Structured stochastic bandits

A finite number K of arms have reward distributions $(\nu_k)_{k \in [K]}$ in a known sub-Gaussian canonical exponential family with one parameter (the mean of the distribution) in an open interval $\Theta \subseteq \mathbb{R}$. The vector of means of these distributions is denoted by $\mu \in \mathbb{R}^K$. The arm with highest mean for vector μ is denoted by $i^*(\mu)$ and we suppose that it is unique. We write μ^* for the value of that highest mean. An arm with mean below μ^* is called sub-optimal. For $x, y \in \Theta$, we denote by $d(x, y)$ the Kullback-Leibler divergence from the distribution parametrised by x to that parametrised by y .

The mean vector μ is known to belong to a set $\mathcal{M} \subseteq \Theta^K$, which represents the structure of the problem: the structural knowledge restricts the set of admissible mean parameters. For example, Lipschitz bandits prescribe that the means

of successive arms cannot be far apart. We make the assumption that there exists a compact set $\mathcal{C} \subseteq \Theta^K$ such that $\mathcal{M} \subseteq \mathcal{C}$, which is convenient for the proof, but most likely not necessary. We also restrict our attention to what we call *regular* structures: those verifying Assumption 7 in appendix C. These include all examples of structures we found in the literature. The simplified Assumption 1 below presents the same idea, is verified by all structures we consider but the Sparse one and is presented here instead of the more detailed one to preserve the brevity of the exposition. For $j \in [K]$ we define $\neg j = \text{cl}(\{\mu \in \mathcal{M} \mid i^*(\mu) \neq j\})$ the closure of the set of alternative structured problems where the best arm is not j . Let also $\neg_x i = \{\lambda \in \mathbb{R}^K : \lambda^i = x\} \cap \neg i$.

Assumption 1 (Regularity, simplified). There exists $c_{\mathcal{M}}$ such that for all $\mu \in \mathcal{M}$, denoting $i^*(\mu) = *$, if $\neg_{\mu^*} * \neq \emptyset$ then for all $\lambda \in \neg_*$ there exists $\xi \in \neg_{\mu^*} *$ such that for all $k \in [K]$, $|\xi^k - \lambda^k| \leq c_{\mathcal{M}} |\mu^* - \lambda^*|$.

The interaction between the algorithm and its environment is the following: at stage $t \geq 1$, **(1)** the algorithm pulls arm $k_t \in [K]$, **(2)** it observes a reward sample $X_t^{k_t} \sim \nu_{k_t}$, **(3)** its total reward is accrued by the mean reward μ^{k_t} (which is unobserved). An algorithm is a sequence of functions, one for each time $t \in \mathbb{N}$, that take $(k_1, X_1^{k_1}, \dots, k_{t-1}, X_{t-1}^{k_{t-1}})$ as input and return $k_t \in [K]$.

We define the gap of an arm $k \in [K]$ as $\Delta^k = \mu^* - \mu^k$. Let N_T^k be the number of pulls of arm k up to time T . The goal of a bandit algorithm is to maximise its cumulated expected reward. Subtracting obtained reward from achievable reward, we arrive at the standard evaluation metric of expected regret,

$$\mathbb{E} R_T = T\mu^* - \sum_{t=1}^T \mathbb{E} \mu^{k_t} = \mathbb{E} \sum_{t=1}^T \Delta^{k_t} = \sum_{k=1}^K \Delta^k \mathbb{E} N_T^k.$$

1.2. Contributions

An algorithm is said to be asymptotically optimal if its regret verifies that for all $\mu \in \mathcal{M}$, $\mathbb{E}_{\mu}[R_t]/\ln(t)$ converges to a constant $V^{\mathcal{M}}(\mu)$, which matches the constant prescribed by the corresponding lower bound (see Section 2.1).

- We introduce a family of algorithms that are asymptotically optimal for all structures \mathcal{M} , while having explicit non-asymptotic regret bounds.
- We exhibit members of that family with computational complexity much lower than that of earlier structure-adaptive algorithms for many structures, since they never solve fully the lower bound minimax problem.
- On experiments, we verify that the proposed algorithms adapt to the structure. Their regret is close to that of the UCB algorithm up to a time depending on the complexity $V^{\mathcal{M}}(\mu)$ of the problem instance, after which the structural information is successfully exploited and the regret is of order $V^{\mathcal{M}}(\mu) \ln T$.

1.3. Related work on structure adaptive methods

Lower bounds for the regret of asymptotically consistent algorithms (i.e. with regret $o(T)$ for all $\mu \in \mathcal{M}$) take the form of a constrained minimisation problem (Lai & Robbins, 1985; Graves & Lai, 1997). The OSSB algorithm (Combes et al., 2017) uses a test to decide whether to exploit or explore (we will use a similar mechanism). If it explores, it solves a plug-in estimate of the lower bound problem, and pulls arms in order to approach the corresponding optimal sampling behaviour. Its analysis is based on the continuity of the optimiser, seen as a function of the bandit instance μ : if the plug-in estimate is close enough to the true mean vector, the estimated pulling proportions are close to optimal. In order to make sure that the empirical means converge to the true means, it pulls all arms $o(\ln T)$ times, a technique that is called “forced exploration”. The OSSB algorithm is claimed to be asymptotically optimal up to a multiplicative factor, which can be as close to 1 as wanted, with the drawback that the asymptotic regime is delayed.

A similar idea of tracking a plug-in estimate of the optimal pulling proportions was used before in the setting of fixed-confidence pure exploration, with the Track-and-Stop algorithm of Garivier & Kaufmann (2016a) (generalised to combinatorial settings by Chen et al. 2017). Track-and-Stop achieves asymptotic optimality for any structure under the above continuity assumption (Kaufmann & Koolen, 2018), and without it (Degenne & Koolen, 2019). In recent work, a game-based algorithm was developed which gets non-asymptotic guarantees for any structure in the pure exploration setting (Degenne et al., 2019). We take inspiration from that line of work.

Our algorithm uses an explore or exploit test as is done in OSSB. If it explores, it treats the lower bound as the value of a zero-sum game between two players. We implement two regret-minimising algorithms, one for each player, and their interaction ensures convergence to the value of the game.

2. Algorithm and Results

Our goal is to design efficient algorithms for the problem of reward maximisation in structured stochastic multi-armed bandit models. We specifically aim to incur little regret. To calibrate our expectations about what can be achieved (at least in principle), we start by reviewing the classic asymptotic lower bound of (Graves & Lai, 1997).

2.1. Asymptotic Lower Bound

Consider any algorithm, and assume that it is reasonable in the sense of *asymptotic consistency*, meaning that for any $\mu \in \mathcal{M}$, any sub-optimal arm $k \notin \text{argmax}_i \mu^i$ is sampled only sub-polynomially often, i.e. $\mathbb{E}_{\mu}[N_T^k] \in \cap_{\alpha>0} o(T^{\alpha})$.

Theorem 1 (Graves & Lai 1997). *An algorithm asymptotically consistent for bandit structure \mathcal{M} must incur regret*

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_{\mu} [R_T]}{\ln T} \geq V^{\mathcal{M}}(\mu) \quad \text{for any instance } \mu \in \mathcal{M}$$

where the instance-dependent asymptotic regret rate $V^{\mathcal{M}}(\mu)$ is the value of the optimisation problem (here $\Lambda := \neg i^*(\mu)$)

$$\min_{N \geq 0} \sum_k N^k \Delta^k \quad \text{s.t.} \quad \inf_{\lambda \in \Lambda} \sum_k N^k d(\mu^k, \lambda^k) \geq 1. \quad (1)$$

Given this insurmountable limit, we take as our goal to construct algorithms that satisfy $\mathbb{E}_{\mu} [R_T] \leq V^{\mathcal{M}}(\mu) \ln T + o(\ln T)$, with explicit finite-time control over the lower-order term.

2.2. Perturbed Game and Saddle Point Problems

In a multi-armed bandit, the optimal arm has zero gap, $\Delta^* = 0$. This creates several technical complications later on, which we choose to avoid by picking a small $\varepsilon > 0$ and defining the ε -perturbed gaps $\Delta_{\varepsilon}^k = \Delta^k \vee \varepsilon$. We call $V_{\varepsilon}^{\mathcal{M}}(\mu)$ the instance-dependent regret rate (1) with these perturbed gaps substituted. We find three useful saddle point expressions.

Lemma 1. *The reciprocal $D_{\varepsilon}^{\mathcal{M}}(\mu) := 1/V_{\varepsilon}^{\mathcal{M}}(\mu)$ satisfies*

$$D_{\varepsilon}^{\mathcal{M}}(\mu) = \max_{\mathbf{w} \in \Delta} \inf_{\lambda \in \Lambda} \frac{\sum_k w^k d(\mu^k, \lambda^k)}{\sum_k w^k \Delta_{\varepsilon}^k} \quad (2a)$$

$$= \max_{\tilde{\mathbf{w}} \in \Delta} \inf_{\lambda \in \Lambda} \sum_k \tilde{w}^k \frac{d(\mu^k, \lambda^k)}{\Delta_{\varepsilon}^k} \quad (2b)$$

$$= \inf_{\mathbf{q} \in \Delta(\Lambda)} \max_k \frac{\mathbb{E}_{\lambda \sim \mathbf{q}} [d(\mu^k, \lambda^k)]}{\Delta_{\varepsilon}^k} \quad (2c)$$

Compared to (1), problem (2a) is parametrised by the **fraction of rounds** $w^k \propto N^k$ spent pulling each arm k . The objective here is quasi-concave in \mathbf{w} but not concave. The rewrite (2b) uses the **fraction of regret** $\tilde{w}^k \propto N^k \Delta_{\varepsilon}^k$ incurred by pulling arm k . This objective is linear (hence concave) in $\tilde{\mathbf{w}}$. From either form, randomising $\lambda \sim \mathbf{q}$ licenses the min-max swap resulting in (2c).

These expressions correspond to a zero-sum game in which a pure strategy for the learner selects an arm $k \in [K]$, while for the opponent it picks a confusing instance $\lambda \in \neg i^*(\mu)$. The resulting payoff is then the information-regret ratio $\frac{d(\mu^k, \lambda^k)}{\Delta_{\varepsilon}^k}$, which quantifies the rate of progress in satisfying the information constraint in (1) per unit of objective value (i.e. regret) invested. In (2) the moves are ordered, upon which the outermost player needs to employ randomisation.

We quantify the cost of perturbation (Proof in Appendix E)

Theorem 2. *Under Assumption 1, there is a $c > 0$ such that for small ε we have $D_{\varepsilon}^{\mathcal{M}} \geq D^{\mathcal{M}} - c\sqrt{\varepsilon D^{\mathcal{M}}}$ and hence $V_{\varepsilon}^{\mathcal{M}} \leq V^{\mathcal{M}} + c\sqrt{\varepsilon V^{\mathcal{M}}}$.*

2.3. Noise-Free Case

In this section we assume that we know the bandit model $\mu \in \mathcal{M}$, and our goal is to compute the perturbed lower-bound value $V_{\varepsilon}^{\mathcal{M}}(\mu)$ and a matching strategy \mathbf{N} from (1) or \mathbf{w} , $\tilde{\mathbf{w}}$ or \mathbf{q} from (2). Our approach will be to pick either form (2b) or (2c), and run an online learner for the outer player against best response for the inner player. We will run this interaction for a carefully selected number of rounds n to get our result.

In the following we will call the maximising player, controlling k , the k -player, even if this player randomises $k \sim \tilde{\mathbf{w}}$. Similarly we will talk about the minimising λ -player. To treat the structure in a modular way, we will make the following computational assumption:

Assumption 2. We are given an *alt-min* oracle computing

$$(\mu, \mathbf{w}, j, k) \mapsto \operatorname{argmin}_{\lambda \in \mathcal{M}: \lambda_k \geq \lambda_j} \sum_i w^i d(\mu^i, \lambda^i) \quad (3)$$

for any vector $\mu \in \Theta^K$, non-negative weights \mathbf{w} and arms $j \neq k$. This implies tractability of argmin over \mathcal{M} and $\neg j$.

This assumption is satisfied (either by a closed-form expression, a binary search or full-blown numerical convex optimisation) for all structures we use for the experiments in Section 3 (unconstrained, sparse, linear, concave, unimodal and categorised).

2.3.1. λ -LEARNER

In this section we will target the saddle point (2c) using an online learner for λ . Each round t , this λ -learner outputs a distribution \mathbf{q}_t on Λ . Given \mathbf{q}_t , we define the k -opponent to play best response, i.e.

$$k_t \in \operatorname{argmin}_k \frac{\mathbb{E}_{\lambda \sim \mathbf{q}_t} [d(\mu^k, \lambda^k)]}{\Delta_{\varepsilon}^k}. \quad (4)$$

We then have the λ -learner update based on the linear loss function $\ell_t(\mathbf{q}) := \mathbb{E}_{\lambda \sim \mathbf{q}} [d(\mu^{k_t}, \lambda^{k_t})]$. A λ -learner regret bound of \mathcal{B}_n^{λ} for n rounds of interaction gives us

$$\begin{aligned} \sum_{t=1}^n \ell_t(\mathbf{q}_t) &\leq \inf_{\lambda \in \Lambda} \sum_{t=1}^n d(\mu^{k_t}, \lambda^{k_t}) + \mathcal{B}_n^{\lambda} \\ &= \inf_{\lambda \in \Lambda} \sum_k N_n^k d(\mu^k, \lambda^k) + \mathcal{B}_n^{\lambda}. \end{aligned} \quad (5)$$

By definition of k_t ,

$$\begin{aligned} \sum_{t=1}^n \ell_t(\mathbf{q}_t) &= \sum_{t=1}^n \Delta_\varepsilon^{k_t} \frac{\mathbb{E}_{\lambda \sim \mathbf{q}_t} [d(\mu^{k_t}, \lambda^{k_t})]}{\Delta_\varepsilon^{k_t}} \\ &\stackrel{(4)}{=} \sum_{t=1}^n \Delta_\varepsilon^{k_t} \max_k \frac{\mathbb{E}_{\lambda \sim \mathbf{q}_t} [d(\mu^k, \lambda^k)]}{\Delta_\varepsilon^k}. \end{aligned}$$

We then have, abbreviating $R_n^\varepsilon = \sum_{t=1}^n \Delta_\varepsilon^{k_t}$,

$$\begin{aligned} \sum_{t=1}^n \ell_t(\mathbf{q}_t) &\geq \max_k \sum_{t=1}^n \Delta_\varepsilon^{k_t} \frac{\mathbb{E}_{\lambda \sim \mathbf{q}_t} [d(\mu^k, \lambda^k)]}{\Delta_\varepsilon^k} \\ &= R_n^\varepsilon \max_k \sum_{t=1}^n \frac{\Delta_\varepsilon^{k_t}}{R_n^\varepsilon} \frac{\mathbb{E}_{\lambda \sim \mathbf{q}_t} [d(\mu^k, \lambda^k)]}{\Delta_\varepsilon^k} \\ &\geq R_n^\varepsilon \inf_{\mathbf{q}} \max_k \frac{\mathbb{E}_{\lambda \sim \mathbf{q}} [d(\mu^k, \lambda^k)]}{\Delta_\varepsilon^k} \\ &\stackrel{(2c)}{=} R_n^\varepsilon D_\varepsilon^{\mathcal{M}}(\boldsymbol{\mu}). \end{aligned} \quad (6)$$

We will choose n adaptively, running the algorithm as long as $\inf_{\lambda \in \Lambda} \sum_k N_n^k d(\mu^k, \lambda^k) \leq \ln T$. Chaining (5) and (6) then yield $R_n^\varepsilon D_\varepsilon^{\mathcal{M}}(\boldsymbol{\mu}) \leq \ln T + \mathcal{B}_n^\lambda$, so that

$$R_n := \sum_{t=1}^n \Delta^{k_t} \leq R_n^\varepsilon \leq V_\varepsilon^{\mathcal{M}}(\boldsymbol{\mu}) (\ln T + \mathcal{B}_n^\lambda).$$

At this point we need \mathcal{B}_n^λ to be $o(\ln T)$ for $V_\varepsilon^{\mathcal{M}}(\boldsymbol{\mu})$ to be the dominant term. This will be feasible, as $n = \mathcal{O}(\ln T)$ and \mathcal{B}_n^λ can be taken to be $\mathcal{O}(\sqrt{n})$.

In particular, because $n\varepsilon \leq R_n^\varepsilon \leq V_\varepsilon^{\mathcal{M}}(\boldsymbol{\mu}) (\ln T + c\sqrt{n})$,

$$\sqrt{n} \leq \frac{cV_\varepsilon^{\mathcal{M}}(\boldsymbol{\mu}) + \sqrt{c^2V_\varepsilon^{\mathcal{M}}(\boldsymbol{\mu})^2 + 4\varepsilon V_\varepsilon^{\mathcal{M}}(\boldsymbol{\mu}) \ln T}}{2\varepsilon}. \quad (7)$$

So that indeed $\sqrt{n} = o(\ln T)$. Moreover, these asymptotics start kicking in once $\ln T \gg \frac{c^2V_\varepsilon^{\mathcal{M}}(\boldsymbol{\mu})}{4\varepsilon}$. Conversely, to ensure asymptotic optimality we need to pick $\varepsilon = \omega(1/\ln T)$. In our next section we will develop an anytime version using a time-decaying ε_t .

The practicality of implementing a λ -learner depends on the geometry of the sets $\neg j$ and the makeup of the loss function. One may always write $\neg j$ as a union of $K-1$ cells, which are intersections of \mathcal{M} with one linear inequality

$$\neg j = \bigcup_{k \neq j} \{\lambda \in \mathcal{M} \mid \lambda_k \geq \lambda_j\}.$$

For the structures considered in the experiments section (except the sparse structure), we find that these cells are in fact convex sets (while $\neg j$ is not). Moreover, for Gaussian and Bernoulli bandits the relative entropy function $d(\cdot, \cdot)$ is strongly convex in its second argument. Hence we can instantiate a sub-learner on each cell. An interesting choice

is Follow-the-Leader, which is tractable by Assumption 2 and incurs $\mathcal{O}(\ln t)$ regret. On the meta-level we aggregate the sub-learner iterates using an experts algorithm (we use AdaHedge by De Rooij et al. (2014) for $\mathcal{O}(\sqrt{t})$ regret).

2.3.2. k -LEARNER

In this section we will target the saddle point (2b). We will work with a k -learner that in each round t outputs an action \tilde{w}_t (recall these are the desired per-arm fractions of regret, not of rounds). Given \tilde{w}_t from the k -learner, we define time proportions w_t and the best-response opponent $\lambda_t \in \Lambda$ by

$$w_t^k \propto \tilde{w}_t^k / \Delta_\varepsilon^k, \quad (8)$$

$$\lambda_t \in \operatorname{argmin}_{\lambda \in \Lambda} \sum_k w_t^k d(\mu^k, \lambda^k). \quad (9)$$

The k -learner then updates using the linear gain function

$$g_t(\tilde{w}) = \left(\sum_k w_t^k \Delta_\varepsilon^k \right) \sum_k \tilde{w}_t^k \frac{d(\mu^k, \lambda_t^k)}{\Delta_\varepsilon^k}. \quad (10)$$

Note that the w_t^k in its leading factor do *not* vary with the argument \tilde{w} ; they are computed from \tilde{w}_t using (8).

A k -learner regret bound of \mathcal{B}_n^k provides

$$\sum_{t=1}^n g_t(\tilde{w}_t) \geq \max_k \sum_{t=1}^n \left(\sum_j w_t^j \Delta_\varepsilon^j \right) \frac{d(\mu^k, \lambda_t^k)}{\Delta_\varepsilon^k} - \mathcal{B}_n^k. \quad (11)$$

Moreover, the total gain satisfies

$$\begin{aligned} \sum_{t=1}^n g_t(\tilde{w}_t) &= \sum_{t=1}^n \left(\sum_k w_t^k \Delta_\varepsilon^k \right) \sum_k \tilde{w}_t^k \frac{d(\mu^k, \lambda_t^k)}{\Delta_\varepsilon^k} \\ &\stackrel{(8)}{=} \sum_{t=1}^n \sum_k w_t^k d(\mu^k, \lambda_t^k) \\ &\stackrel{(9)}{=} \sum_{t=1}^n \inf_{\lambda \in \Lambda} \sum_k w_t^k d(\mu^k, \lambda^k) \\ &\leq \inf_{\lambda \in \Lambda} \sum_{t=1}^n \sum_k w_t^k d(\mu^k, \lambda^k). \end{aligned} \quad (12)$$

Running as long as $\inf_{\lambda \in \Lambda} \sum_{t=1}^n \sum_k w_t^k d(\mu^k, \lambda^k) \leq \ln T$ and chaining (11) with (12) results in

$$\begin{aligned} \ln T &\geq \max_k \sum_{t=1}^n \left(\sum_j w_t^j \Delta_\varepsilon^j \right) \frac{d(\mu^k, \lambda_t^k)}{\Delta_\varepsilon^k} - \mathcal{B}_n^k \\ &= R_n^\varepsilon \max_k \sum_{t=1}^n \frac{\left(\sum_j w_t^j \Delta_\varepsilon^j \right)}{R_n^\varepsilon} \frac{d(\mu^k, \lambda_t^k)}{\Delta_\varepsilon^k} - \mathcal{B}_n^k \end{aligned} \quad (13)$$

where we abbreviated $R_n^\varepsilon = \sum_{t=1}^n \sum_k w_t^k \Delta_\varepsilon^k$. Minimizing over $\mathbf{q} \in \Delta(\Lambda)$,

$$\begin{aligned} \ln T &\geq R_n^\varepsilon \inf_{\mathbf{q} \in \Delta(\Lambda)} \max_k \frac{\mathbb{E}_{\lambda \sim \mathbf{q}} [d(\mu^k, \lambda^k)]}{\Delta_\varepsilon^k} - \mathcal{B}_n^k \\ &= R_n^\varepsilon D_\varepsilon^{\mathcal{M}}(\boldsymbol{\mu}) - \mathcal{B}_n^k, \end{aligned} \quad (14)$$

All in all we showed

$$R_n := \sum_{t=1}^n \sum_k w_t^k \Delta^k \leq R_n^\varepsilon \leq V_\varepsilon^{\mathcal{M}}(\boldsymbol{\mu}) (\ln T + \mathcal{B}_n^k)$$

and again we need to set things up so that $\mathcal{B}_n^k = o(\ln T)$. Now this may work, as we will have $n = \mathcal{O}(\ln T)$ and $\mathcal{B}_n^k = \mathcal{O}(\sqrt{n})$ so that, in total, $\mathcal{B}_n^k = \mathcal{O}(\sqrt{\ln T})$.

Tuning ε is a bit trickier than for (7), since the range of the gain function (10) scales with $1/\varepsilon$ and so hence will \mathcal{B}_n^k . Still $n = \mathcal{O}(V_\varepsilon/\varepsilon \ln T)$ and hence the asymptotics kick in for $\ln T \gg (V_\varepsilon/\varepsilon)^3$, or equivalently $\varepsilon \gg V_\varepsilon/(\ln T)^{1/3}$. We expect that power of T is artificial; the range of practically observed gains (10) may well be constant.

A bandit algorithm cannot pull any \mathbf{w}_t in the simplex, as would be prescribed by the k -learner. It has to pull one arm at each time. We circumvent that difficulty by using a so-called tracking procedure, which ensures that for all times $N_t \approx \sum_{s=1}^t \mathbf{w}_s$. We will use $k_t \in \operatorname{argmin}_{k \in [K]} N_{t-1}^k - \sum_{s=1}^t w_s^k$ (breaking ties arbitrarily). A similar rule was analysed in Garivier & Kaufmann (2016a). Our analysis reveals that it ensures that for all $t \in \mathbb{N}$ and $k \in [K]$,

$$-\ln(K) \leq N_t^k - \sum_{s=1}^t w_s^k \leq 1.$$

The previous result (Garivier & Kaufmann, 2016a, Lemma 15) has $K-1$ instead of our $\ln K$. Our proof makes use of a subtle invariant; it can be found in Appendix B.

In terms of implementation, we need to supply two things. First, a learner for linear losses defined on the simplex. This is a standard experts problem, for which we use AdaHedge (De Rooij et al., 2014). Second, we need to compute the best response $\lambda_t \in \neg i^*(\boldsymbol{\mu})$. This is where the structure dependence of the approach manifests; this step is tractable by Assumption 2. The computational complexity of either the k -learner or λ -learner based iterative saddle point approach is dominated by $K-1$ alt-min oracle calls every round.

2.4. Saddle Point-Based MAB Reward Maximisation

In this section we build atop the noise-free saddle point computation presented above. We obtain algorithms for regret minimisation in structured bandit problems that have finite-time bounds which convey, in particular, asymptotic optimality. The k -learner variant is displayed as Algorithm 1.

Algorithm 1 SP_k Learner

Require: Exploration threshold $f(t)$

Require: Learner \mathcal{A} for linear losses on the simplex

- 1: Start an instance \mathcal{A}_j for each arm j .
 - 2: Pull each arm once and get $\hat{\boldsymbol{\mu}}_K$.
 - 3: Initialise $N_t^k = 1$ and $n_t^k = 0$ for all $k \in [K]$.
 - 4: **for** $t = K + 1, \dots, T$ **do**
 - 5: **if** there is an exploit-worthy i for which (15) **then**
 - 6: $k_t = i$ (pick any if there are several suitable).
 - 7: **else**
 - 8: Estimate best arm: $j_t = \operatorname{argmax}_k \hat{\mu}_{t-1}^k$.
 - 9: **if** $n_t^{j_t}$ is even **then** $k_t = j_t$. **else**
 - 10: Compute confidence interval $[\underline{\mu}_{t-1}^k, \overline{\mu}_{t-1}^k]$ for every arm $k \in [K]$ using (16).
 - 11: Estimate gaps: $\tilde{\Delta}_{\varepsilon_t}^k = \max\{\varepsilon_t, \overline{\mu}_{t-1}^{j_t} - \overline{\mu}_{t-1}^k\}$.
 - 12: Get $\tilde{\mathbf{w}}_t$ from \mathcal{A}_{j_t} , compute $w_t^k \propto \tilde{w}_t^k / \tilde{\Delta}_{\varepsilon_t}^k$.
 - 13: Find the best response in $\neg j_t$ to \mathbf{w}_t given $\hat{\boldsymbol{\mu}}_{t-1}$:

$$\lambda_t = \operatorname{argmin}_{\lambda \in \neg j_t} \sum_k w_t^k d(\hat{\mu}_{t-1}^k, \lambda^k).$$
 - 14: Compute estimates UCB_t^k as in (17)
 - 15: Feed $g_t(\tilde{\mathbf{w}}) = \left(\sum_k w_t^k \tilde{\Delta}_{\varepsilon_t}^k\right) \sum_k \tilde{w}_t^k \text{UCB}_t^k$ to \mathcal{A}_{j_t} .
 - 16: Pull $k_t = \operatorname{argmin}_{k \in [K]} N_{t-1}^k - \sum_{s=1}^t w_s^k$.
 - 17: **end if**
 - 18: Increase $n_t^{j_t}$ by 1.
 - 19: **end if**
 - 20: Access reward $X_t^{k_t}$, update $\hat{\boldsymbol{\mu}}_t$ and N_t .
 - 21: **end for**
-

The main challenge is that we do not know $\boldsymbol{\mu}$ (nor anything derived, including the gaps Δ_ε and the index of the best arm $i^*(\boldsymbol{\mu})$), and that we hence need to estimate these live.

We sketch an outline of the construction. The main distinction made by our algorithm is whether to *explore* or *exploit*. Exploitation occurs when we are certain enough that we have the right best arm for our bandit $\boldsymbol{\mu}$, that is when

$$\exists i \in [K], \min_{\lambda \in \neg i} \sum_k N_{t-1}^k d(\hat{\mu}_{t-1}^k, \lambda^k) > f(t-1), \quad (15)$$

where we will use a threshold $f(t)$ ($\approx \ln t$ plus lower order) that is high enough so that the cumulative contribution to the regret of rounds where (15) fails is bounded by a constant. Note that we cannot afford a failure probability $\geq \frac{1}{t}$, for then the contribution of the failure cases would be of order $\geq \ln T$, voiding asymptotic optimality.

During the other rounds, our algorithm explores. The main idea here is to pick an online learner for either the k or λ side, and adapt the corresponding noise-free pipeline (13) and (14) or (5) and (6). These rounds all happen under

the complement of (15), which now fulfils the role of the stopping condition “running as long as . . .” in Section 2.3.

We now go over the major upgrades one by one, relating back to our Algorithm 1 as we go along.

Estimating $i^*(\mu)$. First, in Section 2.3 we assumed we know the best arm $i^*(\mu)$. Here that is no longer the case, and instead we have to rely on estimates, which can be wrong. Our approach to deal with this is to run not one but K many saddle point computations, one for every candidate best arm. In each exploration round, we produce an estimate j_t of the best arm, and only advance the saddle point interaction corresponding to that estimate. This way, the saddle point iteration for arm $i^*(\mu)$ will simulate a pure trajectory as per Section 2.3. Using concentration of measure and invoking our learners’ regret bounds, we show (in Appendix C.7.2) that the other instances corresponding to incorrect estimates of the best arm will only make a lower-order $\mathcal{O}(\ln \ln T)$ contribution to the regret. With that out of the way, we can reason about exploration rounds in which $j_t = i^*(\mu)$.

Estimating μ . Our second problem is then that we need to estimate the bandit instance μ and its perturbed sub-optimality gaps Δ_ε^k . We will be guided by the following idea. We want to look at the noise-free chain (13) and (14), and relate all occurrences of estimates $\hat{\mu}_t, \tilde{\Delta}_{\varepsilon_t}^k$, to their unknown truth μ and $\Delta_{\varepsilon_t}^k$. We will take inspiration from the classic UCB analysis, and use concentration to bound the instantaneous estimation error after n exploration rounds by $\sqrt{\ln(n)/n}$. These errors then accumulate over n exploration rounds to an order $\sqrt{n \ln n}$ overhead. This is a lower-order term as long as the number n of exploration steps is $o(\ln T)^2$.

More precisely, let E_s^* be the event that we explore at stage s and have the right guess for $i^*(\mu)$ and define $n_t^* = \sum_{s=1}^t \mathbb{1}\{E_s^*\}$. With estimates $\hat{\mu}_t$, we have similarly to equation (12) that if we explore at state $t + 1$,

$$\ln t \geq \inf_{\lambda \in \Lambda} \sum_{s \leq t: E_s^*} \sum_{k=1}^K w_s^k d(\hat{\mu}_t^k, \lambda^k).$$

Our algorithm cannot play at time $s < t$ based on $\hat{\mu}_t$, since it does not know it yet. It will use the available estimate $\hat{\mu}_{s-1}$. We then quantify how far these estimated means are from μ with a concentration event \mathcal{E}_n^{exp} , which states that for all $k \in [K]$ and all times t ,

$$\begin{aligned} \mu^k &\in [\underline{\mu}_t^k(n), \overline{\mu}_t^k(n)], \text{ with } \underline{\mu}_t^k(n), \overline{\mu}_t^k(n) \text{ solution to} \\ N_t^k d(\hat{\mu}_t^k, x) &= \ln(n) + \mathcal{O}(\ln \ln t). \end{aligned} \quad (16)$$

That event is such that apart from a small number of rounds, when we explore event $\mathcal{E}_{n_t^*}^{exp}$ happens. If we can ensure that

n_t^* is less than a power of $\ln(t)$, then these intervals are of order $\ln \ln t$. We get

$$\ln t \geq \inf_{\lambda \in \Lambda} \sum_{s \leq t: E_s^* \cap \mathcal{E}_{n_t^*}^{exp}} \sum_{k=1}^K w_s^k d(\hat{\mu}_{s-1}^k, \lambda^k) - \mathcal{O}(\sqrt{n_t^*}).$$

In order to continue along computations (12), (13) and (14), we need estimates for the gaps. Based on the small confidence intervals $[\underline{\mu}_t^k, \overline{\mu}_t^k]$ (where $\overline{\mu}_t^k = \overline{\mu}_t^k(n_t^*)$), we define $\tilde{\Delta}_t^k = \max\{\varepsilon, \underline{\mu}_t^k - \overline{\mu}_t^k\}$. Then up to a $\mathcal{O}(\sqrt{n_t^*})$ term

$$\begin{aligned} \ln t &\geq \inf_{\lambda \in \Lambda} \sum_{s \leq t: E_s^* \cap \mathcal{E}_{n_t^*}^{exp}} \left(\sum_{j=1}^K w_s^j \tilde{\Delta}_s^j \right) \sum_{k=1}^K \tilde{w}_s^k \frac{d(\hat{\mu}_{s-1}^k, \lambda^k)}{\tilde{\Delta}_s^k} \\ &\geq \sum_{s \leq t: E_s^* \cap \mathcal{E}_{n_t^*}^{exp}} \left(\sum_{j=1}^K w_s^j \tilde{\Delta}_s^j \right) \sum_{k=1}^K \tilde{w}_s^k \frac{d(\hat{\mu}_{s-1}^k, \lambda^k)}{\tilde{\Delta}_s^k}. \end{aligned}$$

Optimism. Using the estimated gains directly could lead to high regret, in the same way as the Follow-The-Leader strategy might accrue linear regret in stochastic bandits. We use the now widely employed Optimism in Face of Uncertainty principle (Agrawal, 1995) and replace the gains by an optimistic estimate. In the hypothetical noiseless case, we would like at time s to feed to the k -learner the gain vector $k \mapsto (\sum_{i=1}^K w_s^i \Delta_{\varepsilon_s}^i) \frac{d(\mu^k, \lambda_s^k)}{\Delta_{\varepsilon_s}^k}$ (see (10)). We define

$$\text{UCB}_t^k = \max_{\xi \in [\underline{\mu}_{t-1}^k, \overline{\mu}_{t-1}^k]} \frac{d(\xi, \lambda_s^k)}{\max\{\varepsilon_t, \mathbb{1}\{k \neq j_t\}(\mu_{t-1}^{j_t} - \xi)\}}. \quad (17)$$

Under the event that $\mu^k \in [\underline{\mu}_{t-1}^k, \overline{\mu}_{t-1}^k]$, this is an upper bound for the factor depending on k in the gain of the k -learner, up to the difference between $\overline{\mu}_{t-1}^{j_t}$ and μ^{j_t} , which will be small since we ensure that j_t is pulled every other round in the exploration phase. Again under that event, $\tilde{\Delta}_{t-1}^k (\text{UCB}_t^k - d(\hat{\mu}_{t-1}^k, \lambda_s^k) / \tilde{\Delta}_{t-1}^k)$ is upper bounded, such that the total cost of the introduction of those upper confidence bounds is $\mathcal{O}(\sqrt{n_t^*})$. We can now finish the computations of equations (13) and (14): up to $\mathcal{O}(\sqrt{n_t^*})$,

$$\begin{aligned} \ln t &\geq \sum_{s \leq t: E_s^* \cap \mathcal{E}_{n_t^*}^{exp}} \left(\sum_{j=1}^K w_s^j \tilde{\Delta}_s^j \right) \sum_{k=1}^K \tilde{w}_s^k \text{UCB}_s^k \\ &\geq \max_{k \in K} \sum_{s \leq t: E_s^* \cap \mathcal{E}_{n_t^*}^{exp}} \left(\sum_{j=1}^K w_s^j \tilde{\Delta}_s^j \right) \text{UCB}_s^k \\ &\geq \max_{k \in K} \sum_{s \leq t: E_s^* \cap \mathcal{E}_{n_t^*}^{exp}} \left(\sum_{j=1}^K w_s^j \tilde{\Delta}_s^j \right) \frac{d(\mu^k, \lambda_s^k)}{\Delta_{\varepsilon_s}^k} \\ &\geq R_t^{\varepsilon, *} D_\varepsilon^{\mathcal{M}}(\mu) - \mathcal{O}(\sqrt{n_t^*}). \end{aligned}$$

Finally, since $n_t^* \leq R_t^{\varepsilon,*} / \varepsilon$ this is an equation which asymptotically gives $R_T^{\varepsilon,*} \leq V^{\mathcal{M}}(\boldsymbol{\mu}) \ln T$ as desired. Since this is the only term in the regret which is not $o(\ln T)$, we have asymptotic optimality if we have $\varepsilon \rightarrow 0$ at a suitable rate.

Effect of the confidence intervals on the sampling behaviour. During the first rounds, the algorithm does not exploit yet. Hence the number of exploration rounds n_t is t . If the best arm is pulled enough to be well estimated, our gap estimates are $\tilde{\Delta}_t^k \approx \mu^{j_t} - \mu^k - \sqrt{2 \ln(n_t) / N_t^k}$ in the Gaussian case. For these estimates to be representative of the true gaps $\mu^{j_t} - \mu^k$, we need $N_t^k \propto \ln(n_t) / (\Delta^k)^2$. This is of order $\ln \ln t$ once n_t is logarithmic in t , but it is $\ln(t) / (\Delta^k)^2$ at the beginning. We indeed verify experimentally that during the first rounds our algorithm pulls similarly to vanilla UCB. Afterwards, the exploitation regime starts and the structure adaptive sampling begins.

How our algorithm relates to previous techniques. Our algorithm employs a explore/exploit test based on a log-likelihood ratio, as is done in the OSSB algorithm (Combes et al., 2017). While the exploration phase of OSSB is akin to the exploration algorithm Track-and-Stop (Garivier & Kaufmann, 2016b), our exploration phase is inspired from the game point of view of (Degenne et al., 2019). The strategy presented above generalizes that exploration algorithm in several key aspects (the pure exploration case is recovered by forcing all gaps Δ^k to 1).

- addressing the regret saddle-point problem with online learning algorithms is non-trivial even for known gaps, and requires several innovations (Section 2.3). For example, one of the players has to play regret proportions instead of pull proportions, as in (8), and gap-weighted regret bounds need to be employed, as in (10).
- the unknown gaps need to be estimated appropriately (line 11 of Alg. 1),
- one of the gaps is 0, which is a problem since we divide by the gaps: we introduce the epsilon-perturbed problem to solve that issue,
- we need to relate the values of the perturbed and original games (Theorem 2).

2.5. Main Result

Theorem 3. *For any k -learner or λ -learner with regret bound of order $\mathcal{O}(\sqrt{n})$ after n steps, the expected regret of our structure adaptive algorithm verifies for all $\boldsymbol{\mu} \in \mathcal{M}$,*

$$\lim_{T \rightarrow +\infty} \frac{\mathbb{E}_{\boldsymbol{\mu}}[R_T]}{\ln T} \leq V^{\mathcal{M}}(\boldsymbol{\mu}).$$

The arguments presented above in fact lead to a finite time bound, but it contains many $o(\ln T)$ terms that we choose not to detail here. Proof in Appendix C.

3. Experiments

As mentioned in Section 2.3, the asymptotics kick in when $\ln T \gg \frac{c^2 V_{\varepsilon_T}^{\mathcal{M}}(\boldsymbol{\mu})}{4\varepsilon_T}$ for the λ -learner and $\ln T \gg (V_{\varepsilon_T}^{\mathcal{M}}(\boldsymbol{\mu}) / \varepsilon_T)^3$ for the k -learner. This indicates that for small time horizons, the asymptotic theoretical guarantees become meaningless. It is worth noting that this is ubiquitous in the asymptotic literature. We still investigate empirically whether and when these algorithms start exploiting the structures and how they perform in small horizons. Implementation details are in Appendix G.


We perform four series of experiments. In all cases, rewards are Gaussian with variance 1. The structures are:

- Figure 1(a): unconstrained. This is the usual stochastic bandit setting, without additional structure.
- Figures 1(b) and 1(c): categorised bandits with strong dominance (Jedor et al., 2019). The arms belong to one of two categories (and this information on each arm is known), and all arms of one category have means higher than all arms of the other (but which category is the “good” one is unknown).
- Figure 1(d) : linear (Auer, 2002). There exists an unknown parameter $\theta \in \mathbb{R}^d$ with $d < K$ and known vectors $x_1, \dots, x_K \in \mathbb{R}^d$ such that $\mu^k = x_k^T \theta$.
- Figures 1(e) and 1(f): unimodal (Combes & Proutiere, 2014). The mean function $k \mapsto \mu^k$ is unimodal, i.e. $\mu^1 \leq \mu^2 \leq \dots \leq \mu^* \geq \dots \geq \mu^{K-1} \geq \mu^K$.
- Figures 1(g) and 1(h): sparse (Kwon et al., 2017). Given sparsity s and value γ , s of K arms have means above γ while others have means γ .


We mainly compare the following four algorithms,

- SP_k , our saddle point algorithm based on a k -learner. The learner used is AdaHedge (De Rooij et al., 2014).
- SP_{λ} , our algorithm based on a λ -learner. In all experiments except the sparse structure, the minimisation over the alternative sets is equivalent to minimising on the union of a small number of convex sets. For each such convex set, we implement one expert running Follow-The-Leader, and the λ -learner aggregates these experts using AdaHedge. For the sparse setting, there is still a decomposition into a finite number of convex sets, but that number is combinatorially large and we did not implement SP_{λ} in that case.
- The OSSB algorithm of (Combes et al., 2017), which solves a plug-in estimate of optimisation problem (1) with empirical estimates and tracks its solution.
- The vanilla UCB algorithm of (Auer et al., 2002; Garivier & Cappé, 2011), which selects the arm with the highest upper confidence bound calculated without additional structural information.


The implementation of our algorithms departs from the




(a) Unconstrained bandits
 $\mu = [0, 0.33, 0.67, 1]$.



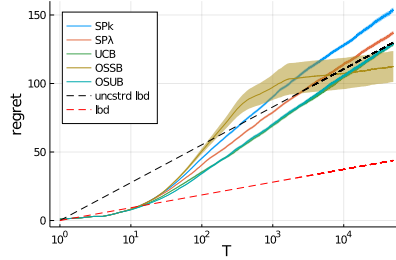
(b) Categorical bandits with $\mu^1 = [2]$ and
 $\mu^2 = [1, 0.96, 0]$.



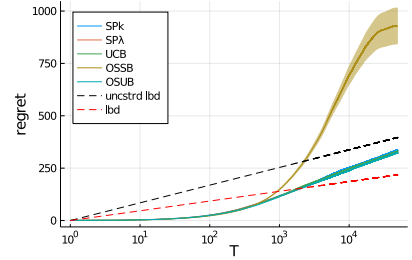
(c) Categorical bandits with $\mu^1 = [3, 1.50]$
 and $\mu^2 = [1.44, 1.44, 0]$.




(d) Linear bandits with arms
 $\mathbf{x} = [(\cos(\rho), \sin(\rho)), \rho \in \{2i\pi/5, 2i\pi/5 + 0.15\}_{i=0}^4]$ and
 $\theta = (0, 2)$.



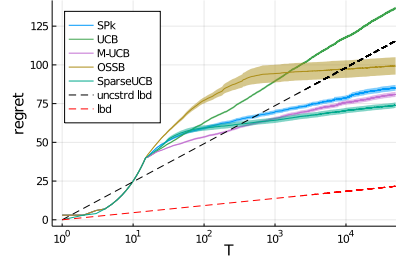
(e) Unimodal bandits with
 $\mu = [0.2, 0.4, 0.9, 0.7, 0.1]$.



(f) Unimodal bandits with
 $\mu = [0, 0, 0, 1.00, 0, 0, 0]$.



(g) Sparse bandits with $s = 1$, $\gamma = 0.3$ and
 $\mu_1 = 0.8$.



(h) Sparse bandits with $s = 2$, $\gamma = 0$ and
 $[\mu_1, \mu_2] = [3.00, 2.00]$.

Figure 1. Regret of algorithms.

theory in two ways. First, we set $\varepsilon_t \propto 1/n_t^{\hat{j}_t}$, where $n_t^{\hat{j}_t}$ is the number of exploration steps with the current estimated best arm. This is a faster decrease than allowed by our proof. Second, the width of the small confidence intervals (16) is smaller: $\ln(n_t)$ instead of $\ln(n_t) + \mathcal{O}(\ln \ln t)$.

We also compare with structure specific algorithms:

- The structural UCB algorithm, denoted by \mathcal{M} -UCB, which looks for the highest upper confidence bound over the intersection of the confidence region and the structure. In the linear case, it uses an adapted estimation of the means, such that it is the same as Lin-UCB (Auer, 2002).
- The CATSE algorithm of (Jedor et al., 2019) for cat-

egorised instances, which eliminates all arms in the “bad” category once we are confident to tell which category is “bad”.

- The “End of Optimism” algorithm of (Lattimore & Szepesvári, 2017) (called “EOO” on figures) for linear instances, which is similar to OSSB with different forced exploration.
- The OSUB algorithm of (Combes & Proutiere, 2014) for unimodal instances, which pulls among the neighbourhood of the empirical optimal arm.
- The SparseUCB algorithm of (Kwon et al., 2017) for sparse instances, which constructs a set of $\geq s$ arms which are estimated to have means larger than γ and then play UCB among this set.

Figure 1 reports the mean regret of these algorithms over 200 repetitions. The shaded areas show an interval around the mean of width twice its empirical standard deviation. We contrast the empirical regret of the algorithms with the unconstrained lower bound (called “unconstrd lbd” in the figures) $\sum_{k \neq i^*(\mu)} \Delta_k / d(\mu^k, \mu^*) \ln(T)$ and the structural lower bound (called “lbd”) $V^{\mathcal{M}}(\mu) \ln(T)$.

We investigate when the theoretically motivated SP_k , SP_λ and OSSB also benefit from good empirical performances. OSSB adapts to the structure in more experiments (the slope of its regret is close to the slope of the structural lower bound, see Figures 1(a), 1(c), 1(e), 1(h)), but it incurs a high initial regret in cases 1(a), 1(b), 1(c), 1(f). In Figure 1(g), OSSB performs poorly and is equivalent to Follow-The-Leader, since its exploration test never triggers on that structure. In contrast, our algorithms are safer, in the sense that their regret is at worst comparable to that of UCB, but they adapt to the structure often only later, sometimes after the chosen horizon. We verify that adaptation to the structure indeed happens, notably in Figures 1(b) and 1(h).

We are not able to provide a characterisation of the problems on which these algorithms manage to adapt to the structure for small horizons. Obtaining such a characterization, or methods that would adapt to the structure for small times, is the most obvious open problem for structured bandits.

4. Open Questions in Structured Bandits

While our algorithms are asymptotically optimal, the dependence of the non-asymptotic bound on problem parameters like the number of arms is definitely not the best possible. Some structures would allow an algorithm to face a problem with a huge number of arms and not suffer from that multiplicity, while our bounds have terms linear in K . Broadly, our algorithm is suited to large times, while adaptivity to structure in a small horizon regime remains to be explored.

On the topic of computational complexity, we remarked that our algorithm never solves the lower bound problem completely, but for example in the case of the k -learner only computes the best-response over an alternative set. If the structure set \mathcal{M} is complicated, that computation can still be expensive, or even infeasible. However one could argue that it does not make sense to compute exactly the best response to a noisy problem, and indeed our algorithm remains asymptotically optimal if that minimization is approximate, with an error at the n^{th} exploration step of order $1/\sqrt{n}$.

A weakness of the analysis of algorithms based on an explore/exploit test like ours is the concentration inequality used to bound the number of times that the exploitation phase is wrongly entered. A concentration result gives a threshold $\beta(t, \delta)$ such that with probability $1 - \delta$, a deviation

is lower than $\beta(t, \delta)$ for all times. That threshold is linear in the number of arms K , while it could be smaller for many particular structures. For example, in the unconstrained case, it can be made to depend only on $\ln(K)$. (see also the discussion on the Pure Exploration Problem Rank by Kaufmann & Koolen (2018)). Adapting the concentration inequality to the structure is in general an open question.

Finally, the behaviour of our algorithms in experiments has an initial phase where the structure information cannot yet be used. The length of this phase depends on the instance complexity (our bounds currently admit an exponential dependence here), and in some of our experiments we have not decidedly left the initial phase. To shorten this phase without changing the high-level approach, one could work on improving confidence regions for the gradients. Our current choice, UCB_t^k in (17), is possibly quite conservative.

Acknowledgements

This work was funded in part by the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute). H. Shao is supported in part by the National Science Foundation under grant CCF-181501. H. Shao would like to thank Emilie Kaufmann for her early discussion on unimodal bandits.

References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pp. 2312–2320, 2011.
- Agrawal, R. Sample mean based index policies by $o(\log n)$ regret for the multi-armed bandit problem. *Advances in Applied Probability*, 27(4):1054–1078, 1995.
- Auer, P. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- Boyd, S. and Vandenberghe, L. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004. ISBN 0521833787.
- Bubeck, S. and Cesa-Bianchi, N. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012.

- Cesa-Bianchi, N. and Lugosi, G. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422, 2012.
- Chen, L., Gupta, A., Li, J., Qiao, M., and Wang, R. Nearly optimal sampling algorithms for combinatorial pure exploration. In *Proceedings of the 2017 Conference on Learning Theory*, volume 65, pp. 482–534. PMLR, July 2017.
- Combes, R. and Proutiere, A. Unimodal bandits: Regret lower bounds and optimal algorithms. In *International Conference on Machine Learning*, pp. 521–529, 2014.
- Combes, R., Magureanu, S., and Proutiere, A. Minimal exploration in structured stochastic bandits. In *Advances in Neural Information Processing Systems*, pp. 1763–1771, 2017.
- Dani, V., Hayes, T. P., and Kakade, S. M. Stochastic linear optimization under bandit feedback. In *21st Annual Conference on Learning Theory - COLT 2008, Helsinki, Finland, July 9-12, 2008*, pp. 355–366. Omnipress, 2008.
- De Rooij, S., Van Erven, T., Grünwald, P. D., and Koolen, W. M. Follow the leader if you can, hedge if you must. *The Journal of Machine Learning Research*, 15(1):1281–1316, 2014.
- Degenne, R. and Koolen, W. M. Pure exploration with multiple correct answers. In *Advances in Neural Information Processing Systems (NeurIPS) 32*, pp. 14564–14573. Curran Associates, Inc., December 2019.
- Degenne, R., Koolen, W. M., and Ménard, P. Non-asymptotic pure exploration by solving games. In *Advances in Neural Information Processing Systems*, 2019.
- Garivier, A. and Cappé, O. The KL-UCB algorithm for bounded stochastic bandits and beyond. In *Proceedings of the 24th annual conference on learning theory*, pp. 359–376, 2011.
- Garivier, A. and Kaufmann, E. Optimal best arm identification with fixed confidence. In *Proceedings of the 29th Conference On Learning Theory (COLT)*, 2016a.
- Garivier, A. and Kaufmann, E. Optimal best arm identification with fixed confidence. In *Conference on Learning Theory*, pp. 998–1027, 2016b.
- Graves, T. L. and Lai, T. L. Asymptotically efficient adaptive choice of control laws in controlled Markov chains. *SIAM Journal on Control and Optimization*, 35(3):715–743, 1997.
- Jedor, M., Perchet, V., and Louedec, J. Categorized bandits. In *Advances in Neural Information Processing Systems*, pp. 14399–14409, 2019.
- Kaufmann, E. and Koolen, W. M. Mixture martingales revisited with applications to sequential tests and confidence intervals. Preprint, October 2018.
- Kveton, B., Wen, Z., Ashkan, A., and Szepesvári, C. Tight regret bounds for stochastic combinatorial semi-bandits. In *Artificial Intelligence and Statistics*, pp. 535–543, 2015.
- Kwon, J., Perchet, V., and Vernade, C. Sparse stochastic bandits. In *Conference On Learning Theory*, 2017.
- Lai, T. L. and Robbins, H. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1): 4–22, 1985.
- Lattimore, T. and Szepesvári, C. The end of optimism? an asymptotic analysis of finite-armed linear bandits. In *Artificial Intelligence and Statistics*, pp. 728–737, 2017.
- Lattimore, T. and Szepesvári, C. Bandit algorithms. *Cambridge University Press*, 2019.
- Magureanu, S., Combes, R., and Proutiere, A. Lipschitz bandits: Regret lower bound and optimal algorithms. In *Conference on Learning Theory*, pp. 975–999, 2014.
- Thompson, W. R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.

A. Concentration Results

We need two concentration results:

- One to bound the probability that $\hat{\theta}_t$ has a different best arm than θ and determine the rule used to stop exploration and switch to exploitation.
- The other to build optimistic estimates for the gains that we feed to one of the learners.

A.1. Exploration Stopping Rule

Our algorithm optimizes the likelihood ratio $\inf_{\lambda \in \neg i_t} \sum_{k=1}^K N_t^k d(\hat{\theta}_t^k, \lambda^k)$. One can then get a concentration inequality of the form: with probability $1 - \delta$, $\sum_{k=1}^K N_t^k d(\hat{\theta}_t^k, \theta^k) \leq \beta(t, \delta)$ for some function $\beta(t, \delta)$. Then if $\inf_{\lambda \in \neg i_t} \sum_{k=1}^K N_t^k d(\hat{\theta}_t^k, \lambda^k) > \beta(t, \delta)$, we can conclude that $\theta \notin \neg i^*(\hat{\theta}_t)$.

Let W_{-1} be the negative branch of the Lambert W function and for $x > 1$, let $\bar{W}(x) = -W_{-1}(-e^{-x})$. It verifies $x + \log(x) \leq \bar{W}(x) \leq x + \log(x) + \min\{\frac{1}{2}, \frac{1}{\sqrt{x}}\}$.

Theorem 4. *Let $t > 1$, $\theta \in \mathcal{M}$. With probability $1 - \delta$,*

$$\sum_{k=1}^K N_t^k d(\hat{\theta}_t^k, \theta^k) \leq \beta(t, \delta) \quad \text{with} \quad \beta(t, \delta) = 2K\bar{W}\left(\frac{1}{2K} \log \frac{e}{\delta} + \frac{1}{2} \log(8eK \log t)\right).$$

Proof. From the proof of Theorem 2 in (Magureanu et al., 2014), if we restrict ourselves to $t > 1$, we obtain for any $x > K + 1$,

$$\mathbb{P}\left(\sum_{k=1}^K N_t^k d(\hat{\theta}_t^k, \theta^k) > x\right) \leq e \left(\frac{2ex^2 \log(t)}{K}\right)^K e^{-x}.$$

We can equivalently write that, for any $t > 1$, with probability $1 - \delta$,

$$\sum_{k=1}^K N_t^k d(\hat{\theta}_t^k, \theta^k) \leq \beta(t, \delta),$$

where $\beta(t, \delta)$ is the solution in x of $\delta = e \left(\frac{2ex^2 \log(t)}{K}\right)^K e^{-x}$. □

A.2. Confidence Intervals

Theorem 5. *For all $t > 1$ and $\eta > 0$ such that $\log_{1+\eta}(t)$ is not an integer, with probability $1 - \delta$, for all $s \leq t$ and all $k \in [K]$,*

$$N_s^k d(\hat{\theta}_s^k, \theta^k) \leq (1 + \eta) \left(\log \frac{2K \log t}{\delta} - \log \log(1 + \eta) \right).$$

Optimizing over η , we obtain that with probability $1 - \delta$, for all $s \leq t$ and all $k \in [K]$,

$$N_s^k d(\hat{\theta}_s^k, \theta^k) \leq \bar{W}\left(\log \frac{2K \log t}{\delta}\right) \exp\left(1/\bar{W}\left(\log \frac{2K \log t}{\delta}\right)\right).$$

For all $\eta > 0$, with probability $1 - \delta$, for all $s \in \mathbb{N}$ and all $k \in [K]$, if $N_s^k > 0$ then

$$N_s^k d(\hat{\theta}_s^k, \theta^k) \leq (1 + \eta) \left(\log \frac{2\zeta(2)K}{\delta} - 2 \log\left(1 + \frac{\log N_s^k}{\log(1 + \eta)}\right) \right).$$

Proof. We prove the inequality for one arm, for one-sided deviations. The final result is then obtained by an union bound over arms and the two sides. We drop the arm superscript in this proof.

For $i \in \{0, \dots, \lceil \log_{1+\eta}(t) \rceil\}$, let θ_i be such that $d(\theta_i, \theta) = \frac{x}{(1+\eta)^i}$. We denote the quantity $\lceil \log_{1+\eta}(t) \rceil$ by m_t .

Let $i_t \in \{0, \dots, m_t - 1\}$ be such that $(1+\eta)^{i_t} \leq N_t < (1+\eta)^{i_t+1}$. Remark that

$$\begin{aligned} \{\hat{\theta}_t \geq \theta, N_t d(\hat{\theta}_t, \theta) \geq x\} &\subseteq \{\hat{\theta}_t \geq \theta, d(\hat{\theta}_t, \theta) \geq \frac{x}{(1+\eta)^{i_t+1}}\} \\ &= \{\hat{\theta}_t \geq \theta, d(\hat{\theta}_t, \theta) - d(\hat{\theta}_t, \theta_{i_t+1}) \geq \frac{x}{(1+\eta)^{i_t+1}}\} \\ &\subseteq \{\hat{\theta}_t \geq \theta, N_t(d(\hat{\theta}_t, \theta) - d(\hat{\theta}_t, \theta_{i_t+1})) \geq \frac{x}{(1+\eta)}\}. \end{aligned}$$

The equality uses that for $\lambda > \theta$, for one-dimensional exponential families, $(\hat{\theta}_t \geq \theta, d(\hat{\theta}_t, \theta) \geq d(\lambda, \theta)) \Leftrightarrow (\hat{\theta}_t \geq \theta, d(\hat{\theta}_t, \theta) - d(\hat{\theta}_t, \lambda) \geq d(\lambda, \theta))$.

For all distributions ρ on \mathbb{R} , we can verify that the following quantity is a martingale with expectation 1:

$$W_t(\rho) = \mathbb{E}_{y \sim \rho} e^{N_t(d(\hat{\theta}_t, \theta) - d(\hat{\theta}_t, y))}.$$

Let ρ be the uniform distribution on $\{\theta_1, \dots, \theta_{m_t}\}$. Summing up the preceding arguments and using Doob's inequality,

$$\begin{aligned} \mathbb{P}(\hat{\theta}_t \geq \theta, N_t d(\hat{\theta}_t, \theta) \geq x) &\leq \mathbb{P}(\hat{\theta}_t \geq \theta, N_t(d(\hat{\theta}_t, \theta) - d(\hat{\theta}_t, \theta_{i_t+1})) \geq \frac{x}{(1+\eta)}) \\ &\leq \mathbb{P}(\hat{\theta}_t \geq \theta, \log \sum_i \frac{1}{m_t - 1} e^{N_t(d(\hat{\theta}_t, \theta) - d(\hat{\theta}_t, \theta_i))} \geq \frac{x}{(1+\eta)} - \log(m_t - 1)) \\ &= \mathbb{P}(\hat{\theta}_t \geq \theta, W_t(\rho) \geq \frac{1}{m_t - 1} e^{\frac{x}{(1+\eta)}}) \\ &\leq \frac{\log t}{\log(1+\eta)} e^{-x/(1+\eta)} \end{aligned}$$

Equivalently, with probability $1 - \delta$, if $\hat{\theta}_t \geq \theta$ then

$$N_t d(\hat{\theta}_t, \theta) \leq (1+\eta) \left(\log \frac{\log t}{\delta} - \log \log(1+\eta) \right).$$

We now prove the third claim.

Extend the definition of θ_i to all $i \in \mathbb{N}$. With ρ the distribution supported on $\cup_{i \geq 1} \{\theta_i\}$ with $\rho(\theta_i) = 1/(\zeta(2)i^2)$, we get

$$\begin{aligned} \{\hat{\theta}_t \geq \theta, N_t d(\hat{\theta}_t, \theta) \geq x\} &\subseteq \{\hat{\theta}_t \geq \theta, N_t(d(\hat{\theta}_t, \theta) - d(\hat{\theta}_t, \theta_{i_t+1})) \geq \frac{x}{(1+\eta)}\} \\ &\subseteq \{\hat{\theta}_t \geq \theta, \log \sum_i \rho(\theta_i) e^{N_t(d(\hat{\theta}_t, \theta) - d(\hat{\theta}_t, \theta_i))} \geq \frac{x}{(1+\eta)} + \log \rho(\theta_{i_t+1})\} \\ &= \{\hat{\theta}_t \geq \theta, W_t(\rho) \geq \rho(\theta_{i_t+1}) e^{\frac{x}{(1+\eta)}}\}. \end{aligned}$$

With probability $1 - \delta$, the martingale $W_t(\rho)$ is smaller than $1/\delta$. That is, with probability $1 - \delta$,

$$\begin{aligned} N_t d(\hat{\theta}_t, \theta) &\leq (1+\eta) \left(\log \frac{1}{\delta} - \log \rho(\theta_{i_t+1}) \right) \\ &= (1+\eta) \left(\log \frac{1}{\delta} + \log(\zeta(2)(i_t+1)^2) \right) \\ &\leq (1+\eta) \left(\log \frac{1}{\delta} + \log(\zeta(2)(\log_{1+\eta}(N_t) + 1)^2) \right). \end{aligned}$$

□

B. Tracking

We call tracking the following interaction. Starting from vectors $W_0 = N_0 = (0, \dots, 0) \in \mathbb{R}^K$, for each stage $t = 1, 2, \dots$

- Nature reveals a vector w_t in the simplex Δ_K and updates $W_t = W_{t-1} + w_t$.
- A tracking rule selects $k_t \in [K]$ based on the sequence (w_1, \dots, w_t) and forms $N_t = N_{t-1} + e_{k_t}$, where $(e_i)_{i \in [K]}$ are the canonical basis.

Note that w_t is known to the tracking rule when choosing k_t . We are interested in tracking rules such that $\|N_t - W_t\|_\infty$ is as small as possible.

Definition 1. We call C-Tracking any rule which for all stages $t \geq 1$ selects $k_t \in \operatorname{argmin}_{k \in [K]} N_{t-1}^k - W_t^k$.

This defines C-tracking up to the choice of k_t when the argmin is not unique.

Theorem 6. *The C-Tracking procedure described above ensures that for all $t \in \mathbb{N}$, for all $k \in [K]$,*

$$-\sum_{j=2}^K \frac{1}{j} \leq N_t^k - W_t^k \leq 1.$$

Proof. The upper bound was proved in (Garivier & Kaufmann, 2016a). We prove the lower bound.

Let $\mathcal{S}_0 = \{v \in \mathbb{R}^K : \sum_{k=1}^K v^k = 0\}$. The tracking procedure is such that for all stages $t \in \mathbb{N}$, $N_t - W_t \in \mathcal{S}_0$. Our proof strategy is to characterize the subset of \mathcal{S}_0 that can be reached during the tracking procedure, starting from $v_0 = N_0 - W_0 = \mathbf{0}$.

We define a move \rightarrow_w as function from \mathcal{S}_0 to itself parametrized by w that maps v to $v - w + e_k$, where $k = \operatorname{argmin}_{j \in [K]} v - w$. If the value of that function at v is u , we write $v \rightarrow_w u$. A vector $u \in \mathcal{S}_0$ is said to be reachable in one move from $v \in \mathcal{S}_0$ if there exists $w \in \Delta_K$ such that $v \rightarrow_w u$. We denote it by $v \rightarrow u$. It is said to be reachable from v if there is a finite sequence of such moves such that $v \rightarrow \dots \rightarrow u$.

A reverse move $\leftarrow_{k,w}$ is a function from \mathcal{S}_0 to itself parametrized by k and w that maps v to $v + w - e_k$. A reverse move is said to be valid at v if $v^k \leq \min_j v^j + 1$. If the value of that function at v is u , we write $u \leftarrow_{k,w} v$. A vector $u \in \mathcal{S}_0$ is said to be reverse-reachable in one move from $v \in \mathcal{S}_0$ if there exists $k \in [K]$ and $w \in \Delta_K$ such that $u \leftarrow_{k,w} v$ and such that this is a valid reverse move at v . We denote it by $u \leftarrow v$.

We now prove that $(u \rightarrow v) \Leftrightarrow (u \leftarrow v)$. First, if $u \rightarrow v$ then let w be the parameter of a move $u \rightarrow_w v$ and let $k = \operatorname{argmin}_{j \in [K]} u - w$. Then $u \leftarrow_{k,w} v$ is a valid reverse move. Second, if $u \leftarrow_{k,w} v$ is a valid reverse move, then $k = \operatorname{argmin}_{j \in [K]} u - w$ and we have $u \rightarrow_w v$.

We characterize the elements v of \mathcal{S}_0 such that $\mathbf{0} \leftarrow \dots \leftarrow v$.

Let $u, v \in \mathcal{S}_0$ be such that $u \leftarrow v$. Let $M_v = \{k \in [K] : v^k \leq \min_j v^j + 1\}$. Then for any set $S \subseteq [K]$ such that $M_v \subseteq S$, $\sum_{i \in S} u^i \leq \sum_{i \in S} v^i$. Indeed, for the reverse move to be valid, one of the coordinates in M_v was decreased by 1, and they were added coordinates of a $w \in \Delta_K$, that sum at most to 1.

Let $S \subseteq [K]$ and $A_S = \{u \in \mathcal{S}_0 : \forall k \notin S, u^k > \frac{1}{|S|} \sum_{i \in S} u^i + 1\}$. We now prove that if $u \leftarrow v$ and $v \in A_S$, then $u \in A_S$ and as a consequence, that if $u \leftarrow \dots \leftarrow v$ and $v \in A_S$, then $u \in A_S$. Indeed,

- Since $v \in A_S$, we have $M_v \subseteq S$, hence the previous remark proves $\frac{1}{|S|} \sum_{i \in S} u^i \leq \frac{1}{|S|} \sum_{i \in S} v^i$.
- For $k \notin S$, then $k \notin M_v$ and $u^k \geq v^k > \frac{1}{|S|} \sum_{i \in S} v^i + 1 \geq \frac{1}{|S|} \sum_{i \in S} u^i + 1$.

Since $\mathbf{0} \notin \bigcup_{S \in \mathcal{P}([K]) \setminus \{\{K\}\}} A_S$, we can now state that if $\mathbf{0} \leftarrow \dots \leftarrow v$, then $v \notin \bigcup_{S \in \mathcal{P}([K]) \setminus \{\{K\}\}} A_S$.

Let $j \in [2 : K]$ and let $v_{(j)} \in \mathcal{S}_0$ be such that $v_{(j)}^1 \geq \dots \geq v_{(j)}^{j-1} > v_{(j)}^j = \dots = v_{(j)}^K$. Then we will prove that one of the two following statements is true:

1. $v_{(j)}^{j-1} > v_{(j)}^j + 1$ and $v_{(j)}$ is not reachable from $\mathbf{0}$,

2. $v_{(j)}^{j-1} \leq v_{(j)}^j + 1$. Then let $v_{j-1:K}$ be the mean of $v_{(j)}^{j-1}, \dots, v_{(j)}^K$ and let $\mathbf{v}_{(j-1)}$ be the vector with $v_{(j-1)}^1 = v_{(j)}^1, \dots, v_{(j-1)}^{j-2} = v_{(j)}^{j-2}$ and $v_{(j-1)}^{j-1} = \dots = v_{(j-1)}^K = v_{j-1:K}$. Then $\mathbf{v}_{(j-1)} \leftarrow \mathbf{v}_{(j)}$.

Case 1: $v_{(j)}^{j-1} > v_{(j)}^j + 1$. Let $S = [j : K]$. then $S \neq [K]$ and $\mathbf{v} \in A_S$. Hence \mathbf{v} is not reachable from $\mathbf{0}$.

Case 2: $v_{(j)}^{j-1} \leq v_{(j)}^j + 1$. Let \mathbf{w} be defined by $w^1 = \dots = w^{j-2} = 0, w^{j-1} = 1 - (K - j + 1)(v_{j-1:K} - v_{(j)}^j)$ and $w^j = \dots = w^K = v_{j-1:K} - v_{(j)}^j$. Since $v_{(j)}^{j-1} \leq v_{(j)}^j + 1$ and $\mathbf{w} \in \Delta_K$, the reverse move $\leftarrow_{j-1, \mathbf{w}}$ is valid at $\mathbf{v}_{(j)}$. Then the image of $\mathbf{v}_{(j)}$ by that reverse move is $\mathbf{v}_{(j-1)}$. Note that $w^{j-1} \geq w^j = \dots = w^K$.

We have now all the tools to state the characterization of the the elements \mathbf{v} of \mathcal{S}_0 such that $\mathbf{0} \leftarrow \dots \leftarrow \mathbf{v}$. By a simple induction using the last case distinction, we have the following: let $\mathbf{v} \in \mathcal{S}_0$ and let $i_1, \dots, i_K \in [K]$ be such that $v^{i_1} \geq \dots \geq v^{i_K}$. If $\mathbf{0} \leftarrow \dots \leftarrow \mathbf{v}$, then there exists $\mathbf{u}_1, \dots, \mathbf{u}_{K-2}$ and $\mathbf{w}_1, \dots, \mathbf{w}_{K-1}$ such that

1. $\mathbf{0} \leftarrow_{i_1, \mathbf{w}_1} \mathbf{u}_1 \leftarrow_{i_2, \mathbf{w}_2} \dots \leftarrow_{i_{K-2}, \mathbf{w}_{K-2}} \mathbf{u}_{K-2} \leftarrow_{i_{K-1}, \mathbf{w}_{K-1}} \mathbf{v}$,
2. for all $j \in [K - 2]$, $u_j^{i_1} \geq \dots \geq u_j^{i_j} \geq u_j^{i_{j+1}} = \dots = u_j^{i_K} = \frac{1}{K-j} \sum_{k=j+1}^K v^{i_k}$,
3. for all $j \in [K - 1]$, $w_j^{i_1} = \dots = w_j^{i_{j-1}} = 0$ and $w_j^{i_j} \geq w_j^{i_{j+1}} = \dots = w_j^{i_K}$.

In order to prove the theorem, we then only need a bound on $v^{i_K} = -\sum_{j=1}^{K-1} w_j^K$. The characterization of \mathbf{w}_j implies that $w_j^K \leq 1/(K - j + 1)$. Hence $v^{i_K} \geq -\sum_{j=2}^K \frac{1}{j}$. □

Lemma 2. For all tracking rules, for all stages t and vectors N_t, W_t , there exists a sequence $w_{t+1}, \dots, w_{t+K-1} \in \Delta_{K-1}$ such that $\min_{k \in [K]} (N_{t+K-1}^k - W_{t+K-1}^k) - (N_t^k - W_t^k) \leq -\sum_{j=2}^K \frac{1}{j}$.

In particular, with $t = 0$ and $N_0 = W_0 = 0$, there exists a sequence $w_1, \dots, w_{K-1} \in \Delta_{K-1}$ such that $\min_{k \in [K]} (N_{K-1}^k - W_{K-1}^k) = -\sum_{j=2}^K \frac{1}{j}$.

Proof. Let w_{t+1} be the vector with all coordinates equal to $1/K$. Define then $w_{t+2}^k = 1/(K - 1)$ for $k \neq k_{t+1}$ and $w_{t+2}^{k_{t+1}} = 0$.

In general, let n_s be the number of different arms chosen between times $t + 1$ and $t + s$ and let $w_{t+s}^k = 1/(K - n_s)$ if k was not yet chosen in that time interval, and $w_{t+s}^k = 0$ otherwise.

There is one arm $i \in [K]$ such that $i \neq k_{t+1}, \dots, i \neq k_{t+K-1}$. For that arm, $w_{t+1}^i + \dots + w_{t+K-1}^i \geq \sum_{j=2}^K \frac{1}{j}$ (with equality if all k_{t+s} are different) and $N_{t+K-1}^i = N_t^i$. □

C. Regret Bound Proof

C.1. Assumptions

We summarize below the assumptions under which our results are proven. We show in Appendix F that all structures used in the paper verify Assumption 7.

Assumption 3. Each arm has a distribution in the same 1-parameter exponential family, with parameter in an open interval $\Theta \subseteq \mathbb{R}$. The distribution of arm k has mean parameter μ^k .

Assumption 4. The arm distributions are σ^2 -sub-Gaussian, meaning that they verify $\log \mathbb{E} e^{t(X - \mathbb{E}[X])} \leq \frac{1}{2} \sigma^2 t^2$ for all $t \in \mathbb{R}$.

Consequence: the Kullback-Leibler divergence between distributions with means x and y verifies $d(x, y) \geq (x - y)^2 / (2\sigma^2)$.

Definition 2. A structure set is a subset \mathcal{M} of Θ^K .

For simplicity we consider only structures with a unique best arm.

Assumption 5. For all $\mu \in \mathcal{M}$, $i^*(\mu) = \operatorname{argmax}_k \mu^k$ is uniquely defined.

Definition 3. Let $\operatorname{cl}(A)$ denote the closure of a set A . Given a structure set $\mathcal{M} \subseteq \Theta^K$, for $i \in [K]$ we define the alternative sets

$$\begin{aligned} \neg i &= \operatorname{cl}(\{\lambda \in \mathcal{M} : i^*(\lambda) \neq i\}), \\ \neg_x i &= \neg i \cap \{\xi \in \Theta^K : \xi^i = x\}. \end{aligned}$$

Our algorithm minimizes over the sets $\neg i$, while the asymptotic complexity is expressed in terms of the sets $\neg_x i$.

Definition 4. For all $\mu, \lambda \in \Theta^K$, we define

$$\begin{aligned} Q_\mu(\lambda) &= \{\xi \in \Theta^K : \forall k \in [K], d(\mu^k, \xi^k) \geq d(\mu^k, \lambda^k)\}, \\ Q_\mu^>(\lambda) &= \{\xi \in \Theta^K : \forall k \in [K], d(\mu^k, \xi^k) \geq d(\mu^k, \lambda^k), \exists j \in [K], d(\mu^j, \xi^j) > d(\mu^j, \lambda^j)\}, \end{aligned}$$

For all $\mu \in \Theta^K$ and $\Lambda \subset \Theta^K$, we define the boundary of Λ relative to μ by

$$\partial_\mu \Lambda = \{\lambda \in \Lambda : \forall \xi \in \Lambda, \lambda \notin Q_\mu^>(\xi)\}.$$

The next lemma states that the lower bound problem can be restricted to that boundary (and justifies the term ‘‘boundary’’).

Lemma 3. For all $\mu \in \mathcal{M}$, $v \in (\mathbb{R}^+)^K$, there exists $\lambda \in \partial_\mu \neg i^*(\mu)$ such that $\lambda \in \operatorname{argmin}_{\xi \in \neg i^*(\mu)} \sum_{k=1}^K v^k d(\mu^k, \xi^k)$.

The main idea that we will use is that the lower bound problem can be restricted to $\neg i^*(\mu)$ minus any union of $Q_\mu(\lambda) \setminus \{\lambda\}$ for some $\lambda \in \neg i^*(\mu)$. We introduce the boundary because the Sparse structure (see Section F) is such that $\neg i^*(\mu)$ does not verify the property described in Assumption 7 (see below), but $\partial_\mu \neg i^*(\mu)$ does.

Assumption 6. There exists a compact set $\mathcal{C} \subseteq \Theta^K$ such that $\mathcal{M} \subseteq \mathcal{C}$.

As a consequence of that assumption, we have the following lemma.

Lemma 4. There exists $C > 0$ such that for all $\theta, \lambda \in \mathcal{M}$, for all $k \in [K]$, $d(\theta^k, \lambda^k) \leq C/2$. There exists $L > 0$ such that for all $\lambda \in \mathcal{M}$, for all $k \in [K]$, the function $x \mapsto d(x, \lambda^k)$ is L -Lipschitz on \mathcal{M} . There exists $L' > 0$ such that for all $\lambda \in \mathcal{M}$, for all $k \in [K]$, the function $x \mapsto d(\lambda^k, x)$ is L' -Lipschitz on \mathcal{M} .

The next assumption, which describes a sort of Lipschitz property for the structure, is key to linking the value of the perturbed lower bound problem and the unperturbed one together.

Assumption 7. There exists $c_{\mathcal{M}}$ such that for all $\mu \in \mathcal{M}$, denoting $i^*(\mu) = *$, if $\neg_{\mu^*} * \neq \emptyset$ then

$$\forall \lambda \in \partial_\mu \neg_{\mu^*} *, \exists \xi \in \neg_{\mu^*} *, \forall k \in [K], |\xi^k - \lambda^k| \leq c_{\mathcal{M}} |\mu^* - \lambda^*|.$$

C.2. Events and Counts

We first define random variables and events used in the analysis.

Events and random variables based on the algorithm definition:

- E_t : the algorithm explores at stage t .
- $E_t^j = E_t \cap \{j_t = j\}$.
- $E_t^{j\%2} = E_t^j \cap \{n_{t-1}^j \bmod 2 = 1\}$.
- $n_t = \sum_{s \leq t} \mathbb{I}(E_s)$.
- $n_t^j = \sum_{s \leq t} \mathbb{I}(E_s^j)$.

Events and random variables based on concentration:

- $\mathcal{E}_t = \{\sum_{k=1}^K N_t^k d(\hat{\theta}_t^k, \theta^k) \leq f(t)\}$ where $f(t) = \beta(t, 1/(t \log(t)))$.
- $c_t = \sum_{s \leq t} \mathbb{I}(\mathcal{E}_s)$, number of stages without concentration at that level.
- $\mathcal{E}_n^{exp} = \{\forall k \in [K], \forall t \in \mathbb{N}^*, N_{t-1}^k d(\hat{\theta}_{t-1}^k, \theta^k) \leq g_{t-1}(n)\}$ with $g_t(n) = (1 + \eta)[\log(n) + \log(\frac{2K^2 \log^2(t)}{\log(1+\eta)})]$ for a fixed $\eta > 0$.

- $m_t = \sum_{s \leq t} \mathbb{I}(E_s \cap \overline{\mathcal{E}_{n_{s-1}^{j_s}}^{exp}})$.
- $m_t^j = \sum_{s \leq t} \mathbb{I}(E_s^j \cap \overline{\mathcal{E}_{n_{s-1}^j}^{exp}})$.

For $k \in [K]$ and $t \in \mathbb{N}$, we define

$$\begin{aligned} \overline{\theta_{t-1}^k} &= \sup\{\xi \in \mathbb{R} : N_{t-1}^k d(\hat{\theta}_{t-1}^k, \xi) \leq g_{t-1}(n_{t-1}^{j_t})\}, \\ \underline{\theta_{t-1}^k} &= \inf\{\xi \in \mathbb{R} : N_{t-1}^k d(\hat{\theta}_{t-1}^k, \xi) \leq g_{t-1}(n_{t-1}^{j_t})\}, \end{aligned}$$

such that $\mathcal{E}_{n_{t-1}^{j_t}}^{exp} = \{\forall k \in [K], \forall t \in \mathbb{N}^*, \hat{\theta}_{t-1} \in [\underline{\theta_{t-1}^k}, \overline{\theta_{t-1}^k}]\}$. Due to the sub-Gaussian assumption, we have

$$\hat{\theta}_{t-1}^k - \sqrt{2\sigma^2 \frac{g_{t-1}(n_{t-1}^{j_t})}{N_{t-1}^k}} \leq \underline{\theta_{t-1}^k} \leq \hat{\theta}_{t-1}^k \leq \overline{\theta_{t-1}^k} \leq \hat{\theta}_{t-1}^k + \sqrt{2\sigma^2 \frac{g_{t-1}(n_{t-1}^{j_t})}{N_{t-1}^k}}.$$

Expectations of some of these random variables.

$$\begin{aligned} \mathbb{E}[c_t] &\leq \sum_{s=1}^t \mathbb{P}(\overline{\mathcal{E}}_s) \leq 1 + \sum_{s=2}^t \frac{1}{s \log s} \leq 1 + \log \log(t). \\ \mathbb{E}[m_t^j] &\leq \sum_{n=1}^t \mathbb{P}(\overline{\mathcal{E}}_n^{exp}) \leq 1 + \sum_{n=2}^t \frac{1}{Kn \log n} \leq 1 + \frac{1}{K} \log \log(t). \\ \mathbb{E}[m_t] &= \sum_{j=1}^K \mathbb{E}[m_t^j] \leq K + \log \log(t). \end{aligned}$$

C.3. Estimation of the Gaps

Once a candidate best arm j_s is selected at stage s , we form estimates of the gaps for each arm. We use an optimistic estimator for the value of the best arm and define for all $k \in [K]$, $\tilde{\Delta}_s^k = \max\{\varepsilon_s, \overline{\theta_{s-1}^{j_s}} - \overline{\theta_{s-1}^k}\}$. This quantity estimates the gap in a worst case fashion, with an added perturbation of ε_s . Note that $\tilde{\Delta}_s^{j_s} = \varepsilon_s$. Obviously for all $k \in [K]$, $\tilde{\Delta}_s^k \geq \varepsilon_s$. Under $\mathcal{E}_{n_{s-1}^{j_s}}^{exp}$, for $k \neq j_s$,

$$\tilde{\Delta}_s^k \geq \max\left\{\varepsilon_s, \theta^{j_s} - \theta^k - (\overline{\theta_{s-1}^k} - \underline{\theta_{s-1}^k})\right\} \geq \max\left\{\varepsilon_s, \theta^{j_s} - \theta^k - 2\sqrt{2\sigma^2 \frac{g_{s-1}(n_{s-1}^{j_s})}{N_{s-1}^k}}\right\}. \quad (18)$$

$$\tilde{\Delta}_s^k \leq \max\left\{\varepsilon_s, \theta^{j_s} - \theta^k + (\overline{\theta_{s-1}^k} - \underline{\theta_{s-1}^k})\right\} \leq \max\left\{\varepsilon_s, \theta^{j_s} - \theta^k + 2\sqrt{2\sigma^2 \frac{g_{s-1}(n_{s-1}^{j_s})}{N_{s-1}^k}}\right\}. \quad (19)$$

C.4. UCBs

We introduce upper confidence bounds on the ratio of Kullback-Leibler divergence and gap.

$$\text{UCB}_s^k = \sup_{\xi \in [\underline{\theta_{s-1}^k}, \overline{\theta_{s-1}^k}]} \frac{\mathbb{E}_{\lambda \sim q_s} d(\xi, \lambda^k)}{\max\{\varepsilon_s, \mathbb{1}\{k \neq j_s\}(\overline{\theta_{s-1}^{j_s}} - \xi)\}}.$$

Note that the supremum in θ can only be achieved at either end of the range, or at the kink induced by the max in the denominator, as a non-negative convex (since divergence in first argument) over positive concave (since linear/constant) function is quasi-convex (Boyd & Vandenberghe, 2004, Example 3.38).

Upper bound.

$$\text{UCB}_s^k \leq \frac{\sup_{\xi \in [\underline{\theta_{s-1}^k}, \overline{\theta_{s-1}^k}]} \mathbb{E}_{\lambda \sim q_s} d(\xi, \lambda^k)}{\inf_{\xi \in [\underline{\theta_{s-1}^k}, \overline{\theta_{s-1}^k}]} \max\{\varepsilon_s, \mathbb{1}\{k \neq j_s\}[\overline{\theta_{s-1}^{j_s}} - \xi]\}} = \frac{\sup_{\xi \in [\underline{\theta_{s-1}^k}, \overline{\theta_{s-1}^k}]} \mathbb{E}_{\lambda \sim q_s} d(\xi, \lambda^k)}{\tilde{\Delta}_s^k}. \quad (20)$$

Lower bound. We consider only the case $j_s = *$.

We now relate UCB_s^k and $\frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^k, \lambda^k)}{\Delta_s^k}$ where $\Delta_s^k = \max\{\varepsilon_s, \Delta^k\}$. Under $\mathcal{E}_{n_{s-1}^*}^{exp}$, θ^k belongs to the interval over which the maximization is performed in the definition of UCB_s^k for all $k \in [K]$, hence

$$\text{UCB}_s^k \geq \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^k, \lambda^k)}{\max\{\varepsilon_s, \mathbb{1}\{k \neq *\}\}[\bar{\theta}_{s-1}^{j_s} - \theta^k]} \geq \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^k, \lambda^k)}{\max\left\{\varepsilon_s, \mathbb{1}\{k \neq *\}\left[\Delta^k + 2\sqrt{2\sigma^2 \frac{g_{s-1}(n_{s-1}^*)}{N_{s-1}^*}}\right]\right\}}.$$

Hence $\text{UCB}_s^* \geq \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^*, \lambda^*)}{\varepsilon_s}$ and for $k \neq *$, using $1/(1+x) \geq 1-x$,

$$\text{UCB}_s^k \geq \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^k, \lambda^k)}{\Delta_s^k} \frac{\Delta_s^k}{\Delta_s^k + 2\sqrt{2\sigma^2 \frac{g_{s-1}(n_{s-1}^*)}{N_{s-1}^*}}} \geq \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^k, \lambda^k)}{\Delta_s^k} \left(1 - \sqrt{\frac{8\sigma^2 g_{s-1}(n_{s-1}^*)}{(\Delta_s^k)^2 N_{s-1}^*}}\right). \quad (21)$$

C.5. Decomposition of the regret

The regret at the end of stage t is

$$R_t = \sum_{s \leq t} \Delta_{k_s} = \underbrace{\sum_{s \leq t, E_s \cap \bar{\mathcal{E}}_s} \Delta_{k_s}}_{R_t^{(1)}} + \underbrace{\sum_{s \leq t, E_s \cap \mathcal{E}_s} \Delta_{k_s}}_{R_t^{(2)}} + \underbrace{\sum_{s \leq t, E_s \cap \bar{\mathcal{E}}_{n_{s-1}^{exp}}^{j_s}} \Delta_{k_s}}_{R_t^{(3)}} + \sum_{j=1}^K \underbrace{\sum_{s \leq t, E_s^j \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \Delta_{k_s}}_{R_t^{(4j)}}.$$

C.6. Bounding the three first terms of the regret

First, by definition of c_t , $\mathbb{E} R_t^{(1)} \leq \Delta_{\max} \mathbb{E} c_t$.

$R_t^{(2)}$ corresponds to the event that concentration holds and we enter the exploitation phase. That is, we have

$$\min_{\lambda \in \neg k_t} \sum_k N_{t-1}^k d(\hat{\theta}_{t-1}^k, \lambda^k) \stackrel{\text{since } \bar{E}_t}{>} f(t-1) \stackrel{\text{since } \mathcal{E}_t}{\geq} \sum_k N_{t-1}^k d(\hat{\theta}_{t-1}^k, \theta^k).$$

In that case, $k_t = *$ (or we would find a contradiction since then $\theta \in \neg k_t$) and there is no regret. $R_t^{(2)} = 0$.

Finally, by definition of m_t , $\mathbb{E} R_t^{(3)} = \mathbb{E} \sum_{s \leq t, E_s \cap \bar{\mathcal{E}}_{n_{s-1}^{exp}}^{j_s}} \Delta_{k_s} \leq \Delta_{\max} \mathbb{E} m_t$.

We proved at that point that

$$\mathbb{E} R_t \leq \Delta_{\max} (\mathbb{E} c_t + \mathbb{E} m_t) + \mathbb{E} \sum_{j=1}^K \sum_{s \leq t, E_s^j \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \Delta_{k_s} = \Delta_{\max} (\mathbb{E} c_t + \mathbb{E} m_t) + \sum_{j=1}^K \mathbb{E} R_t^{(4j)}. \quad (22)$$

C.7. Bounding $R_t^{(4j)}$

$R_t^{(4j)}$ is the regret incurred when the algorithm explores and chooses the alternative set $\neg j$, and when the concentration event $\mathcal{E}_{n_{t-1}^j}^{exp}$ holds.

C.7.1. BOUNDING $R_t^{(4*)}$

For the special case of $*$, there is no regret when $n_{s-1}^* \bmod 2 = 0$ (event $E_t^* \setminus E_t^{* \% 2}$). Hence $R_t^{(4*)} = \sum_{s \leq t, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \Delta_{k_s}$.

Using the exploration criterion. Let $t + 1$ be a stage in which there is exploration (event E_{t+1}). Remember that by hypothesis, all divergences are bounded by C . Then, using the tracking Theorem 6, $N_t \geq \sum_{s=1}^t \mathbf{w}_s - \log(K)$ and

$$\begin{aligned} f(t) &\geq \inf_{\lambda \in \neg^*} \sum_{k=1}^K N_t^k d(\hat{\theta}_t^k, \lambda^k) \geq \inf_{\lambda \in \neg^*} \sum_{k=1}^K \sum_{s=1}^t w_s^k d(\hat{\theta}_t^k, \lambda^k) - CK \log K \\ &= \inf_{\lambda \in \neg^*} \sum_{s=1}^t \sum_{k=1}^K w_s^k d(\hat{\theta}_t^k, \lambda^k) - CK \log K. \end{aligned}$$

Concentration to go to an online learning formulation.

$$\begin{aligned} \inf_{\lambda \in \neg^*} \sum_{s=1}^t \sum_{k=1}^K w_s^k d(\hat{\theta}_t^k, \lambda^k) &\geq \inf_{\lambda \in \neg^*} \sum_{s \leq t, E_s^*} \sum_{k=1}^K w_s^k d(\hat{\theta}_t^k, \lambda^k) \\ &= \inf_{\lambda \in \neg^*} \sum_{s \leq t, E_s^*} \sum_{k=1}^K w_s^k d(\hat{\theta}_{s-1}^k, \lambda^k) - \sum_{s \leq t, E_s^*} \sum_{k=1}^K w_s^k [d(\hat{\theta}_{s-1}^k, \lambda^k) - d(\hat{\theta}_t^k, \lambda^k)]. \end{aligned}$$

For all $\lambda \in \neg^*$, if $\mathcal{E}_{n_{s-1}^*}^{exp}$ then for all $k \in [K]$, $d(\hat{\theta}_{s-1}^k, \lambda^k) - d(\hat{\theta}_t^k, \lambda^k) \leq 2L \sqrt{2\sigma^2 \frac{g_{s-1}(n_{s-1}^*)}{N_{s-1}^k}}$ by the Lipschitz assumption. If concentration does not hold (which happens for m_t^* steps), it is smaller than C . For $j \in [K]$, let $c_0^j(t) = \sum_{s \leq t, E_s^j \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \sum_{k=1}^K w_s^k \sqrt{\frac{g_{s-1}(n_{s-1}^j)}{N_{s-1}^k}}$. This quantity, of order $\sqrt{K n_t^j \log n_t^j}$, will appear several times in the following computations. We proved

$$\begin{aligned} \inf_{\lambda \in \neg^*} \sum_{s=1}^t \sum_{k=1}^K w_s^k d(\hat{\theta}_t^k, \lambda^k) &\geq \inf_{\lambda \in \neg^*} \sum_{s \leq t, E_s^*} \sum_{k=1}^K w_s^k d(\hat{\theta}_{s-1}^k, \lambda^k) - 2Lc_0^*(t) - Cm_t^* \\ &\geq \inf_{\lambda \in \neg^*} \sum_{s \leq t, E_s^* \% 2} \sum_{k=1}^K w_s^k d(\hat{\theta}_{s-1}^k, \lambda^k) - 2Lc_0^*(t) - Cm_t^*. \end{aligned}$$

The λ -player. By the regret property of the λ -player algorithm,

$$\inf_{\lambda \in \neg^*} \sum_{s \leq t, E_s^* \% 2} \sum_{k=1}^K w_s^k d(\hat{\theta}_{s-1}^k, \lambda^k) \geq \sum_{s \leq t, E_s^* \% 2} \sum_{k=1}^K w_s^k \mathbb{E}_{\lambda \sim q_s^*} d(\hat{\theta}_{s-1}^k, \lambda^k) - R_{n_t^*}^\lambda / 2.$$

Let $C_2(t) = 2Lc_0^*(t) + Cm_t^* + CK \log K + R_{n_t^*}^\lambda / 2$. We now drop the rounds in which the concentration event does not hold from the sum.

$$f(t) \geq \sum_{s \leq t, E_s^* \% 2 \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \mathbb{E}_{\lambda \sim q_s^*} d(\hat{\theta}_{s-1}^k, \lambda^k) - C_2(t).$$

UCBs. We introduce the gap estimates $\tilde{\Delta}_s^k$ discussed in Section C.3.

$$f(t) \geq \sum_{s \leq t, E_s^* \% 2 \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \mathbb{E}_{\lambda \sim q_s} d(\hat{\theta}_{s-1}^k, \lambda^k) - C_2(t) = \sum_{s \leq t, E_s^* \% 2 \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \frac{\mathbb{E}_{\lambda \sim q_s} d(\hat{\theta}_{s-1}^k, \lambda^k)}{\tilde{\Delta}_s^k} - C_2(t).$$

We now introduce upper confidence bounds for the ratios in the expression above. We use UCB_s^k as discussed in Section C.4.

$$f(t) \geq \sum_{s \leq t, E_s^* \% 2 \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \text{UCB}_s^k - \sum_{s \leq t, E_s^* \% 2 \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \left[\text{UCB}_s^k - \frac{\mathbb{E}_{\lambda \sim q_s} d(\hat{\theta}_{s-1}^k, \lambda^k)}{\tilde{\Delta}_s^k} \right] - C_2(t).$$

By the upper bound on UCB_s^k of equation (20), and the Lipschitz property of d , for all $j \in [K]$,

$$\begin{aligned}
 & \sum_{s \leq t, E_s^{j\%2} \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \left[\text{UCB}_s^k - \frac{\mathbb{E}_{\lambda \sim q_s} d(\hat{\theta}_{s-1}^k, \lambda^k)}{\tilde{\Delta}_s^k} \right] \\
 & \leq \sum_{s \leq t, E_s^{j\%2} \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \sum_{k=1}^K w_s^k \left[\sup_{\xi \in [\underline{\theta}_{s-1}^k, \bar{\theta}_{s-1}^k]} \mathbb{E}_{\lambda \sim q_s} d(\xi, \lambda^k) - \mathbb{E}_{\lambda \sim q_s} d(\hat{\theta}_{s-1}^k, \lambda^k) \right] \\
 & \leq L \sum_{s \leq t, E_s^{j\%2} \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \sum_{k=1}^K w_s^k \sup_{\xi \in [\underline{\theta}_{s-1}^k, \bar{\theta}_{s-1}^k]} |\xi - \hat{\theta}_{s-1}^k| \\
 & \leq Lc_0^j(t).
 \end{aligned}$$

Regret of the k -player. Let $C_3(t) = C_2(t) + Lc_0^*(t)$. We introduce $\tilde{w}_s^k = \frac{w_s^k \tilde{\Delta}_s^k}{\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k}$, the proportion of (ϵ -perturbed) regret allocated to arm k .

$$f(t) \geq \sum_{s \leq t, E_s^{*\%2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \text{UCB}_s^k - C_3(t) = \sum_{s \leq t, E_s^{*\%2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \left(\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \right) \sum_{k=1}^K \tilde{w}_s^k \text{UCB}_s^k - C_3(t).$$

We now use a regret-minimizing learner for that \tilde{w} action, where the loss of vector $\mathbf{v} \in \Delta_K$ at time s is $\ell_s(\mathbf{v}) = -(\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k) \sum_{k=1}^K v^k \text{UCB}_s^k$. Let $R_{n_t^*}^w$ be the regret of that learner after $n_t^*/2$ steps, such that

$$f(t) \geq \max_{i \in [K]} \sum_{s \leq t, E_s^{*\%2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \left(\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \right) \text{UCB}_s^i - R_{n_t^*}^w - C_3(t).$$

Using optimism. We use the lower bound on UCB_s^i of equation (21). Under $\mathcal{E}_{n_{s-1}^*}^{exp}$, with $\Delta_s^i = \max\{\varepsilon_s, \Delta^i\}$,

$$\text{UCB}_s^i \geq \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^i, \lambda^i)}{\Delta_s^i} \left(1 - \sqrt{\frac{8\sigma^2 g_{s-1}(n_{s-1}^*)}{(\Delta_s^i)^2 N_{s-1}^*}} \right)$$

Let $c_4(t) = \max_{i \neq *} \sum_{s \leq t, E_s^{*\%2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \left(\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \right) \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^i, \lambda^i)}{\Delta_s^i} \sqrt{\frac{8\sigma^2 g_{s-1}(n_{s-1}^*)}{(\Delta_s^i)^2 N_{s-1}^*}}$ and let $C_4(t) = C_3(t) + R_{n_t^*}^w + c_4(t)$.

Then

$$f(t) \geq \max_{i \in [K]} \sum_{s \leq t, E_s^{*\%2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \left(\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \right) \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^i, \lambda^i)}{\Delta_s^i} - C_4(t). \quad (23)$$

An inequality on the ε -perturbed regret. For $s \in [t]$, let $\alpha_s = \frac{\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k}{\sum_{s \leq t, E_s^{*\%2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \tilde{\Delta}_s^k}$. These are positive and sum to 1.

$$f(t) \geq \left(\sum_{s \leq t, E_s^{*\%2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \right) \max_k \sum_{s \leq t, E_s^{*\%2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \alpha_s \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^k, \lambda^k)}{\Delta_s^k} - C_4(t).$$

We use that $\tilde{\Delta}_s^k \geq \Delta_s^k - \mathbb{1}\{k \neq *\} 2\sqrt{2\sigma^2 \frac{g_{s-1}(n_{s-1}^*)}{N_{s-1}^k}}$. Remark that $2 \sum_{s \leq t, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}, k_s \neq *} \sum_{k=1}^K w_s^k \sqrt{2\sigma^2 \frac{g_{s-1}(n_{s-1}^*)}{N_{s-1}^k}} \leq 2c_0^*(t)$.

$$f(t) \geq \left(\sum_{s \leq t, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \Delta_s^k - 2c_0^*(t) \right) \max_k \sum_{s \leq t, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \alpha_s \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^k, \lambda^k)}{\Delta_s^k} - C_4(t).$$

If all $(\varepsilon_s)_{s \geq 1}$ are equal, let $\Delta_\varepsilon^k = \max\{\varepsilon, \Delta^k\}$. The sum of expectations $\sum_s \alpha_s \mathbb{E}_{q_s}$ is the expectation under another distribution over λ . Hence

$$f(t) \geq \left(\sum_{s \leq t, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \Delta_\varepsilon^k - 2c_0^*(t) \right) \inf_{\mathbf{q}} \max_k \frac{\mathbb{E}_{\lambda \sim \mathbf{q}} d(\theta^k, \lambda^k)}{\Delta_\varepsilon^k} - C_4(t).$$

Let $D_\varepsilon = \inf_{\mathbf{q}} \max_k \frac{\mathbb{E}_{\lambda \sim \mathbf{q}} d(\theta^k, \lambda^k)}{\Delta_\varepsilon^k}$. This is smaller than $D = \inf_{\mathbf{q}} \max_k \frac{\mathbb{E}_{\lambda \sim \mathbf{q}} d(\theta^k, \lambda^k)}{\Delta^k}$. Let $C_5(t) = \frac{1}{D} C_4(t) + 2c_0^*(t)$. The bound on the pertubed regret (perturbed since function of the gaps Δ_ε^k) is

$$\frac{D_\varepsilon}{D} \sum_{s \leq t, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \Delta_\varepsilon^k \leq \frac{f(t)}{D} + C_5(t).$$

In general, if $(\varepsilon_s)_{s \geq 1, j_s = *}$ is a non-increasing sequence, we remove the first terms of the sum in (23),

$$f(t) \geq \max_{i \in [K]} \sum_{s \leq t, n_s^* \geq \sqrt{n_t^*}, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \left(\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \right) \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^i, \lambda^i)}{\Delta_s^i} - C_4(t).$$

We introduce positive coefficients α that sum to 1 as above and follow the same steps, such that

$$\begin{aligned} f(t) &\geq \left(\sum_{s \leq t, n_s^* \geq \sqrt{n_t^*}, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \Delta_s^k - 2c_0^*(t) \right) \max_k \sum_{s \leq t, n_s^* \geq \sqrt{n_t^*}, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \alpha_s \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^k, \lambda^k)}{\Delta_s^k} - C_4(t) \\ &\geq \left(\sum_{s \leq t, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \Delta_s^k - \max\{\varepsilon_0, \Delta_{\max}\} \sqrt{n_t^*} - 2c_0^*(t) \right) \max_k \sum_{s \leq t, n_s^* \geq \sqrt{n_t^*}, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \alpha_s \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^k, \lambda^k)}{\Delta_s^k} - C_4(t). \end{aligned}$$

Now let $\varepsilon(n_t^*)$ be the maximum of all ε_s for s such that $s \leq t, n_s^* \geq \sqrt{n_t^*}, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}$. Then

$$f(t) \geq \left(\sum_{s \leq t, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \Delta_s^k - \max\{\varepsilon_0, \Delta_{\max}\} \sqrt{n_t^*} - 2c_0^*(t) \right) \inf_{\mathbf{q}} \max_k \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^k, \lambda^k)}{\max\{\varepsilon(n_t^*), \Delta^k\}} - C_4(t).$$

We have

$$\frac{D_{\varepsilon(n_t^*)}}{D} \sum_{s \leq t, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \Delta_s^k \leq \frac{f(t)}{D} + C_5(t) + \max\{\varepsilon_0, \Delta_{\max}\} \sqrt{n_t^*}.$$

By Theorem 2, there exists $c > 0$ such that $D_{\varepsilon(n_t^*)} \geq D - c\sqrt{\varepsilon(n_t^*)D}$, hence

$$\left(1 - c\sqrt{\frac{\varepsilon(n_t^*)}{D}} \right) \sum_{s \leq t, E_s^{* \% 2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \Delta_s^k \leq \frac{f(t)}{D} + C_5(t) + \max\{\varepsilon_0, \Delta_{\max}\} \sqrt{n_t^*}.$$

Inequality on the regret. Let $R_{t,\varepsilon}^{(4*)}$ be the perturbed regret. Then with ε_t^* the smallest ε_s for $s \leq t$ such that $E_s^{*\%2} \cap \mathcal{E}_{n_{s-1}^*}$,

$$\varepsilon_t^*(n_t^*/2 - m_t^*) \leq \sum_{s \leq t, E_s^{*\%2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \varepsilon_s \leq \sum_{s \leq t, E_s^{*\%2} \cap \mathcal{E}_{n_{s-1}^*}^{exp}} \sum_{k=1}^K w_s^k \Delta_s^k = R_{t,\varepsilon}^{(4*)}.$$

From Lemma 7 and with H as defined in that Lemma, $C_5(t) \leq H\sqrt{n_t^* g_t(n_t^*)} + \frac{1}{D} [R_{n_t^*/2}^w + R_{n_t^*/2}^\lambda + Cm_t^*]$, unless $n_t^* < K$ or $g_t(n_t^*) > n_t^*/\log^2(n_t^*)$.

If $n_t^* < K$ then $R_t^{(4*)} < K\Delta_{\max}$. If $g_t(n_t^*) > n_t^*/\log^2(n_t^*)$ then there exists a constant G such that $n_t^* \leq G \log \log t$ and the regret is again bounded by that number times Δ_{\max} . If none of these two inequalities are true, let R be such that $R_{n_t^*/2}^w + R_{n_t^*/2}^\lambda \leq R\sqrt{n_t^*}$. We have

$$\begin{aligned} \left(1 - c\sqrt{\frac{\varepsilon(n_t^*)}{D}}\right) R_{t,\varepsilon}^{(4*)} &\leq \frac{f(t)}{D} + H\sqrt{n_t^* g_t(n_t^*)} + \frac{1}{D} [R\sqrt{n_t^*} + Cm_t^*] \\ &\leq \frac{f(t)}{D} + \left(H + \frac{R}{D}\right) \sqrt{\left(2\frac{R_{t,\varepsilon}^{(4*)}}{\varepsilon_t^*} + m_t^*\right) g_t\left(2\frac{R_{t,\varepsilon}^{(4*)}}{\varepsilon_t^*} + m_t^*\right)} + \frac{C}{D} m_t^*. \end{aligned}$$

Let $y = 2\frac{R_{t,\varepsilon}^{(4*)}}{\varepsilon_t^*} + m_t^*$, such that for a constant ε sequence, the equation above implies

$$\left(1 - c\sqrt{\frac{\varepsilon}{D}}\right) y \leq \frac{2f(t)}{\varepsilon D} + \frac{2}{\varepsilon} \left(H + \frac{R}{D}\right) \sqrt{y g_t(y)} + \left(\frac{2C}{\varepsilon D} + 1\right) m_t^*.$$

We need to solve an equation of the form $y \leq A + B\sqrt{y g_t(y)}$. By concavity of $y \mapsto \sqrt{y g_t(y)}$ (which comes from the concavity of $y \mapsto \sqrt{y \log(y)}$), we have the same inequality in expectation: $\mathbb{E}y \leq \mathbb{E}A + B\sqrt{\mathbb{E}y g_t(\mathbb{E}y)}$.

For $A \gg B$ we get approximately $\mathbb{E}y \leq A$, i.e. $\mathbb{E}R_{t,\varepsilon}^{(4*)} \leq \frac{f(t) + C\mathbb{E}m_t^*}{D(1 - c\sqrt{\frac{\varepsilon}{D}})} \approx \frac{f(t)}{D(1 - c\sqrt{\frac{\varepsilon}{D}})}$ since $\mathbb{E}m_t^*$ is of order $\log \log t$ and $f(t) \approx \log(t)$.

For $\varepsilon_s = 1/(n_s^{j_s})^{1/a}$ with a big enough, we have by construction $\varepsilon(n_t^*) \leq 1/(n_t^*)^{1/(2a)}$, $\varepsilon_t^* \geq 1/(n_t^*)^{1/a}$ and we get an equation on the regret that implies asymptotic optimality.

C.7.2. BOUNDING $R_t^{(4j)}$ FOR $j \neq *$

We prove an upper bound on n_t^j . Then $R_t^{(4j)} \leq \Delta_{\max} n_t^j$. We first show that there exists k for which UCB_s^k is bigger than some constant.

Lemma 5. For all $s \leq t$ such that $j_s \neq i^*(\theta)$ and the event $\mathcal{E}_{n_{s-1}^{j_s}}^{exp}$ holds, there exists an arm $k \in [K]$ such that

$$\text{UCB}_s^k \geq \frac{\Delta_{\min}^2}{8\sigma^2} \frac{1}{\max\left\{\varepsilon_s, \Delta_{\max} + \Delta_{\min} + 2\sqrt{2\sigma^2 g_t(n_{s-1}^{j_s})}\right\}}.$$

Proof. Set a time s and let $j = j_s \neq i^*(\theta)$. Since $i^*(\hat{\theta}_{s-1}^j) \neq i^*(\theta)$, there exists k such that $|\hat{\theta}_{s-1}^k - \theta^k| > \Delta_{\min}/2$. Then by the sub-Gaussian property of the arm distributions, for that arm,

$$d(\theta^k, \hat{\theta}_{s-1}^k) \geq \frac{1}{2\sigma^2} (\hat{\theta}_{s-1}^k - \theta^k)^2 \geq \frac{\Delta_{\min}^2}{8\sigma^2}.$$

Under $\mathcal{E}_{n_{s-1}^j}^{exp}$,

$$\max_{\xi \in [\underline{\theta}_{s-1}^k, \overline{\theta}_{s-1}^k]} \mathbb{E}_{\lambda \sim q_s} d(\xi, \lambda^k) \geq \max_{\xi \in [\underline{\theta}_{s-1}^k, \overline{\theta}_{s-1}^k]} d(\xi, \hat{\theta}_{s-1}^k) \geq d(\theta^k, \hat{\theta}_{s-1}^k).$$

We obtain that for the arm k discussed above,

$$\begin{aligned} \text{UCB}_s^k &= \max_{\xi \in [\underline{\theta}_{s-1}^k, \overline{\theta}_{s-1}^k]} \left[\mathbb{E}_{\lambda \sim q_s} \frac{d(\xi, \lambda^k)}{\max\{\varepsilon_s, \overline{\theta}_{s-1}^j - \xi\}} \right] \geq \frac{\max_{\xi \in [\underline{\theta}_{s-1}^k, \overline{\theta}_{s-1}^k]} \mathbb{E}_{\lambda \sim q_s} d(\xi, \lambda^k)}{\max\{\varepsilon_s, \overline{\theta}_{s-1}^j - \underline{\theta}_{s-1}^k\}} \\ &= \frac{\max_{\xi \in [\underline{\theta}_{s-1}^k, \overline{\theta}_{s-1}^k]} \mathbb{E}_{\lambda \sim q_s} d(\xi, \lambda^k)}{\max\{\varepsilon_s, \overline{\theta}_{s-1}^j - \theta^k + (\overline{\theta}_{s-1}^j - \theta^j) + (\theta^k - \underline{\theta}_{s-1}^k)\}}. \end{aligned}$$

Under $\mathcal{E}_{n_{s-1}^j}^{exp}$,

$$\begin{aligned} \text{UCB}_s^k &\geq \frac{\max_{\xi \in [\underline{\theta}_{s-1}^k, \overline{\theta}_{s-1}^k]} \mathbb{E}_{\lambda \sim q_s} d(\xi, \lambda^k)}{\max\{\varepsilon_s, \overline{\theta}_{s-1}^j - \theta^k + (\overline{\theta}_{s-1}^j - \theta_{s-1}^j) + (\overline{\theta}_{s-1}^k - \underline{\theta}_{s-1}^k)\}} \\ &\geq \frac{\Delta_{\min}}{8\sigma^2} \frac{1}{\max\{\varepsilon_s, \overline{\theta}_{s-1}^j - \theta^k + (\overline{\theta}_{s-1}^j - \theta_{s-1}^j) + (\overline{\theta}_{s-1}^k - \underline{\theta}_{s-1}^k)\}}. \end{aligned}$$

Then if $\overline{\theta}_{s-1}^j - \underline{\theta}_{s-1}^j \leq \Delta_{\min}$, under $\mathcal{E}_{n_{s-1}^j}^{exp}$,

$$\begin{aligned} \text{UCB}_s^k &\geq \frac{\Delta_{\min}^2}{8\sigma^2} \frac{1}{\max\left\{\varepsilon_s, \overline{\theta}_{s-1}^j - \theta^k + \Delta_{\min} + 2\sqrt{2\sigma^2 \frac{g_{s-1}(n_{s-1}^j)}{N_{s-1}^k}}\right\}} \\ &\geq \frac{\Delta_{\min}^2}{8\sigma^2} \frac{1}{\max\left\{\varepsilon_s, \Delta_{\max} + \Delta_{\min} + 2\sqrt{2\sigma^2 g_t(n_{s-1}^j)}\right\}}. \end{aligned}$$

If $\overline{\theta}_{s-1}^j - \underline{\theta}_{s-1}^j > \Delta_{\min}$, then

$$\max_{\xi \in [\underline{\theta}_{s-1}^j, \overline{\theta}_{s-1}^j]} d(\xi, \hat{\theta}_{s-1}^j) \geq \frac{1}{2\sigma^2} \max_{\xi \in [\underline{\theta}_{s-1}^j, \overline{\theta}_{s-1}^j]} (\xi - \hat{\theta}_{s-1}^j)^2 \geq \frac{1}{8\sigma^2} (\overline{\theta}_{s-1}^j - \underline{\theta}_{s-1}^j)^2 \geq \frac{\Delta_{\min}^2}{8\sigma^2},$$

and we get that $\text{UCB}_s^j \geq \frac{\Delta_{\min}^2}{8\sigma^2 \varepsilon_s}$. □

Suppose that $(\varepsilon_s)_{s \in \mathbb{N}, j_s=j}$ is non-increasing. We proved that for all s such that $E_s^j \cap \mathcal{E}_{n_{s-1}^j}^{exp}$, there exists $k \in [K]$ with $(\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k) \text{UCB}_s^k \geq (\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k) \frac{\Delta_{\min}^2}{16\sigma^2(\max\{\varepsilon_0, \Delta_{\max}\} + \sqrt{2\sigma^2 g_t(n_t^j)})}$. Let $\Delta_{\max,0} = \max\{\varepsilon_0, \Delta_{\max}\}$. Then there exists also a $k' \in [K]$ for which

$$\sum_{s \leq t, E_s^j \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \left(\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \right) \text{UCB}_s^{k'} \geq \frac{1}{K} \sum_{s \leq t, E_s^j \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \left(\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \right) \frac{\Delta_{\min}^2}{16\sigma^2 \left(\Delta_{\max,0} + \sqrt{2\sigma^2 g_t(n_t^j)} \right)}.$$

Since $\theta \in \neg j$,

$$\sum_{s \leq t, E_s^j \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \sum_{k=1}^K w_s^k d(\hat{\theta}_{s-1}^k, \theta^k) \geq \inf_{\lambda \in \neg j} \sum_{s \leq t, E_s^j \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \sum_{k=1}^K w_s^k d(\hat{\theta}_{s-1}^k, \lambda^k).$$

Left hand side:

$$\sum_{s \leq t, E_s^j \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \sum_{k=1}^K w_s^k d(\hat{\theta}_{s-1}^k, \theta^k) \leq \sum_{s \leq t, E_s^j \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \sum_{k=1}^K \frac{w_s^k}{N_{s-1}^k} g_{s-1}(n_{s-1}^j) \leq 2K \log(n_t^j) g_t(n_t^j).$$

Right hand side: we apply the same proof steps as for the bounding of $R_t^{(4*)}$.

$$\begin{aligned} \inf_{\lambda \in \neg j} \sum_{s \leq t, E_s^j \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \sum_{k=1}^K w_s^k d(\hat{\theta}_{s-1}^k, \lambda^{k_s}) &\geq \inf_{\lambda \in \neg j} \sum_{s \leq t, E_s^{j\%2} \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \sum_{k=1}^K w_s^k d(\hat{\theta}_{s-1}^k, \lambda^{k_s}) \\ &\geq \inf_{\lambda \in \neg j} \sum_{s \leq t, E_s^{j\%2}} \sum_{k=1}^K w_s^k d(\hat{\theta}_{s-1}^k, \lambda^{k_s}) - C m_t^j \\ &\geq \sum_{s \leq t, E_s^{j\%2}} \sum_{k=1}^K w_s^k \mathbb{E}_{\lambda \sim q_s} d(\hat{\theta}_{s-1}^k, \lambda^{k_s}) - C m_t^j - R_{n_t^j}^\lambda. \\ \sum_{s \leq t, E_s^{j\%2}} \sum_{k=1}^K w_s^k \mathbb{E}_{\lambda \sim q_s} d(\hat{\theta}_{s-1}^k, \lambda^{k_s}) &\geq \sum_{s \leq t, E_s^{j\%2} \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \mathbb{E}_{\lambda \sim q_s} \frac{d(\hat{\theta}_{s-1}^k, \lambda^k)}{\tilde{\Delta}_s^k} \\ &\geq \sum_{s \leq t, E_s^{j\%2} \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \text{UCB}_s^k - L c_0^j(t) \\ &\geq \max_{i \in [K]} \sum_{s \leq t, E_s^{j\%2} \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \left(\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \right) \text{UCB}_s^i - R_{n_t^j/2}^w - L c_0^j(t). \end{aligned}$$

As we argued above, that maximum is such that

$$\begin{aligned} \max_{i \in [K]} \sum_{s \leq t, E_s^{j\%2} \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \left(\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \right) \text{UCB}_s^i &\geq \frac{1}{K} \sum_{s \leq t, E_s^{j\%2} \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \left(\sum_{k=1}^K w_s^k \tilde{\Delta}_s^k \right) \frac{\Delta_{\min}^2}{16\sigma^2 \left(\Delta_{\max,0} + \sqrt{2\sigma^2 g_t(n_t^j)} \right)} \\ &\geq (n_t^j/2 - m_t^j) \frac{1}{K} \varepsilon_t^j \frac{\Delta_{\min}^2}{16\sigma^2 \left(\Delta_{\max,0} + \sqrt{2\sigma^2 g_t(n_t^j)} \right)}. \end{aligned}$$

In the last line, $\varepsilon_t^j = \varepsilon_{t'}$ where t' is the last time before t with $j_t = j$.

We obtain that n_t^j verifies

$$2K \log(n_t^j) g_t(n_t^j) \geq (n_t^j/2 - m_t^j) \frac{\varepsilon_t^j}{K} \frac{\Delta_{\min}^2}{16\sigma^2 \left(\Delta_{\max,0} + \sqrt{2\sigma^2 g_t(n_t^j)} \right)} - R_{n_t^j/2}^w - R_{n_t^j/2}^\lambda - C m_t^j - L c_0^j(t).$$

Let R be such that $R_{n_t^j/2}^w + R_{n_t^j/2}^\lambda \leq R \sqrt{n_t^j}$. From Lemma 7, either $n_t^j \leq K$ or $c_0^j(t) \leq K(\log K + \sqrt{8}) \sqrt{2\sigma^2 n_t^j g_t(n_t^j)}$.

Then either $n_t^j \leq G \log \log t$ for a constant G (that depends on K), or $\log(n_t^j) \sqrt{g_t(n_t^j)} \leq \sqrt{n_t^j}$. When $n_t^j >$

$\max\{K, G \log \log t\}$, we have

$$\begin{aligned}
 n_t^j/2 &\leq m_t + \frac{16K\sigma^2}{\Delta_{\min}^2 \varepsilon_t^j} \left[2K\Delta_{\max,0} \log(n_t^j) g_t(n_t^j) + R\Delta_{\max,0} \sqrt{n_t^j} + LK\Delta_{\max,0} (\log(K) + \sqrt{8}) \sqrt{2\sigma^2 n_t^j g_t(n_t^j)} \right. \\
 &\quad + 2K \log(n_t^j) g_t(n_t^j) \sqrt{2\sigma^2 g_t(n_t^j)} + R\sqrt{2\sigma^2 n_t^j g_t(n_t^j)} + LK(\log(K) + \sqrt{8}) 2\sigma^2 \sqrt{n_t^j g_t(n_t^j)} \\
 &\quad \left. + Cm_t \left(\Delta_{\max,0} + \sqrt{2\sigma^2 g_t(n_t^j)} \right) \right] \\
 &\leq m_t + \frac{16K\sigma^2}{\Delta_{\min}^2 \varepsilon_t^j} \sqrt{n_t^j g(n_t^j)} \left[2K\Delta_{\max,0} + R\Delta_{\max,0} + LK\Delta_{\max,0} (\log(K) + \sqrt{8}) \sqrt{2\sigma^2} \right. \\
 &\quad \left. + 2K\sqrt{2\sigma^2} + R\sqrt{2\sigma^2} + LK(\log(K) + \sqrt{8}) 2\sigma^2 \right] \\
 &\quad + \frac{16K\sigma^2}{\Delta_{\min}^2 \varepsilon_t^j} \left[Cm_t \left(\Delta_{\max,0} + \sqrt{2\sigma^2 g_t(n_t^j)} \right) \right].
 \end{aligned}$$

Bounding the last $\sqrt{g_t(n_t^j)}$ by $\sqrt{g_t(t)}$, we obtain an equation of the form $\varepsilon_t^j n_t^j \leq \alpha_1 + \alpha_2 \sqrt{n_t^j g_t(n_t^j)} + m_t(\alpha_3 + \alpha_4 \sqrt{g_t(t)})$. The quantities $\alpha_1, \alpha_2, \alpha_3, \alpha_4$ in that inequality are deterministic. By concavity of the function of n_t^j present on the right hand side, we have

$$\mathbb{E}[\varepsilon_t^j n_t^j] \leq \alpha_1 + \alpha_2 \sqrt{\mathbb{E} n_t^j g_t(\mathbb{E} n_t^j)} + \mathbb{E} m_t (\alpha_3 + \alpha_4 \sqrt{g_t(t)}).$$

If ε_t^j is constant, it only remains to solve that inequality to get a bound of the form $\mathbb{E} n_t^j = \mathcal{O}(\sqrt{\log t})$ where the dominant term is $\sqrt{g_t(t)} \approx \sqrt{\log t}$. Hence $\mathbb{E} R_t^{(4j)} = \mathcal{O}(\sqrt{\log t})$ as well since $\mathbb{E} R_t^{(4j)} \leq \Delta_{\max} n_t^j$.

If we take $\varepsilon_s = 1/(n_s^{j_s})^{1/a}$ for $a > 2$, we also get a bound $\mathcal{O}((\log t)^{a/(2(a-1))})$.

C.8. Regret bound: summary

Recall that we proved (22) in Section C.6:

$$\mathbb{E} R_t \leq \Delta_{\max} (\mathbb{E} c_t + \mathbb{E} m_t) + \sum_{j=1}^K \mathbb{E} R_t^{(4j)}.$$

The two first terms are $\mathcal{O}(\log \log t)$. For the last one, we proved that for $j \neq *$, $\mathbb{E} R_t^{(4j)} = \mathcal{O}(\sqrt{\log t})$ and that $\mathbb{E} R_t^{(4*)}$ is asymptotically equivalent to $f(t)/D_\varepsilon$ (for fixed ε), with $f(t) \approx \log(t)$. Hence that ratio is also the dominant term of the whole regret. By Theorem 2, $D_\varepsilon \geq D - c\sqrt{\varepsilon D}$ for a constant c .

In order to get asymptotic optimality, we choose $\varepsilon_s = 1/(n_s^{j_s})^{1/a}$ for $a > 2$ big enough.

It is of course possible to get a concrete, explicit bound for the regret from the results in this section. As we foster hardly any hope that the resulting lower order terms are either representative or necessary, we instead focused on the asymptotically dominant term (while providing some hints on the time at which it becomes dominant).

D. Technical Lemmas

Lemma 6. *C-Tracking verifies*

$$\begin{aligned}
 \sum_{s=1}^t \sum_{k=1}^K \frac{w_s^k}{\sqrt{N_{s-1}^k}} &\leq K \log K + 2\sqrt{2Kt}. \\
 \sum_{s=1}^t \sum_{k=1}^K \frac{w_s^k}{N_{s-1}^k} &\leq 2K \log(t).
 \end{aligned}$$

The proof is the same as for Lemma 9 of (Degenne et al., 2019), except that a factor K is replaced by $\log K$ due to our improved tracking result (see Theorem 6).

Lemma 7. *We have the inequalities*

$$\begin{aligned} c_0^j(t) &\leq \sqrt{2\sigma^2 g_t(n_t^j)} (K \log K + 2\sqrt{2Kn_t^j}), \\ c_4(t) &\leq 4C \frac{\max\{\varepsilon_1, \Delta_{\max}\}}{\Delta_{\min}^2} \sqrt{2\sigma^2 n_t^* g_t(n_t^*)} + \frac{8\sigma^2 C}{\Delta_{\min}} \log(n_t) g_t(n_t^*), \end{aligned}$$

where these quantities are defined in the sample complexity proof. Furthermore, if $g(n_t^*) \leq n_t^*/\log^2(n_t^*)$ and $n_t^* \geq K$,

$$\begin{aligned} c_0^*(t) &\leq \sqrt{K}(\log K + \sqrt{8}) \sqrt{2\sigma^2 n_t^* g_t(n_t^*)}, \\ c_4(t) &\leq \frac{4C}{\Delta_{\min}} \left(\frac{\max\{\varepsilon_1, \Delta_{\max}\}}{\Delta_{\min}} + \sqrt{2\sigma^2} \right) \sqrt{2\sigma^2 n_t^* g_t(n_t^*)}, \\ C_5(t) &\leq \left(\frac{4C}{\Delta_{\min} D} \left(\frac{\max\{\varepsilon_1, \Delta_{\max}\}}{\Delta_{\min}} + \sqrt{2\sigma^2} \right) + (2 + \frac{3L}{D}) \sqrt{K}(\log K + \sqrt{8}) \right) \sqrt{2\sigma^2 n_t^* g_t(n_t^*)} \\ &\quad + \frac{1}{D} \left[R_{n_t^*/2}^w + R_{n_t^*/2}^\lambda + C m_t^* \right]. \end{aligned}$$

We define H as the term such that this last inequality reads $C_5(t) \leq H \sqrt{n_t^* g_t(n_t^*)} + \frac{1}{D_\varepsilon} \left[R_{n_t^*/2}^w + R_{n_t^*/2}^\lambda + C m_t^* \right]$.

Proof. Let $c_0^j(t) = \sum_{s \leq t, E_s^j \cap \mathcal{E}_{n_{s-1}^j}^{exp}} \sum_{k=1}^k w_s^k \sqrt{\frac{2\sigma^2 g_{s-1}(n_{s-1}^j)}{N_{s-1}^k}}$. Then

$$c_0^j(t) \leq \sqrt{2\sigma^2 g_t(n_t^j)} \sum_{s \leq t, E_s^j} \sum_{k=1}^K \frac{w_s^k}{\sqrt{N_{s-1}^k}} \leq \sqrt{2\sigma^2 g_t(n_t^j)} (K \log K + 2\sqrt{2Kn_t^j}).$$

We now bound $c_4(t)$.

$$c_4(t) = \max_{i \neq * } \sum_{s \leq t, E_s^{*2} \cap \mathcal{E}_{n_{s-1}^{*2}}^{exp}} \sum_{k=1}^k w_s^k \tilde{\Delta}_s^k \frac{\mathbb{E}_{\lambda \sim q_s} d(\theta^i, \lambda^i)}{\Delta_s^i} \sqrt{\frac{8\sigma^2 g_{s-1}(n_{s-1}^*)}{(\Delta_s^i)^2 N_{s-1}^*}}.$$

By equation (19),

$$\tilde{\Delta}_s^k \leq \Delta_s^k + 2\sqrt{2\sigma^2 \frac{g_{s-1}(n_t^*)}{N_{s-1}^*}} \leq \Delta_s^k + 2\sqrt{2\sigma^2 \frac{g_t(n_t^*)}{N_{s-1}^*}} \leq \max\{\varepsilon_s, \Delta_{\max}\} + 2\sqrt{2\sigma^2 \frac{g_t(n_t^*)}{N_{s-1}^*}}.$$

We bound the divergence by C and get

$$\begin{aligned} c_4(t) &\leq \max_{k \neq * } \sum_{s \leq t, E_s^{*2} \cap \mathcal{E}_{n_{s-1}^{*2}}^{exp}} C \left(\frac{\max\{\varepsilon_1, \Delta_{\max}\}}{\Delta_{\min}} + \sqrt{\frac{8\sigma^2 g_t(n_t^*)}{(\Delta_s^k)^2 N_{s-1}^*}} \right) \sqrt{\frac{8\sigma^2 g_t(n_t^*)}{(\Delta_s^k)^2 N_{s-1}^*}} \\ &\leq \sum_{s \leq t, E_s^{*2} \cap \mathcal{E}_{n_{s-1}^{*2}}^{exp}} C \left(\frac{\max\{\varepsilon_1, \Delta_{\max}\}}{\Delta_{\min}} + \sqrt{\frac{8\sigma^2 g_t(n_t^*)}{\Delta_{\min}^2 N_{s-1}^*}} \right) \sqrt{\frac{8\sigma^2 g_t(n_t^*)}{\Delta_{\min}^2 N_{s-1}^*}}. \end{aligned}$$

By design of the algorithm, $N_{s-1}^* \geq n_{s-1}^*/2$.

$$\begin{aligned} c_4(t) &\leq 2C \frac{\max\{\varepsilon_1, \Delta_{\max}\}}{\Delta_{\min}^2} \sum_{s \leq t, E_s^{*2} \cap \mathcal{E}_{n_{s-1}^{*2}}^{exp}} \sqrt{2\sigma^2 \frac{g_t(n_t^*)}{n_{s-1}^*}} + \frac{8\sigma^2 C}{\Delta_{\min}} \sum_{s \leq t, E_s^{*2} \cap \mathcal{E}_{n_{s-1}^{*2}}^{exp}} \frac{g_t(n_t^*)}{n_{s-1}^*} \\ &\leq 4C \frac{\max\{\varepsilon_1, \Delta_{\max}\}}{\Delta_{\min}^2} \sqrt{2\sigma^2 (n_t^* - m_t^*) g_t(n_t^*)} + \frac{8\sigma^2 C}{\Delta_{\min}} \log(n_t^* - m_t^*) g_t(n_t^*). \end{aligned}$$

The definition of $C_5(t)$ is

$$C_5(t) = 2c_0^*(t) + \frac{1}{D} \left[c_4(t) + R_{n_t^*/2}^w + R_{n_t^*/2}^\lambda + 3Lc_0^*(t) + Cm_t^* \right].$$

□

Lemma 8. Let $a, b > 0$. Let $x > 0$ be such that $x \leq a + \sqrt{x \log(bx)}$. Then for all $\eta \in (0, 1)$

$$x \leq \max \left\{ \frac{a}{1-\eta}, \frac{1}{\eta^2} \overline{W}(\log \frac{b}{\eta^2}) \right\}.$$

Proof. The two terms of the max correspond to the two possibilities $\log(bx) < \eta^2 x$ and $\log(bx) \geq \eta^2 x$. □

E. Perturbation

Proof of Theorem 2. Fix $\mu, c > 0$ and a small $\epsilon \leq \Delta_{\min}$. Let $*$ denote the best arm for μ , and let \mathbf{q} be an optimal solution for problem $D_\epsilon^{\mathcal{M}}(\mu)$ from (2c) which is supported on $\partial_{\mu^{-i^*}}(\mu)$. Then in particular, with the sub-Gaussian property (Assumption 4),

$$\frac{\mathbb{E}_{\lambda \sim \mathbf{q}}(\mu^* - \lambda^*)^2}{2\sigma^2\epsilon} \leq \frac{\mathbb{E}_{\lambda \sim \mathbf{q}}[d(\mu^*, \lambda^*)]}{\epsilon} \leq D_\epsilon^{\mathcal{M}}(\mu).$$

Consider any $\lambda \in \partial_{\mu^{-i^*}}(\mu)$. By Assumption 7, we know that unless $\neg_{\mu^*} = \emptyset$ there is another $\tilde{\lambda} \in \neg_{\mu^*}$ such that $\tilde{\lambda}^* = \mu^*$, and moreover for all $k \in [K]$, $|\tilde{\lambda}^k - \lambda^k| \leq c_{\mathcal{M}}|\lambda^* - \mu^*|$. If $\neg_{\mu^*} = \emptyset$ then $D^{\mathcal{M}}(\mu) = 0$ and the inequality of the theorem is verified.

Let $L' > 0$ be such that for all $k \in [K]$, $x \mapsto d(\mu^k, x)$ is L' -Lipschitz (Lemma 4). Then

$$\begin{aligned} d(\mu^k, \tilde{\lambda}^k) &\leq d(\mu^k, \lambda^k) + L'|\lambda^k - \tilde{\lambda}^k| \\ &\leq d(\mu^k, \lambda^k) + L'c_{\mathcal{M}}|\lambda^* - \mu^*|, \end{aligned}$$

Let $\tilde{\mathbf{q}}$ be any distribution obtained from \mathbf{q} by replacing each λ in its support by such a $\tilde{\lambda}$. Then

$$\begin{aligned} \mathbb{E}_{\tilde{\lambda} \sim \tilde{\mathbf{q}}} d(\mu^k, \tilde{\lambda}^k) &\leq \mathbb{E}_{\lambda \sim \mathbf{q}} d(\mu^k, \lambda^k) + L'c_{\mathcal{M}}\mathbb{E}_{\lambda \sim \mathbf{q}}|\lambda^* - \mu^*| \\ &\leq \mathbb{E}_{\lambda \sim \mathbf{q}} d(\mu^k, \lambda^k) + L'c_{\mathcal{M}}\sqrt{\mathbb{E}_{\lambda \sim \mathbf{q}}(\lambda^* - \mu^*)^2} \\ &\leq \mathbb{E}_{\lambda \sim \mathbf{q}} d(\mu^k, \lambda^k) + L'c_{\mathcal{M}}\sqrt{2\sigma^2\epsilon D_\epsilon^{\mathcal{M}}(\mu)}. \end{aligned}$$

This implies in particular that,

$$\begin{aligned} D^{\mathcal{M}}(\mu) &\leq \max_{k \neq *} \frac{\mathbb{E}_{\tilde{\lambda} \sim \tilde{\mathbf{q}}} [d(\mu^k, \tilde{\lambda}^k)]}{\Delta^k} \\ &\leq \max_{k \neq *} \frac{\mathbb{E}_{\lambda \sim \mathbf{q}} d(\mu^k, \lambda^k) + L'c_{\mathcal{M}}\sqrt{2\sigma^2\epsilon D_\epsilon^{\mathcal{M}}(\mu)}}{\Delta^k} \\ &\leq D_\epsilon^{\mathcal{M}}(\mu) + \frac{L'c_{\mathcal{M}}\sqrt{2\sigma^2 D_\epsilon^{\mathcal{M}}(\mu)}}{\Delta_{\min}} \sqrt{\epsilon}. \end{aligned}$$

We may invert this using that $x \leq y + c\sqrt{y}$ with $c \geq 0$ implies that $y \geq x - \frac{c}{2}(\sqrt{c^2 + 4x} - c) \geq x - c\sqrt{x}$, and hence we find

$$D_\epsilon^{\mathcal{M}}(\mu) \geq D^{\mathcal{M}}(\mu) - \frac{L'c_{\mathcal{M}}\sqrt{2\sigma^2 D^{\mathcal{M}}(\mu)}}{\Delta_{\min}} \sqrt{\epsilon}. \quad \square$$

F. Example of Structures: Verifying Assumption 7

In the following examples, we will describe possibly unbounded structure sets $\mathcal{U} \subseteq \Theta^K$. Then our algorithm will apply to $\mathcal{M} = \mathcal{U} \cap \{\Delta_{\min} > 0\} \cap [a, b]^K$ for an interval $[a, b] \subseteq \Theta$ such that \mathcal{M} verifies Assumptions 5 and 6.

Categorised, Lipschitz, Unconstrained or Unimodal. Unconstrained: $\mathcal{U} = \Theta^K$.

Unimodal: $\mathcal{U} = \{\theta \in \Theta^K : \exists k_0 \in [K], \theta^1 \leq \theta^2 \leq \dots \leq \theta^{k_0} \geq \dots \geq \theta^K\}$.

Lipschitz: $\mathcal{U} = \{\theta \in \Theta^K : \forall k \in [K-1], |\theta^k - \theta^{k+1}| \leq L\}$.

Categorised: the arms are known to each belong to one of C categories, and one of the categories c^* verifies that $\min_{k \in c^*} \theta^k$ is larger than $\max_{k \in c} \theta^k$ for all other categories.

We verify that Assumption 7 holds with $c_{\mathcal{M}} = 1$, using the observation that for these structures, translating all arms by the same amount keeps the structure. Set $\mu \in \mathcal{M}$ and let $* = i^*(\mu)$. Let $\lambda \in \neg*$. Let ξ be defined by $\xi^k = \max\{a, \min\{\lambda^k + (\mu^* - \lambda^*), b\}\}$. Then $\xi \in \neg_{\mu^*}$ and for all $k \in [K]$, $|\xi^k - \lambda^k| \leq |\mu^* - \lambda^*|$. QED.

Linear. A set of arm vectors $a_1, \dots, a_K \in \mathbb{R}^d$ is given. let $A \in \mathbb{R}^{K \times d}$ be the matrix with rows $a_1^\top, \dots, a_K^\top$. The structure is $\mathcal{U} = \{\lambda \in \Theta^K : \exists \eta \in \mathbb{R}^d, \lambda = A\eta\}$.

We verify that Assumption 7 holds. Set $\mu \in \mathcal{M}$ and let $* = i^*(\mu)$. Let $\lambda \in \neg*$. Then let $\eta \in \mathbb{R}^d$ be such that $\lambda = A\eta$ and ζ be a solution of the following problem, if there exists one:

$$\begin{aligned} & \inf_{\zeta \in \mathbb{R}^d} \max_k |a_k^\top (\eta - \zeta)| \\ & \text{s.t. } a_*^\top \zeta = \mu^*, \\ & (a_j - a_*)^\top \zeta \geq 0. \end{aligned}$$

Remark that it verifies

$$\left\{ \begin{array}{l} \inf_{\zeta \in \mathbb{R}^d} \max_k |a_k^\top (\eta - \zeta)| \\ \text{s.t. } a_*^\top \zeta = \mu^*, \\ (a_j - a_*)^\top \zeta \geq 0 \end{array} \right\} \leq \left\{ \begin{array}{l} \inf_{\mathbf{u} \in \mathbb{R}^d} \max_k |a_k^\top \mathbf{u}| \\ \text{s.t. } a_*^\top \mathbf{u} = \mu^* - \lambda^*, \\ (a_j - a_*)^\top \mathbf{u} \geq 0 \end{array} \right\} \leq |\mu^* - \lambda^*| \left\{ \begin{array}{l} \inf_{\mathbf{u} \in \mathbb{R}^d} \max_k |a_k^\top \mathbf{u}| \\ \text{s.t. } a_*^\top \mathbf{u} = 1, \\ a_j^\top \mathbf{u} \geq 1 \text{ if } \mu^* > \lambda^* \\ a_j^\top \mathbf{u} \leq 1 \text{ if } \mu^* < \lambda^* \end{array} \right.$$

and a solution \mathbf{u} to the final optimisation problem gives a feasible $\zeta = \eta + (\mu^* - \lambda^*)\mathbf{u}$ for the first problem. If a_* and a_j are not collinear, a solution exists. If all arms are collinear, then $\neg_{\mu^*} = \emptyset$. Otherwise, there exists at least one suitable j .

For a_j not collinear with a_* , let $c_{j,*}$ be the maximal value of the last optimisation problem for the two possibilities $\mu^* > \lambda^*$ and $\mu^* < \lambda^*$. Then $\|\lambda - A\zeta\|_\infty \leq |\lambda^* - \mu^*| c_{j,*}$.

Assumption 7 is verified for $c_{\mathcal{M}} = \max_{i,j \in [K]} c_{i,j}$ where $c_{i,j}$ is taken to be 0 if a_i and a_j are collinear.

Sparse. Among K arms, s arms (s is known) have means greater than a known level $\gamma \in \Theta$. The $K - s$ other arms have mean γ . Let S_μ be the sparse support of μ , i.e. the set of s arms with values greater than γ .

Lemma 9. For all $\mu \in \mathcal{M}$,

$$\begin{aligned} \partial_{\mu} \neg_{i^*}(\mu) \subseteq & (\cup_{j \in S_{\mu}, j \neq *}\{\xi \in \Theta^K : \xi^j = \xi^*, \forall k \notin \{*, j\}, \xi^k = \mu^k\}) \\ & \cup (\cup_{j \notin S_{\mu}} (\cup_{i \in S_{\mu}, i \neq *}\{\xi \in \Theta^K : \xi^i = \gamma, \xi^j = \xi^*; \forall k \notin \{*, i, j\}, \xi^k = \mu^k\} \\ & \cup \{\xi \in \Theta^K : \xi^* = \gamma; \forall k \neq *, \xi^k = \mu^k\})). \end{aligned}$$

Proof. Let $\xi \in \neg_{i^*}$ with best arm j (or one of the best arms being j). If j is in the sparse support of μ , then define λ equal to μ except that $\lambda^* = \xi^*$ and $\lambda^j = \xi^j$. Then $\lambda \in \neg_{i^*}$ since it is s -sparse and $\lambda^j \geq \lambda^*$. Remark then that $\xi \in Q_\mu(\lambda)$. If $\lambda^j < \mu^*$, let λ_2 be equal to λ except that $\lambda_2^j = \lambda^j$, otherwise let λ_2 be equal to λ except that $\lambda_2^j = \lambda_2^* = \mu^*$. Then λ_2 belongs to the described set and $\lambda \in Q_\mu(\lambda_2)$.

If j is not in the sparse support of μ , let i be in that support such that $\xi^i = \gamma$ (it exists since ξ is also s -sparse). Define λ with coordinates equal to those of μ except for $*, i, j$, for which they are equal to the coordinates of ξ . Then $\lambda \in \neg_{i^*}$ since it is s -sparse and $\lambda^j \geq \lambda^*$. Then $\xi \in Q_\mu(\lambda)$.

If $i = *$, then let λ_2 be the same as λ except that $\lambda_2^j = \mu^j$. $\lambda \in Q_\mu(\lambda_2)$ and λ_2 belongs to the prescribed set. Otherwise let λ_2 be equal to λ except that if $\lambda^j < \mu^*$, $\lambda_2^j = \lambda^j$ and if $\lambda^j \geq \mu^*$, $\lambda_2^j = \lambda_2^* = \mu^*$. Then $\lambda_2 \in \neg_{i^*}$ and belong to one of the sets in the decomposition and $\lambda \in Q_\mu(\lambda_2)$. \square

Lemma 10. *Assumption 7 holds for the Sparse structure with $c_{\mathcal{M}} = 1$.*

We note that the assumption holds trivially for $s = 1$, since in that case \neg_{μ^*} is empty for all $\mu \in \mathcal{M}$.

Proof. We show that the assumption holds for each of the sets in the decomposition of Lemma 9.

Case 1: $j \in S_{\mu^*}$. Let $\lambda \in \neg_* \cap \{\xi \in \Theta^K : \xi^j = \xi^*, \forall k \notin \{*, j\}, \xi^k = \mu^k\}$. Then $\lambda + (\mu^* - \lambda^*)(e_j + e_*) \in \neg_{\mu^*}$ and it verifies the conditions of the assumption with $c_{\mathcal{M}} = 1$.

Case 2: $j \notin S_{\mu^*}$. Let $i \in S_{\mu^*}, i \neq *$ and $\lambda \in \neg_* \cap \{\xi \in \Theta^K : \xi^i = \gamma, \xi^j = \xi^*; \forall k \notin \{*, i, j\}, \xi^k = \mu^k\}$. Then $\lambda + (\mu^* - \lambda^*)(e_j + e_*) \in \neg_{\mu^*}$ and it verifies the conditions of the assumption with $c_{\mathcal{M}} = 1$.

Let $\lambda \in \{\xi \in \Theta^K : \xi^* = \gamma; \forall k \neq *, \xi^k = \mu^k\}$. If $\neg_{\mu^*} \neq \emptyset$, then $s > 1$. Take $l, m \neq *$ with $\lambda^m > \gamma$, and remark that $\lambda + (\mu^* - \gamma)e_* + (\mu^* - \lambda^l)e_l - (\lambda^m - \gamma)e_m \in \neg_{\mu^*}$ and it verifies the conditions of the assumption with $c_{\mathcal{M}} = 1$. \square

G. Experiments Notes

The experiment code is available in the repository https://bitbucket.org/wmkoolen/tidnabbil/src/master/regret_games_paper/.

G.1. Parameters used for the Competition

We use OSSB with forced exploration parameter $\epsilon = 0.02/\sqrt{\ln T}$ and concentration threshold $f(t) = (1 + \gamma) \ln(1 + t)$ with explore-exploit threshold parameter $\gamma = 0$. We run CATSE with $\delta = 1/T$.

G.2. Implementation Notes

We describe how to implement the required alt-min oracle (3) for the structures used in our experiments. Recall that k is the index of the best arm in bandit instance μ , and j is the index of another arm, which we are required to make better than k in λ .

- For the unconstrained case, we set $\lambda_k = \lambda_j$ equal to the weighted average $\frac{N_j \mu_j + N_k \mu_k}{N_j + N_k}$, and we set $\lambda_i = \mu_i$ for all $i \notin \{k, j\}$.
- For the s -sparse case, we reason separately about the three possible combinations of k and j being sparse or not. If both are not sparse, we equalise them as in the unconstrained case. For the other arms $i \notin \{k, j\}$, we set $\lambda_i = \gamma$ for the arms minimising $N_i d(\mu_i, \gamma)$, and we set $\lambda_i = \mu_i$ for the remaining arms.
- For the categorised case, we distinguish the case where k and j are in the same category, and the case where they are not. We have only implemented the case of two categories. Our approach is to find a separating level between the two categories by binary search.
- For the Lipschitz, Simplex and Concave case we use the general-purpose convex quadratic optimiser OSQP through the JuMP interface.
- For the unimodal case we use the PAVA algorithm for isotonic regression. We perform left-to-right increasing regressions and right-to-left decreasing regression, storing all the partial optima. This allows us to answer for any $k \neq j$ the cost of increasing up to arm k and then decreasing from $k + 1$ to the end.
- For the linear case we first try the unconstrained projection, which has a linear algebra closed form. If this leaves λ_k and λ_j in the wrong order, we add the equality constraint and perform the projection in one dimension less.