
Unimodal Bandits: Regret Lower Bounds and Optimal Algorithms

Richard Combes

KTH, Royal Institute of technology, Stockholm, Sweden

RCOMBES@KTH.SE

Alexandre Proutiere

KTH, Royal Institute of technology, Stockholm, Sweden

ALEPRO@KTH.SE

Abstract

We consider stochastic multi-armed bandits where the expected reward is a unimodal function over partially ordered arms. This important class of problems has been recently investigated in (Cope, 2009; Yu & Mannor, 2011). The set of arms is either discrete, in which case arms correspond to the vertices of a finite graph whose structure represents similarity in rewards, or continuous, in which case arms belong to a bounded interval. For discrete unimodal bandits, we derive asymptotic lower bounds for the regret achieved under any algorithm, and propose OSUB, an algorithm whose regret matches this lower bound. Our algorithm optimally exploits the unimodal structure of the problem, and surprisingly, its asymptotic regret does not depend on the number of arms. We also provide a regret upper bound for OSUB in non-stationary environments where the expected rewards smoothly evolve over time. The analytical results are supported by numerical experiments showing that OSUB performs significantly better than the state-of-the-art algorithms. For continuous sets of arms, we provide a brief discussion. We show that combining an appropriate discretization of the set of arms with the UCB algorithm yields an order-optimal regret, and in practice, outperforms recently proposed algorithms designed to exploit the unimodal structure.

1. Introduction

Stochastic Multi-Armed Bandits (MAB) (Robbins, 1952; Gittins, 1989) constitute the most fundamental sequential decision problems with an exploration vs. exploitation trade-off. In such problems, the decision maker selects an arm in each round, and observes a realization of the corresponding unknown reward distribution. Each decision is based on past decisions and observed rewards. The objective is to maximize the expected cumulative reward over some time horizon by balancing exploitation (arms with higher observed rewards should be selected often) and exploration (all arms should be explored to learn their average rewards). Equivalently, the performance of a decision rule or algorithm can be measured through its expected *regret*, defined as the gap between the expected reward achieved by the algorithm and that achieved by an oracle algorithm always selecting the best arm. MAB problems have found many fields of application, including sequential clinical trials, communication systems, economics, see e.g. (Cesa-Bianchi & Lugosi, 2006; Bubeck & Cesa-Bianchi, 2012).

In their seminal paper (Lai & Robbins, 1985), Lai and Robbins solve MAB problems where the successive rewards of a given arm are i.i.d., and where the expected rewards of the various arms are not related. They derive an asymptotic (when the time horizon grows large) lower bound of the regret satisfied by any algorithm, and present an algorithm whose regret matches this lower bound. This initial algorithm was quite involved, and many researchers have tried to devise simpler and yet efficient algorithms. The most popular of these algorithms are UCB (Auer et al., 2002) and its extensions, e.g. KL-UCB (Garivier & Cappé, 2011; Cappé et al., 2013) (note that KL-UCB algorithm was initially proposed in (Lai, 1987), see (2.6)). When the expected rewards of the various arms are not related (Lai & Robbins, 1985), the regret of the best algorithm is essentially of the order $O(K \log(T))$ where K denotes the number of arms, and T is the time horizon. When

K is very large or even infinite, MAB problems become more challenging. Fortunately, in such scenarios, the expected rewards often exhibit some structural properties that the decision maker can exploit to design efficient algorithms. Various structures have been investigated in the literature, e.g., Lipschitz (Agrawal, 1995; Kleinberg et al., 2008; Bubeck et al., 2008), linear (Dani et al., 2008), convex (Flaxman et al., 2005).

We consider bandit problems where the expected reward is a unimodal function over partially ordered arms as in (Yu & Mannor, 2011). The set of arms is either discrete, in which case arms correspond to the vertices of a finite graph whose structure represents similarity in rewards, or continuous, in which case arms belong to a bounded interval. This unimodal structure occurs naturally in many practical decision problems, such as sequential pricing (Yu & Mannor, 2011) and bidding in online sponsored search auctions (B., 2005).

Our contributions. We mainly investigate unimodal bandits with finite sets of arms, and are primarily interested in cases where the time horizon T is much larger than the number of arms K .

(a) For these problems, we derive an asymptotic regret lower bound satisfied by any algorithm. This lower bound does not depend on the structure of the graph, nor on its size: it actually corresponds to the regret lower bound in a classical bandit problem (Lai & Robbins, 1985), where the set of arms is just a neighborhood of the best arm in the graph.

(b) We propose OSUB (Optimal Sampling for Unimodal Bandits), a simple algorithm whose regret matches our lower bound, i.e., it optimally exploits the unimodal structure. The asymptotic regret of OSUB does not depend on the number of arms. This contrasts with LSE (Line Search Elimination), the algorithm proposed in (Yu & Mannor, 2011) whose regret scales as $O(\gamma D \log(T))$ where γ is the maximum degree of vertices in the graph and D is its diameter. We present a finite-time analysis of OSUB, and derive a regret upper bound that scales as $O(\gamma \log(T) + K)$. Hence OSUB offers better performance guarantees than LSE as soon as the time horizon satisfies $T \geq \exp(K/\gamma D)$. Although this is not explicitly mentioned in (Yu & Mannor, 2011), we believe that LSE was meant to address bandits where the number of arms is not negligible compared to the time horizon.

(c) We further investigate OSUB performance in non-stationary environments where the expected rewards smoothly evolve over time but keep their unimodal structure.

(d) We conduct numerical experiments and show that OSUB significantly outperforms LSE and other classi-

cal bandit algorithms when the number of arms is much smaller than the time horizon.

(e) Finally, we briefly discuss systems with a continuous set of arms. We show that using a simple discretization of the set of arms, UCB-like algorithms are order-optimal, and actually outperform more advanced algorithms such as those proposed in (Yu & Mannor, 2011). This result suggests that in discrete unimodal bandits with a very large number of arms, it is wise to first prune the set of arms, so as to reduce its size to a number of the order of $\sqrt{T}/\log(T)$.

2. Related work

Unimodal bandits have received relatively little attention in the literature. They are specific instances of bandits in metric spaces (Kleinberg, 2004; Kleinberg et al., 2008; Bubeck et al., 2008). In this paper, we add unimodality and show how this structure can be optimally exploited. Unimodal bandits have been specifically addressed in (Cope, 2009; Yu & Mannor, 2011). In (Cope, 2009), bandits with a continuous set of arms are studied, and the author shows that the Kiefer-Wolfowitz stochastic approximation algorithm achieves a regret of the order of $O(\sqrt{T})$ under some strong regularity assumptions on the reward function. In (Yu & Mannor, 2011), for the same problem, the authors present LSE, an algorithm whose regret scales as $O(\sqrt{T} \log(T))$ without the need for a strong regularity assumption. The LSE algorithm is based on Kiefer’s golden section search algorithm. It iteratively eliminates subsets of arms based on PAC-bounds derived after appropriate sampling. By design, under LSE, the sequence of parameters used for the PAC bounds is pre-defined, and in particular does not depend of the observed rewards. As a consequence, LSE may explore too much sub-optimal parts of the set of arms. For bandits with a continuum set of arms, we actually show that combining an appropriate discretization of the decision space (i.e., reducing the number of arms to $\sqrt{T}/\log(T)$ arms) and the UCB algorithm can outperform LSE in practice (this is due to the adaptive nature of UCB). Note that the parameters used in LSE to get a regret of the order $O(\sqrt{T} \log(T))$ depend on the time horizon T .

In (Yu & Mannor, 2011), the authors also present an extension of the LSE algorithm to problems with discrete sets of arms, and provide regret upper bounds of this algorithm. These bounds depends on the structure of the graph defining unimodal structure, and on the number of arms as mentioned previously. LSE performs better than classical bandit algorithms only when the number of arms is very large, and actually becomes comparable to the time horizon. Here we are interested in bandits with relatively small number of arms.

Non-stationary bandits have been studied in

(Hartland et al., 2007; Garivier & Moulines, 2008; Slivkins & Upfal, 2008; Yu & Mannor, 2011). Except for (Slivkins & Upfal, 2008), these papers deal with environments where the expected rewards and the best arm change abruptly. This ensures that arms are always well separated, and in turn, simplifies the analysis. In (Slivkins & Upfal, 2008), the expected rewards evolve according to independent brownian motions. We consider a different, but more general class of dynamic environments: here the rewards smoothly evolve over time. The challenge for such environments stems from the fact that, at some time instants, arms can have expected rewards arbitrarily close to each other.

Finally, we should mention that bandit problems with structural properties such as those we address here can often be seen as specific instances of problems in the control of Markov chains, see (Graves & Lai, 1997). We leverage this observation to derive regret lower bounds. However, algorithms developed for the control of generic Markov chains are often too complex to implement in practice. Our algorithm, OSUB, is optimal and straightforward to implement.

3. Model and Objectives

We consider a stochastic multi-armed bandit problem with $K \geq 2$ arms. We discuss problems where the set of arms is continuous in Section 6. Time proceeds in rounds indexed by $n = 1, 2, \dots$. Let $X_k(n)$ be the reward obtained at time n if arm k is selected. For any k , the sequence of rewards $(X_k(n))_{n \geq 1}$ is i.i.d. with distribution and expectation denoted by ν_k and μ_k respectively. Rewards are independent across arms. Let $\mu = (\mu_1, \dots, \mu_K)$ represent the expected rewards of the various arms. At each round, a decision rule or algorithm selects an arm depending on the arms chosen in earlier rounds and their observed rewards. We denote by $k^\pi(n)$ the arm selected under π in round n . The set Π of all possible decision rules consists of policies π satisfying: for any $n \geq 1$, if \mathcal{F}_n^π is the σ -algebra generated by $(k^\pi(t), X_{k^\pi(t)}(t))_{1 \leq t \leq n}$, then $k^\pi(n+1)$ is \mathcal{F}_n^π -measurable.

3.1. Unimodal Structure

The expected rewards exhibit a *unimodal* structure, similar to that considered in (Yu & Mannor, 2011). More precisely, there exists an undirected graph $G = (V, E)$ whose vertices correspond to arms, i.e., $V = \{1, \dots, K\}$, and whose edges characterize a partial order (initially unknown to the decision maker) among expected rewards. We assume that there exists a unique arm k^* with maximum expected reward μ^* , and that from any sub-optimal arm $k \neq k^*$, there exists a path $p = (k_1 = k, \dots, k_m = k^*)$ of length m (depending on k) such that for all $i = 1, \dots, m-1$, $(k_i, k_{i+1}) \in E$ and $\mu_{k_i} < \mu_{k_{i+1}}$. We denote by \mathcal{U}_G the set of vectors μ satisfying this unimodal structure.

This notion of unimodality is quite general, and includes, as a special case, classical unimodality (where G is just a line). Note that we assume that the decision maker knows the graph G , but ignores the best arm, and hence the partial order induced by the edges of G .

3.2. Stationary and non-stationary environments

The model presented above concerns stationary environments, where the expected rewards for the various arms do not evolve over time. In this paper, we also consider non-stationary environments where these expected rewards could evolve over time according to some deterministic dynamics. In such scenarios, we denote by $\mu_k(n)$ the expected reward of arm k at time n , i.e., $\mathbb{E}[X_k(n)] = \mu_k(n)$, and $(X_k(n))_{n \geq 1}$ constitutes a sequence of independent random variables with evolving mean. In non-stationary environments, the sequences of rewards are still assumed to be independent across arms. Moreover, at any time n , $\mu(n) = (\mu_1(n), \dots, \mu_K(n))$ is unimodal with respect to some fixed graph G , i.e., $\mu(n) \in \mathcal{U}_G$ (note however that the partial order satisfied by the expected rewards may evolve over time).

3.3. Regrets

The performance of an algorithm $\pi \in \Pi$ is characterized by its *regret* up to time T (where T is typically large). The way regret is defined differs depending on the type of environment.

Stationary Environments. In such environments, the regret $R^\pi(T)$ of algorithm $\pi \in \Pi$ is simply defined through the number of times $t_k^\pi(T) = \sum_{1 \leq n \leq T} \mathbf{1}\{k^\pi(n) = k\}$ that arm k has been selected up to time T : $R^\pi(T) = \sum_{k=1}^K (\mu^* - \mu_k) \mathbb{E}[t_k^\pi(T)]$. Our objectives are (1) to identify an asymptotic (when $T \rightarrow \infty$) regret lower bound satisfied by *any* algorithm in Π , and (2) to devise an algorithm that achieves this lower bound.

Non-stationary Environments. In such environments, the regret of an algorithm $\pi \in \Pi$ quantifies how well π tracks the best arm over time. Let $k^*(n)$ denote the optimal arm with expected reward $\mu^*(n)$ at time n . The regret of π up to time T is hence defined as: $R^\pi(T) = \sum_{n=1}^T (\mu^*(n) - \mathbb{E}[\mu_{k^\pi(n)}(n)])$.

4. Stationary environments

In this section, we consider unimodal bandit problems in stationary environments. We derive an asymptotic lower bound of regret when the reward distributions belong to a parametrized family of distributions, and propose OSUB, an algorithm whose regret matches this lower bound.

4.1. Lower bound on regret

To simplify the presentation, we assume here that the reward distributions belong to a parametrized family of distributions. More precisely, we define a set of distributions $\mathcal{V} = \{\nu(\theta)\}_{\theta \in [0,1]}$ parametrized by $\theta \in [0, 1]$. The expectation of $\nu(\theta)$ is denoted by $\mu(\theta)$ for any $\theta \in [0, 1]$. $\nu(\theta)$ is absolutely continuous with respect to some positive measure m on \mathbb{R} , and we denote by $p(x, \theta)$ its density. The Kullback-Leibler (KL) divergence number between $\nu(\theta)$ and $\nu(\theta')$ is: $KL(\theta, \theta') = \int_{\mathbb{R}} \log(p(x, \theta)/p(x, \theta'))p(x, \theta)m(dx)$. We denote by θ^* a parameter (it might not be unique) such that $\mu(\theta^*) = \mu^*$, and we define the minimal divergence number between $\nu(\theta)$ and $\nu(\theta^*)$ as: $I_{\min}(\theta, \theta^*) = \inf_{\theta \in [0,1]: \mu(\theta) \geq \mu^*} KL(\theta, \theta')$.

Finally, we say that arm k has parameter θ_k if $\nu_k = \nu(\theta_k)$, and we denote by Θ_G the set of all parameters $\theta = (\theta_1, \dots, \theta_K) \in [0, 1]^K$ such that the corresponding expected rewards are unimodal with respect to graph G : $\mu = (\mu_1, \dots, \mu_K) \in \mathcal{U}_G$. Of particular interest is the family of Bernoulli distributions: the support of m is $\{0, 1\}$, $\mu(\theta) = \theta$, and $I_{\min}(\theta, \theta^*) = I(\theta, \theta^*)$ where $I(\theta, \theta^*) = \theta \log(\frac{\theta}{\theta^*}) + (1 - \theta) \log(\frac{1-\theta}{1-\theta^*})$ is KL divergence number between Bernoulli distributions of respective means θ and θ^* .

We are now ready to derive an asymptotic regret lower in parametrized unimodal bandit problems as defined above. Without loss of generality, we restrict our attention to so-called uniformly good algorithms, as defined in (Lai & Robbins, 1985) (uniformly good algorithms exist as shown later on). We say that $\pi \in \Pi$ is uniformly good if for all $\theta \in \Theta_G$, we have that $R^\pi(T) = o(T^a)$ for all $a > 0$.

Theorem 4.1 *Let $\pi \in \Pi$ be a uniformly good algorithm, and assume that $\nu_k = \nu(\theta_k) \in \mathcal{V}$ for all k . Then for any $\theta \in \Theta_G$,*

$$\liminf_{T \rightarrow +\infty} \frac{R^\pi(T)}{\log(T)} \geq c(\theta) = \sum_{(k, k^*) \in E} \frac{\mu^* - \mu_k}{I_{\min}(\theta_k, \theta^*)}. \quad (1)$$

The above theorem is a consequence of results in optimal control of Markov chains (Graves & Lai, 1997). All proofs are presented in (Combes & Proutiere, 2014). As in classical discrete bandit problems, the regret scales at least logarithmically with time (the regret lower bound derived in (Lai & Robbins, 1985) is obtained from Theorem 4.1 assuming that G is the complete graph). We also observe that the unimodal structure, if optimally exploited, can bring significant performance improvements: the regret lower bound does not depend on the size K of the decision space. Indeed $c(\theta)$ includes only terms corresponding to arms that are neighbors in G of the optimal arm (as if one could learn without regret that all other arms are sub-optimal).

In the case of Bernoulli rewards, the lower regret bound becomes $\log(T) \sum_{(k, k^*) \in E} \frac{\mu^* - \mu_k}{I(\theta_k, \theta^*)}$. Note that LSE and GLSE, the algorithms proposed in (Yu & Mannor, 2011), have performance guarantees that do not match our lower bound: when G is a line, LSE achieves a regret bounded by $41/\Delta^2 \log(T)$, whereas in the general case, GLSE incurs a regret of the order of $O(\gamma D \log(T))$ where γ is the maximal degree of vertices in G , and D is its diameter. The performance of LSE critically depends on the graph structure, and the number of arms. Hence there is an important gap between the performance of existing algorithms and the lower bound derived in Theorem 4.1. In the next section, we close this gap and propose an asymptotically optimal algorithm.

4.2. The OSUB Algorithm

We now describe OSUB, a simple algorithm whose regret matches the lower bound derived in Theorem of 4.1 for Bernoulli rewards, i.e., OSUB is asymptotically optimal. The algorithm is based on KL-UCB proposed in (Lai, 1987; Cappé et al., 2013), and uses KL-divergence upper confidence bounds to define an *index* for each arm. OSUB can be readily extended to systems where reward distributions are within one-parameter exponential families by simply modifying the definition of arm indices as done in (Cappé et al., 2013). In OSUB, each arm is attached an index that resembles the KL-UCB index, but the arm selected at a given time is the arm with maximal index within the neighborhood in G of the arm that yielded the highest empirical reward. Note that since the sequential choices of arms are restricted to some neighborhoods in the graph, OSUB is not an index policy. To formally describe OSUB, we need the following notation. For $p \in [0, 1]$, $s \in \mathbb{N}$, and $n \in \mathbb{N}$, we define:

$$F(p, s, n) = \sup\{q \geq p : sI(p, q) \leq \log(n) + c \log(\log(n))\}, \quad (2)$$

with the convention that $F(p, 0, n) = 1$, and $F(1, s, n) = 1$, and where $c > 0$ is a constant. Let $k(n)$ be the arm selected under OSUB at time n , and let $t_k(n)$ denote the number of times arm k has been selected up to time n . The empirical reward of arm k at time n is $\hat{\mu}_k(n) = \frac{1}{t_k(n)} \sum_{t=1}^n \mathbf{1}\{k(t) = k\} X_k(t)$, if $t_k(n) > 0$ and $\hat{\mu}_k(n) = 0$ otherwise. We denote by $L(n) = \arg \max_{1 \leq k \leq K} \hat{\mu}_k(n)$ the index of the arm with the highest empirical reward (ties are broken arbitrarily). Arm $L(n)$ is referred to as the *leader* at time n . Further define $l_k(n) = \sum_{t=1}^n \mathbf{1}\{L(t) = k\}$ the number of times arm k has been the leader up to time n . Now the index of arm k at time n is defined as:

$$b_k(n) = F(\hat{\mu}_k(n), t_k(n), l_k(L(n))).$$

Finally for any k , let $N(k) = \{k' : (k', k) \in E\} \cup \{k\}$ be the neighborhood of k in G . The pseudo-code of OSUB is

presented below.

Algorithm OSUB

Input: graph $G = (V, E)$

For $n \geq 1$, select the arm $k(n)$ where:

$$k(n) = \begin{cases} L(n) & \text{if } \frac{l_{L(n)}(n)-1}{\gamma+1} \in \mathbb{N}, \\ \arg \max_{k \in N(L(n))} b_k(n) & \text{otherwise,} \end{cases}$$

where γ is the maximal degree of nodes in G and ties are broken arbitrarily.

Note that OSUB forces us to select the current leader often: $L(n)$ is chosen when $l_{L(n)}(n) - 1$ is a multiple of $\gamma + 1$. This ensures that the number of times an arm has been selected is at least proportional to the number of times this arm has been the leader. This property significantly simplifies the regret analysis, but it could be removed.

4.3. Finite-time analysis of OSUB

Next we provide a finite time analysis of the regret achieved under OSUB. Let Δ denote the minimal separation between an arm and its best adjacent arm: $\Delta = \min_{1 \leq k \leq K} \max_{k': (k, k') \in E} \mu_{k'} - \mu_k$. Note that Δ is not known a priori.

Theorem 4.2 *Assume that the rewards lie in $[0, 1]$ (i.e., the support of ν_k is included in $[0, 1]$, for all k), and that $(\mu_1, \dots, \mu_K) \in \mathcal{U}_G$. The number of times suboptimal arm k is selected under OSUB satisfies: for all $\epsilon > 0$ and all $T \geq 3$,*

$$\mathbb{E}[t_k(T)] \leq \begin{cases} (1 + \epsilon) \frac{\log(T) + c \log(\log(T))}{I(\mu_k, \mu^*)} & \text{if } (k, k^*) \in E, \\ + C_1 \log \log(T) + \frac{C_2}{T^{\beta(\epsilon)}} & \\ \frac{C_3}{\Delta^2} & \text{otherwise,} \end{cases}$$

where $\beta(\epsilon) > 0$, and $0 < C_1 < 7$, $C_2 > 0$, $C_3 > 0$ are constants.

To prove this upper bound, we analyze the regret accumulated (i) when the best arm k^* is the leader, and (ii) when the leader is arm $k \neq k^*$. (i) When k^* is the leader, the algorithm behaves like KL-UCB restricted to the arms around k^* , and the regret at these rounds can be analyzed as in (Cappé et al., 2013). (ii) Bounding the number of rounds where $k \neq k^*$ is not the leader is more involved. To do this, we decompose this set of rounds into further subsets (such as the time instants where k is the leader and its mean is not well estimated), and control their expected cardinalities using concentration inequalities. Along the way, we establish Lemma 4.3, a new concentration inequality of independent interest.

Lemma 4.3 *Let $\{Z_t\}_{t \in \mathbb{Z}}$ be a sequence of independent random variables with values in $[0, B]$. Define \mathcal{F}_n the σ -algebra generated by $\{Z_t\}_{t \leq n}$ and the filtration $\mathcal{F} = (\mathcal{F}_n)_{n \in \mathbb{Z}}$. Consider $s \in \mathbb{N}$, $n_0 \in \mathbb{Z}$ and $T \geq n_0$. We define $S_n = \sum_{t=n_0}^n B_t(Z_t - \mathbb{E}[Z_t])$, where $B_t \in \{0, 1\}$ is a \mathcal{F}_{t-1} -measurable random variable. Further define $t_n = \sum_{t=n_0}^n B_t$. Define $\phi \in \{n_0, \dots, T+1\}$ a \mathcal{F} -stopping time such that either $t_\phi \geq s$ or $\phi = T+1$.*

Then we have that: $\mathbb{P}[S_\phi \geq t_\phi \delta, \phi \leq T] \leq \exp(-2s\delta^2 B^{-2})$. As a consequence: $\mathbb{P}[|S_\phi| \geq t_\phi \delta, \phi \leq T] \leq 2 \exp(-2s\delta^2 B^{-2})$.

Lemma 4.3 concerns the sum of products of i.i.d. random variables and of a previsible sequence, evaluated at a stopping time (for the natural filtration). We believe that concentration results for such sums can be instrumental in bandit problems, where typically, we need information about the empirical rewards at some specific random time epochs (that often are stopping times). Refer to (Combes & Proutiere, 2014) for a proof. A direct consequence of Theorem 4.2 is the asymptotic optimality of OSUB in the case of Bernoulli rewards:

Corollary 4.4 *Assume that rewards distributions are Bernoulli (i.e for any k , $\nu_k \sim \text{Bernoulli}(\theta_k)$), and that $\theta \in \Theta_G$. Then the regret achieved under $\pi = \text{OSUB}$ satisfies: $\limsup_{T \rightarrow +\infty} R^\pi(T) / \log(T) \leq c(\theta)$.*

5. Non-stationary environments

We now consider time-varying environments. We assume that the expected reward of each arm varies smoothly over time, i.e., it is Lipschitz continuous: for all $n, n' \geq 1$ and $1 \leq k \leq K$: $|\mu_k(n) - \mu_k(n')| \leq \sigma |n - n'|$.

We further assume that the unimodal structure is preserved (with respect to the same graph G): for all $n \geq 1$, $\mu(n) \in \mathcal{U}_G$. Considering smoothly varying rewards is more challenging than scenarios where the environment is abruptly changing. The difficulty stems from the fact that the rewards of two or more arms may become arbitrarily close to each other (this happens each time the optimal arm changes), and in such situations, regret is difficult to control. To get a chance to design an algorithm that efficiently tracks the best arm, we need to make some assumption to limit the proportion of time when the separation of arms becomes too small. Define for $T \in \mathbb{N}$, and $\Delta > 0$:

$$H(\Delta, T) = \sum_{n=1}^T \sum_{(k, k') \in E} \mathbf{1}\{|\mu_k(n) - \mu_{k'}(n)| < \Delta\}.$$

Assumption 1 *There exists a function Φ and Δ_0 such that for all $\Delta < \Delta_0$: $\limsup_{T \rightarrow +\infty} H(\Delta, T) / T \leq \Phi(K)\Delta$.*

5.1. OSUB with a Sliding Window

To cope with the changing environment, we modify the OSUB algorithm, so that decisions are based on past choices and observations over a time-window of fixed duration equal to $\tau + 1$ rounds. The idea of adding a sliding window to algorithms initially designed for stationary environments is not novel (Garivier & Moulines, 2008); but here, the unimodal structure and the smooth evolution of rewards make the regret analysis more challenging.

Define: $t_k^\tau(n) = \sum_{t=n-\tau}^n \mathbf{1}\{k(t) = k\}$; $\hat{\mu}_k^\tau(n) = (1/t_k^\tau(n)) \sum_{t=n-\tau}^n \mathbf{1}\{k(t) = k\} X_k(t)$ if $t_k^\tau(n) > 0$ and $\hat{\mu}_k^\tau(n) = 0$ otherwise; $L^\tau(n) = \arg \max_{1 \leq k \leq K} \hat{\mu}_k^\tau(n)$; $l_k^\tau(n) = \sum_{t=n-\tau}^n \mathbf{1}\{L^\tau(t) = k\}$. The index of arm k at time n then becomes: $b_k^\tau(n) = F(\hat{\mu}_k^\tau(n), t_k^\tau(n), l_k^\tau(L^\tau(n)))$. The pseudo-code of SW-OSUB is presented below.

Algorithm SW-OSUB

Input: graph $G = (V, E)$, window size $\tau + 1$

For $n \geq 1$, select the arm $k(n)$ where:

$$k(n) = \begin{cases} L^\tau(n) & \text{if } \frac{l_{L^\tau(n)}^\tau(n)-1}{\gamma+1} \in \mathbb{N}, \\ \arg \max_{k \in N(L^\tau(n))} b_k^\tau(n) & \text{otherwise.} \end{cases}$$

5.2. Regret Analysis

In non-stationary environments, achieving sublinear regrets is often not possible. In (Garivier & Moulines, 2008), the environment is subject to abrupt changes or breakpoints. It is shown that if the density of breakpoints is strictly positive, which typically holds in practice, then the regret of any algorithm has to scale linearly with time. We are interested in similar scenarios, and consider smoothly varying environments where the number of times the optimal arm changes has a positive density. The next theorem provides an upper bound of the regret per unit of time achieved under SW-OSUB. This bound holds for any non-stationary environment with σ -Lipschitz rewards.

Theorem 5.1 *Let $\Delta: 2\tau\sigma < \Delta < \Delta_0$. Assume that for any $n \geq 1$, $\mu(n) \in \mathcal{U}_G$ and $\mu^*(n) \in [a, 1 - a]$ for some $a > 0$. Further suppose that $\mu_k(\cdot)$ is σ -Lipschitz for any k . The regret per unit time under $\pi = \text{SW-OSUB}$ with a sliding window of size $\tau + 1$ satisfies: if $a > \sigma\tau$, then for any $T \geq 1$,*

$$\begin{aligned} \frac{R^\pi(T)}{T} &\leq \frac{H(\Delta, T)}{T} (1 + \Delta) + \frac{C_1 K \log(\tau)}{\tau(\Delta - 4\tau\sigma)^2} \\ &\quad + \gamma \left(1 + g_0^{-1/2}\right) \frac{\log(\tau) + c \log(\log(\tau)) + C_2}{2\tau(\Delta - 2\tau\sigma)^2}, \end{aligned}$$

where C_1, C_2 are positive constants and $g_0 = (a - \sigma\tau)(1 - a + \sigma\tau)/2$.

Corollary 5.2 *Assume that for any $n \geq 1$, $\mu(n) \in \mathcal{U}_G$ and $\mu^*(n) \in [a, 1 - a]$ for some $a > 0$, and that $\mu_k(\cdot)$ is σ -Lipschitz for any k . Set $\tau = \sigma^{-3/4} \log(1/\sigma)/8$. The regret per unit of time of $\pi = \text{SW-OSUB}$ with window size $\tau + 1$ satisfies:*

$$\limsup_{T \rightarrow \infty} \frac{R^\pi(T)}{T} \leq C \Phi(K) \sigma^{1/4} \log\left(\frac{1}{\sigma}\right) (1 + Kj(\sigma)),$$

for some constant $C > 0$, and some function j such that $\lim_{\sigma \rightarrow 0^+} j(\sigma) = 0$.

These results state that the regret per unit of time achieved under SW-OSUB decreases and actually vanishes when the speed at which expected rewards evolve decreases to 0. Also observe that the dependence of this regret bound in the number of arms is typically mild (in many practical scenarios, $\Phi(K)$ may actually not depend on K).

The proof of Theorem 5.1 relies on the same types of arguments as those used in stationary environments. To establish the regret upper bound, we need to evaluate the performance of the KL-UCB algorithm in non-stationary environments (the result and the corresponding analysis are presented in (Combes & Proutiere, 2014)).

6. Continuous Set of Arms

In this section, we briefly discuss the case where the decision space is continuous. The set of arms is $[0, 1]$, and the expected reward function $\mu : [0, 1] \rightarrow \mathbb{R}$ is assumed to be Lipschitz continuous, and unimodal: there exists $x^* \in [0, 1]$ such that $\mu(x') \geq \mu(x)$ if $x' \in [x, x^*]$ or $x' \in [x^*, x]$. Let $\mu^* = \mu(x^*)$ denote the highest expected reward. A decision rule selects at any round $n \geq 1$ an arm x and observes the corresponding reward $X(x, n)$. For any $x \in [0, 1]$, $(X(x, n))_{n \geq 1}$ is an i.i.d. sequence. We make the following additional assumption on function μ .

Assumption 2 *There exists $\delta_0 > 0$ such that (i) for all x, y in $[x^*, x^* + \delta_0]$ (or in $[x^* - \delta_0, x^*]$), $C_1 |x - y|^\alpha \leq |\mu(x) - \mu(y)|$; (ii) for $\delta \leq \delta_0$, if $|x - x^*| \leq \delta$, then $|\mu(x^*) - \mu(x)| \leq C_2 \delta^\alpha$.*

This assumption is more general than that used in (Yu & Mannor, 2011). In particular it holds for functions with a plateau and a peak: $\mu(x) = \max(1 - |x - x^*|/\epsilon, 0)$. Now as for the case of a discrete set of arms, we denote by Π the set of possible decision rules, and the regret achieved under rule $\pi \in \Pi$ up to time T is: $R^\pi(T) = T\mu^* - \sum_{n=1}^T \mathbb{E}[\mu(x^\pi(n))]$, where $x^\pi(n)$ is the arm selected under π at time n .

There is no known precise asymptotic lower bound for continuous bandits. However, we know that for our problem, the regret must be at least of the order of $O(\sqrt{T})$ up to

logarithmic factor. In (Yu & Mannor, 2011), the authors show that the LSE algorithm achieves a regret scaling as $O(\sqrt{T} \log(T))$, under more restrictive assumptions. We show that combining discretization and the UCB algorithm as initially proposed in (Kleinberg, 2004) yields lower regrets than LSE in practice (see Section 7), and is order-optimal, i.e., the regret grows as $O(\sqrt{T} \log(T))$.

For $\delta > 0$, we define a discrete bandit problem with $K = \lceil 1/\delta \rceil$ arms, and where the rewards of k -th arm are distributed as $X((k-1)/\delta, n)$. The expected reward of the k -th arm is $\mu_k = \mu((k-1)/\delta)$. Let π be an algorithm running on this discrete bandit problem. The regret of π for the initial continuous bandit problem is at time T :

$R^\pi(T) = T\mu^* - \sum_{k=1}^{\lceil 1/\delta \rceil} \mu_k \mathbb{E}[t_k^\pi(T)]$. We denote by $\text{UCB}(\delta)$ the UCB algorithm (Auer et al., 2002) applied to the discretized bandit. In the following proposition, we show that when $\delta = (\log(T)/\sqrt{T})^{1/\alpha}$, $\text{UCB}(\delta)$ is order-optimal. In practice, one may not know the time horizon T in advance. In this case, using the “doubling trick” (see e.g. (Cesa-Bianchi & Lugosi, 2006)) would incur an additional logarithmic multiplicative factor in the regret.

Proposition 1 Consider a unimodal bandit on $[0, 1]$ with rewards in $[0, 1]$ and satisfying Assumption 2. Set $\delta = (\log(T)/\sqrt{T})^{1/\alpha}$. The regret under $\text{UCB}(\delta)$ satisfies:

$$\limsup_{T \rightarrow \infty} \frac{R^\pi(T)}{\sqrt{T} \log(T)} \leq C_2 3^\alpha + 16/C_1.$$

7. Numerical experiments

7.1. Discrete bandits

We compare the performance of our algorithm to that of KL-UCB (Cappé et al., 2013), LSE (Yu & Mannor, 2011), UCB (Auer et al., 2002), and UCB-U. The latter algorithm is obtained by applying UCB restricted to the arms which are adjacent to the current leader as in OSUB. We add the prefix “SW” to refer to Sliding Window versions of these algorithms.

Stationary environments. In our first experiment, we consider $K = 17$ arms with Bernoulli rewards of respective averages $\mu = (0.1, 0.2, \dots, 0.9, 0.8, \dots, 0.1)$. The rewards are unimodal (the graph G is simply a line). The regret achieved under the various algorithms is presented in Figure 1 and Table 1. The parameters in LSE algorithm are chosen as suggested in Proposition 4.5 (Yu & Mannor, 2011). Regrets are calculated averaging over 50 independent runs. OSUB significantly outperforms all other algorithms. The regret achieved under LSE is not presented in Figure 1, because it is typically much larger than that of other algorithms. This poor performance can be explained by the non-adaptive nature of LSE, as already discussed earlier. LSE can beat UCB when the number of arms is

T	1000	10000	100000
UCB	30.1	35.1	39
KL-UCB	18.8	21.4	23
UCB-U	8.5	11.7	13.9
OSUB	5.8	5.9	6
LSE	36.3	271.5	999.1

Table 1. $R^\pi(T)/\log(T)$ for different algorithms – 17 arms.




Figure 1. Regret vs. time in stationary environments – $K = 17$ arms.

not negligible compared to the time horizon (e.g. in Figure 4 in (Yu & Mannor, 2011), $K = 250.000$ and the time horizon is less than $3K$): in such scenarios, UCB-like algorithms perform poorly because of their initialization phase (all arms have to be tested once).

In Figure 2, the number of arms is 129, and the expected rewards form a triangular shape as in the previous example, with minimum and maximum equal to 0.1 and 0.9, respectively. Similar observations as in the case of 17 arms can be made. We deliberately restrict the plot to small time horizons: this corresponds to scenarios where LSE can perform well.

Non-stationary environments. We now investigate the per-




Figure 2. Regret vs. time in stationary environments – $K = 129$ arms.




Figure 3. Regret vs. time in a slowly varying environment – $K = 10$ arms, $\sigma = 10^{-3}$.




Figure 4. Regret per unit of time $R^\pi(T)/T$ vs. speed σ – $K = 10$ arms.

formance of SW-OSUB in a slowly varying environment. There are $K = 10$ arms whose expected rewards form a moving triangle: for $k = 1, \dots, K$, $\mu_k(n) = (K-1)/K - |w(n) - k|/K$, where $w(n) = 1 + (K-1)(1 + \sin(n\sigma))/2$. Figure 3 presents the regret as a function of time under various algorithms when the speed at which the environment evolves is $\sigma = 10^{-3}$. The window size are set as follows for the various algorithms: $\tau = \sigma^{-4/5}$ for SW-UCB and SW-KL-UCB (the rationale for this choice is explained in (Combes & Proutiere, 2014)), $\tau = \sigma^{-3/4} \log(1/\sigma)/8$ for SW-UCB-U and OSUB. In Figure 4, we show how the speed σ impacts the regret per time unit. SW-OSUB provides the most efficient way of tracking the optimal arm.

7.2. Continuous bandits

In Figure 5, we compare the performance of the LSE and $\text{UCB}(\delta)$ algorithms when the set of arms is continuous. The expected rewards form a triangle: $\mu(x) = 1/2 - |x - 1/2|$ so that $\mu^* = 1/2$ and $x^* = 1/2$. The parameters used in LSE are those given in (Yu & Mannor, 2011), whereas the discretization parameter δ in $\text{UCB}(\delta)$ is set to $\delta = \log(T)/\sqrt{T}$. $\text{UCB}(\delta)$ significantly outperforms LSE at any time: an appropriate discretization of continuous bandit problems might actually be more efficient than other meth-




Figure 5. Regret vs. time for a continuous set of arms.




Figure 6. Normalized regret vs. K/T , $T = 5 \cdot 10^4$ for a continuous set of arms.

ods based on ideas taken from classical optimization theory.

Figure 6 compares the regret of the discrete version of LSE (with optimized parameters), and of OSUB as the number of arms K grows large, $T = 50,000$. The average rewards of arms are extracted from the triangle used in the continuous bandit, and we also provide the regret achieved under $\text{UCB}(\delta)$. OSUB outperforms $\text{UCB}(\delta)$ even if the number of arms gets as large as 7500! OSUB also beats LSE unless the number of arms gets bigger than $0.6 \times T$.

8. Conclusion

In this paper, we address stochastic bandit problems with a unimodal structure, and a finite set of arms. We provide asymptotic regret lower bounds for these problems and design an algorithm that asymptotically achieves the lowest regret possible. Hence our algorithm optimally exploits the unimodal structure of the problem. Our preliminary analysis of the continuous version of this bandit problem suggests that when the number of arms become very large and comparable to the time horizon, it might be wiser to prune the set of arms before actually running any algorithm.

References

- Agrawal, R. The continuum-armed bandit problem. *SIAM J. Control and Optimization*, 33(6):1926–1951, November 1995.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. Finite time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002.
- B., Edelman. Strategic bidder behavior in sponsored search auctions. In *Proc. of Workshop on Sponsored Search Auctions, ACM Electronic Commerce*, pp. 192–198, 2005.
- 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.
- Bubeck, S., Munos, R., Stoltz, G., and Szepesvári, C. Online optimization in x-armed bandits. In *Advances in Neural Information Processing Systems 22*, 2008.
- Cappé, O., Garivier, A., Maillard, O., Munos, R., and Stoltz, G. Kullback-leibler upper confidence bounds for optimal sequential allocation. *Annals of Statistics*, 41(3): 516–541, June 2013.
- Cesa-Bianchi, N. and Lugosi, G. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- Combes, R. and Proutiere, A. Unimodal bandits: Regret lower bounds and optimal algorithms. Technical Report, people.kth.se/~alepro/pdf/tr-icml2014.pdf, 2014.
- Cope, E. W. Regret and convergence bounds for a class of continuum-armed bandit problems. *IEEE Trans. Automat. Contr.*, 54(6):1243–1253, 2009.
- Dani, V., Hayes, T. P., and Kakade, S. M. Stochastic linear optimization under bandit feedback. In *Proc. of Conference On Learning Theory (COLT)*, pp. 355–366, 2008.
- Flaxman, A., Kalai, A. T., and McMahan, H. B. Online convex optimization in the bandit setting: gradient descent without a gradient. In *Proc. of ACM/SIAM symposium on Discrete Algorithms (SODA)*, pp. 385–394, 2005.
- Garivier, A. and Cappé, O. The KL-UCB algorithm for bounded stochastic bandits and beyond. In *Proc. of Conference On Learning Theory (COLT)*, 2011.
- Garivier, A. and Moulines, E. On upper-confidence bound policies for non-stationary bandit problems. In *Proc. of Algorithmic Learning Theory (ALT)*, 2008.
- Gittins, J.C. *Bandit Processes and Dynamic Allocation Indices*. John Wiley, 1989.
- Graves, T. L. and Lai, T. L. Asymptotically efficient adaptive choice of control laws in controlled markov chains. *SIAM J. Control and Optimization*, 35(3):715–743, 1997.
- Hartland, C., Baskiotis, N., Gelly, S., Teytaud, O., and Sebag, M. Change point detection and meta-bandits for online learning in dynamic environments. In *Proc. of conférence francophone sur l'apprentissage automatique (CAp07)*, 2007.
- Kleinberg, R., Slivkins, A., and Upfal, E. Multi-armed bandits in metric spaces. In *Proc. of the 40th annual ACM Symposium on Theory of Computing (STOC)*, pp. 681–690, 2008.
- Kleinberg, R. D. Nearly tight bounds for the continuum-armed bandit problem. In *Proc. of the conference on Neural Information Processing Systems (NIPS)*, 2004.
- Lai, T. L. Adaptive treatment allocation and the multi-armed bandit problem. *The Annals of Statistics*, 15(3): 1091–1114, 09 1987.
- Lai, T.L. and Robbins, H. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1): 4–2, 1985.
- Robbins, H. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952.
- Slivkins, A. and Upfal, E. Adapting to a changing environment: the brownian restless bandits. In *Proc. of Conference On Learning Theory (COLT)*, pp. 343–354, 2008.
- Yu, J. and Mannor, S. Unimodal bandits. In *Proc. of International Conference on Machine Learning (ICML)*, pp. 41–48, 2011.

Unimodal Bandits: Regret Lower Bounds and Optimal Algorithms

Supplementary File

To avoid ambiguity, when referring to theorems and lemmas from the main article, we add suffix "(a)". For instance the theorem stating the regret lower bound for stationary unimodal bandits is referred to as Theorem 4.1(a).

This document is organized as follows. In Section 1, we prove Theorem 4.1(a). In Section 2, we state and prove several concentration inequalities which are the cornerstone of our regret analysis of the OSUB algorithm for both stationary and non-stationary environments. In Section 3 we prove Theorem 4.2(a). In Section 4, we prove Theorem 5.1(a). Finally, Section 5 is devoted to the proof of Proposition 6.1(a).

1. Proof of Theorem 4.1(a)

We derive here a regret lower bound for the unimodal bandit problem. To this aim, we apply the techniques used by Graves and Lai (Graves & Lai, 1997) to investigate efficient adaptive decision rules in controlled Markov chains. We recall here their general framework. Consider a controlled Markov chain $(X_t)_{t \geq 0}$ on a finite state space \mathcal{S} with a control set U . The transition probabilities given control $u \in U$ are parametrized by θ taking values in a compact metric space Θ : the probability to move from state x to state y given the control u and the parameter θ is $p(x, y; u, \theta)$. The parameter θ is not known. The decision maker is provided with a finite set of stationary control laws $G = \{g_1, \dots, g_K\}$ where each control law g_j is a mapping from \mathcal{S} to U : when control law g_j is applied in state x , the applied control is $u = g_j(x)$. It is assumed that if the decision maker always selects the same control law g the Markov chain is then irreducible with stationary distribution π_θ^g . Now the expected reward obtained when applying control u in state x is denoted by $r(x, u)$, so that the expected reward achieved under control law g is: $\mu_\theta(g) = \sum_x r(x, g(x))\pi_\theta^g(x)$. There is an optimal control law given θ whose expected reward is denoted $\mu_\theta^* \in \arg \max_{g \in G} \mu_\theta(g)$. Now the objective of the decision maker is to sequentially select control laws so as to maximize the expected reward up to a given time horizon T . As for MAB problems, the performance of a decision scheme can be quantified through the notion of regret

which compares the expected reward to that obtained by always applying the optimal control law.

We now apply the above framework to our unimodal bandit problem, and we consider $\theta \in \Theta_G$. The Markov chain has values in $\mathcal{S} = \mathbb{R}$. The set of control laws is $G = \{1, \dots, K\}$. These laws are constant, in the sense that the control applied by control law k does not depend on the state of the Markov chain, and corresponds to selecting arm k . The transition probabilities are: $p(x, y; k, \theta) = p(y, \theta_k)$. Finally, the reward $r(x, k)$ does not depend on the state and is equal to $\mu(\theta_k)$, which is also the expected reward obtained by always using control law k .

We now fix $\theta \in \Theta_G$. Define $KL^k(\theta, \lambda) = KL(\theta_k, \lambda_k)$ for any k . Further define the set $B(\theta)$ consisting of all *bad* parameters $\lambda \in \Theta_G$ such that k^* is not optimal under parameter λ , but which are statistically *indistinguishable* from θ :

$$B(\theta) = \{\lambda \in \Theta_G : \lambda_{k^*} = \theta_{k^*} \text{ and } \max_k \mu(\lambda_k) > \mu(\lambda_{k^*})\},$$

$B(\theta)$ can be written as the union of sets $B_k(\theta)$, $(k, k^*) \in E$ defined as:

$$B_k(\theta) = \{\lambda \in B(\theta) : \mu(\lambda_k) > \mu(\lambda_{k^*})\}.$$

We have that $B(\theta) = \cup_{(k, k^*) \in E} B_k(\theta)$, because if $\mu(\lambda_{k^*}) < \max_k \mu(\lambda_k)$, then there must exist k such that $(k, k^*) \in E$, and $\mu(\lambda_k) > \mu(\lambda_{k^*})$. By applying Theorem 1 in (Graves & Lai, 1997), we know that $c(\theta)$ is the minimal value of the following LP:

$$\begin{aligned} \min \quad & \sum_k c_k (\mu(\theta_{k^*}) - \mu(\theta_k)) \\ \text{s.t.} \quad & \inf_{\lambda \in B_k(\theta)} \sum_{l \neq k^*} c_l KL^l(\theta, \lambda) \geq 1, (k, k^*) \in E \\ & c_k \geq 0, \quad \forall k. \end{aligned} \tag{1}$$

Next we show that the constraints (2) on the c_k 's are equivalent to:

$$\min_{(k, k^*) \in E} c_k I_{\min}(\theta_k, \mu(\theta_{k^*})) \geq 1. \tag{2}$$

Consider k fixed with $(k, k^*) \in E$. We prove that:

$$\inf_{\lambda \in B_k(\theta)} \sum_{l \neq k^*} c_l KL^l(\theta, \lambda) = c_k I_{\min}(\theta_k, \mu(\theta_{k^*})). \tag{3}$$

This is simply due to the following two observations:

- 110 • Since $\mu(\lambda_k) > \mu(\theta_{k^*})$ and the KL divergence is positive:
111

$$112 \sum_{l \neq k^*} c_l KL^l(\theta, \lambda) \geq c_k KL^k(\theta, \lambda) \\ 113 \\ 114 \geq c_k I_{min}(\theta_k, \mu(\theta_{k^*})).$$

- 115 • For $\epsilon > 0$, define λ_ϵ as follows: $\mu(\lambda_k) > \mu(\theta_{k^*})$ and
116 $KL(\theta_k, \lambda_k) \leq I_{min}(\theta_k, \mu(\theta_{k^*})) + \epsilon$ and $\lambda_l = \theta_l$ for
117 $l \neq k$. By construction, $\lambda_\epsilon \in B_k(\theta)$, and

$$118 \lim_{\epsilon \rightarrow 0} \sum_{l \neq k^*} c_l KL^l(\theta, \lambda_\epsilon) = c_k I_{min}(\theta_k, \mu(\theta_{k^*})).$$

119 From (5), we deduce that constraints (2) are equivalent
120 to (4) (indeed, for $(k, k^*) \in E$, (2) is equivalent to
121 $c_k I_{min}(\theta_k, \mu(\theta_{k^*})) \geq 1$). With the constraints (4), the op-
122 timization problem becomes straightforward to solve, and
123 its solution yields:

$$124 c(\theta) = \sum_{(k, k^*) \in E} \frac{\mu(\theta_{k^*}) - \mu(\theta_k)}{I_{min}(\theta_k, \mu(\theta_{k^*}))}.$$

125 □

126 2. Concentration inequalities and Preliminaries

127 We state and prove Lemma 2.1, a new concentration in-
128 equality which extends Hoeffding's inequality, and is used
129 for the regret analysis in subsequent sections. We believe
130 that Lemma 2.1 could be useful in a variety of bandit prob-
131 lems, where an upper bound on the deviation of the empiri-
132 cal mean *sampled at a stopping time* is needed. An example
133 would be the probability that the empirical reward of the k -
134 th arm deviates from its expectation, when it is sampled for
135 the s -th time.

136 **Lemma 2.1** *Let $\{Z_t\}_{t \in \mathbb{Z}}$ be a sequence of independent*
137 *random variables with values in $[0, B]$. Define \mathcal{F}_n the*
138 *σ -algebra generated by $\{Z_t\}_{t \leq n}$ and the filtration $\mathcal{F} =$*
139 *$(\mathcal{F}_n)_{n \in \mathbb{Z}}$. Consider $s \in \mathbb{N}$, $n_0 \in \mathbb{Z}$ and $T \geq n_0$. We*
140 *define $S_n = \sum_{t=n_0}^n B_t(Z_t - \mathbb{E}[Z_t])$, where $B_t \in \{0, 1\}$*
141 *is a \mathcal{F}_{t-1} -measurable random variable. Further define*
142 *$t_n = \sum_{t=n_0}^n B_t$. Define $\phi \in \{n_0, \dots, T+1\}$ a \mathcal{F} -stopping*
143 *time such that either $t_\phi \geq s$ or $\phi = T+1$.*

144 Then we have that:

$$145 \mathbb{P}[S_\phi \geq t_\phi \delta, \phi \leq T] \leq \exp\left(-\frac{2s\delta^2}{B^2}\right).$$

146 As a consequence:

$$147 \mathbb{P}[|S_\phi| \geq t_\phi \delta, \phi \leq T] \leq 2 \exp\left(-\frac{2s\delta^2}{B^2}\right).$$

148 **Proof.** Let $\lambda > 0$, and define $G_n = \exp(\lambda(S_n - \delta t_n)) \mathbf{1}\{n \leq T\}$. We have that:

$$149 \mathbb{P}[S_\phi \geq t_\phi \delta, \phi \leq T] \\ 150 = \mathbb{P}[\exp(\lambda(S_\phi - \delta t_\phi)) \mathbf{1}\{\phi \leq T\} \geq 1] \\ 151 = \mathbb{P}[G_\phi \geq 1] \leq \mathbb{E}[G_\phi].$$

152 Next we provide an upper bound for $\mathbb{E}[G_\phi]$. We define the
153 following quantities:

$$154 Y_t = B_t[\lambda(Z_t - \mathbb{E}[Z_t]) - \lambda^2 B^2/8]$$

$$155 \tilde{G}_n = \exp\left(\sum_{t=n_0}^n Y_t\right) \mathbf{1}\{n \leq T\}.$$

156 So that G can be written:

$$157 G_n = \tilde{G}_n \exp(-t_n(\lambda\delta - \lambda^2 B^2/8)).$$

158 Setting $\lambda = 4\delta/B^2$:

$$159 G_n = \tilde{G}_n \exp(-2t_n\delta^2/B^2).$$

160 Using the fact that $t_\phi \geq s$ if $\phi \leq T$, we can upper bound
161 G_ϕ by:

$$162 G_\phi = \tilde{G}_\phi \exp(-2t_\phi\delta^2/B^2) \leq \tilde{G}_\phi \exp(-2s\delta^2/B^2).$$

163 It is noted that the above inequality holds even when $\phi =$
164 $T+1$, since $G_{T+1} = \tilde{G}_{T+1} = 0$. Hence:

$$165 \mathbb{E}[G_\phi] \leq \mathbb{E}[\tilde{G}_\phi] \exp(-2s\delta^2/B^2).$$

166 We prove that $(\tilde{G}_n)_n$ is a super-martingale. We have that
167 $\mathbb{E}[\tilde{G}_{T+1}|\mathcal{F}_T] = 0 \leq \tilde{G}_T$. For $n \leq T-1$, since B_{n+1} is
168 \mathcal{F}_n measurable:

$$169 \mathbb{E}[\tilde{G}_{n+1}|\mathcal{F}_n] = \tilde{G}_n((1 - B_{n+1}) + B_{n+1}\mathbb{E}[\exp(Y_{n+1})]).$$

170 As proven by Hoeffding (Hoeffding, 1963)[eq. 4.16] since
171 $Z_{n+1} \in [0, B]$:

$$172 \mathbb{E}[\exp(\lambda(Z_{n+1} - \mathbb{E}[Z_{n+1}]))] \leq \exp(\lambda^2 B^2/8),$$

173 so $\mathbb{E}[\exp(Y_{n+1})] \leq 1$ and $(\tilde{G}_n)_n$ is indeed a supermartin-
174 gale: $\mathbb{E}[\tilde{G}_{n+1}|\mathcal{F}_n] \leq \tilde{G}_n$. Since $\phi \leq T+1$ almost surely,
175 and $(\tilde{G}_n)_n$ is a supermartingale, Doob's optional stopping
176 theorem yields: $\mathbb{E}[\tilde{G}_\phi] \leq \mathbb{E}[\tilde{G}_{n_0-1}] = 1$, and so

$$177 \mathbb{P}[S_\phi \geq t_\phi \delta, \phi \leq T] \leq \mathbb{E}[G_\phi] \\ 178 \leq \mathbb{E}[\tilde{G}_\phi] \exp(-2s\delta^2/B^2) \leq \exp(-2s\delta^2/B^2).$$

179 which concludes the proof. The second inequality is ob-
180 tained by symmetry. □

181 Lemma 2.2 states that if a set of instants Λ can be decom-
182 posed into a family of subsets $(\Lambda(s))_{s \geq 1}$ of instants (each
183 subset has at most one instant) where k is tried sufficiently
184 many times ($t_k(n) \geq \epsilon s$, for $n \in \Lambda(s)$), then the expected
185 number of instants in Λ at which the average reward of k is
186 badly estimated is finite.

165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219

Lemma 2.2 Let $k \in \{1, \dots, K\}$, and $\epsilon > 0$. Define \mathcal{F}_n the σ -algebra generated by $(X_k(t))_{1 \leq t \leq n, 1 \leq k \leq K}$. Let $\Lambda \subset \mathbb{N}$ be a (random) set of instants. Assume that there exists a sequence of (random) sets $(\Lambda(s))_{s \geq 1}$ such that (i) $\Lambda \subset \cup_{s \geq 1} \Lambda(s)$, (ii) for all $s \geq 1$ and all $n \in \Lambda(s)$, $t_k(n) \geq \epsilon s$, (iii) $|\Lambda(s)| \leq 1$, and (iv) the event $n \in \Lambda(s)$ is \mathcal{F}_n -measurable. Then for all $\delta > 0$:

$$\mathbb{E}\left[\sum_{n \geq 1} \mathbf{1}\{n \in \Lambda, |\hat{\mu}_k(n) - \mathbb{E}[\hat{\mu}_k(n)]| > \delta\}\right] \leq \frac{1}{\epsilon \delta^2}. \quad (6)$$

Proof. Let $T \geq 1$. For all $s \geq 1$, since $\Lambda(s)$ has at most one element, define $\phi_s = T + 1$ if $\Lambda(s) \cap \{1, \dots, T\}$ is empty and $\{\phi_s\} = \Lambda(s)$ otherwise. Since $\Lambda \subset \cup_{s \geq 1} \Lambda(s)$, we have:

$$\begin{aligned} & \sum_{n=1}^T \mathbf{1}\{n \in \Lambda, |\hat{\mu}_k(n) - \mathbb{E}[\hat{\mu}_k(n)]| > \delta\} \\ & \leq \sum_{s \geq 1} \mathbf{1}\{|\hat{\mu}_k(\phi_s) - \mathbb{E}[\hat{\mu}_k(\phi_s)]| > \delta, \phi_s \leq T\}. \end{aligned}$$

Taking expectations:

$$\begin{aligned} & \mathbb{E}\left[\sum_{n=1}^T \mathbf{1}\{n \in \Lambda, |\hat{\mu}_k(n) - \mathbb{E}[\hat{\mu}_k(n)]| > \delta\}\right] \\ & \leq \sum_{s \geq 1} \mathbb{P}[|\hat{\mu}_k(\phi_s) - \mathbb{E}[\hat{\mu}_k(\phi_s)]| > \delta, \phi_s \leq T]. \end{aligned}$$

Since ϕ_s is a stopping time upper bounded by $T + 1$, and that $t_k(\phi_s) \geq \epsilon s$ we can apply Lemma 2.1 to obtain:

$$\begin{aligned} & \mathbb{E}\left[\sum_{n=1}^T \mathbf{1}\{n \in \Lambda, |\hat{\mu}_k(n) - \mathbb{E}[\hat{\mu}_k(n)]| > \delta\}\right] \\ & \leq \sum_{s \geq 1} 2 \exp(-2s\epsilon\delta^2) \leq \frac{1}{\epsilon\delta^2}. \end{aligned}$$

We have used the inequality: $\sum_{s \geq 1} e^{-s} \leq \int_0^{+\infty} e^{-uw} du = 1/w$. Since the above reasoning is valid for all T , we obtain the claim (6). \square

A useful corollary of Lemma 2.2 is obtained by choosing $\delta = \Delta_{k,k'}/2$, when arms k and k' are separated by at least $\Delta_{k,k'}$.

Lemma 2.3 Let $k, k' \in \{1, \dots, K\}$ with $k \neq k'$ and $\epsilon > 0$. Define \mathcal{F}_n the σ -algebra generated by $(X_k(t))_{1 \leq t \leq n, 1 \leq k \leq K}$. Let $\Lambda \subset \mathbb{N}$ be a (random) set of instants. Assume that there exists a sequence of (random) sets $(\Lambda(s))_{s \geq 1}$ such that (i) $\Lambda \subset \cup_{s \geq 1} \Lambda(s)$, (ii) for all $s \geq 1$ and all $n \in \Lambda(s)$, $t_k(n) \geq \epsilon s$ and $t_{k'}(n) \geq \epsilon s$, (iii) for all s we have $|\Lambda(s)| \leq 1$ almost surely and (iv) for all $n \in \Lambda$, we have $\mathbb{E}[\hat{\mu}_k(n)] \leq \mathbb{E}[\hat{\mu}_{k'}(n)] - \Delta_{k,k'}$ (v) the

event $n \in \Lambda(s)$ is \mathcal{F}_n -measurable. Then:

$$\mathbb{E}\left[\sum_{n \geq 1} \mathbf{1}\{n \in \Lambda, \hat{\mu}_k(n) > \hat{\mu}_{k'}(n)\}\right] \leq \frac{8}{\epsilon \Delta_{k,k'}^2}. \quad (7)$$

Lemma 2.4 is straightforward from (Garivier & Cappé, 2011)[Theorem 10]. It should be observed that this result is not a direct application of Sanov's theorem; Lemma 2.4 provides sharper bounds in certain cases, and it is also valid for non-Bernoulli distributed random variables.

Lemma 2.4 For $1 \leq t_k(n) \leq \tau$ and $\delta > 0$, if $\{X_k(i)\}_{1 \leq i \leq \tau}$ are independent random variables with mean μ_k , we have that:

$$\begin{aligned} \mathbb{P}\left[t_k(n) I\left(\frac{1}{t_k(n)} \sum_{i=1}^{t_k(n)} X_k(i), \mu_k\right) \geq \delta\right] \\ \leq 2e[\delta \log(\tau)] \exp(-\delta). \end{aligned}$$

We present results related to the KL divergence that will be instrumental when manipulating indexes $b_k(n)$. Lemma 2.5 gives an upper and a lower bound for the KL divergence. The lower bound is Pinsker's inequality. The upper bound is due to the fact that $I(p, q)$ is convex in its second argument.

Lemma 2.5 For all $p, q \in [0, 1]^2$, $p \leq q$:

$$2(p - q)^2 \leq I(p, q) \leq \frac{(p - q)^2}{q(1 - q)}. \quad (8)$$

and

$$I(p, q) \sim \frac{(p - q)^2}{q(1 - q)}, \quad q \rightarrow p^+ \quad (9)$$

Proof. The lower bound is Pinsker's inequality. For the upper bound, we have:

$$\frac{\partial I}{\partial q}(p, q) = \frac{q - p}{q(1 - q)}.$$

Since $q \mapsto \frac{\partial I}{\partial q}(p, q)$ is increasing, the fundamental theorem of calculus gives the announced result:

$$I(p, q) \leq \int_p^q \frac{\partial I}{\partial u}(p, u) du \leq \frac{(p - q)^2}{q(1 - q)}.$$

The equivalence comes from a Taylor development of $q \rightarrow I(p, q)$ at p , since:

$$\frac{\partial I}{\partial q}(p, q)|_{q=p} = 0,$$

$$\frac{\partial^2 I}{\partial q^2}(p, q)|_{q=p} = \frac{1}{q(1 - q)}.$$

\square

We prove a deviation bound similar to that of Lemma 2.2 for non-stationary environments.

Lemma 2.6 Let $k \in \{1, \dots, K\}$, $n_0 \in \mathbb{N}$ and $\epsilon > 0$. Let $\Lambda \subset \mathbb{N}$ be a (random) set of instants. Assume that there exists a sequence of (random) sets $(\Lambda(s))_{s \geq 1}$ such that (i) $\Lambda \subset \cup_{s \geq 1} \Lambda(s)$, (ii) for all $s \geq 1$ and all $n \in \Lambda(s)$, $t_k(n) \geq \epsilon s$, and (iii) for all $s \geq 1$ $|\Lambda(s) \cap [n_0, n_0 + \tau]| \leq 1$. Then for all $\delta > 0$:

$$\mathbb{E}\left[\sum_{n=n_0}^{n_0+\tau} \mathbf{1}\{n \in \Lambda, |\hat{\mu}_k(n) - \mathbb{E}[\hat{\mu}_k(n)]| > \delta\}\right] \leq \frac{\log(\tau)}{2\epsilon\delta^2} + 2.$$

Proof. Fix $s_0 \geq 1$. We use the following decomposition, depending on the value of s with respect to s_0 :

$$\{n \in \Lambda, |\hat{\mu}_k(n) - \mathbb{E}[\hat{\mu}_k(n)]| > \delta\} \subset A \cup B,$$

where

$$A = \{n_0, \dots, n_0 + \tau\} \cap (\cup_{1 \leq s \leq s_0} \Lambda(s)),$$

$$B = \{n_0, \dots, n_0 + \tau\}$$

$$\cap \{n \in \cup_{s \geq s_0} \Lambda(s) : |\hat{\mu}_k(n) - \mathbb{E}[\hat{\mu}_k(n)]| > \delta\}.$$

Since for all s , $|\Lambda(s) \cap \{n_0, \dots, n_0 + \tau\}| \leq 1$, we have $|A| \leq s_0$. The expected size of B is upper bounded by:

$$\begin{aligned} E[|B|] &\leq \sum_{n=n_0}^{n_0+\tau} \mathbb{P}[n \in \cup_{s \geq s_0} \Lambda(s), |\hat{\mu}_k(n) - \mathbb{E}[\hat{\mu}_k(n)]| > \delta] \\ &\leq \sum_{n=n_0}^{n_0+\tau} \mathbb{P}[|\hat{\mu}_k(n) - \mathbb{E}[\hat{\mu}_k(n)]| > \delta, t_k(n) \geq \epsilon s_0]. \end{aligned}$$

For a given n , we apply Lemma 2.1 with $n - \tau$ in place of n_0 , and $\phi = n$ if $t_k(n) \geq \epsilon s_0$ and $\phi = T + 1$ otherwise. It is noted that ϕ is indeed a stopping time. We get:

$$\begin{aligned} \mathbb{P}[|\hat{\mu}_k(n) - \mathbb{E}[\hat{\mu}_k(n)]| > \delta, t_k(n) \geq \epsilon s_0] \\ \leq 2 \exp(-2s_0\epsilon\delta^2). \end{aligned}$$

Therefore, setting $s_0 = \log(\tau)/(2\epsilon\delta^2)$,

$$E[|B|] \leq 2\tau \exp(-2s_0\epsilon\delta^2) = 2.$$

Finally we obtain the announced result:

$$\mathbb{E}\left[\sum_{n=n_0}^{n_0+\tau} \mathbf{1}\{n \in \Lambda, |\hat{\mu}_k(n) - \mathbb{E}[\hat{\mu}_k(n)]| > \delta\}\right] \leq \frac{\log(\tau)}{2\epsilon\delta^2} + 2.$$

□

Lemma 2.7 Consider $k, k' \in \{1, \dots, K\}$, $n_0 \in \mathbb{N}$ and $\epsilon > 0$. Let $\Lambda \subset \mathbb{N}$ be a (random) set of instants. Assume that there exists a sequence of (random) sets $(\Lambda(s))_{s \geq 1}$ such that (i) $\Lambda \subset \cup_{s \geq 1} \Lambda(s)$, and (ii) for all $s \geq 1$ and all $n \in \Lambda(s)$, $t_k(n) \geq \epsilon s$, $t_{k'}(n) \geq \epsilon s$ and (iii) for all $s \geq 1$ $|\Lambda(s) \cap [n_0, n_0 + \tau]| \leq 1$ and (iv) for all $n \in \Lambda$, we have $\mathbb{E}[\hat{\mu}_k(n)] \leq \mathbb{E}[\hat{\mu}_{k'}(n)] - \Delta_{k,k'}$.

Then for all $\delta > 0$:

$$\mathbb{E}\left[\sum_{n=n_0}^{n_0+\tau} \mathbf{1}\{n \in \Lambda, \hat{\mu}_k(n) > \hat{\mu}_{k'}(n)\}\right] \leq \frac{4 \log(\tau)}{\epsilon \Delta_{k,k'}^2} + 4.$$

3. Proofs for stationary environments

3.1. Proof of Theorem 4.2(a)

Notations. Throughout the proof, by a slight abuse of notation, we omit the floor/ceiling functions when it does not create ambiguity. Consider a suboptimal arm $k \neq k^*$. Define the difference between the average reward of k and k' : $\Delta_{k,k'} = |\mu_{k'} - \mu_k| > 0$. We use the notation:

$$t_{k,k'}(n) = \sum_{t=1}^n \mathbf{1}\{L(t) = k, k(t) = k'\}.$$

$t_{k,k'}(n)$ is the number of times up time n that k' has been selected given that k was the leader.

Proof. Let $T > 0$. The regret $R^{OSUB}(T)$ of OSUB algorithm up to time T is:

$$R^{OSUB}(T) = \sum_{k \neq k^*} (\mu_{k^*} - \mu_k) \mathbb{E}\left[\sum_{n=1}^T \mathbf{1}\{k(n) = k\}\right].$$

We use the following decomposition:

$$\begin{aligned} \mathbf{1}\{k(n) = k\} &= \mathbf{1}\{L(n) = k^*, k(n) = k\} \\ &\quad + \mathbf{1}\{L(n) \neq k^*, k(n) = k\}. \end{aligned}$$

Now

$$\begin{aligned} \sum_{k \neq k^*} (\mu_{k^*} - \mu_k) \mathbb{E}\left[\sum_{n=1}^T \mathbf{1}\{L(n) \neq k^*, k(n) = k\}\right] \\ \leq \sum_{k \neq k^*} \mathbb{E}\left[\sum_{n=1}^T \mathbf{1}\{L(n) \neq k^*, k(n) = k\}\right] \\ \leq \sum_{k \neq k^*} \mathbb{E}[l_k(T)]. \end{aligned}$$

Observing that when $L(n) = k^*$, the algorithm selects a decision $(k, k^*) \in E$, we deduce that:

$$\begin{aligned} R^{OSUB}(T) &\leq \sum_{k \neq k^*} \mathbb{E}[l_k(T)] \\ &\quad + \sum_{(k,k^*) \in E} (\mu_{k^*} - \mu_k) \mathbb{E}\left[\sum_{n=1}^T \mathbf{1}\{L(n) = k^*, k(n) = k\}\right] \end{aligned}$$

Then we analyze the two terms in the r.h.s. in the above inequality. The first term corresponds to the average number of times where k^* is not the leader, while the second term represents the accumulated regret when the leader is k^* . The following result states that the first term is $O(\log(\log(T)))$:

Theorem 3.1 For $k \neq k^*$, $\mathbb{E}[l_k(T)] = O(\log(\log(T)))$.

From the above theorem, we conclude that the leader is k^* except for a negligible number of instants (in expectation). When k^* is the leader, OSUB behaves as KL-UCB restricted to the set $N(k^*)$ of possible decisions. Following the same analysis as in (Garivier & Cappé, 2011) (the analysis of KL-UCB), we can show that for all $\epsilon > 0$ there are constants $C_1 \leq 7$, $C_2(\epsilon)$ and $\beta(\epsilon) > 0$ such that:

$$\begin{aligned} & \mathbb{E}\left[\sum_{n=1}^T \mathbf{1}\{L(n) = k^*, k(n) = k\}\right] \\ & \leq \mathbb{E}\left[\sum_{n=1}^T \mathbf{1}\{b_k(n) \geq b_{k^*}(n)\}\right] \\ & \leq (1 + \epsilon) \frac{\log(T)}{I(\mu_k, \mu_{k^*})} + C_1 \log(\log(T)) + \frac{C_2(\epsilon)}{T^{\beta(\epsilon)}}. \end{aligned} \quad (10)$$

Combining the above bound with Theorem 3.1, we get:

$$R^{OSUB}(T) \leq (1 + \epsilon)c(\theta) \log(T) + O(\log(\log(T))), \quad (11)$$

which concludes the proof of Theorem 4.2(a). \square

It remains to show that Theorem 3.1 holds, which is done in the next section. The proof of Theorem 3.1 is technical, and requires the concentration inequalities presented in section 2. The theorem itself is proved in 3.2.

3.2. Proof of Theorem 3.1

Let k be the index of a suboptimal arm. Let $\delta > 0$, $\epsilon > 0$ small enough (we provide a more precise definition later on). We define $k_2 = \arg \max_{k' : (k, k') \in E} \mu_{k'}$ the best neighbor of k . To derive an upper bound of $\mathbb{E}[l_k(T)]$, we decompose the set of times where k is the leader into the following sets:

$$\{n \leq T : L(n) = k\} \subset A_\epsilon \cup B_\epsilon^T,$$

where

$$\begin{aligned} A_\epsilon &= \{n : L(n) = k, t_{k_2}(n) \geq \epsilon l_k(n)\} \\ B_\epsilon^T &= \{n \leq T : L(n) = k, t_{k_2}(n) \leq \epsilon l_k(n)\}. \end{aligned}$$

Hence we have:

$$\mathbb{E}[l_k(T)] \leq \mathbb{E}[|A_\epsilon| + |B_\epsilon^T|],$$

Next we provide upper bounds of $\mathbb{E}[|A_\epsilon|]$ and $\mathbb{E}[|B_\epsilon^T|]$.

Bound on $\mathbb{E}[|A_\epsilon|]$. Let $n \in A_\epsilon$ and assume that $l_k(n) = s$. By design of the algorithm, $t_k(n) \geq s/(\gamma + 1)$. Also $t_{k_2}(n) \geq \epsilon l_k(n) = \epsilon s$. We apply Lemma 2.3 with $\Lambda(s) = \{n \in A_\epsilon, l_k(n) = s\}$, $\Lambda = \cup_{s \geq 1} \Lambda(s)$. Of course, for any s , $|\Lambda(s)| \leq 1$. We have: $A_\epsilon = \{n \in \Lambda : \hat{\mu}_k(n) \geq \hat{\mu}_{k_2}(n)\}$, since when $n \in A_\epsilon$, k is the leader. Lemma 2.3 can be applied with $k' = k_2$. We get: $\mathbb{E}[|A_\epsilon|] < \infty$.

Bound on $\mathbb{E}[|B_\epsilon^T|]$. We introduce the following sets:

- C_δ is the set of instants at which the average reward of the leader k is badly estimated:

$$C_\delta = \{n : L(n) = k, |\hat{\mu}_k(n) - \mu_k| > \delta\}.$$

- $D_\delta = \cup_{k' \in N(k) \setminus \{k_2\}} D_{\delta, k'}$ where $D_{\delta, k'} = \{n : L(n) = k, k(n) = k', |\hat{\mu}_{k'}(n) - \mu_{k'}| > \delta\}$ is the set of instants at which k is the leader, k' is selected and the average reward of k' is badly estimated.
- $E^T = \{n \leq T : L(n) = k, b_{k_2}(n) \leq \mu_{k_2}\}$, is the set of instants at which k is the leader, and the upper confidence index $b_{k_2}(n)$ underestimates the average reward μ_{k_2} .

We first prove that $|B_\epsilon^T| \leq 2\gamma(1 + \gamma)(|C_\delta| + |D_\delta| + |E^T|) + O(1)$ as T grows large, and then provide upper bounds on $\mathbb{E}[|C_\delta|]$, $\mathbb{E}[|D_\delta|]$, and $\mathbb{E}[|E^T|]$. Let $n \in B_\epsilon^T$. When k is the leader, the selected decision is in $N(k)$:

$$l_k(n) = t_{k, k_2}(n) + \sum_{k' \in N(k) \setminus \{k_2\}} t_{k, k'}(n).$$

We recall that $t_{k, k'}(n)$ denotes the number of times up to time n when k is the leader and k' is selected. Since $n \in B_\epsilon^T$, $t_{k, k_2}(n) \leq \epsilon l_k(n)$, from which we deduce that:

$$(1 - \epsilon)l_k(n) \leq \sum_{k' \in N(k) \setminus \{k_2\}} t_{k, k'}(n).$$

Choose $\epsilon < 1/(2(\gamma + 1))$. With this choice, from the previous inequality, we must have that either (a) there exists $k_1 \in N(k) \setminus \{k, k_2\}$, $t_{k, k_1}(n) \geq l_k(n)/(\gamma + 1)$ or (b) $t_{k, k}(n) \geq (3/2)l_k(n)/(\gamma + 1) + 1$.

(a) Assume that $t_{k, k_1}(n) \geq l_k(n)/(\gamma + 1)$. Since $t_{k, k_1}(n)$ is only incremented when k_1 is selected and k is the leader, and since $n \mapsto l_k(n)$ is increasing, there exists a unique $\phi(n) < n$ such that $L(\phi(n)) = k$, $k(\phi(n)) = k_1$, $t_{k, k_1}(\phi(n)) = \lfloor l_k(n)/(2(\gamma + 1)) \rfloor$. $\phi(n)$ is indeed unique because $t_{k, k_1}(\phi(n))$ is incremented at time $\phi(n)$.

Next we prove by contradiction that for $l_k(n) \geq l_0$ large enough and δ small enough, we must have $\phi(n) \in C_\delta \cup D_\delta \cup E^T$. Assume that $\phi(n) \notin C_\delta \cup D_\delta \cup E^T$. Then $b_{k_2}(\phi(n)) \geq \mu_{k_2}$, $\hat{\mu}_{k_1}(\phi(n)) \leq \mu_{k_1} + \delta$. Using Pinsker's inequality and the fact that $t_{k_1}(\phi(n)) \geq t_{k, k_1}(\phi(n))$:

$$\begin{aligned} b_{k_1}(\phi(n)) &\leq \hat{\mu}_{k_1}(\phi(n)) \\ &+ \sqrt{\frac{\log(l_k(\phi(n))) + c \log(\log(l_k(\phi(n))))}{2t_{k_1}(\phi(n))}} \\ &\leq \mu_{k_1} + \delta + \sqrt{\frac{\log(l_k(n)) + c \log(\log(l_k(n)))}{2\lfloor l_k(n)/(2(\gamma + 1)) \rfloor}}. \end{aligned}$$

Now select $\delta < (\mu_{k_2} - \mu_k)/2$ and l_0 such that $\sqrt{(\log(l_0) + c \log(\log(l_0)))/2} \lceil l_0/(2(\gamma + 1)) \rceil \leq \delta$. If $l_k(n) \geq l_0$:

$$b_{k_1}(\phi(n)) \leq \mu_{k_1} + 2\delta < \mu_{k_2} \leq b_{k_2}(\phi(n)),$$

which implies that k_1 cannot be selected at time $\phi(n)$ (because $b_{k_1}(\phi(n)) < b_{k_2}(\phi(n))$), a contradiction.

(b) Assume that $t_{k,k}(n) \geq (3/2)l_k(n)/(\gamma + 1) + 1 = l_k(n)/(\gamma + 1) + l_k(n)/(2(\gamma + 1)) + 1$. There are at least $l_k(n)/(2(\gamma + 1)) + 1$ instants \tilde{n} such that $l_k(\tilde{n}) - 1$ is not a multiple of $1/(\gamma + 1)$, $L(\tilde{n}) = k$ and $k(\tilde{n}) = k$. By the same reasoning as in (a) there exists a unique $\phi(n) < n$ such that $L(\phi(n)) = k$, $k(\phi(n)) = k$, $t_{k,k}(\phi(n)) = \lfloor l_k(n)/(2(\gamma + 1)) \rfloor$ and $(l_k(\phi(n)) - 1)$ is not a multiple of $1/(\gamma + 1)$. So $b_k(\phi(n)) \geq b_{k_2}(\phi(n))$. The same reasoning as that applied in (a) (replacing k_1 by k) yields $\phi(n) \in C_\delta \cup D_\delta \cup E^T$.

We define $B_{\epsilon, l_0}^T = \{n : n \in B_\epsilon^T, l_k(n) \geq l_0\}$, and we have that $|B_\epsilon^T| \leq l_0 + |B_{\epsilon, l_0}^T|$. We have defined a mapping ϕ from B_{ϵ, l_0}^T to $C_\delta \cup D_\delta \cup E^T$. To bound the size of B_{ϵ, l_0}^T , we use the following decomposition:

$$\begin{aligned} & \{n : n \in B_{\epsilon, l_0}^T, l_k(n) \geq l_0\} \\ & \subset \cup_{n' \in C_\delta \cup D_\delta \cup E^T} \{n : n \in B_{\epsilon, l_0}^T, \phi(n) = n'\}. \end{aligned}$$

Let us fix n' . If $n \in B_{\epsilon, l_0}^T$ and $\phi(n) = n'$, then $\lfloor l_k(n)/(2(\gamma + 1)) \rfloor \in \cup_{k' \in N(k) \setminus \{k_2\}} \{t_{k, k'}(n')\}$ and $l_k(n)$ is incremented at time n because $L(n) = k$. Therefore:

$$|\{n : n \in B_{\epsilon, l_0}^T, \phi(n) = n'\}| \leq 2\gamma(\gamma + 1).$$

Using union bound, we obtain the desired result:

$$|B_\epsilon^T| \leq l_0 + |B_{\epsilon, l_0}^T| \leq O(1) + 2\gamma(\gamma + 1)(|C_\delta| + |D_\delta| + |E^T|).$$

Bound on $\mathbb{E}|C_\delta|$. We apply Lemma 2.2 with $\Lambda(s) = \{n : L(n) = k, l_k(n) = s\}$, and $\Lambda = \cup_{s \geq 1} \Lambda(s)$. Then of course, $|\Lambda(s)| \leq 1$ for all s . Moreover by design, $t_k(n) \geq s/(\gamma + 1)$ when $n \in \Lambda(s)$, so we can choose any $\epsilon < 1/(\gamma + 1)$ in Lemma 2.2. Now $C_\delta = \{n \in \Lambda : |\hat{\mu}_k(n) - \mu_k| > \delta\}$. From (6), we get $\mathbb{E}|C_\delta| < \infty$.

Bound on $\mathbb{E}|D_\delta|$. Let $k' \in N(k) \setminus \{k_2\}$. Define for any s , $\Lambda(s) = \{n : L(n) = k, k(n) = k', t_{k'}(n) = s\}$, and $\Lambda = \cup_{s \geq 1} \Lambda(s)$. We have $|\Lambda(s)| \leq 1$, and for any $n \in \Lambda(s)$, $t_{k'}(n) = s \geq \epsilon s$ for any $\epsilon < 1$. We can now apply Lemma 2.2 (where k is replaced by k'). Note that $D_{\delta, k'} = \{n \in \Lambda : |\hat{\mu}_{k'}(n) - \mu_{k'}| > \delta\}$, and hence (6) leads to $\mathbb{E}|D_{\delta, k'}| < \infty$, and thus $\mathbb{E}|D_\delta| < \infty$.

Bound on $\mathbb{E}|E^T|$. We can show as in (Garivier & Cappé, 2011) (the analysis of KL-UCB) that $\mathbb{E}|E^T| =$

$O(\log(\log(T)))$ (more precisely, this result is a simple application of Theorem 10 in (Garivier & Cappé, 2011)).

We have shown that $\mathbb{E}|B_\epsilon^T| = O(\log(\log(T)))$, and hence $\mathbb{E}[l_k(T)] = O(\log(\log(T)))$, which concludes the proof of Theorem 3.1. \square

4. Proofs for non-stationary environments

To simplify the notation, we remove the superscript τ throughout the proofs, e.g. $t_k^\tau(n)$ and $l_k^\tau(n)$ are denoted by $t_k(n)$ and $l_k(n)$.

4.1. A lemma for sums over a sliding window

We will use Lemma 4.1 repeatedly to bound the number of times some events occur over a sliding window of size τ .

Lemma 4.1 *Let $A \subset \mathbb{N}$, and $\tau \in \mathbb{N}$ fixed. Define $a(n) = \sum_{t=n-\tau}^{n-1} \mathbf{1}\{t \in A\}$. Then for all $T \in \mathbb{N}$ and $s \in \mathbb{N}$ we have the inequality:*

$$\sum_{n=1}^T \mathbf{1}\{n \in A, a(n) \leq s\} \leq s \lceil T/\tau \rceil. \quad (12)$$

As a consequence, for all $k \in \{1, \dots, K\}$, we have:

$$\begin{aligned} & \sum_{n=1}^T \mathbf{1}\{k(n) = k, t_k(n) \leq s\} \leq s \lceil T/\tau \rceil, \quad (13) \\ & \sum_{n=1}^T \mathbf{1}\{L(n) = k, l_k(n) \leq s\} \leq s \lceil T/\tau \rceil. \end{aligned}$$

These inequalities are obtained by choosing $A = \{n : k(n) = k\}$ and $A = \{n : L(n) = k\}$ in (12).

Proof. We decompose $\{1, \dots, T\}$ into intervals of size τ : $\{1, \dots, \tau\}$, $\{\tau + 1, \dots, 2\tau\}$ etc. We have:

$$\begin{aligned} & \sum_{n=1}^T \mathbf{1}\{n \in A, a(n) \leq s\} \\ & \leq \sum_{i=0}^{\lceil T/\tau \rceil - 1} \sum_{n=1}^{\tau} \mathbf{1}\{n + i\tau \in A, a(n + i\tau) \leq s\}. \quad (14) \end{aligned}$$

Fix i and assume that $\sum_{n=1}^{\tau} \mathbf{1}\{n + i\tau \in A, a(n + i\tau) \leq s\} > s$. Then there must exist $n' < \tau$ such that $n' \in A$ and $\sum_{n=1}^{n'} \mathbf{1}\{n + i\tau \in A, a(n + i\tau) \leq s\} = s$. Since $a(n' + i\tau) \geq \sum_{n=1}^{n'} \mathbf{1}\{n + i\tau \in A, a(n + i\tau) \leq s\}$, we have $a(n' + i\tau) \geq s$. As $n' \in A$, we must have $a(n'' + i\tau) \geq$

($s + 1$) for all $n'' > n'$ such that $n'' \in A$. So

$$\begin{aligned} & \sum_{n=1}^{\tau} \mathbf{1}\{n + i\tau \in A, a(n + i\tau) \leq s\} \\ &= \sum_{n=1}^{n'} \mathbf{1}\{n + i\tau \in A, a(n + i\tau) \leq s\} = s, \end{aligned}$$

which is a contradiction. Hence, for all i :

$$\sum_{n=1}^{\tau} \mathbf{1}\{n + i\tau \in A, a(n + i\tau) \leq s\} \leq s,$$

and substituting in (14) gives the desired result:

$$\sum_{n=1}^T \mathbf{1}\{n \in A, a(n) \leq s\} \leq \sum_{i=0}^{\lceil T/\tau \rceil - 1} s = s \lceil T/\tau \rceil.$$

□

4.2. Regret of SW-KL-UCB

In order to analyze the regret of SW-OSUB, we first have to analyze the regret SW-KL-UCB on which SW-OSUB is based.

Theorem 4.2 *Let $\Delta: 2\tau\sigma < \Delta < \Delta_0$. Assume that for any $n \geq 1$, $\mu^*(n) \in [a, 1 - a]$ for some $a > 0$. Further suppose that $\mu_k(\cdot)$ is σ -Lipschitz for any k . The regret per unit time under $\pi = \text{SW-KL-UCB}$ with a sliding window of size τ satisfies: if $a > \sigma\tau$, then for any $T \geq 1$,*

$$\begin{aligned} \frac{R^\pi(T)}{T} &\leq \frac{H(\Delta, T)}{T} \Delta \\ &+ K \left(1 + g_0^{-1/2}\right) \frac{\log(\tau) + c \log(\log(\tau)) + C_1}{2\tau(\Delta - 2\tau\sigma)^2}, \end{aligned}$$

where C_1 is a positive constant and $g_0 = (a - \sigma\tau)(1 - a + \sigma\tau)/2$.

Recall that due to the changing environment and the use of a sliding window, the empirical reward is a biased estimator of the average reward, and that its bias is upper bounded by $\sigma\tau$.

To ease the regret analysis, we first provide bounds on the empirical reward. Unlike in the stationary case, the empirical reward $\hat{\mu}_k(n)$ is not a sum of $t_k(n)$ i.i.d. variables. We define $\bar{X}_k(n', n) = X_k(n') + (\mu_k(n) + \sigma|n' - n| - \mu_k(n'))$, $\underline{X}_k(n', n) = X_k(n') + (\mu_k(n) - \sigma|n' - n| - \mu_k(n'))$ and:

$$\begin{aligned} \hat{\mu}_k(n) &= \frac{1}{t_k(n)} \sum_{n'=n-\tau}^n \underline{X}_k(n', n) \mathbf{1}\{k(n') = k\}, \\ \bar{\mu}_k(n) &= \frac{1}{t_k(n)} \sum_{n'=n-\tau}^n \bar{X}_k(n', n) \mathbf{1}\{k(n') = k\}. \end{aligned}$$

Then of course, $\hat{\mu}_k(n) \leq \hat{\mu}_k(n) \leq \bar{\mu}_k(n)$.

Now the regret under $\pi = \text{SW-OSUB}$ is given by:

$$R^\pi(T) = \sum_{n=1}^T \sum_{k=1}^K (\mu_{k^*}(n) - \mu_k(n)) \mathbb{P}[k(n) = k].$$

We define $I_{\min} = 2(\Delta - 2\tau\sigma)^2$. Let $\epsilon > 0$ and $\mathcal{K}^\tau = (1 + \epsilon) \frac{\log(\tau) + c \log(\log(\tau))}{I_{\min}}$. We introduce the following sets of events:

(i) $A = \cup_{k=1}^K A_k$, where

$$A_k = \{1 \leq n \leq T : k(n) = k, |\mu_k(n) - \mu_{k^*}(n)| < \Delta\},$$

A_k is the set of times at which k is chosen, and k is "close" to the optimal decision. Note that, by definition, $|A| \leq H(\Delta, T)$.

(ii) $B = \{1 \leq n \leq T : b_{k^*}(n) \leq \mu_{k^*}(n) - \tau\sigma\}$. B is the set of times at which the index $b_{k^*}(n)$ underestimates the average reward of the optimal decision (with an error greater than the bias $\tau\sigma$).

(iii) $C = \cup_{k=1}^K C_k$, $C_k = \{1 \leq n \leq T : k(n) = k, t_k(n) \leq \mathcal{K}^\tau\}$. C_k is the set of times at which k is selected and it has been tried less than \mathcal{K}^τ times.

(iv) $D = \cup_{k=1}^K D_k$, $D_k = \{1 \leq n \leq T : k(n) = k, n \notin (A \cup B \cup C)\}$. D_k is the set of times where (a) k is chosen, (b) k has been tried more than \mathcal{K}^τ times, (c) k is not close to the optimal decision, and (d) the average reward of the optimal decision is not underestimated.

We will show that:

$$\sum_{n \in A} (\mu^*(n) - \mu_{k(n)}(n)) \leq \Delta H(\Delta, T). \quad (15)$$

and the following inequalities

$$\begin{aligned} \mathbb{E}[|B|] &\leq O(T/\tau), \quad \mathbb{E}[|C_k|] \leq \mathcal{K}^\tau \lceil T/\tau \rceil, \\ \mathbb{E}[|D_k|] &\leq \frac{T}{(\tau \log(\tau)^c)^{g_0 \epsilon^2}}. \end{aligned}$$

We deduce that:

$$\begin{aligned} R^\pi(T) &\leq \Delta H(\Delta, T) + O(T/\tau) \\ &+ K \mathcal{K}^\tau \lceil T/\tau \rceil + \frac{KT}{(\tau \log(\tau)^c)^{g_0 \epsilon^2}}, \end{aligned}$$

which proves Theorem 4.2.

Proof of (15). Let $n \in A_k$. If $n \in A_k$, by definition we have $|\mu_{k^*}(n) - \mu_k(n)| < \Delta$. Then if $k(n) = k$, we have that $\mu^*(n) - \mu_{k(n)}(n) \leq \Delta$ so that:

$$\sum_{n \in A} (\mu^*(n) - \mu_{k(n)}(n)) \leq \Delta |A| \leq \Delta H(\Delta, T),$$

which completes the proof of (15).

Bound on $\mathbb{E}[|B|]$. Let $n \in B$. Note that $\hat{\mu}_{k^*}(n) \leq \bar{\mu}_{k^*}(n) \leq b_{k^*}(n)$. Since $b_{k^*}(n) \leq \mu_{k^*}(n) - \sigma\tau$, we deduce that: $\hat{\mu}_{k^*}(n) \leq \mu_{k^*}(n) - \sigma\tau$. Now we have:

$$\begin{aligned} \mathbb{P}[n \in B] &= \mathbb{P}[b_{k^*}(n) \leq \mu_{k^*}(n) - \sigma\tau] \\ &= \mathbb{P}[t_{k^*}(n)I(\hat{\mu}_{k^*}(n), \mu_{k^*}(n) - \sigma\tau) \\ &\quad \geq \log(\tau) + c \log(\log(\tau))] \\ &\stackrel{(a)}{\leq} \mathbb{P}[t_{k^*}(n)I(\hat{\mu}_{k^*}(n), \mu_{k^*}(n) - \sigma\tau) \\ &\quad \geq \log(\tau) + c \log(\log(\tau))] \\ &\stackrel{(b)}{\leq} \frac{2e}{\tau(\log(\tau))^{c-2}}, \end{aligned}$$

where (a) is due to the fact that $\hat{\mu}_{k^*}(n) \leq \bar{\mu}_{k^*}(n)$, and (b) is obtained applying Lemma 2.4. Hence: $\mathbb{E}[|B|] \leq O(T/\tau)$.

Bound on $\mathbb{E}[|C_k|]$. Using Lemma 4.1, we get $|C_k| \leq \mathcal{K}^\tau \lceil T/\tau \rceil$, and hence $|C| \leq K\mathcal{K}^\tau \lceil T/\tau \rceil$.

Bound on $\mathbb{E}[|D_k|]$. We will prove that $n \in D_k$ implies that $\bar{\mu}_k(n)$ deviates from its expectation by at least $f(\epsilon, I_{\min}) > 0$ so that:

$$\mathbb{P}[n \in D_k] \leq \mathbb{P}[\bar{\mu}_k(n) - \mathbb{E}[\bar{\mu}_k(n)] > f(\epsilon, I_{\min})].$$

Let $n \in D_k$. Since $k(n) = k$ and $b_{k^*}(n) \geq \mu_{k^*}(n) - \sigma\tau$, we have $b_k(n) \geq \mu_{k^*}(n) - \sigma\tau$. We decompose D_k as follows:

$$\begin{aligned} D_k &= D_{k,1} \cup D_{k,2} \\ D_{k,1} &= \{n \in D_k : \bar{\mu}_k(n) \geq \mu_{k^*}(n) - \sigma\tau\} \\ D_{k,2} &= \{n \in D_k : \bar{\mu}_k(n) \leq \mu_{k^*}(n) - \sigma\tau\} \end{aligned}$$

If $n \in D_{k,1}$, $\bar{\mu}_k(n) - \mathbb{E}[\bar{\mu}_k(n)] \geq \mu_{k^*}(n) - \mu_k(n) - 2\sigma\tau > 0$ so that $\bar{\mu}_k(n)$ indeed deviates from its expectation. Now let $n \in D_{k,2}$. We have:

$$\begin{aligned} \mathbb{P}[n \in D_{k,2}] &\leq \mathbb{P}[b_k(n) \geq \mu_{k^*}(n) - \sigma\tau, n \in D_{k,2}] \\ &= \mathbb{P}[t_k(n)I(\bar{\mu}_k(n), \mu_{k^*}(n) - \sigma\tau) \\ &\quad \leq \log(\tau) + c \log(\log(\tau)), n \in D_{k,2}] \\ &\stackrel{(a)}{\leq} \mathbb{P}[\mathcal{K}^\tau I(\bar{\mu}_k(n), \mu_{k^*}(n) - \sigma\tau) \\ &\quad \leq \log(\tau) + c \log(\log(\tau)), t_k(n) \geq \mathcal{K}^\tau] \\ &= \mathbb{P}\left[I(\bar{\mu}_k(n), \mu_{k^*}(n) - \sigma\tau) \leq \frac{I_{\min}}{1+\epsilon}, t_k(n) \geq \mathcal{K}^\tau\right], \end{aligned}$$

where in (a), we used the facts that: $\bar{\mu}_k(n) \leq \mu_{k^*}(n) - \sigma\tau$, $\bar{\mu}_k(n) \geq \hat{\mu}_k(n)$, and $t_k(n) \geq \mathcal{K}^\tau$ ($n \notin C$). It is noted

that since $n \notin A_k$, by Pinsker's inequality we have that: $I(\mu_k(n) + \tau\sigma, \mu_{k^*}(n) - \tau\sigma) \geq 2(\mu_{k^*}(n) - \mu_k(n) - 2\tau\sigma)^2 \geq 2(\Delta - 2\tau\sigma)^2 = I_{\min}$. By continuity and monotonicity of the KL divergence, there exists a unique positive function f such that:

$$\begin{aligned} I(\mu_k(n) + \sigma\tau + f(\epsilon, I_{\min}), \mu_{k^*}(n) - \sigma\tau) &= \frac{I_{\min}}{1+\epsilon}, \\ \mu_k(n) + \sigma\tau + f(\epsilon, I_{\min}) &\leq \mu_{k^*}(n) - \sigma\tau. \end{aligned}$$

We are interested in the asymptotic behavior of f when ϵ , I_{\min} both tend to 0. Define μ' , μ'' and μ_0 such that

$$\mu_k(n) + \sigma\tau \leq \mu' \leq \mu'' \leq \mu_0 = \mu_{k^*}(n) - \sigma\tau.$$

and

$$I(\mu', \mu_0) = I_{\min}, \quad I(\mu'', \mu_0) = \frac{I_{\min}}{1+\epsilon}.$$

Using the equivalent (9) given in Lemma 2.5, there exists a function a such that:

$$\begin{aligned} \frac{(\mu_0 - \mu')^2}{\mu_0(1 - \mu_0)}(1 + a(\mu_0 - \mu')) &= I_{\min}, \\ \frac{(\mu_0 - \mu'')^2}{\mu_0(1 - \mu_0)}(1 + a(\mu_0 - \mu'')) &= \frac{I_{\min}}{1+\epsilon}. \end{aligned}$$

with $a(\delta) \rightarrow 0$ when $\delta \rightarrow 0^+$. It is noted that $0 \leq \mu_0 - \mu'' \leq \mu_0 - \mu' = o(1)$ when $I_{\min} \rightarrow 0^+$ by continuity of the KL divergence. Hence:

$$\mu'' - \mu' = \left(\frac{\epsilon}{2} + o(1)\right) \sqrt{\mu_0(1 - \mu_0)I_{\min}}.$$

Using the inequality

$$\begin{aligned} f(\epsilon, I_{\min}) &= \mu'' - (\mu_k(n) + \sigma\tau) \\ &\geq \mu'' - \mu' = \frac{\epsilon}{2} \sqrt{\mu_0(1 - \mu_0)I_{\min}}, \end{aligned}$$

we have proved that:

$$2f(\epsilon, I_{\min})^2 \geq \epsilon^2 g_0 I_{\min} + o(\epsilon^2)$$

with

$$g_0 = (a - \sigma\tau)(1 - a + \sigma\tau)/2.$$

Therefore, since $\mathbb{E}[\bar{\mu}_k(n)] \leq \mu_k(n) + \sigma\tau$, as claimed, we have

$$\begin{aligned} \mathbb{P}[n \in D_k] &\leq \mathbb{P}[\bar{\mu}_k(n) - \mathbb{E}[\bar{\mu}_k(n)] \geq f(\epsilon, I_{\min}), t_k(n) \geq \mathcal{K}^\tau]. \end{aligned}$$

We now apply Lemma 2.1 with $n - \tau$ in place of n_0 , \mathcal{K}^τ in place of s and $\phi = n$ if $t_k(n) \geq \mathcal{K}^\tau$ and $\phi = T + 1$ otherwise. We obtain, for all n :

$$\begin{aligned} \mathbb{P}[n \in D_k] &\leq \mathbb{P}[\bar{\mu}_k(n) - \mathbb{E}[\bar{\mu}_k(n)] \geq f(\epsilon, I_{\min}), t_k(n) \geq \mathcal{K}^\tau] \\ &\leq \exp(-2\mathcal{K}^\tau f(\epsilon, I_{\min})^2) \leq \frac{1}{(\tau \log(\tau)^c) g_0 \epsilon^2}, \end{aligned}$$

825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879

and we get the desired bound by summing over n :

$$\mathbb{E}[|D_k|] = \sum_{n=1}^T \mathbb{P}[n \in D_k] \leq \frac{T}{(\tau \log(\tau))^c g_0 \epsilon^2}.$$

4.3. Proof of Theorem 5.1(a)

We first introduce some notations. For any set A of instants, we use the notation: $A[n_0, n] = A \cap \{n_0, \dots, n_0 + \tau\}$. Let $n_0 \leq n$. We define $t_k(n_0, n)$ the number of times k has been chosen during interval $\{n_0, \dots, n_0 + \tau\}$, $l_k(n_0, n)$ the number of times k has been the leader, and $t_{k,k'}(n_0, n)$ the number of times k' has been chosen while k was the leader:

$$\begin{aligned} t_k(n_0, n) &= \sum_{n'=n_0}^n \mathbf{1}\{k(n') = k\}, \\ l_k(n_0, n) &= \sum_{n'=n_0}^n \mathbf{1}\{L(n') = k\}, \\ t_{k,k'}(n_0, n) &= \sum_{n'=n_0}^n \mathbf{1}\{L(n') = k, k(n') = k'\}. \end{aligned}$$

Note that $l_k(n - \tau, n) = l_k(n)$, $t_k(n - \tau, n) = t_k(n)$ and $t_{k,k'}(n - \tau, n) = t_{k,k'}(n)$. Given $\Delta > 0$, we define the set of instants at which the average reward of k is separated from the average reward of its neighbours by at least Δ :

$$\mathcal{N}_k(\Delta) = \cap_{(k',k) \in E} \{n : |\mu_k(n) - \mu_{k'}(n)| > \Delta\}.$$

We further define the amount of time that k is suboptimal, k is the leader, and it is well separated from its neighbors:

$$\mathcal{L}_k(\Delta) = \{n : L(n) = k \neq k^*(n), n \in \mathcal{N}_k(\Delta)\}.$$

By definition of the regret under $\pi = \text{SW-OSUB}$:

$$R^\pi(T) = \sum_{n=1}^T \sum_{k \neq k^*(n)} (\mu_{k^*}(n) - \mu_k(n)) \mathbb{P}[k(n) = k].$$

To bound the regret, as in the stationary case, we split the regret into two components: the regret accumulated when the leader is the optimal arm, and the regret generated when the leader is not the optimal arm. The regret when the

leader is suboptimal satisfies:

$$\begin{aligned} & \sum_{n=1}^T \sum_{k \neq k^*(n)} (\mu_{k^*}(n) - \mu_k) \mathbf{1}\{k(n) = k, L(n) \neq k^*(n)\} \\ & \leq \sum_{n=1}^T \mathbf{1}\{L(n) \neq k^*(n)\} \\ & \leq \sum_{n=1}^T \sum_{k \neq k^*(n)} \mathbf{1}\{L(n) = k \neq k^*(n)\} \\ & \leq \sum_{n=1}^T \sum_{k \neq k^*(n)} \mathbf{1}\{n \in \mathcal{L}_k(\Delta)\} \\ & \quad + \mathbf{1}\{\exists k' : (k, k') \in E : |\mu_k(n) - \mu_{k'}(n)| \leq \Delta\} \\ & \leq \left(\sum_{k=1}^K |\mathcal{L}_k(\Delta)[0, T]| + H(\Delta, T) \right). \end{aligned}$$

Therefore the regret satisfies:

$$\begin{aligned} R^\pi(T) & \leq \left(H(\Delta, T) + \sum_{k=1}^K \mathbb{E}[|\mathcal{L}_k(\Delta)[0, T]|] \right) \\ & \quad + \sum_{n=1}^T \sum_{(k, k^*(n)) \in E} (\mu_{k^*}(n) - \mu_k(n)) \mathbb{P}[k(n) = k]. \end{aligned} \tag{16}$$

The second term of the r.h.s in (16) is the regret of SW-OSUB when $k^*(n)$ is the leader. This term can be analyzed using the same techniques as those used for the analysis of SW-KL-UCB and is upper bounded by the regret of SW-KL-UCB. It remains to bound the first term of the r.h.s in (16).

Theorem 4.3 Consider $\Delta > 4\tau\sigma$. Then for all k :

$$\mathbb{E}[|\mathcal{L}_k(\Delta)[0, T]|] \leq C_1 \times \frac{T \log(\tau)}{\tau(\Delta - 4\tau\sigma)^2}, \tag{17}$$

where $C_1 > 0$ does not depend on T , τ , σ and Δ .

Substituting (17) in (16), we obtain the announced result. \square

4.4. Proof of Theorem 4.3

It remains to prove Theorem 4.3. Define $\delta = (\Delta - 4\tau\sigma)/2$. We can decompose $\{1, \dots, T\}$ into at most $\lceil T/\tau \rceil$ intervals of size τ . Therefore, to prove the theorem, it is sufficient to prove that for all $n_0 \in \mathcal{L}_k(\Delta)$ we have:

$$\mathbb{E}[|\mathcal{L}_k(\Delta)[n_0, n_0 + \tau]|] \leq O\left(\frac{\log(\tau)}{\delta^2}\right).$$

In the remaining of the proof, we consider an interval $\{n_0, \dots, n_0 + \tau\}$, with $n_0 \in \mathcal{L}_k(\Delta)$ fixed. It is noted that the best neighbour of k changes with time. We define $k_2(n)$ the best neighbor of k at time n . From the Lipschitz assumption and the fact that $\Delta > 4\tau\sigma$, we have that for all $n \in \{n_0, \dots, n_0 + \tau\}$, $k_2(n) = k_2(n_0)$. Indeed for all $n \in \{n_0, \dots, n_0 + \tau\}$:

$$\begin{aligned} & \mu_{k_2(n_0)}(n) - \mu_k(n) \\ & \geq \mu_{k_2(n_0)}(n_0) - \mu_k(n_0) - 2(n - n_0)\sigma \\ & \geq \Delta - 2\tau\sigma \geq 2\tau\sigma > 0. \end{aligned}$$

We write $k_2 = k_2(n_0) = k_2(n)$ when this does not create ambiguity. We will use the fact that, for all $n \in \{n_0, \dots, n_0 + \tau\}$:

$$\begin{aligned} \mathbb{E}[\hat{\mu}_{k_2}(n)] - \mathbb{E}[\hat{\mu}_k(n)] & \geq \mu_{k_2}(n) - \mu_k(n) - 2\tau\sigma, \\ & \geq \mu_{k_2}(n_0) - \mu_k(n_0) - 4\tau\sigma, \\ & \geq \Delta - 4\tau\sigma = 2\delta > 0. \end{aligned}$$

We decompose $\mathcal{L}_k(\Delta)[n_0, n_0 + \tau] = A_\epsilon^{n_0} \cup B_\epsilon^{n_0}$, with:

$A_\epsilon^{n_0} = \{n \in \mathcal{L}_k(\Delta)[n_0, n_0 + \tau], t_{k_2}(n) \geq \epsilon l_k(n_0, n)\}$ the set of times where k is the leader, k is not the optimal arm, and its best neighbor k_2 has been tried sufficiently many times during interval $\{n_0, \dots, n_0 + \tau\}$,

$B_\epsilon^{n_0} = \{n \in \mathcal{L}_k(\Delta)[n_0, n_0 + \tau], t_{k_2}(n) \leq \epsilon l_k(n_0, n)\}$ the set of times where k is the leader, k is not the optimal arm, and its best neighbor k_2 has been little tried during interval $\{n_0, \dots, n_0 + \tau\}$.

Bound on $\mathbb{E}[A_\epsilon^{n_0}]$. Let $n \in A_\epsilon^{n_0}$. We recall that $\mathbb{E}[\hat{\mu}_{k_2}(n)] - \mathbb{E}[\hat{\mu}_k(n)] \geq 2\delta$, so that the reward of k or k_2 must be badly estimated at time n :

$$\begin{aligned} \mathbb{P}[n \in A_\epsilon^{n_0}] & \leq \mathbb{P}[|\hat{\mu}_k(n) - \mathbb{E}[\hat{\mu}_k(n)]| > \delta] \\ & \quad + \mathbb{P}[|\hat{\mu}_{k_2}(n) - \mathbb{E}[\hat{\mu}_{k_2}(n)]| > \delta]. \end{aligned}$$

We apply Lemma 2.7, with $k' = k_2$, $\Delta_{k,k'} = 2\delta$, $\Lambda(s) = \{n \in A_\epsilon^{n_0}, l_k(n_0, n) = s\}$, $t_{k_2}(n) \geq \epsilon l_k(n_0, n) = \epsilon s$. By design of SW-OSUB : $t_k(n) \geq l_k(n_0, n)/(\gamma + 1) = s/(\gamma + 1)$. Using the fact that $|\Lambda(s)| \leq 1$ for all s , we have that:

$$\mathbb{E}[A_\epsilon^{n_0}] \leq O\left(\frac{\log(\tau)}{\epsilon\delta^2}\right).$$

Bound on $\mathbb{E}[B_\epsilon^{n_0}]$. Define l_0 such that

$$\sqrt{\frac{\log(l_0) + c \log(\log(l_0))}{2\lfloor l_0/(2(\gamma + 1)) \rfloor}} \leq \delta.$$

In particular we can choose $l_0 = 2(\gamma + 1)(\log(1/\delta)/\delta^2)$. Indeed, with such a choice we have that

$$\sqrt{\frac{\log(l_0) + c \log(\log(l_0))}{2\lfloor l_0/(2(\gamma + 1)) \rfloor}} \sim \delta/2, \delta \rightarrow 0^+.$$

Let $\epsilon < 1/(2(\gamma + 1))$, and define the following sets:

$C_\delta^{n_0}$ is the set of instants at which the average reward of the leader k is badly estimated:

$$\begin{aligned} C_\delta^{n_0} & = \{n \in \{n_0, \dots, n_0 + \tau\} \\ & \quad : L(n) = k \neq k^*(n), |\hat{\mu}_k(n) - \mathbb{E}[\hat{\mu}_k(n)]| > \delta\}; \end{aligned}$$

$D_\delta^{n_0} = \cup_{k' \in N(k) \setminus \{k_2\}} D_{\delta,k'}^{n_0}$, where $D_{\delta,k'}^{n_0} = \{n : L(n) = k \neq k^*(n), k(n) = k', |\hat{\mu}_{k'}(n) - \mathbb{E}[\hat{\mu}_{k'}(n)]| > \delta\}$. $D_\delta^{n_0}$ is the set of instants at which k is the leader, k' is selected and the average reward of k' is badly estimated.

$E^{n_0} = \{n \leq T : L(n) = k \neq k^*(n), b_{k_2}(n) \leq \mathbb{E}[\hat{\mu}_{k_2}(n)]\}$ is the set of instants at which k is the leader, and the upper confidence index $b_{k_2}(n)$ underestimates the average reward $\mathbb{E}[\hat{\mu}_{k_2}(n)]$.

Let $n \in B_\epsilon^{n_0}$. Write $s = l_k(n_0, n)$, and we assume that $s \geq l_0$. Since $t_{k_2}(n_0, n) \leq \epsilon l_k(n_0, n)$ and the fact that $l_k(n_0, n) = t_{k_2}(n_0, n) + \sum_{k' \in N(k) \setminus \{k_2\}} t_{k'}(n_0, n)$, we must have (a) there exists $k_1 \in N(k) \setminus \{k, k_2\}$ such that $t_{k_1}(n_0, n) \geq s/(\gamma + 1)$ or (b) $t_{k_1}(n_0, n) \geq (3/2)s/(\gamma + 1) + 1$. Since $t_{k,k}(n)$ and $t_{k,k_2}(n)$ are incremented only at times when $k(n) = k$ and $k(n) = k_2$ respectively, there must exist a unique index $\phi(n) \in \{n_0, \dots, n_0 + \tau\}$ such that either: (a) $t_{k,k_1}(\phi(n)) = \lfloor s/(2(\gamma + 1)) \rfloor$ and $k(\phi(n)) = k_1$; or (b) $t_{k,k_2}(\phi(n)) = \lfloor (3/2)s/(\gamma + 1) \rfloor$ and $k(n) = k$ and $l_k(\phi(n))$ is not a multiple of 3. In both cases, as in the proof of theorem 3.1, we must have that $\phi(n) \in C_\delta^{n_0} \cup D_\delta^{n_0} \cup E^{n_0}$.

We now upper bound the number of instants n which are associated to the same $\phi(n)$. Let $n, n' \in B_\epsilon^{n_0}$ and $s = l_k(n_0, n)$. We see that $\phi(n') = \phi(n)$ implies either $\lfloor l_k(n_0, n')/(2(\gamma + 1)) \rfloor = \lfloor l_k(n_0, n)/(2(\gamma + 1)) \rfloor$ or $\lfloor (3/2)l_k(n_0, n')/(\gamma + 1) \rfloor = \lfloor (3/2)l_k(n_0, n)/(\gamma + 1) \rfloor$. Furthermore, $n' \mapsto l_k(n_0, n')$ is incremented at time n' . Hence for all $n \in B_\epsilon^{n_0}$:

$$|n' \in B_\epsilon^{n_0}, \phi(n') = \phi(n)| \leq 2\gamma(\gamma + 1).$$

We have established that:

$$\begin{aligned} |B_\epsilon^{n_0}| & \leq l_0 + 2\gamma(\gamma + 1)(|C_\delta^{n_0}| + |D_\delta^{n_0}| + |E^{n_0}|) \\ & = 2(\gamma + 1)\log(1/\delta)/\delta^2 \\ & \quad + 2\gamma(\gamma + 1)(|C_\delta^{n_0}| + |D_\delta^{n_0}| + |E^{n_0}|). \end{aligned}$$

We complete the proof by providing bounds of the expected sizes of sets $C_\delta^{n_0}$, $D_\delta^{n_0}$ and E^{n_0} .

Bound of $\mathbb{E}[C_\delta^{n_0}]$: Using Lemma 2.6 with $\Lambda(s) = \{n \in C_\delta^{n_0}, l_k(n_0, n) = s\}$, and by design of SW-OSUB : $t_k(n) \geq l_k(n_0, n)/(\gamma + 1) = s/(\gamma + 1)$. Since $|\Lambda(s)| \leq 1$ for all s , we have that:

$$\mathbb{E}[|C_\delta^{n_0}|] \leq O\left(\frac{\log(\tau)}{\delta^2}\right).$$

Bound of $\mathbb{E}[D_\delta^{n_0}]$: Using Lemma 2.6 with $\Lambda(s) = \{n \in D_\delta^{n_0}, t_{k,k'}(n_0, n) = s\}$, and $|\Lambda(s)| \leq 1$ for all s , we have that:

$$\mathbb{E}[|D_{\delta,k'}^{n_0}|] \leq O\left(\frac{\log(\tau)}{\delta^2}\right).$$

Bound of $\mathbb{E}[E^{n_0}]$: By Lemma 2.4 since $l_k(n) \leq \tau$:

$$\begin{aligned} \mathbb{P}[n \in E^{n_0}] &\leq 2e[\log(\tau)(\log(\tau) + c \log(\log(\tau)))] \\ &\quad \exp(-\log(\tau) + c \log(\log(\tau))) \\ &\leq \frac{4e}{\tau \log(\tau)^{c-2}}. \end{aligned}$$

Thus

$$\mathbb{E}[|E^{n_0}|] \leq \frac{4e}{(\log \tau)^{c-2}}.$$

Putting the various bounds all together, we have:

$$\mathbb{E}[|\mathcal{L}_k(\Delta)[n_0, n_0 + \tau]|] \leq O\left(\frac{\log(\tau)}{\delta^2}\right),$$

for all $n_0 \in \mathcal{L}_k(\Delta)$, uniformly in δ , which concludes the proof. \square

5. Proof of Proposition 6.1(a)

The regret of UCB(δ) is defined as:

$$R^\pi(T) \leq \sum_{k=1}^{\lceil 1/\delta \rceil} \mathbb{E}[t_k(T)](\mu^* - \mu_k).$$

We separate the arms into three different sets. $\{1, \dots, \lceil 1/\delta \rceil\} = A \cup B \cup C$, with: $A = \{k^* - 1, k^*, k^* + 1\}$ the optimal arm and its neighbors, $B = \{k : k \notin A, (k-1)\delta \in [x^* - \delta_0, x^* + \delta_0]\}$ the arms which are not neighbors of the optimal arm, but are in $[x^* - \delta_0, x^* + \delta_0]$, and $C = \{k : (k-1)\delta \notin [x^* - \delta_0, x^* + \delta_0]\}$ the rest of the arms.

We consider $\delta < \delta_0/3$, so that $A \subset [x^* - \delta_0, x^* + \delta_0]$. By our assumption on the reward function, if $k \in A$, $|x^* -$

$\delta(k-1)| \leq 2\delta$ then $|\mu^* - \mu_k| \leq C_2(2\delta)^\alpha$. The regret is upper bounded by:

$$R^\pi(T) \leq TC_2(2\delta)^\alpha + \sum_{k \in B \cup C} \mathbb{E}[t_k(T)](\mu^* - \mu_k).$$

Using the fact that $\mu^* - \mu_{k^*} \leq C_2\delta^\alpha$ and $\sum_{k=1}^{\lceil 1/\delta \rceil} \mathbb{E}[t_k(T)] \leq T$, the bound becomes:

$$R^\pi(T) \leq TC_2(3\delta)^\alpha + \sum_{k \in B \cup C} \mathbb{E}[t_k(T)](\mu_{k^*} - \mu_k).$$

By (Auer et al., 2002) (the analysis of UCB), for all k , $\mathbb{E}[t_k(T)] \leq 8 \log(T)/(\mu_{k^*} - \mu_k)^2$. Replacing in the regret upper bound:

$$R^\pi(T) \leq TC_2(3\delta)^\alpha + \sum_{k \in B \cup C} 8 \log(T)/(\mu_{k^*} - \mu_k).$$

If $k \in B$, $|\delta(k^* - 1) - \delta(k-1)| \geq \delta(|k^* - k| - 1)$, so $\mu_{k^*} - \mu_k \geq C_1\delta^\alpha(|k^* - k| - 1)^\alpha$. If $k \in C$, then $|\delta(k^* - 1) - \delta(k-1)| \geq \delta_0/2$, so $\mu_{k^*} - \mu_k \geq C_1(\delta_0/2)^\alpha$. So the regret for arms in $B \cup C$ reduces to:

$$R^\pi(T) \leq TC_2(3\delta)^\alpha + \frac{8 \log(T) \lceil 1/\delta \rceil}{C_1(\delta_0/2)^\alpha} + 2 \sum_{k=1}^{\lceil 1/\delta \rceil} \frac{8 \log(T)}{C_1(\delta k)^\alpha}.$$

Using a sum-integral comparison: $\sum_{k=1}^{\lceil 1/\delta \rceil} k^{-\alpha} \leq \sum_{k=1}^{\lceil 1/\delta \rceil} k^{-1} \leq 1 + \log(\lceil 1/\delta \rceil)$, so that:

$$\begin{aligned} R^\pi(T) &\leq TC_2(3\delta)^\alpha \\ &\quad + 8 \log(T) \left(\frac{\lceil 1/\delta \rceil}{C_1(\delta_0/2)^\alpha} + \frac{2(1 + \log(\lceil 1/\delta \rceil))}{C_1\delta^\alpha} \right). \end{aligned}$$

Setting $\delta = (\log(T)/\sqrt{T})^{1/\alpha}$, the regret becomes:

$$\begin{aligned} R^\pi(T) &\leq TC_2(3^\alpha)(\log(T)/\sqrt{T}) + \\ &\quad 8 \log(T) \left(\frac{\lceil (\sqrt{T}/\log(T))^{1/\alpha} \rceil}{C_1(\delta_0/2)^\alpha} + \frac{2(1 + \log(T))}{C_1 \log(T)/\sqrt{T}} \right). \end{aligned}$$

we have used the fact that $\lceil 1/\delta \rceil \leq T$.

$$\begin{aligned} R^\pi(T) &\leq C_2(3^\alpha) \log(T) \sqrt{T} \\ &\quad + 8 \left(\frac{\sqrt{T} + 1}{C_1(\delta_0/2)^\alpha} + \frac{2\sqrt{T}(1 + \log(T))}{C_1} \right) \end{aligned}$$

Letting $T \rightarrow \infty$ gives the result:

$$\limsup_T R^\pi(T)/(\sqrt{T} \log(T)) \leq C_2 3^\alpha + 16/C_1.$$

1210	References	1265
1211		1266
1212	Auer, P., Cesa-Bianchi, N., and Fischer, P. Finite time analysis of the multiarmed bandit problem. <i>Machine Learning</i> , 47(2-3):235–256, 2002.	1267
1213		1268
1214		1269
1215	Garivier, A. and Cappé, O. The KL-UCB algorithm for bounded stochastic bandits and beyond. In <i>Proceedings of Conference On Learning Theory (COLT)</i> , 2011.	1270
1216		1271
1217		1272
1218		1273
1219	Graves, T. L. and Lai, T. L. Asymptotically efficient adaptive choice of control laws in controlled markov chains. <i>SIAM J. Control and Optimization</i> , 35(3):715–743, 1997.	1274
1220		1275
1221		1276
1222		1277
1223	Hoeffding, Wassily. Probability inequalities for sums of bounded random variables. <i>Journal of the American Statistical Association</i> , 58(301):pp. 13–30, 1963.	1278
1224		1279
1225		1280
1226		1281
1227		1282
1228		1283
1229		1284
1230		1285
1231		1286
1232		1287
1233		1288
1234		1289
1235		1290
1236		1291
1237		1292
1238		1293
1239		1294
1240		1295
1241		1296
1242		1297
1243		1298
1244		1299
1245		1300
1246		1301
1247		1302
1248		1303
1249		1304
1250		1305
1251		1306
1252		1307
1253		1308
1254		1309
1255		1310
1256		1311
1257		1312
1258		1313
1259		1314
1260		1315
1261		1316
1262		1317
1263		1318
1264		1319