



Bounded Regret for Finitely Parameterized Multi-Armed Bandits

Kishan Panaganti, Dileep Kalathil, and Pravin Varaiya

Abstract We consider multi-armed bandits where the model of the underlying stochastic environment is characterized by a common unknown parameter. The true parameter is unknown to the learning agent. However, the set of possible parameters, which is finite, is known a priori. We propose an algorithm that is simple and easy to implement, which we call Finitely Parameterized Upper Confidence Bound (FP-UCB) algorithm, which uses the information about the underlying parameter set for faster learning. In particular, we show that the FP-UCB algorithm achieves a bounded regret under a structural condition on the underlying parameter set. We also show that, if the underlying parameter set does not satisfy this structural condition, the FP-UCB algorithm achieves a logarithmic regret, but with a smaller preceding constant compared to the standard UCB algorithm. We also validate the superior performance of the FP-UCB algorithm through extensive numerical simulations.

1 Introduction

The Multi-Armed Bandit (MAB) problem is a canonical formalism for studying how an agent learns to take optimal actions through repeated interactions with a stochastic environment. The learning agent receives a reward at each time step which will depend on the action of the agent as well as the stochastic uncertainty of the environment. The goal of the agent is to act so as to maximize the cumulative reward. When the model of the environment is known, computing the optimal action is a standard optimization problem. The challenge in MAB is that the agent does not know the stochastic model of environment a priori. The agent needs to *explore*, i.e., take ac-

Kishan Panaganti and Dileep Kalathil are with the Department of Electrical and Computer Engineering at Texas A&M University, College Station, Texas, USA, e-mail: kpb,dileep.kalathil@tamu.edu · Pravin Varaiya is with the Department of Electrical Engineering and Computer Sciences at University of California, Berkeley, California, USA, e-mail: varaiya@berkeley.edu

tions to gather information and estimate the model of the system. At the same time the agent must *exploit* the available information to maximize the cumulative reward. This *exploration vs. exploitation* trade-off is at the core of the MAB problem.

Lai and Robbins in their seminal paper [19] formulated the non-Bayesian stochastic MAB problem and characterized the performance of a learning algorithm using the metric of *regret*. They showed that no learning algorithm can achieve a regret better than $\mathcal{O}(\log T)$. They also proposed a learning algorithm that achieves an asymptotic logarithmic regret, matching the fundamental lower bound. A simple index-based algorithm called UCB algorithm was introduced in [5] which achieves the order optimal regret in a non-asymptotic manner. This approach led to a number of interesting algorithms, among them linear bandits [13], contextual bandits [11], combinatorial bandits [10], and decentralized and multi-player bandits [15].

Thompson (Posterior) Sampling is another class of algorithms that give superior numerical performance for MAB problems. The posterior sampling heuristic was first introduced by Thompson [25], but the first rigorous performance guarantee, an $\mathcal{O}(\log T)$ regret, was given in [2]. The Thompson sampling idea has been used in algorithms for bandits with multiple plays [17], contextual bandits [3], general online learning problem [14], and reinforcement learning [23]. Both classes of algorithms have been used in a number of practical applications, like communication networks [24], smart grids [16], and recommendation systems [29].

Our contribution: We consider a class of multi-armed bandits problems where the reward corresponding to each arm is characterized by a common unknown parameter with a finite set of possible values. This restriction is inspired by real-world applications. For example, in recommendation systems and e-commerce applications (Amazon, Netflix), it is typical to assume that each user has a certain ‘type’ parameter (denoted by θ in our formulation), and the set of possible parameter values is finite. The preferences of the user is characterized by her type (for example, prefer science books over fiction books). The set of all possible types and the preferences of each type may be known a priori, but the type of a new user may be unknown. So, instead of learning the preferences of this user over all possible choices, it may be easier to learn the type parameter of this user from a few observations. In this work, we propose an algorithm that explicitly uses the availability of such structural information about the underlying parameter set which enables faster learning.

We propose an algorithm that is simple and easy to implement, which we call FP-UCB algorithm, which uses the structural information for faster learning. We show that the proposed FP-UCB algorithm can achieve a bounded regret ($\mathcal{O}(1)$) under some structural condition on the underlying parameter set. This is in sharp contrast to the increasing ($\mathcal{O}(\log T)$) regret of standard multi-armed bandits algorithms. We also show that, if the underlying parameter set does not satisfy the structural condition, the FP-UCB algorithm achieves a regret of $\mathcal{O}(\log T)$, but with a smaller preceding constant compared to the standard UCB algorithm. The regret achieved by our algorithm also matches with the fundamental lower bound given by [1]. One remarkable aspect of our algorithm is that, it is oblivious to whether the underlying parameter set satisfies the necessary condition or not, thereby avoiding re-tuning

of the algorithm depending on the problem instance. Instead, it achieves the best possible performance given the problem instance.

Related work: Finitely parameterized multi-armed bandits problem were first studied by Agrawal et al. [1]. They proposed an algorithm for this setting, and proved that it achieves a bounded regret when the parameter set satisfies some necessary condition, and logarithmic regret otherwise. However, their algorithm is rather complicated, which limits practical implementations and extension to other settings. The regret analysis is also involved and asymptotic in nature, different from the recent simpler index-based bandits algorithms and their finite time analysis. [1] also provided a fundamental lower bound for this class of problems. Compared to this work, our FP-UCB algorithm is simple, easy to implement, and easy to analyze, while providing non-asymptotic performance guarantees that match the lower bound.

Some recent works exploit the available structure of the MAB problem to get tighter regret bounds. In particular, [4] [20] [22] [12] consider the problem setting similar to our paper where the mean reward of each arm is characterized by a single unknown parameter. [4] assumes that the reward functions are continuous in the global parameter and gives a bounded regret result. [20] gives specific conditions on the mean reward to achieve a bounded regret. [22] considers a latent bandit problem where the reward distributions are partitioned into a number of clusters and indexed by a latent parameter corresponding to the cluster. [12] characterizes the minimal rates at which sub-optimal arms have to be explored depending on the structural information, and proposes an algorithm that achieves these rates. [8] [7] [26] exploit a different structural information where it is shown that if the mean value of the best arm and the second best arm (but not the identity of the arms) are known, a bounded regret can be achieved. There also are bandit algorithms that exploit side information [28] [9], and recently in the context of contextual bandits [6]. Our problem formulation, algorithm, and analysis are different from these works. We also note that our problem formulation is fundamentally different from the system identification problems [21] [18] because the goal here is to learn an optimal policy online.

2 Problem Formulation

We consider the following sequential decision-making problem. In each time step $t \in \{1, 2, \dots, T\}$, the agent selects an arm (action) from the set of L possible arms, denoted $a(t) \in [L] = \{1, \dots, L\}$. Each arm i , when selected, yields a random real-valued reward. Let $X_i(\tau)$ be the random reward from arm i in its τ th selection. We assume that $X_i(\tau)$ is drawn according to a probability distribution $P_i(\cdot; \theta^o)$ with mean $\mu_i(\theta^o)$. Here θ^o is the (true) parameter that determines the distribution of the stochastic rewards. The agent does not know θ^o or the corresponding mean value $\mu_i(\theta^o)$. The random rewards obtained from playing an arm repeatedly are i.i.d. and independent of the plays of the other arms. The rewards are bounded with support

in $[0, 1]$. The goal of the agent is to select a sequence of actions that maximizes the expected cumulative reward, $\mathbb{E}[\sum_{t=1}^T \mu_{a(t)}(\theta^o)]$. The action $a(t)$ depends on the history of observations available to the agent until time t . So, $a(t)$ is stochastic and the expectation is with respect to its randomness.

Clearly, the optimal choice is to select the best arm (the arm with the highest mean value) all the time, i.e., $a(t) = a^*(\theta^o), \forall t$, where $a^*(\theta^o) = \arg \max_{i \in [L]} \mu_i(\theta^o)$. However, the agent will be able to make this optimal decision only if she knows the parameter θ^o or the corresponding mean values $\mu_i(\theta^o)$ for all i . The goal of a MAB algorithm is to learn to make the optimal sequence of decisions without knowing the true parameter θ^o .

We consider the setting where the agent knows the set of possible parameters Θ . We assume that Θ is finite. If the true parameter were $\theta \in \Theta$, then agent selecting arm i will get a random reward drawn according to a distribution $P_i(\cdot; \theta)$ with mean $\mu_i(\theta)$. We assume that for each $\theta \in \Theta$, the agent knows $P_i(\cdot; \theta)$ and $\mu_i(\theta)$ for all $i \in [L]$. The optimal arm corresponding to the parameter θ is $a^*(\theta) = \arg \max_{i \in [L]} \mu_i(\theta)$. We emphasize that the agent does not know the true parameter θ^o (and hence the optimal action $a^*(\theta^o)$) except that it is in the finite set Θ .

In the multi-armed bandits literature, it is standard to characterize the performance of an online learning algorithm using the metric of regret. Regret is defined as the performance loss of an algorithm as compared to the optimal algorithm with complete information. Since this is $b(t) = a^*(\theta^o)$, the expected cumulative regret of a multi-armed bandits algorithm after T time steps is defined as

$$\mathbb{E}[R(T)] := \mathbb{E} \left[\sum_{t=1}^T (\mu_{a^*(\theta^o)}(\theta^o) - \mu_{a(t)}(\theta^o)) \right]. \quad (1)$$

The goal of a MAB learning algorithm is to select actions sequentially in order to minimize $\mathbb{E}[R(T)]$.

3 UCB Algorithm for Finitely Parameterized Multi-Armed Bandits

In this section, we present our algorithm for finitely parameterized multi-armed bandits and the main theorem. We first introduce a few notations for presenting the algorithm and the results succinctly.

Let $n_i(t)$ be the number of times arm i has been selected by the algorithm until time t , i.e., $n_i(t) = \sum_{\tau=1}^t \mathbb{1}\{a(\tau) = i\}$. Here $\mathbb{1}\{\cdot\}$ is an indicator function. Define the empirical mean corresponding to arm i at time t as,

$$\hat{\mu}_i(t) := \frac{1}{n_i(t)} \sum_{\tau=1}^{n_i(t)} X_i(\tau). \quad (2)$$

Define the set $A := \{a^*(\theta) : \theta \in \Theta\}$, which is the collection of optimal arms corresponding to all parameters in Θ . Intuitively, a learning agent can restrict selection to arms from the set A . Clearly, $A \subset [L]$ and this reduction can be useful when $|A|$ is much smaller than L .

Our FP-UCB Algorithm is given in Algorithm 1. Figure 1 gives an illustration of the episodes and time slots of the FP-UCB algorithm.

For stating the main result, we introduce a few more notations. We define the *confusion set* $B(\theta^o)$ and $C(\theta^o)$ as,

$$\begin{aligned} B(\theta^o) &:= \{\theta \in \Theta : a^*(\theta) \neq a^*(\theta^o) \text{ and } \mu_{a^*(\theta^o)}(\theta^o) = \mu_{a^*(\theta^o)}(\theta)\}, \\ C(\theta^o) &:= \{a^*(\theta) : \theta \in B(\theta^o)\}. \end{aligned}$$

Intuitively, $B(\theta^o)$ is the set of parameters that can be confused with the true parameter θ^o . If $B(\theta^o)$ is non-empty, selecting $a^*(\theta^o)$ and estimating the empirical mean is not sufficient to identify the true parameter because the same mean reward can result from other parameters in $B(\theta^o)$. So, if $B(\theta^o)$ is non-empty, more exploration (i.e., selecting sub-optimal arms other than $a^*(\theta^o)$) is necessary to identify the true parameter. This exploration will contribute to the regret. On the other hand, if $B(\theta^o)$ is empty, the optimal parameter can be identified with much less exploration, which results in a bounded regret. $C(\theta^o)$ is the corresponding set of arms that needs to be explored sufficiently to identify the optimal parameter. So, whether $B(\theta^o)$ is empty or not is the structural condition that decides the performance of the algorithm.

We make the following assumption.

Assumption (Unique best action) For all $\theta \in \Theta$, the optimal action, $a^*(\theta)$, is unique.

We note that this is a standard assumption in the literature. This assumption can be removed at the expense of more notations. We define Δ_i as,

$$\Delta_i := \mu_{a^*(\theta^o)}(\theta^o) - \mu_i(\theta^o), \quad (3)$$

which is the difference between the mean value of the optimal arm and the mean value of arm i for the true parameter θ^o . This is the standard optimality gap notion used in the MAB literature [5]. Without loss of generality assume natural logarithms.

For each arm in $i \in C(\theta^o)$, we define,

$$\beta_i := \min_{\theta: \theta \in B(\theta^o), a^*(\theta)=i} |\mu_i(\theta^o) - \mu_i(\theta)|. \quad (4)$$

We use the following Lemma to compare our result with classical MAB result. The proof for this lemma is given in the appendix.

Lemma 1 Let Δ_i and β_i be as defined in (3) and (4) respectively. Then, for each $i \in C(\theta^o)$, $\beta_i > 0$. Moreover, $\beta_i > \Delta_i$.

We now present the finite time performance guarantee for our FP-UCB algorithm.

Algorithm 1 FP-UCB

-
- 1: Initialization: Select each arm in the set A once
 - 2: Initialize episode number $k = 1$, time step $t = |A| + 1$
 - 3: **while** $t \leq T$ **do**
 - 4: $t_k = t - 1$
 - 5: Compute the set

$$A_k = \left\{ a^*(\theta), \theta \in \Theta : \forall i \in A, |\hat{\mu}_i(t_k) - \mu_i(\theta)| \leq \sqrt{\frac{3 \log(k)}{n_i(t_k)}} \right\}$$

- 6: **if** $|A_k| \neq 0$ **then**
 - 7: Select each arm in the set A_k once
 - 8: $t \leftarrow t + |A_k|$
 - 9: **else**
 - 10: Select each arm in the set A once
 - 11: $t \leftarrow t + |A|$
 - 12: **end if**
 - 13: $k \leftarrow k + 1$
 - 14: **end while**
-




Fig. 1: An illustration of the episodes and time slots of the FP-UCB algorithm.

Theorem 1 Under the FP-UCB algorithm,

$$\begin{aligned} \mathbb{E}[R(T)] &\leq D_1, \text{ if } B(\theta^o) \text{ empty, and} \\ \mathbb{E}[R(T)] &\leq D_2 + 12 \log(T) \sum_{i \in C(\theta^o)} \frac{\Delta_i}{\beta_i^2}, \text{ if } B(\theta^o) \text{ non-empty,} \end{aligned} \quad (5)$$

where D_1 and D_2 are problem dependent constants that depend only on the problem parameters $|A|$ and $(\mu_i(\theta), \theta \in \Theta)$, but do not depend on T .

Remark 1 (Comparison with the classical MAB results) Both UCB type algorithms and Thompson Sampling type algorithms give a problem dependent regret bound $\mathcal{O}(\log T)$. More precisely, assuming that the optimal arm is arm 1, the regret of the UCB algorithm, $\mathbb{E}[R_{\text{UCB}}(T)]$, is given by [5]

$$\mathbb{E}[R_{\text{UCB}}(T)] = \mathcal{O} \left(\sum_{i=2}^L \frac{1}{\Delta_i} \log T \right).$$

On the other hand, the FP-UCB algorithm achieves the regret, $\mathbb{E}[R_{\text{FP-UCB}}(T)]$,

$$\mathcal{O}(1), \text{ if } B(\theta^o) \text{ empty, and } \mathcal{O}\left(\sum_{i \in C(\theta^o)} \frac{\Delta_i}{\beta_i^2} \log T\right), \text{ if } B(\theta^o) \text{ non-empty.}$$

Clearly, for some MAB problems, FP-UCB algorithm achieves a bounded regret ($\mathcal{O}(1)$) as opposed to the increasing regret ($\mathcal{O}(\log T)$) of the standard UCB algorithm. Even in the cases where FP-UCB algorithm incurs an increasing regret ($\mathcal{O}(\log T)$), the preceding constant (Δ_i/β_i^2) is smaller than the preceding constant ($1/\Delta_i$) of the standard UCB algorithm because $\beta_i > \Delta_i$.

We now give the asymptotic lower bound for the finitely parameterized multi-armed bandits problem from [1], for comparing the performance of our FP-UCB algorithm.

Theorem 2 (Lower bound [1])

For any uniformly good control scheme under the parameter θ^o ,

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}[R(T)]}{\log(T)} \geq \min_{h \in H} \max_{\theta \in B(\theta^o)} \frac{\sum_{u \in A \setminus \{a^*(\theta^o)\}} h_u(\mu_{a^*(\theta^o)}(\theta^o) - \mu_u(\theta^o))}{\sum_{u \in A \setminus \{a^*(\theta^o)\}} h_u D_u(\theta^o \| \theta)}.$$

where H is a probability simplex with $|A| - 1$ vertices and, for any $u \in A \setminus \{a^*(\theta^o)\}$, $D_u(\theta^o \| \theta) = \int P_u(x; \theta^o) \log(P_u(x; \theta^o)/P_u(x; \theta)) dx$ is the KL-divergence between the probability distributions $P_u(\cdot; \theta^o)$ and $P_u(\cdot; \theta)$.

Remark 2 (Optimality of the FP-UCB algorithm) From Theorem 2, the achievable regret of any multi-armed bandits learning algorithm is lower bounded by $\Omega(1)$ when $B(\theta^o)$ is empty, and $\Omega(\log T)$ when $B(\theta^o)$ is non-empty. Our FP-UCB algorithm achieves these bounds and hence achieves the order optimal performance.

4 Analysis of the FP-UCB Algorithm

In this section, we give the proof of Theorem 1. For reducing the notation, without loss of generality we assume that the true optimal arm is arm 1, i.e., $a^* = a^*(\theta^o) = 1$. We will also denote $\mu_j(\theta^o)$ as μ_j^o , for any $j \in A$.

Now, we can rewrite the expected regret from (1) as

$$\begin{aligned} \mathbb{E}[R(T)] &= \mathbb{E}\left[\sum_{t=1}^T (\mu_1^o - \mu_{a(t)}^o)\right] \\ &= \sum_{i=2}^L \Delta_i \mathbb{E}\left[\sum_{t=1}^T \mathbb{1}\{a(t) = i\}\right] = \sum_{i=2}^L \Delta_i \mathbb{E}[n_i(T)]. \end{aligned}$$

Since the algorithm selects arms only from the set A , this can be written as

$$\mathbb{E}[R(T)] = \sum_{i \in A} \Delta_i \mathbb{E}[n_i(T)]. \quad (6)$$

We first prove the following important propositions.

Proposition 1 For all $i \in A \setminus C(\theta^o), i \neq 1$, under FP-UCB algorithm,

$$\mathbb{E}[n_i(T)] \leq C_i, \quad (7)$$

where C_i is a problem dependent constant that does not depend on T .

Proof Consider an arm $i \in A \setminus C(\theta^o), i \neq 1$. Then, by definition, there exists a $\theta \in \Theta$ such that $a^*(\theta) = i$. Fix a θ which satisfies this condition. Define

$$\alpha_1(\theta) := |\mu_1(\theta^o) - \mu_1(\theta)|.$$

It is straightforward to note that when $i \in A \setminus C(\theta^o)$, then the θ which we considered above is not in $B(\theta^o)$. Hence, by definition, $\alpha_1(\theta) > 0$.

For notational convenience, we will denote $\mu_j(\theta)$ simply as μ_j , for any $j \in A$. Notice that the algorithm picks i^{th} arm once in $t \in \{1, \dots, |A|\}$. Define K_T (note that this is a random variable) to be the total number of episodes in time horizon T for the FP-UCB algorithm. It is straightforward that $K_T \leq T$. Now,

$$\begin{aligned} \mathbb{E}[n_i(T)] &= 1 + \mathbb{E} \left[\sum_{t=|A|+1}^T \mathbb{1}\{a(t) = i\} \right] \\ &\stackrel{(a)}{=} 1 + \mathbb{E} \left[\sum_{k=1}^{K_T} (\mathbb{1}\{i \in A_k\} + \mathbb{1}\{A_k = \emptyset\}) \right] \\ &\leq 1 + \sum_{k=1}^T [\mathbb{P}(\{i \in A_k\}) + \mathbb{P}(\{A_k = \emptyset\})] \end{aligned} \quad (8)$$

$$\begin{aligned} &= 1 + \sum_{k=1}^T [\mathbb{P}(\{i \in A_k, 1 \in A_k\}) + \mathbb{P}(\{i \in A_k, 1 \notin A_k\}) + \mathbb{P}(\{A_k = \emptyset\})] \\ &\leq 1 + \sum_{k=1}^T [\mathbb{P}(\{i \in A_k, 1 \in A_k\}) + \mathbb{P}(\{i \in A_k, 1 \notin A_k\}) + \mathbb{P}(\{i \notin A_k, 1 \notin A_k\})] \\ &\leq 1 + \sum_{k=1}^T [\mathbb{P}(\{i \in A_k, 1 \in A_k\}) + \mathbb{P}(\{1 \notin A_k\})]. \end{aligned} \quad (9)$$

Here (a) follows from the algorithm definition.

We will first analyze the second summation term in (9). First observe that, we can write $n_j(t_k) = 1 + \sum_{\tau=1}^{k-1} (\mathbb{1}\{j \in A_\tau\} + \mathbb{1}\{A_\tau = \emptyset\})$ for any $j \in A$ and episode k . Thus, $n_j(t_k)$ lies between 1 and k . Now,

$$\begin{aligned}
 \sum_{k=1}^T \mathbb{P}(\{1 \notin A_k\}) &\stackrel{(a)}{=} \sum_{k=1}^T \mathbb{P} \left(\bigcup_{j \in A} \left\{ |\hat{\mu}_j(t_k) - \mu_j^o| > \sqrt{\frac{3 \log k}{n_j(t_k)}} \right\} \right) \\
 &\stackrel{(b)}{\leq} \sum_{k=1}^T \sum_{j \in A} \mathbb{P} \left(|\hat{\mu}_j(t_k) - \mu_j^o| > \sqrt{\frac{3 \log k}{n_j(t_k)}} \right) \\
 &\stackrel{(c)}{=} \sum_{k=1}^T \sum_{j \in A} \mathbb{P} \left(\left| \frac{1}{n_j(t_k)} \sum_{\tau=1}^{n_j(t_k)} X_j(\tau) - \mu_j^o \right| > \sqrt{\frac{3 \log k}{n_j(t_k)}} \right) \\
 &\stackrel{(d)}{\leq} \sum_{k=1}^T \sum_{j \in A} \sum_{m=1}^k \mathbb{P} \left(\left| \frac{1}{m} \sum_{\tau=1}^m X_j(\tau) - \mu_j^o \right| > \sqrt{\frac{3 \log k}{m}} \right) \\
 &\stackrel{(e)}{\leq} \sum_{k=1}^T \sum_{j \in A} \sum_{m=1}^k 2 \exp \left(-2m \frac{3 \log k}{m} \right) = \sum_{k=1}^T \sum_{j \in A} 2k^{-5} \leq 4|A|. \quad (10)
 \end{aligned}$$

Here (a) follows from algorithm definition, (b) from the union bound, and (c) from the definition in (2). Inequality (d) follows by conditioning the random variable $n_j(t_k)$ that lies between 1 and k for any $j \in A$ and episode k . Inequality (e) follows from Hoeffding’s inequality [27, Theorem 2.2.6].

For analyzing the first summation term in (9), define the event $E_k := \{n_1(t_k) < 12 \log k / \alpha_1^2(\theta)\}$. Denote the complement of this event as E_k^c . Now the first summation term in (9) can be written as

$$\begin{aligned}
 \sum_{k=1}^T \mathbb{P}(\{i \in A_k, 1 \in A_k\}) &= \underbrace{\sum_{k=1}^T \mathbb{P}(\{i \in A_k, 1 \in A_k, E_k^c\})}_{= Term_1} + \underbrace{\sum_{k=1}^T \mathbb{P}(\{i \in A_k, 1 \in A_k, E_k\})}_{= Term_2}. \quad (11)
 \end{aligned}$$

Analyzing $Term_1$ in (11), we get,

$$\begin{aligned}
 &\mathbb{P}(\{i \in A_k, 1 \in A_k, E_k^c\}) \\
 &= \mathbb{P} \left(\bigcap_{j \in A} \{|\hat{\mu}_j(t_k) - \mu_j^o| < \sqrt{\frac{3 \log k}{n_j(t_k)}}\} \cap \bigcap_{j \in A} \{|\hat{\mu}_j(t_k) - \mu_j| < \sqrt{\frac{3 \log k}{n_j(t_k)}}\} \cap E_k^c \right) \\
 &\leq \mathbb{P} \left(\{|\hat{\mu}_1(t_k) - \mu_1^o| < \sqrt{\frac{3 \log k}{n_1(t_k)}}\}, \{|\hat{\mu}_1(t_k) - \mu_1| < \sqrt{\frac{3 \log k}{n_1(t_k)}}\}, E_k^c \right) = 0. \quad (12)
 \end{aligned}$$

This is because the events $\{|\hat{\mu}_1(t_k) - \mu_1^o| < \sqrt{\frac{3 \log k}{n_1(t_k)}}\}$ and $\{|\hat{\mu}_1(t_k) - \mu_1| < \sqrt{\frac{3 \log k}{n_1(t_k)}}\}$ are disjoint under E_k^c , that is, when $n_1(t_k) \geq 12 \log(k) / \alpha_1^2(\theta)$. To see this, notice that

$$\left\{ |\hat{\mu}_1(t_k) - \mu_1^o| < \sqrt{\frac{3 \log k}{n_1(t_k)}} \right\} \subseteq \left\{ |\hat{\mu}_1(t_k) - \mu_1^o| < \frac{\alpha_1(\theta)}{2} \right\},$$

$$\left\{ |\hat{\mu}_1(t_k) - \mu_1| < \sqrt{\frac{3 \log k}{n_1(t_k)}} \right\} \subseteq \left\{ |\hat{\mu}_1(t_k) - \mu_1| < \frac{\alpha_1(\theta)}{2} \right\},$$

for $n_1(t_k) \geq 12 \log k / \alpha_1^2(\theta)$. Moreover, since $|\mu_1^o - \mu_1| = \alpha_1(\theta)$, $\{|\hat{\mu}_1(t_k) - \mu_1^o| < \alpha_1(\theta)/2\}$ and $\{|\hat{\mu}_1(t_k) - \mu_1| < \alpha_1(\theta)/2\}$ are disjoint sets. Hence, their subsets are also disjoint.

For analyzing $Term_2$ in (11), we start by setting up few notations. Define $n'_1(t_k) := 1 + \sum_{\tau=1}^{k-1} \mathbb{1}\{1 \in A_\tau\}$. Note that, according to the FP-UCB algorithm, arm 1 can be selected if A_τ is empty as well, so $n'_1(t_k) \leq n_1(t_k)$. Define $k_i(\theta)$ and $m(k)$ as,

$$k_i(\theta) := \min \{k : k \geq 3, k > \lceil 12 \log(k) / \alpha_1^2(\theta) \rceil\}, \quad (13)$$

$$m(k) := \max \{1, k - \lceil 12 \log(k) / \alpha_1^2(\theta) \rceil\}. \quad (14)$$

Note that $k_i(\theta)$ is a problem dependent constant and does not depend on T . Also, $m(k) = k - \lceil 12 \log(k) / \alpha_1^2(\theta) \rceil$ for all $k \geq k_i(\theta)$. We claim that for all $k \geq k_i(\theta)$,

$$\{n'_1(t_k) < 12 \log(k) / \alpha_1^2(\theta)\} \subseteq \{1 \notin A_\tau, \text{ for some } \tau, m(k) \leq \tau \leq k-1\}. \quad (15)$$

To see this, suppose there exists no τ , $m(k) \leq \tau \leq k-1$, such that $1 \notin A_\tau$. Then, $1 \in A_\tau$ for all τ , where $m(k) \leq \tau \leq k-1$. So, by definition $n'_1(t_k) \geq (k - m(k)) = \lceil 12 \log(k) / \alpha_1^2(\theta) \rceil$ for $k \geq k_i(\theta)$. So, the complement of the RHS of (15) is a subset of the complement of the LHS of (15). Hence the claim follows.

Now,

$$\begin{aligned} & \sum_{k=1}^T \mathbb{P}(\{i \in A_k, 1 \in A_k, E_k\}) \leq \sum_{k=1}^T \mathbb{P}(E_k) \\ & \stackrel{(a)}{\leq} \sum_{k=1}^T \mathbb{P}(n'_1(t_k) < 12 \log(k) / \alpha_1^2(\theta)) \\ & \stackrel{(b)}{\leq} k_i(\theta) + \sum_{k=k_i(\theta)}^T \mathbb{P}(n'_1(t_k) < 12 \log(k) / \alpha_1^2(\theta)) \\ & \stackrel{(c)}{\leq} k_i(\theta) + \sum_{k=k_i(\theta)}^T \mathbb{P}(\{1 \notin A_\tau, \text{ for some } \tau, m(k) \leq \tau \leq k-1\}) \\ & \stackrel{(d)}{=} k_i(\theta) + \sum_{k=k_i(\theta)}^T \mathbb{P}\left(\bigcup_{\tau=m(k)}^{k-1} \bigcup_{j \in A} |\hat{\mu}_j(\tau) - \mu_j^o| > \sqrt{\frac{3 \log \tau}{n_j(t_\tau)}}\right) \end{aligned}$$

$$\begin{aligned} &\leq k_i(\theta) + \sum_{k=k_i(\theta)}^T \sum_{\tau=m(k)}^{k-1} \sum_{j \in A} \mathbb{P} \left(|\hat{\mu}_j(\tau) - \mu_j^o| > \sqrt{\frac{3 \log \tau}{n_j(t_\tau)}} \right) \\ &\stackrel{(e)}{\leq} k_i(\theta) + \sum_{k=k_i(\theta)}^T \sum_{\tau=m(k)}^{k-1} \frac{2|A|}{\tau^5} \end{aligned} \tag{16}$$

$$\begin{aligned} &\leq k_i(\theta) + \sum_{k=k_i(\theta)}^T \frac{2|A|k}{(m(k))^5} \\ &= k_i(\theta) + \sum_{k=k_i(\theta)}^T \frac{2|A|k}{\left(k - \left\lceil \frac{12 \log(k)}{\alpha_1^2(\theta)} \right\rceil\right)^5} \stackrel{(f)}{=} k_i(\theta) + K_i(\theta), \end{aligned} \tag{17}$$

where $K_i(\theta)$ is a problem dependent constant that does not depend on T .

In the above analysis, (a) follows from the definition of E_k and the observation that $n'_1(t_k) \leq n_1(t_k)$. Considering T to be greater than or equal to $k_i(\theta)|A|$, equality (b) follows; note that this is an artifact of the proof technique and does not affect the theorem statement since $\mathbb{E}[n_i(T')]$, for any T' less than $k_i(\theta)|A|$, can be trivially upper bounded by $\mathbb{E}[n_i(T)]$. Inequality (c) follows from (15), (d) by the FP-UCB algorithm, (e) is similar to the analysis in (10), and (f) follows from the fact that $k > \lceil 12 \log(k) / \alpha_1^2(\theta) \rceil$ for all $k \geq k_i(\theta)$.

Now, using (17) and (12) in (11), we get,

$$\sum_{k=1}^T \mathbb{P}(\{i \in A_k, 1 \in A_k\}) \leq k_i(\theta) + K_i(\theta). \tag{18}$$

Using (18) and (10) in (9), we get,

$$\mathbb{E}[n_i(T)] \leq C_i,$$

where $C_i = 1 + 4|A| + \min_{\theta: a^*(\theta)=i} (k_i(\theta) + K_i(\theta))$, which is a problem dependent constant that does not depend on T . This concludes the proof. \square

Proposition 2 For any $i \in C(\theta^o)$, under the FP-UCB algorithm,

$$\mathbb{E}[n_i(T)] \leq 2 + 4|A| + \frac{12 \log(T)}{\beta_i^2}. \tag{19}$$

Proof Fix an $i \in C(\theta^o)$. Then there exists $\theta \in B(\theta^o)$ such that $a^*(\theta) = i$. Fix θ which satisfies this condition. Define the event $F(t) := \{n_i(t-1) < 12 \log T / \beta_i^2\}$. Now,

$$\begin{aligned}
\mathbb{E}[n_i(T)] &= 1 + \mathbb{E} \left[\sum_{t=|A|+1}^T \mathbb{1}\{a(t) = i\} \right] \\
&= 1 + \mathbb{E} \left[\sum_{t=|A|+1}^T \mathbb{1}\{a(t) = i, F(t)\} \right] + \mathbb{E} \left[\sum_{t=|A|+1}^T \mathbb{1}\{a(t) = i, F^c(t)\} \right].
\end{aligned} \tag{20}$$

Analyzing the first summation term in (20) we get,

$$\begin{aligned}
\mathbb{E} \left[\sum_{t=|A|+1}^T \mathbb{1}\{a(t) = i, F(t)\} \right] &= \mathbb{E} \left[\sum_{t=|A|+1}^T \mathbb{1}\{a(t) = i\} \mathbb{1}\{n_i(t-1) < 12 \log T / \beta_i^2\} \right] \\
&\leq 1 + 12 \log T / \beta_i^2.
\end{aligned} \tag{21}$$

We use the same decomposition as in the proof of Proposition 1 for the second summation term in (20). Thus we get,

$$\begin{aligned}
\mathbb{E} \left[\sum_{t=|A|+1}^T \mathbb{1}\{a(t) = i, F^c(t)\} \right] &= \\
&\mathbb{E} \left[\sum_{k=1}^{K_T} \mathbb{1}\{i \in A_k, F^c(t_k + 1)\} + \mathbb{1}\{A_k = \emptyset, F^c(t_k + 1)\} \right] \\
&\leq \sum_{k=1}^T \mathbb{P}(\{i \in A_k, 1 \in A_k, F^c(t_k + 1)\})
\end{aligned} \tag{22}$$

$$+ \sum_{k=1}^T \mathbb{P}(\{1 \notin A_k, F^c(t_k + 1)\}), \tag{23}$$

following the analysis in (9). First, consider (23). From the analysis in (10) we have

$$\sum_{k=1}^T \mathbb{P}(\{1 \notin A_k, F^c(t_k + 1)\}) \leq \sum_{k=1}^T \mathbb{P}(\{1 \notin A_k\}) \leq 4|A|. \tag{24}$$

For any $i \in A$ and episode k under event $F^c(t_k + 1)$, we have

$$n_i(t_k) \geq \frac{12 \log T}{\beta_i^2} \geq \frac{12 \log t_k}{\beta_i^2} \geq \frac{12 \log k}{\beta_i^2}$$

since t_k satisfies $k \leq t_k \leq T$. From (4), it further follows that

$$\sqrt{\frac{3 \log k}{n_j(t_k)}} \leq \frac{\beta_i}{2} \leq \frac{|\mu_i(\theta^o) - \mu_i(\theta)|}{2}.$$

So, following the analysis in (12) for (22), we get

$$\begin{aligned}
 & \mathbb{P}(\{i \in A_k, 1 \in A_k, F^c(t_k + 1)\}) \\
 &= \mathbb{P} \left(\begin{aligned} & \bigcap_{j \in A} \{|\hat{\mu}_j(t_k) - \mu_j(\theta^o)| < \sqrt{\frac{3 \log k}{n_j(t_k)}}\}, \\ & \bigcap_{j \in A} \{|\hat{\mu}_j(t_k) - \mu_j(\theta)\} < \sqrt{\frac{3 \log k}{n_j(t_k)}}\}, F^c(t_k + 1) \end{aligned} \right) \\
 &\leq \mathbb{P} \left(\begin{aligned} & \{|\hat{\mu}_i(t_k) - \mu_i(\theta^o)| < \sqrt{\frac{3 \log k}{n_i(t_k)}}\}, \\ & \{|\hat{\mu}_i(t_k) - \mu_i(\theta)\} < \sqrt{\frac{3 \log k}{n_i(t_k)}}\}, F^c(t_k + 1) \end{aligned} \right) = 0. \quad (25)
 \end{aligned}$$

Using equations (21), (24), and (25) in (20), we get

$$\mathbb{E}[n_i(T)] \leq 2 + 4|A| + \frac{12 \log(T)}{\beta_i^2}.$$

This completes the proof. □

We now give the proof of our main theorem.

Proof (of Theorem 1)

From (6),

$$\mathbb{E}[R(T)] = \sum_{i \in A} \Delta_i \mathbb{E}[n_i(T)] = \sum_{i \in A \setminus C(\theta^o)} \Delta_i \mathbb{E}[n_i(T)] + \sum_{i \in C(\theta^o)} \Delta_i \mathbb{E}[n_i(T)]. \quad (26)$$

Whenever $B(\theta^o)$ is empty, notice that $C(\theta^o)$ is empty. So, using Proposition 1, (26) becomes

$$\mathbb{E}[R(T)] = \sum_{i \in A} \Delta_i \mathbb{E}[n_i(T)] \leq \sum_{i \in A} \Delta_i C_i \leq |A| \max_{i \in A} \Delta_i C_i.$$

Whenever $B(\theta^o)$ is non-empty, $C(\theta^o)$ is non-empty. Analyzing (26), we get,

$$\begin{aligned}
 \mathbb{E}[R(T)] &= \sum_{i \in A \setminus C(\theta^o)} \Delta_i \mathbb{E}[n_i(T)] + \sum_{i \in C(\theta^o)} \Delta_i \mathbb{E}[n_i(T)] \\
 &\stackrel{(a)}{\leq} \sum_{i \in A \setminus C(\theta^o)} \Delta_i C_i + \sum_{i \in C(\theta^o)} \Delta_i \mathbb{E}[n_i(T)] \\
 &\stackrel{(b)}{\leq} \sum_{i \in A \setminus C(\theta^o)} \Delta_i C_i + \sum_{i \in C(\theta^o)} \Delta_i \left(2 + 4|A| + \frac{12 \log(T)}{\beta_i^2} \right) \\
 &\leq |A| \max_{i \in A} \Delta_i (2 + C_i + 4|A|) + 12 \log(T) \sum_{i \in C(\theta^o)} \frac{\Delta_i}{\beta_i^2}.
 \end{aligned}$$

Here (a) follows from Proposition 1 and (b) from Proposition 2. Setting

$$D_1 := |A| \max_{i \in A} \Delta_i C_i \text{ and } D_2 := |A| \max_{i \in A} \Delta_i (2 + C_i + 4|A|) \quad (27)$$

proves the regret bounds in (5) of the theorem. \square

We now provide the following lemma to characterize the problem dependent constants C_i given in Proposition 1. The proof for this lemma is given in the appendix.

Lemma 2 *Under the hypotheses in Proposition 1, we have*

$$C_i \leq 1 + 4|A| + \min_{\theta: a^*(\theta)=i} (2E_i(\theta)(E_i(\theta) + 1)|A| + 4|A|\alpha_1^{10}(\theta)),$$

where $E_i(\theta) = \max\{3, \lceil 144/\alpha_1^4(\theta) \rceil\}$ and $\alpha_1(\theta) = |\mu_1(\theta^o) - \mu_1(\theta)|$.

Now, using the above lemma with (27), we have a characterization of the problem dependent constants in Theorem 1.

5 Simulations

In this section, we present detailed numerical simulation to illustrate the performance of FP-UCB algorithm compared to the other standard multi-armed bandits algorithms.

We first consider a simple setting to illustrate intuition behind FP-UCB algorithm. Consider $\Theta = \{\theta^1, \theta^2\}$ with $[\mu_1(\theta^1), \mu_2(\theta^1)] = [0.9, 0.5]$ and $[\mu_1(\theta^2), \mu_2(\theta^2)] = [0.2, 0.5]$. Consider the reward distributions $P_i, i = 1, 2$ to be Bernoulli. Clearly, $a^*(\theta^1) = 1$ and $a^*(\theta^2) = 2$.

Suppose the true parameter is θ^1 , i.e., $\theta^o = \theta^1$. Then, it is easy to note that, in this case $B(\theta^o)$ is empty, and hence $C(\theta^o)$ is empty. So, according to Theorem 1, FP-UCB will achieve an $\mathcal{O}(1)$ regret. The performance of the algorithm for this setting is shown in Fig. 2. Indeed, the regret doesn't increase after some time steps, which shows the bounded regret property. We note that in all the figures, the regret is averaged over 10 runs, with the thick line showing the average regret and the band around shows the ± 1 standard deviation.

Now, suppose the true parameter is θ^2 , i.e., $\theta^o = \theta^2$. In this case $B(\theta^o)$ is non-empty. In fact, $B(\theta^o) = \theta^1$ and $C(\theta^o) = 1$. So, according to Theorem 1, FP-UCB will achieve an $\mathcal{O}(\log T)$ regret. The performance of the algorithm shown in Fig. 3 suggests the same. Fig. 4 plots the regret scaled by $\log t$, and the curve converges to a constant value, confirming the $\mathcal{O}(\log T)$ regret performance.

We consider a problem with 4 arms where the mean values for the arms (corresponding to the true parameter θ^o) are $\mu(\theta^o) = [0.6, 0.4, 0.3, 0.2]$. Consider the parameter set Θ such that $\mu(\theta)$ for any θ is a permutation of $\mu(\theta^o)$. Note that the cardinality of the parameter set, $|\Theta| = 24$, in this case. It is straightforward to show that $B(\theta^o)$ is empty for this case. We compare the performance of FP-UCB algorithm for this case with two standard multi-armed bandits algorithms. Fig. 5 shows the performance of standard UCB algorithm and that of FP-UCB algorithm. Fig. 6 compares the performance of standard Thompson sampling algorithm with that of FP-UCB algorithm. The standard bandits algorithm incurs an increasing regret,

while FP-UCB achieves a bounded regret. For $\mu(\theta') = [0.4, 0.6, 0.3, 0.2]$, we have $a^*(\theta') = 2$. Now we give a typical value for the $k_2(\theta')$, defined in (13), used in the proof. For this θ' we have $k_2(\theta') = \min \{k : k \geq 3, k > \lceil 12 \log(k) / \alpha_1^2(\theta') \rceil\} = \min \{k : k \geq 3, k > \lceil 12 \log(k) / 0.2^2 \rceil\} = 2326$ since $\alpha_1(\theta') = 0.2$. When the reward distributions are not necessarily Bernoulli, note that $k_i(\theta)$ is 3 for any θ with $a^*(\theta) = i$ satisfying $\alpha_1(\theta) > 2\sqrt{3/e}$.

As before assume that $\mu(\theta^o) = [0.6, 0.4, 0.3, 0.2]$. But consider a larger parameter set Θ such that for any $\theta \in \Theta$, $\mu(\theta) \in \{0.6, 0.4, 0.3, 0.2\}^4$. Note that, due to repetitions in the mean rewards for the arms, definition of $a^*(\theta)$ needs to be updated, and the algorithmic way is to pick the minimum arm index out of which are having the same mean rewards. For example, consider $\mu(\theta) = [0.5, 0.6, 0.6, 0.2]$, and so as per our new definition, $a^*(\theta) = 2$. Even in this scenario, we have $B(\theta^o)$ to be empty. Thus, FP-UCB achieves an $\mathcal{O}(1)$ regret rather than $\mathcal{O}(\log(T))$ as opposed to standard UCB algorithm and Thompson sampling algorithm.

We now consider a case where FP-UCB incurs an increasing regret. We again consider a problem with 4 arms where the mean values for the arms are $\mu(\theta^o) = [0.4, 0.3, 0.2, 0.2]$. But consider a larger parameter set Θ such that for any $\theta \in \Theta$, $\mu(\theta) \in \{0.6, 0.4, 0.3, 0.2\}^4$. Note that the cardinality of Θ , $|\Theta| = 4^4$ in this case. It is easy to observe that $B(\theta^o)$ is non-empty, for instance θ with mean arm values $[0.4, 0.6, 0.3, 0.2]$ is in $B(\theta^o)$. Fig. 7 compares the performance of standard UCB and FP-UCB algorithms for this case. We see FP-UCB incurring $\mathcal{O}(\log(T))$ regret here. Also note that the performance of the FP-UCB in this case also is superior to the standard UCB algorithm.




Fig. 2




Fig. 3




Fig. 4




Fig. 5




Fig. 6




Fig. 7

6 Conclusion and Future Work

We proposed an algorithm for finitely parameterized multi-armed bandits. Our FP-UCB algorithm achieves bounded regret if the parameter set satisfies some necessary condition and logarithmic regret in other cases. In both cases, the theoretical performance guarantees for our algorithm are superior to the standard UCB algorithm for multi-armed bandits. Our algorithm also shows superior numerical performance.

In the future, we will extend this approach to linear bandits and contextual bandits. Reinforcement learning problems where the underlying MDP is finitely parameterized is another research direction we plan to explore. We will also develop similar algorithms using Thompson sampling approaches.

References

1. Rajeev Agrawal, Demosthenis Teneketzis, and Venkatachalam Anantharam. Asymptotically efficient adaptive allocation schemes for controlled iid processes: Finite parameter space. *IEEE Transactions on Automatic Control*, 34(3):258–267, 1989.
2. Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *Proceedings of the 25th Annual Conference on Learning Theory*, volume 23, pages 39.1–39.26. PMLR, 2012.
3. Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, pages 127–135, 2013.
4. Onur Atan, Cem Tekin, and Mihaela Schaar. Global multi-armed bandits with h older continuity. In *Artificial Intelligence and Statistics*, pages 28–36, 2015.
5. Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
6. Hamsa Bastani, Mohsen Bayati, and Khashayar Khosravi. Mostly exploration-free algorithms for contextual bandits. *arXiv:1704.09011*, 2017.
7. S bastien Bubeck and Che-Yu Liu. Prior-free and prior-dependent regret bounds for thompson sampling. In *Advances in Neural Information Processing Systems*, pages 638–646, 2013.
8. S bastien Bubeck, Vianney Perchet, and Philippe Rigollet. Bounded regret in stochastic multi-armed bandits. In *Conference on Learning Theory*, pages 122–134, 2013.
9. St phane Caron, Branislav Kveton, Marc Lelarge, and Smriti Bhagat. Leveraging side observations in stochastic bandits. *Conference on Uncertainty in Artificial Intelligence*, 2012.
10. Nicolo Cesa-Bianchi and G bor Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422, 2012.
11. Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 208–214, 2011.
12. Richard Combes, Stefan Magureanu, and Alexandre Proutiere. Minimal exploration in structured stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 1763–1771, 2017.
13. Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. In *COLT*, 2008.
14. Aditya Gopalan, Shie Mannor, and Yishay Mansour. Thompson sampling for complex online problems. In *International Conference on Machine Learning*, pages 100–108, 2014.
15. Dileep Kalathil, Naumaan Nayyar, and Rahul Jain. Decentralized learning for multiplayer multiarmed bandits. *IEEE Transactions on Information Theory*, 60(4):2331–2345, 2014.

16. Dileep Kalathil and Ram Rajagopal. Online learning for demand response. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 218–222. IEEE, 2015.
17. Junpei Komiyama, Junya Honda, and Hiroshi Nakagawa. Optimal regret analysis of thompson sampling in stochastic multi-armed bandit problem with multiple plays. In *International Conference on Machine Learning*, pages 1152–1161, 2015.
18. Panqanamala Ramana Kumar and Pravin Varaiya. *Stochastic systems: Estimation, identification, and adaptive control*, volume 75. SIAM, 2015.
19. Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
20. Tor Lattimore and Rémi Munos. Bounded regret for finite-armed structured bandits. In *Advances in Neural Information Processing Systems*, pages 550–558, 2014.
21. Lennart Ljung. *System Identification: Theory for the User*. Prentice Hall, New Jersey, 1998.
22. Odalric-Ambrym Maillard and Shie Mannor. Latent bandits. In *International Conference on Machine Learning*, pages 136–144, 2014.
23. Ian Osband, Daniel Russo, and Benjamin Van Roy. (more) efficient reinforcement learning via posterior sampling. In *Advances in Neural Information Processing Systems*, pages 3003–3011, 2013.
24. Cem Tekin and Mingyan Liu. Approximately optimal adaptive learning in opportunistic spectrum access. In *2012 Proceedings IEEE INFOCOM*, pages 1548–1556. IEEE, 2012.
25. William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
26. Sattar Vakili and Qing Zhao. Achieving complete learning in multi-armed bandit problems. In *2013 Asilomar Conference on Signals, Systems and Computers*, pages 1778–1782. IEEE, 2013.
27. Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018.
28. Chih-Chun Wang, Sanjeev R Kulkarni, and H Vincent Poor. Bandit problems with side observations. *IEEE Transactions on Automatic Control*, 50(3):338–355, 2005.
29. Shi Zong, Hao Ni, Kenny Sung, Nan Rosemary Ke, Zheng Wen, and Branislav Kveton. Cascading bandits for large-scale recommendation problems. In *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence*, pages 835–844. AUAI Press, 2016.

Appendix

6.1 Proof of Lemma 1

Proof Fix an $i \in C(\theta^o)$. Then there exists a $\theta \in B(\theta^o)$ such that $a^*(\theta) = i$. For this θ , by the definition of $B(\theta^o)$, we have

$$\mu_1(\theta^o) = \mu_1(\theta). \quad (28)$$

Using Assumption 1, it follows that

$$\mu_i(\theta) = \mu_{a^*(\theta)}(\theta) > \mu_1(\theta) = \mu_1(\theta^o) = \mu_{a^*(\theta^o)}(\theta^o) > \mu_i(\theta^o).$$

Thus, $\beta_i = \min_{\theta: \theta \in B(\theta^o), a^*(\theta) = i} |\mu_i(\theta^o) - \mu_i(\theta)| > 0$.

Now, for any given θ considered above, suppose $|\mu_i(\theta) - \mu_i(\theta^o)| \leq \Delta_i$. Since $\Delta_i > 0$ by definition, this implies that

$$\mu_{a^*(\theta)}(\theta) = \mu_i(\theta) \leq \Delta_i + \mu_i(\theta^o) \stackrel{(a)}{=} \mu_1(\theta^o) - \mu_i(\theta^o) + \mu_i(\theta^o) = \mu_1(\theta^o) \stackrel{(b)}{=} \mu_1(\theta),$$

where (a) follows from definition of Δ_i and (b) from (28). This is a contradiction because $\mu_{a^*(\theta)}(\theta) > \mu_1(\theta)$.

Thus, $|\mu_i(\theta) - \mu_i(\theta^o)| > \Delta_i$ for any $\theta \in B(\theta^o)$ such that $a^*(\theta) = i$. So, $\beta_i > \Delta_i$. \square

6.2 Proof of Lemma 2

Proof We have $C_i = 1 + 4|A| + \min_{\theta: a^*(\theta)=i} (k_i(\theta) + K_i(\theta))$.

First recall that $k_i(\theta) := \min\{k : k \geq 3, k > \lceil 12 \log(k) / \alpha_1^2(\theta) \rceil\}$. Since $\log(x) \leq (x-1)/\sqrt{x}$ for all $1 \leq x < \infty$, we have

$$\left\{ k : k \geq 3, k > \frac{12(k-1)}{\alpha_1^2(\theta)\sqrt{k}} + 1 \right\} \subseteq \{k : k \geq 3, k > \lceil 12 \log(k) / \alpha_1^2(\theta) \rceil\}.$$

The LHS of the above equation simplifies to $\{k : k \geq 3, k > 144/\alpha_1^4(\theta)\}$. Thus, we have $k_i(\theta) \leq \max\{3, \lceil 144/\alpha_1^4(\theta) \rceil\}$.

Now, recall that $K_i(\theta)$ is defined as

$$\begin{aligned} K_i(\theta) &= \sum_{k=k_i(\theta)}^T \frac{2|A|k}{(k - \lceil \frac{12 \log(k)}{\alpha_1^2(\theta)} \rceil)^5} \\ &\leq \sum_{k=k_i(\theta)}^{\infty} \frac{2|A|k}{(k - \lceil \frac{12 \log(k)}{\alpha_1^2(\theta)} \rceil)^5} \\ &= \sum_{k=k_i(\theta)}^{E_i(\theta)} \frac{2|A|k}{(k - \lceil \frac{12 \log(k)}{\alpha_1^2(\theta)} \rceil)^5} + \sum_{k=E_i(\theta)+1}^{\infty} \frac{2|A|k}{(k - \lceil \frac{12 \log(k)}{\alpha_1^2(\theta)} \rceil)^5}. \end{aligned} \quad (29)$$

We analyze the first summation in (29). Thus, we get,

$$\sum_{k=k_i(\theta)}^{E_i(\theta)} \frac{2|A|k}{(k - \lceil \frac{12 \log(k)}{\alpha_1^2(\theta)} \rceil)^5} \leq \sum_{k=k_i(\theta)}^{E_i(\theta)} 2|A|k \leq \sum_{k=1}^{E_i(\theta)} 2|A|k = E_i(\theta)(E_i(\theta) + 1)|A|. \quad (30)$$

Since $\log(x) \leq (x-1)/\sqrt{x}$ for all $1 \leq x < \infty$, we have

$$k - \left\lceil \frac{12 \log(k)}{\alpha_1^2(\theta)} \right\rceil \geq k - \frac{12 \log(k)}{\alpha_1^2(\theta)} - 1 \geq \frac{(k-1)(\alpha_1^2(\theta)\sqrt{k} - 12)}{\alpha_1^2(\theta)\sqrt{k}}.$$

Using this, the second summation in (29) can be bounded as

$$\begin{aligned}
 \sum_{k=E_i(\theta)+1}^{\infty} \frac{2|A|k}{\left(k - \left\lceil \frac{12 \log(k)}{\alpha_1^2(\theta)} \right\rceil\right)^5} &\leq \sum_{k=E_i(\theta)+1}^{\infty} \frac{2|A|k^{7/2}\alpha_1^{10}(\theta)}{\left((k-1)(\alpha_1^2(\theta)\sqrt{k}-12)\right)^5} \\
 &\stackrel{(a)}{\leq} \sum_{k=E_i(\theta)+1}^{\infty} \frac{2|A|k^{7/2}\alpha_1^{10}(\theta)}{(k-1)^5} \\
 &\leq 2|A|\alpha_1^{10}(\theta) \sum_{k=4}^{\infty} \frac{k^{7/2}}{(k-1)^5} \stackrel{(b)}{\leq} 4|A|\alpha_1^{10}(\theta) \quad (31)
 \end{aligned}$$

where (a) follows from the observation that $(\alpha_1^2(\theta)\sqrt{k}-12) > 1$ for $k \geq E_i(\theta) + 1$ and (b) follows from calculus (an integral bound).

Thus using (30) and (31) in (29), we get $K_i(\theta) \leq E_i(\theta)(E_i(\theta) + 1)|A| + 4|A|\alpha_1^{10}(\theta)$. This concludes the proof of this lemma. \square