

Distributed Learning in Multi-Armed Bandit With Multiple Players

Keqin Liu and Qing Zhao

Abstract—We formulate and study a decentralized multi-armed bandit (MAB) problem. There are M distributed players competing for N independent arms. Each arm, when played, offers i.i.d. reward according to a distribution with an unknown parameter. At each time, each player chooses one arm to play without exchanging observations or any information with other players. Players choosing the same arm collide, and, depending on the collision model, either no one receives reward or the colliding players share the reward in an arbitrary way. We show that the minimum system regret of the decentralized MAB grows with time at the same logarithmic order as in the centralized counterpart where players act collectively as a single entity by exchanging observations and making decisions jointly. A decentralized policy is constructed to achieve this optimal order while ensuring fairness among players and without assuming any pre-agreement or information exchange among players. Based on a time-division fair sharing (TDFS) of the M best arms, the proposed policy is constructed and its order optimality is proven under a general reward model. Furthermore, the basic structure of the TDFS policy can be used with any order-optimal single-player policy to achieve order optimality in the decentralized setting. We also establish a lower bound on the system regret for a general class of decentralized policies, to which the proposed policy belongs. This problem finds potential applications in cognitive radio networks, multi-channel communication systems, multi-agent systems, web search and advertising, and social networks.

Index Terms—Cognitive radio, decentralized multi-armed bandit, distributed learning, multi-agent systems, system regret, Web search and advertising.

I. INTRODUCTION

A. The Classic MAB With a Single Player

IN the classic multi-armed bandit (MAB) problem, there are N independent arms and a single player. Playing arm i ($i = 1, \dots, N$) yields i.i.d. random rewards with a distribution parameterized by an unknown θ_i . At each time, the player chooses one arm to play, aiming to maximize the total expected reward in the long run. Had the reward model of each arm been known to the player, the player would have always chosen the arm that

offers the maximum expected reward. Under an unknown reward model, the essence of the problem lies in the well-known tradeoff between exploitation (aiming at gaining immediate reward by choosing the arm that is suggested to be the best by past observations) and exploration (aiming at learning the unknown reward model of all arms to minimize the mistake of choosing an inferior arm in the future).

Under a non-Bayesian formulation where the unknown parameters $\Theta \triangleq (\theta_1, \dots, \theta_N)$ are considered deterministic, a commonly used performance measure of an arm selection policy is the so-called *regret* or *the cost of learning* defined as the reward loss with respect to the case with known reward models. Since the best arm can never be perfectly identified from finite observations (except certain trivial scenarios), the player can never stop learning and will always make mistakes. Consequently, the regret of any policy grows with time.

An interesting question posed by Lai and Robbins in 1985 [1] is on the minimum rate at which the regret grows with time. They showed in [1] that the minimum regret grows at a logarithmic order under certain regularity conditions. The best leading constant was also obtained, and an optimal policy was constructed under a general reward model to achieve the minimum regret (both the logarithmic order and the best leading constant). In 1987, Anantharam *et al.* extended Lai and Robbins's results to MAB with multiple plays: exactly M ($M < N$) arms can be played simultaneously at each time [2]. They showed that allowing multiple plays changes only the leading constant but not the logarithmic order of the regret. They also extended Lai–Robbins policy to achieve the optimal regret under multiple plays.

B. Decentralized MAB With Distributed Multiple Players

In this paper, we formulate and study a decentralized version of the classic MAB, where we consider M ($M < N$) distributed players. At each time, a player chooses one arm to play based on its *local* observation and decision history. Players do not exchange information on their decisions and observations. Collisions occur when multiple players choose the same arm, and, depending on the collision model, either no one receives reward or the colliding players share the reward in an arbitrary way. The objective is to design distributed policies for each player in order to minimize, under any unknown parameter set Θ , the rate at which the system regret grows with time. Here the system regret is defined as the reward loss with respect to the maximum system reward obtained under a known reward model and with centralized scheduling of the players.

The single-player MAB with multiple plays considered in [2] is equivalent to a *centralized* MAB with multiple players. If all M players can exchange observations and make decisions

Manuscript received December 19, 2009; accepted July 13, 2010. Date of publication August 03, 2010; date of current version October 13, 2010. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Mark Coates. This work was supported by the Army Research Laboratory NS-CTA under Grant W911NF-09-2-0053 and by the National Science Foundation under Grant CCF-0830685. Part of this result was presented at ITA, San Diego, CA, January 2010, and IEEE ICASSP, Dallas, TX, March 2010.

The authors are with the Department of Electrical and Computer Engineering, University of California, Davis, CA, 95616 USA (e-mail: kqliu@ucdavis.edu; qzhao@ucdavis.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2010.2062509

jointly, they act collectively as a single player who has the freedom of choosing M arms simultaneously. As a direct consequence, the logarithmic order of the minimum regret established in [2] provides a lower bound on the minimum regret in a *decentralized* MAB where players cannot exchange observations and must make decisions independently based on their individual local observations.¹

C. Main Results

In this paper, we show that in a decentralized MAB where players can only learn from their individual observations and collisions are bound to happen, the system can achieve the same logarithmic order of the regret as in the centralized case. Furthermore, we show that this optimal order can be achieved under a fairness constraint that requires all players accrue reward at the same rate.

A decentralized policy is constructed to achieve the optimal order under the fairness constraint. The proposed policy is based on a time-division fair sharing (TDFS) of the M best arms where no pre-agreement on the time sharing schedule is needed. It is constructed and its order optimality is proven under a general reward model. The TDFS decentralized policy thus applies to and maintains its order optimality in a wide range of potential applications as detailed in Section I-D. Furthermore, the basic structure of TDFS is not tied with a specific single-player policy. It can be used with any single-player policy to achieve an efficient and fair sharing of the M best arms among M distributed players. More specifically, if the single-player policy achieves the optimal logarithmic order in the centralized setting, then the corresponding TDFS policy achieves the optimal logarithmic order in the decentralized setting. The order optimality of the TDFS policy is also preserved when player's local policies are built upon different order-optimal single-player policies.

We also establish a lower bound on the leading constant of the logarithmic order of the regret for a general class of decentralized policies, to which the proposed TDFS policy belongs. This lower bound is tighter than the trivial bound provided by the centralized MAB considered in [2], which indicates, as one would have expected, that decentralized MAB is likely to incur a larger leading constant than its centralized counterpart.

D. Applications

With its general reward model, the TDFS policy finds a wide area of potential applications. We give a few examples below.

Consider first a cognitive radio network where secondary users independently search for idle channels temporarily unused by primary users [3]. Assume that the primary system adopts a slotted transmission structure and the state (busy/idle) of each channel can be modeled by an i.i.d. Bernoulli process. At the beginning of each slot, multiple distributed secondary users need to decide which channel to sense (and subsequently transmit if the chosen channel is idle) without knowing the channel occupancy statistics (i.e., the mean θ_i of the Bernoulli process). An idle channel offers one unit of reward. When multiple secondary users choose the same channel, none or

only one receives reward depending on whether carrier sensing is implemented.

Another potential application is opportunistic transmission over wireless fading channels [4]. In each slot, each user senses the fading realization of a selected channel and chooses its transmission power or data rate accordingly. The reward can model energy efficiency (for fixed-rate transmission) or throughput. The objective is to design distributed channel selection policies under unknown fading statistics.

Consider next a multi-agent system, where each agent is assigned to collect targets among multiple locations (for example, ore mining). When multiple agents choose the same location, they share the reward according to a certain rule. The log-Gaussian distribution with an unknown mean may be used as the reward model when the target is fish [5] or ore [6].

Another potential application is Internet advertising where multiple competing products select websites to post advertisements. The reward obtained from the selected website is measured by the number of interested viewers whose distribution can be modeled as Poisson [7].

E. Related Work

In the context of the classic MAB, there have been several attempts at developing index-type policies that are simpler than Lai-Robbins policy by using a single sample mean based statistic [8], [9]. However, such a simpler form of the policy was obtained at the price of a more restrictive reward model and/or a larger leading constant in the logarithmic order. In particular, the index policy proposed in [8] mainly targets at several members of the exponential family of reward distributions. Auer *et al.* proposed in [9] several index-type policies that achieve order optimality for reward distributions with a known finite support.

Under a Bernoulli reward model in the context of cognitive radio, MAB with multiple players was considered in [10] and [11]. In [10], a heuristic policy based on histogram estimation of the unknown parameters was proposed. This policy provides a linear order of the system regret, thus cannot achieve the maximum average reward. In [11], Anandkumar *et al.* have independently established order-optimal distributed policies by extending the index-type single-user policies proposed in [9]. Compared to [11], the TDFS policy proposed here applies to more general reward models (for example, Gaussian and Poisson reward distributions that have infinite support). It thus has a wider range of potential applications as discussed in Section I-D. Furthermore, the policies proposed in [11] are specific to the single-player policies proposed in [9], whereas the TDFS policy can be used with any order-optimal single-player policy to achieve order optimality in the decentralized setting. Another difference between the policies proposed in [11] and the TDFS policy is on user fairness. The policies in [11] orthogonalize users into different channels that offer different throughput, whereas the TDFS policy ensures that each player achieves the same time-average reward at the same rate. One policy given in [11] does offer probabilistic fairness in the sense that all users have the same chance of settling in the best channel. However, given that the policy operates over an infinite horizon and the order optimality is asymptotic, each

¹While intuitive, this equivalence requires the condition that the M best arms have nonnegative means (see Section IV for details).

user only sees one realization in its lifetime, leading to different throughput among users. A lower bound on the achievable system regret is also given in [11], which is identical to the bound developed here. The derivation of the lower bound in this work, however, applies to a more general class of policies and, first given in [12], precedes that in [11]. In terms of using collision history to orthogonalizing players without pre-agreement, the basic idea used in this work is similar to that in [11]. The difference is that in the TDFS policy, the players are orthogonalized via settling at different offsets in their time sharing schedule, while in [11], players are orthogonalized to different channels. Recently, results obtained in [11] have been extended to incorporate unknown number of users for Bernoulli reward model in the context of cognitive radio [13].

A variation of centralized MAB in the context of cognitive radio has been considered in [14] where a channel offers independent Bernoulli rewards with different (unknown) means for different users. This more general model captures contention among secondary users that are in the communication range of different sets of primary users. A centralized policy that assumes full information exchange and cooperation among users is proposed which achieves the logarithmic order of the regret. We point out that the logarithmic order of the TDFS policy is preserved in the decentralized setting when we allow players to experience different reward distributions on the same arm provided that players have the common set of the M best arms and each of the M best arms has the same mean across players. It would be interesting to investigate whether the same holds when players experience different means on each arm.

F. Notations

Let $|\mathcal{A}|$ denote the cardinality of set \mathcal{A} . For two sets \mathcal{A} and \mathcal{B} , let $\mathcal{A} \setminus \mathcal{B}$ denote the set consisting of all elements in \mathcal{A} that do not belong to \mathcal{B} . For two positive integers k and l , define $k \oslash l \triangleq ((k-1) \bmod l) + 1$, which is an integer taking values from $1, 2, \dots, l$. Let $\Pr_{\theta}\{\cdot\}$ denote the probability measure when the unknown parameter in the associated distribution equals to θ .

II. CLASSIC RESULTS ON SINGLE-PLAYER MAB

In this section, we give a brief review of the main results established in [1], [2], [8], and [9] on the classic MAB with a single player.

A. System Regret

Consider an N -arm bandit with a single player. At each time t , the player chooses exactly M ($1 \leq M < N$) arms to play. Playing arm i yields i.i.d. random reward $S_i(t)$ drawn from a univariate density function $f(s; \theta_i)$ parameterized by θ_i . The parameter set $\Theta \triangleq (\theta_1, \dots, \theta_N)$ is unknown to the player. Let $\mu(\theta_i)$ denote the mean of $S_i(t)$ under the density function $f(s; \theta_i)$. Let $I(\theta, \theta') = \mathbb{E}_{\theta}[\log((f(y, \theta))/(f(y, \theta')))]$ be the Kullback–Liebler distance that measures the dissimilarity between two distributions parameterized by θ and θ' , respectively.

An arm selection policy $\pi = \{\pi(t)\}_{t=1}^{\infty}$ is a series of functions, where $\pi(t)$ maps the previous observations of rewards to the current action that specifies the set of M arms to play at time t . The system performance under policy π is measured by

the system regret $R_T^{\pi}(\Theta)$ defined as the expected total reward loss up to time T under policy π compared to the ideal scenario that Θ is known to the player (thus the M best arms are played at each time). Let σ be a permutation of $\{1, \dots, N\}$ such that $\mu(\theta_{\sigma(1)}) \geq \mu(\theta_{\sigma(2)}) \geq \dots \geq \mu(\theta_{\sigma(N)})$. We have

$$R_T^{\pi}(\Theta) \triangleq T \sum_{j=1}^M \mu(\theta_{\sigma(j)}) - \mathbb{E}_{\pi}[\sum_{t=1}^T S_{\pi(t)}(t)]$$

where $S_{\pi(t)}(t)$ is the random reward obtained at time t under action $\pi(t)$, and $\mathbb{E}_{\pi}[\cdot]$ denotes the expectation with respect to policy π . The objective is to minimize the rate at which $R_T(\Theta)$ grows with T under any parameter set Θ by choosing an optimal policy π^* .

We point out that the system regret is a finer performance measure than the long-term average reward. All policies with a sublinear regret would achieve the maximum long-term average reward. However, the difference in their performance measured by the total expected reward accrued over a time horizon of length T can be arbitrarily large as T increases. It is thus of great interest to characterize the minimum regret and construct policies optimal under this finer performance measure.

A policy is called *uniformly good* if for any parameter set Θ , we have $R_T^{\pi}(\Theta) = o(T^b)$ for any $b > 0$. Note that a uniformly good policy implies the sublinear system regret and achieves the maximum long-term average reward $\sum_{j=1}^M \mu(\theta_{\sigma(j)})$ which is the same as in the case with perfect knowledge of Θ .

B. The Logarithmic Order and the Optimal Policy

We present in the theorem below the result established in [1] and [2] on the logarithmic order as well as the leading constant of the minimum regret of the single-player MAB.

Theorem [1], [2]: Under the regularity conditions (conditions C1–C4 in Appendix A), we have, for any uniformly good policy π ,

$$\liminf_{T \rightarrow \infty} \frac{R_T^{\pi}(\Theta)}{\log T} \geq \sum_{j: \mu(\theta_j) < \mu(\theta_{\sigma(M)})} \frac{\mu(\theta_{\sigma(M)}) - \mu(\theta_j)}{I(\theta_j, \theta_{\sigma(M)})}. \quad (1)$$

Lai and Robbins also constructed a policy that achieves the lower bound on the regret given in (1) under single play ($M = 1$) [1] (which was extended by Anantharam *et al.* to $M > 1$ in [2]). Under this policy, two statistics of the past observations are maintained for each arm. Referred to as the point estimate, the first statistic is an estimate of the mean $\mu(\theta_i)$ given by a function h_k of the past observations on this arm (k denotes the total number of observations). The second statistic is the so-called confidence upper bound which represents the *potential* of an arm: the less frequently an arm is played, the less confident we are about the point estimate, and the higher the potential of this arm. The confidence upper bound, denoted by $g_{t,k}$, thus depends on not only the number k of observations on the arm but also the current time t in order to measure how frequently this arm has been played.

Based on these two statistics, Lai–Robbins policy operates as follows. At each time t , among all “well-sampled” arms, the one with the largest point estimate is selected as the leader denoted as l_t . The player then chooses between the leader l_t and a round-robin candidate $r_t = t \oslash N$ to play. The leader l_t is played if and only if its point estimate exceeds the confidence upper bound of

Lai-Robbins Policy for Single-Player MAB [1]

- Notations and Inputs: let $\tau_{n,t}$ denote the number of times that arm n has been played up to (but excluding) time t and $S_{n,1}, \dots, S_{n,\tau_{n,t}}$ the past observations obtained from arm n . Fix $\delta \in (0, 1/N)$.
- Initializations: in the first N steps, play each arm once.
- At time t ($t > N$), among all arms that have been played at least $(t-1)\delta$ times, let l_t denote the arm with the largest point estimate (referred to as the leader):

$$l_t \triangleq \arg \max_{n: \tau_{n,t} \geq (t-1)\delta} h_{\tau_{n,t}}(S_{n,1}, \dots, S_{n,\tau_{n,t}}).$$

Let $r_t = t \oslash N$ be the round robin candidate at time t . The player plays the leader l_t if $h_{\tau_{l_t,t}}(S_{l_t,1}, \dots, S_{l_t,\tau_{l_t,t}}) > g_{t,\tau_{r_t,t}}(S_{r_t,1}, \dots, S_{r_t,\tau_{r_t,t}})$ and the round-robin candidate r_t otherwise.

Fig. 1. Lai-Robbins policy for single-player MAB [1].

the round-robin candidate r_t . A detailed implementation of this policy is given in Fig. 1.

Lai and Robbins [1] have shown that for point estimates h_k and confidence upper bounds $g_{t,k}$ satisfying certain conditions (condition C5 in Appendix A), the above policy is optimal, i.e., it achieves the minimum regret given in (1). For Gaussian, Bernoulli, and Poisson reward models, h_k and $g_{t,k}$ satisfying condition C5 are as follows [1]:

$$h_k = \frac{\sum_{j=1}^k S_{n,j}}{k}, \quad (2)$$

$$g_{t,k} = \inf \left\{ \lambda : \lambda \geq h_k \text{ and } I(h_k, \lambda) \geq \frac{\log(t-1)}{k} \right\}. \quad (3)$$

While $g_{t,k}$ given in (3) is not in closed-form, it does not affect the implementation of the policy. Specifically, the comparison between the point estimate of the leader l_t and the confidence upper bound of the round-robin candidate r_t (see Fig. 1) is shown to be equivalent to the following two conditions [1]:

$$\begin{aligned} h_{\tau_{r_t,t}}(S_{r_t,1}, \dots, S_{r_t,\tau_{r_t,t}}) &< h_{\tau_{l_t,t}}(S_{l_t,1}, \dots, S_{l_t,\tau_{l_t,t}}) \\ I(h_{\tau_{r_t,t}}, h_{\tau_{l_t,t}}) &> \frac{\log(t-1)}{\tau_{r_t,t}}. \end{aligned}$$

Consequently, we only require the point estimate h_k to implement the policy.

C. Order-Optimal Index Policies

Since Lai and Robbins's seminal work, researchers have developed several index-type policies that are simpler than Lai-Robbins policy by using a single sample mean based statistic [8], [9]. Specifically, under such an index policy, each arm is assigned with an index that is a function of its sample mean and the arm with the greatest index will be played in each slot. To obtain an initial index, each arm is played once in the first N slots.

The indexes proposed in [8] for Gaussian, Bernoulli, Poisson, and exponential distributed reward models are given in (4), shown at the bottom of the page. Except for Gaussian distribution, this index policy only achieves the optimal logarithmic order of the regret but not the best leading constant,² where $I(\cdot)$ is the index function, x is the sample mean of the arm, and a, b are upper bounds of all possible values of θ in the Poisson and exponential reward models, respectively.

²The optimal indexes for Bernoulli, Poisson and exponential distributions are also developed in [8]. However, the indexes are not given in closed-form and are difficult to implement.

$$I(x) = \begin{cases} x + \left(2 \frac{\log(t-1) + 2 \log \log(t-1)}{\tau_{n,t}} \right)^{1/2}, & \text{Gaussian} \\ x + \min \left\{ \left(2 \frac{\log(t-1) + 2 \log \log(t-1)}{\tau_{n,t}} \right)^{1/2} / 2, 1 \right\}, & \text{Bernoulli} \\ x + \min \left\{ \left(2a \frac{\log(t-1) + 2 \log \log(t-1)}{\tau_{n,t}} \right)^{1/2} / 2, a \right\}, & \text{Poisson} \\ x + b \min \left\{ \left(2 \frac{\log(t-1) + 2 \log \log(t-1)}{\tau_{n,t}} \right)^{1/2}, 1 \right\}, & \text{Exponential} \end{cases} \quad (4)$$

Based on [8], a simpler order-optimal sample mean based index policy was established in [9] for reward distributions with a known finite support:

$$I(x) = x + \left(2 \frac{\log(t-1)}{\tau_{n,t}}\right)^{1/2}. \quad (5)$$

While the simple index policy given in [9] is only order-optimal, it allows arms to have different distribution functions $f(\cdot)$.

III. DECENTRALIZED MAB: PROBLEM FORMULATION

In this section, we formulate the decentralized MAB with M distributed players. In addition to conditions C1–C5 required by the centralized MAB, we assume that the M best arms have distinct nonnegative means.³

In the decentralized setting, players may collide and may not receive reward that the arm can potentially offer. We thus refer to $S_i(t)$ as the *state* of arm i at time t (for example, the busy/idle state of a communication channel in the context of cognitive radio). At time t , player i ($1 \leq i \leq M$) chooses an action $a_i(t) \in \{1, \dots, N\}$ that specifies the arm to play and observes its state $S_{a_i(t)}(t)$. The action $a_i(t)$ is based on the player's *local* observation and decision history. Note that the local observation history of each player also includes the collision history. As shown in Section V-C, the observation of a collision is used to adjust the local offset of a player's time sharing schedule of the M best arms to avoid excessive future collisions.

We define a local policy π_i for player i as a sequence of functions $\pi_i = \{\pi_i(t)\}_{t \geq 1}$, where $\pi_i(t)$ maps player i 's past observations and decisions to action $a_i(t)$ at time t . The decentralized policy π is thus given by the concatenation of the local policies for all players:

$$\pi \triangleq [\pi_1, \dots, \pi_M].$$

Define reward $Y(t)$ as the total reward accrued by all players at time t , which depends on the system collision model as given below.

Collision Model 1: When multiple players choose the same arm to play, they share the reward in an arbitrary way. Since how players share the reward has no effect on the total system regret, without loss of generality, we assume that only one of the colliding players obtains a random reward given by the current state of the arm. Under this model, we have

$$Y(t) = \sum_{j=1}^N \mathbb{I}_j(t) S_j(t)$$

where $\mathbb{I}_j(t)$ is the indicator function that equals to 1 if arm j is played by *at least* one player, and 0 otherwise.

Collision model 2: When multiple players choose the same arm to play, no one obtains reward. Under this model, we have

$$Y(t) = \sum_{j=1}^N \hat{\mathbb{I}}_j(t) S_j(t)$$

³This condition can be relaxed to the case that a tie occurs at the M th largest mean.

where $\hat{\mathbb{I}}_j(t)$ is the indicator function that equals to 1 if arm j is played by *exactly* one player, and 0 otherwise.

The system regret of policy π is thus given by

$$R_T^\pi(\Theta) = T \sum_{j=1}^M \mu(\theta_{\sigma(j)}) - \mathbb{E}_\pi \left[\sum_{t=1}^T Y(t) \right].$$

Note that the system regret in the decentralized MAB is defined with respect to the same best possible reward $T \sum_{j=1}^M \mu(\theta_{\sigma(j)})$ as in its centralized counterpart. The objective is to minimize the rate at which $R_T(\Theta)$ grows with time T under any parameter set Θ by choosing an optimal decentralized policy. Similarly, we say a decentralized policy is *uniformly good* if for any parameter set Θ , we have $R_T(\Theta) = o(T^b)$ for any $b > 0$. To address the optimal order of the regret, it is sufficient to focus on uniformly good decentralized policies provided that such policies exist.

We point out that all results developed in this work apply to a more general observation model. Specifically, the arm state observed by different players can be drawn from different distributions, as long as players have the common set of the M best arms and each of the M best arms has the same mean across players. This relaxation in the observation model is particularly important in the application of opportunistic transmission in fading channels, where different users experience different fading environments in the same channel.

IV. THE OPTIMAL ORDER OF THE SYSTEM REGRET

In this section, we show that the optimal order of the system regret of the decentralized MAB is logarithmic, the same as its centralized counterpart as given in Section II.

Theorem 1: Under both collision models, the optimal order of the system regret of the decentralized MAB is logarithmic, i.e., for an optimal decentralized policy π^* , we have

$$L(\Theta) \leq \liminf_{T \rightarrow \infty} \frac{R_T^{\pi^*}(\Theta)}{\log T} \leq \limsup_{T \rightarrow \infty} \frac{R_T^{\pi^*}(\Theta)}{\log T} \leq U(\Theta) \quad (6)$$

for some constants $L(\Theta)$ and $U(\Theta)$ that depend on Θ .

Proof: The proof consists of two parts. First, we prove that the lower bound for the centralized MAB given in (1) is also a lower bound for the decentralized MAB. Second, we construct a decentralized policy (see Section V) that achieves the logarithmic order of the regret.

It appears to be intuitive that the lower bound for the centralized MAB provides a lower bound for the decentralized MAB. This, however, may not hold when some of the M best arms have negative means (modeling the punishment for playing certain arms). The reason is that the centralized MAB considered in [2] requires exactly M arms to be played at each time, while in the decentralized setting, fewer than M arms may be played at each time when players choose the same arm. When the M best arms have nonnegative means, it is easy to show that playing up to M arms cannot achieve a better performance than the optimal performance achieved by playing exactly M arms. The lower bound for the centralized MAB given in (1) thus applies to the decentralized MAB. A detailed proof can be done along the line of [1, Theorem 1]. ■




Fig. 2. The structure of player 1's local policy under π_F^* ($M = 2$, $N = 3$). Let $\hat{\alpha}(1)$ denote the arm considered as the best by player 1 in the first subsequence. In this example, player 1 divides the second subsequence (i.e., all the even slots) into three mini-sequences, each associated with a subset of 2 arms after removing the arm considered to be the best in the first subsequence.).

V. AN ORDER-OPTIMAL DECENTRALIZED POLICY

In this section, we construct a decentralized policy π_F^* that achieves the optimal logarithmic order of the system regret under the fairness constraint.

A. Basic Structure of the Decentralized TDFS Policy π_F^*

The basic structure of the proposed policy π_F^* is a time division structure at each player for selecting the M best arms. For the ease of the presentation, we first assume that there is a pre-agreement among players so that they use different phases (offsets) in their time division schedule. For example, the offset in each player's time division schedule can be predetermined based on the player's ID. In Section V-C, we show that this pre-agreement can be eliminated while maintaining the order-optimality and fairness of the TDFS policy, which leads to a complete decentralization among players.

Consider, for example, the case of $M = 2$. The time sequence is divided into two disjoint subsequences, where the first subsequence consists of all odd slots and the second one consists of all even slots. The pre-agreement is such that player 1 targets at the best arm during the first subsequence and the second best arm during the second subsequence, and player 2 does the opposite.

Without loss of generality, consider player 1. In the first subsequence, player 1 applies a single-player policy, say Lai-Robbins policy, to efficiently learn and select the best arm. In the second subsequence, the second best arm is learned and identified by removing the arm that is considered to be the best and applying the Lai-Robbins policy to identify the best arm in the remaining $N - 1$ arms (which is the second best among all N arms). Note that since the best arm cannot be perfectly identified, which arm is considered as the best in an odd slot t (in the first subsequence) is a random variable determined by the realization of past observations. We thus partition the second subsequence into multiple mini-sequences depending on which arm was considered to be the best in the preceding odd slot and thus should be removed from consideration when identifying the second best. Specifically, as illustrated in Fig. 2, the second subsequence is divided into N disjoint mini-sequences, where the i th mini-sequence consists of all slots that follow a slot in which arm i was played (i.e., arm i was considered the best arm in the preceding slot that belongs to the first subsequence). In the i th mini-sequence, the player applies Lai-Robbins policy to arms $\{1, \dots, i - 1, i + 1, \dots, N\}$ after removing arm i .

In summary, the local policy for each player consists of $N + 1$ parallel Lai-Robbins procedures: one is applied in the subsequence that targets at the best arm and the rest N are applied in the N mini-sequences that target at the second best arm. These parallel procedures, however, are coupled through the common observation history since in each slot, regardless of which subsequence or mini-sequence it belongs to, all the past observations are used in the decision making. We point out that making each mini-sequence use only its own observations is sufficient for the order optimality of the TDFS policy and simplifies the optimality analysis. However, we expect that using all available observations leads to a better constant as demonstrated by simulation results in Section VII.

The basic structure of the TDFS policy π_F^* is the same for the general case of $M > 2$. Specifically, the time sequence is divided into M subsequences, in which each player targets at the M best arms in a round-robin fashion with a different offset. Suppose that player 1 has offset 0, i.e., it targets at the k th ($1 \leq k \leq M$) best arm in the k th subsequence. To player 1, the k th ($k > 1$) subsequence is then divided into C_{k-1}^N mini-sequences, each associated with a subset of $N - k + 1$ arms after removing those $k - 1$ arms that are considered to have a higher rank than k . In each of the C_{k-1}^N mini-sequences, Lai-Robbins single-player policy is applied to the subset of $N - k + 1$ arms associated with this mini-sequence. Note that while the subsequences are deterministic, each mini-sequence is random. Specifically, for a given slot t in the k th subsequence, the mini-sequence which it belongs to is determined by the specific actions of the player in the previous $k - 1$ slots. For example, if arms $1, \dots, k - 1$ are played in the previous $k - 1$ slots, then slot t belongs to the mini-sequence associated with the arm set of $\{k, \dots, N\}$. A detailed implementation of the proposed policy for a general M is given in Fig. 3.

Note that the basic structure of the TDFS policy π_F^* is general. It can be used with any single-player policy to achieve an efficient and fair sharing of the M best arms among M distributed players. Furthermore, players are not restricted to using the same single-player policy. Details are given in Section V-B.

B. Order-Optimality Under Fairness Constraint

Compared to the single-player counterpart given in [1] and [2], the difficulties in establishing the logarithmic regret of π_F^* are twofold. First, compared to the centralized problem where a single player observes M different arms simultaneously, each

player here can only observe one arm at each time. Each player is thus learning the entire rank of the M best arms with fewer observations. Furthermore, since the rank of any arm can never be perfectly identified, the mistakes in identifying the i th ($1 \leq i < M$) best arm will propagate into the learning process for identifying the $(i + 1)$ th, up to the M th best arm. Second, without centralized scheduling, collisions are bound to happen even with pre-agreed offset on sharing the M best arms since players do not always agree on the rank of the arms. Such issues need to be carefully dealt with in establishing the logarithmic order of the regret.

Theorem 2: Under the TDFS policy π_F^* , we have, for some constant $C(\Theta)$,

$$\limsup_{T \rightarrow \infty} \frac{R_T^{\pi_F^*}(\Theta)}{\log T} \leq C(\Theta). \quad (7)$$

Let

$$x_k \triangleq \sum_{j=1}^k \sum_{j: \mu(\theta_j) < \mu(\theta_{\sigma(i)})} \frac{1}{I(\theta_j, \theta_{\sigma(i)})}.$$

Under the collision model 1,

$$C(\Theta) = M \left(\sum_{i=1}^M x_k \mu(\theta_{\sigma(i)}) - \sum_{n: \mu(\theta_n) < \mu(\theta_{\sigma(M)})} \frac{\mu(\theta_n)}{I(\theta_n, \theta_{\sigma(M)})} \right).$$

Under the collision model 2,

$$C(\Theta) = M \left(\sum_{i=1}^M \sum_{k=1}^M x_k \mu(\theta_{\sigma(i)}) - \sum_{n: \mu(\theta_n) < \mu(\theta_{\sigma(M)})} \mu(\theta_n) \max \left\{ \frac{1}{I(\theta_n, \theta_{\sigma(M)})} - \sum_{i=1}^{M-1} \frac{1}{I(\theta_n, \theta_{\sigma(i)})}, 0 \right\} \right).$$

Proof: Note that the system regret is given by the sum of the regrets in the M subsequences. By symmetry in the structure of π_F^* , all subsequences experience the same regret. Consequently, the system regret is equal to M times the regret in each subsequence. We thus focus on a particular subsequence and show that the regret in this subsequence is at most logarithmic. Without loss of generality, consider the first subsequence in which the i th player targets at the i th ($1 \leq i \leq M$) best arm.

To show the logarithmic order of the regret, The basic approach is to show that the number of slots in which the i th player does not play the i th best arm is at most logarithmic with time. This is done by establishing in Lemma 2 a lower bound on the number of slots in which the i th player plays the i th best arm in the first subsequence.

To show Lemma 2, we first establish Lemma 1 by focusing on the *dominant* mini-sequence of the first subsequence. Without loss of generality, consider player i . Define the dominant mini-sequence as the one associated with arm set $\{\sigma(i), \dots, \sigma(N)\}$ (i.e., the $i - 1$ best arms are correctly identified and removed in this mini-sequence). Lemma 1 shows that, in the dominant mini-sequence, the number of slots in which player i does not play the i th best arm is at most logarithmic with time.

Lemma 1: Let $\tau_{n,T}^D$ denote the number of slots up to time T in which arm n is played in the dominant mini-sequence associated with arm set $\{\sigma(i), \dots, \sigma(N)\}$. Then, for any arm n with $\mu(\theta_n) < \mu(\theta_{\sigma(i)})$, we have

$$\limsup_{T \rightarrow \infty} \frac{\mathbb{E}[\tau_{n,T}^D]}{\log T} \leq \frac{1}{I(\theta_n, \theta_{\sigma(i)})}. \quad (8)$$

Proof: Note that this lemma is an extension of [1, Theorem 3]. The proof of this lemma is more complicated since the mini-sequence is random and the decisions made in this mini-sequence depend on all past observations (no matter to which mini-sequence they belong). See Appendix B for details. ■

Next, we establish Lemma 2. The basic approach is to show that the length of the dominant mini-sequence dominates the lengths of all other mini-sequences in the first subsequences. Specifically, we show that the number of slots that do not belong to the dominant mini-sequence is at most logarithmic with time. This, together with Lemma 1 that characterizes the dominant mini-sequence, leads to Lemma 2 below.

Lemma 2: Let $\hat{\tau}_{\sigma(i),T}$ denote the number of slots in which player i plays the i th best arm in the first subsequence up to time T , we have

$$\limsup_{T \rightarrow \infty} \frac{T/M - \mathbb{E}[\hat{\tau}_{\sigma(i),T}]}{\log T} \leq x_i \quad (9)$$

where

$$x_i = \sum_{j=1}^i \sum_{k: \mu(\theta_k) < \mu(\theta_{\sigma(j)})} \frac{1}{I(\theta_{\sigma(k)}, \theta_{\sigma(j)})}. \quad (10)$$

Proof: The proof is based on an induction argument on i . See Appendix C for details. ■

From Lemma 2, for all $1 \leq i \leq M$, the number of slots that the i th best arm is not played by player i is at most logarithmic with time. Consequently, for all $j \neq i$, the number of slots that player j plays the i th best arm is also at most logarithmic with time, i.e., the number of collisions on the i th best arm is at most logarithmic with time. Since a reward loss on the i th ($1 \leq i \leq M$) best arm can only happen when it is not played or a collision happens, the reward loss on the i th best is at most logarithmic order of time, leading to the logarithmic order of the regret.

To establish the upper bound $C(\Theta)$ on the constant of the logarithmic order of the system regret, we consider the worst-case collisions on each arm. See Appendix D for details. ■

From Theorem 1 and Theorem 2, the decentralized policy is order-optimal. Furthermore, as given in Theorem 3 below, the decentralized policy ensures fairness among players under a fair collision model that offers all colliding players the same *expected* reward. For example, if only one colliding player can receive reward, then all players have equal probability to be lucky.

Theorem 3: Define the local regret for player i under π_F^* as

$$R_{T,i}^{\pi_F^*}(\Theta) \triangleq \frac{1}{M} T \sum_{j=1}^M \mu(\theta_{\sigma(j)}) - \mathbb{E}_{\pi_F^*} [\sum_{t=1}^T Y_i(t)]$$

where $Y_i(t)$ is the immediate reward obtained by player i at time t . Under a fair collision model, we have

$$\limsup_{T \rightarrow \infty} \frac{R_{T,i}^{\pi_F^*}(\Theta)}{\log T} = \frac{1}{M} \limsup_{T \rightarrow \infty} \frac{R_T^{\pi_F^*}(\Theta)}{\log T}. \quad (11)$$

Theorem 3 follows directly from the symmetry among players under π_F^* . It shows that each player achieves the same time average reward $(1/M)T \sum_{j=1}^M \mu(\theta_{\sigma(j)})$ at the same rate. We point out that without knowing the reward rate that each arm can offer, ensuring fairness requires that each player identify the entire set of the M best arms and share each of these M arms evenly with other players. As a consequence, each player needs to learn which of the C_M^N possibilities is the correct choice. This is in stark difference from policies that make each player target at a single arm with a specific rank (for example, the i th player targets solely at the i th best arm). In this case, each player only needs to distinguish one arm (with a specific rank) from the rest. The uncertainty facing each player, consequently, the amount of learning required, is reduced from C_M^N to N . Unfortunately, fairness among players is lost.

As mentioned in Section V-A, the proposed policy π_F^* can be used with any single-player policy, which can also be different for different players. More importantly, the order optimality of the TDFS policy is preserved as long as each player's single-player policy achieves the optimal logarithmic order in the single player setting. This statement can be proven along the same line as Theorem 2 by establishing results similar to Lemma 1 and Lemma 2.

C. Eliminating the Pre-Agreement

In this subsection, we show that a pre-agreement among players on the offsets of the time division schedule can be eliminated in the TDFS policy π_F^* while maintaining its order-optimality and fairness.

Specifically, when a player joins the system, it randomly generates a local offset uniformly drawn from $\{0, 1, \dots, M-1\}$ and plays one round of the M arms considered to be the best. For example, if the random offset is 1, the player targets at the second, the third, \dots , the M th, and then the best arms in the M subsequent slots, respectively. If no collision is observed in this round of M plays, the player keeps the same offset; otherwise the player randomly generates a new offset. Over the sequence of slots where the same offset is used, the player implements the local policy of π_F^* (given in Fig. 3) with this offset. To summarize, each player implements M parallel local procedures of π_F^* corresponding to M different offsets. These parallel procedures are coupled through the observation history, i.e., each player uses its entire local observation history in learning no matter which offset is being used.

Note that players can join the system at different time. We also allow each player to leave the system for an arbitrary finite number of slots.

Theorem 4: The decentralized TDFS policy π_F^* without pre-agreement is order-optimal and fair.

Proof: See Appendix E. ■

VI. A LOWER BOUND FOR A CLASS OF DECENTRALIZED POLICES

In this section, we establish a lower bound on the system regret for a general class of decentralized policies, to which the proposed policy π_F^* belongs. This lower bound provides a

tighter performance benchmark compared to the one defined by the centralized MAB. The definition of this class of decentralized policies is given below.

Definition 1: Time Division Selection Policies The class of time division selection (TDS) policies consists of all decentralized policies $\pi = [\pi_1, \dots, \pi_M]$ that satisfy the following property: under local policy π_i , there exists $0 \leq a_{ij} \leq 1$ ($1 \leq i, j \leq M$, $\sum_{j=1}^M a_{ij} = 1 \forall i$) independent of the parameter set Θ such that the expected number of times that player i plays the j th ($1 \leq j \leq M$) best arm up to time T is $a_{ij}T - o(T^b)$ for all $b > 0$.

A policy in the TDS class essentially allows a player to efficiently select each of the M best arms according to a fixed time portion that does not depend on the parameter set Θ . It is easy to see that the TDFS policy π_F^* (with or without pre-agreement) belongs to the TDS class with $a_{ij} = 1/M$ for all $1 \leq i, j \leq M$.

Theorem 5: For any uniformly good decentralized policy π in the TDS class, we have

$$\liminf_{T \rightarrow \infty} \frac{R_T^{\pi}(\Theta)}{\log T} \geq \sum_{i=1}^M \sum_{j: \mu(\theta_j) < \mu(\theta_{\sigma(M)})} \frac{\mu(\theta_{\sigma(M)}) - \mu(\theta_j)}{I(\theta_j, \theta_{\sigma(i)})}. \quad (12)$$

Proof: The basic approach is to establish a lower bound on the number of slots in which each player plays an arm that does not belong to the M best arms. By considering the best case that they do not collide, we arrive at the lower bound on the regret given in (12). The proof is based on the following lemma, which generalizes [1, Theorem 2]. To simplify the presentation, we assume that the means of all arms are distinct. However, Theorem 5 and Lemma 3 apply without this assumption.

Lemma 3: Consider a local policy π_i . If for any parameter set Θ and $b > 0$, there exists a Θ -independent positive increasing function $v(T)$ satisfying $v(T) \rightarrow \infty$ as $T \rightarrow \infty$ such that

$$\mathbb{E}[v(T) - \tau_{\sigma(j), T}] \leq o((v(T))^b) \quad (13)$$

then we have, $\forall k > j$,

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}[\tau_{\sigma(k), T}]}{\log v(T)} \geq \frac{1}{I(\theta_{\sigma(k)}, \theta_{\sigma(j)})}. \quad (14)$$

Note that Lemma 3 is more general than [1, Theorem 2] that assumes $v(T) = T$ and $j = 1$. The proof of Lemma 3 is given in Appendix F.

Consider a uniformly good decentralized policy π in the TDS class. There exists a player, say player u_1 , that plays the best arm for at least $T/M - o(T^b)$ times. Since at these times, u_1 cannot play other arms, there must exist another player, say player u_2 , that plays the second best arm for at least $T/(M(M-1)) - o(T^b)$ times. It thus follows that there exist M different players u_1, \dots, u_M such that under any parameter set Θ , the expected time player u_i ($1 \leq i \leq M$) plays the i th best arm is at least $T/\prod_{j=1}^i (M-j+1) - o(T^b)$ times. Based on Lemma 3, for any arm j with $\mu(\theta_j) < \mu(\theta_{\sigma(M)})$, the expected time that player u_i plays arm j is at least $1/I(\theta_j, \theta_{\sigma(i)}) \log(T/\prod_{j=1}^i (M-j+1)) - o(\log T)$. By considering the best case that players do not collide, we arrive at (12). ■

The Decentralized TDFS Policy π_F^*

Without loss of generality, consider player i .

- Notations and Inputs: In addition to the notations and inputs of Lai-Robbins single-player policy (see Fig. 1), let $m_k(t)$ denote the number of slots in the k th subsequence up to (and including) t , and let $m_{\mathcal{N}}(t)$ denote the number of slots in the mini-sequence associated with arm set \mathcal{N} up to (and including) t .
- At time t , player i does the following.
 1. If t belongs to the i th subsequence (i.e., $t \oslash M = i$), player i targets at the best arm by carrying out the following procedure.
 If $m_i(t) \leq N$, play arm $m_i(t)$. Otherwise, the player chooses between a leader l_t and a round-robin candidate $r_t = m_i(t) \oslash N$, where the leader l_t is the arm with the largest point estimate among all arms that have been played for at least $(m_i(t) - 1)\delta$ times. The player plays the leader l_t if its point estimate is larger than the confidence upper bound of r_t . Otherwise, the player plays the round-robin candidate r_t .
 2. If t belongs to the k th ($k \neq i$) subsequence (i.e., $t \oslash M = k$), the player targets at the j th best arm where $j = (k - i + M + 1) \oslash M$ by carrying out the following procedure.
 If $m_k(t) \leq N$, play arm $m_k(t)$. Otherwise, the player targets at the j th best arm. Let \mathcal{A}_t denote the set of $j - 1$ arms played in the previous $j - 1$ slots. Slot t thus belongs to the mini-sequence associated with the subset $\mathcal{N}_t = \{1, \dots, N\} \setminus \mathcal{A}_t$ of arms. The player chooses between a leader and a round-robin candidate defined within \mathcal{N}_t . Specifically, among all arms that have been played for at least $(m_{\mathcal{N}_t}(t) - 1)\delta$ times, let l_t denote the leader. Let $r_t = m_{\mathcal{N}_t}(t) \oslash (N - j + 1)$ be the round-robin candidate where, for simplicity, we have assumed that arms in \mathcal{N}_t are indexed by $1, \dots, N - j + 1$. The player plays the leader l_t if its point estimate is larger than the confidence upper bound of r_t . Otherwise, the player plays the round-robin candidate r_t .

Fig. 3. The decentralized TDFS policy π_F^* .

VII. SIMULATION EXAMPLES

In this section, we study the performance (i.e., the leading constant of the logarithmic order) of the decentralized TDFS policy in different applications through simulation examples.

A. Cognitive Radio Networks: Bernoulli Reward Model

We first consider a Bernoulli reward model using cognitive radio as an example of application. There are M secondary users independently search for idle channels among N channels. In each slot, the busy/idle state of each channel is drawn from Bernoulli distribution with unknown mean, i.e., $f(s; \theta) = \theta^s(1 - \theta)^{1-s}$ for $s \in \{0, 1\}$. When multiple secondary users choose the same channel for transmission, they will collide and no one gets reward (collision model 2).

In Fig. 4, we plot the performance of π_F^* (with pre-agreement) using different single-player policies. The leading constant of the logarithmic order is plotted as a function of N with fixed M . From Fig. 4, we observe that adopting the optimal Lai-Robbins single-player policy in π_F^* achieves the best performance. This is expected given that Lai-Robbins policy

achieves the best leading constant in the single-player case. We also observe that coupling the parallel single-player procedures through the common observation history in π_F^* leads to a better performance compared to the one without coupling.

In Fig. 5, we study the impact of eliminating pre-agreement on the performance of π_F^* . We plot the leading constant of the logarithmic order as a function of M with fixed N . We observe that eliminating pre-agreement comes with a price in performance in this case. We also observe that the system performance degrades as M increases. One potential cause is the fairness property of the policy that requires each player to learn the entire rank of the M best arms.

B. Multichannel Communications Under Unknown Fading: Exponential Reward Model

In this example, we consider opportunistic transmission over wireless fading channels with unknown fading statistics. In each slot, each user senses the fading condition of a selected channel and then transmits data with a fixed power. The reward obtained from a chosen channel is measured by its capacity (maximum




Fig. 4. The performance of π_F^* built upon different single-player policies (Bernoulli distributions, global horizon length $T = 5000$, $M = 2$, $\Theta = [0.1, 0.2, \dots, (N/10)]$).




Fig. 5. The performance of π_F^* based on Lai-Robbins policy (Bernoulli distributions, global horizon length $T = 5000$, $N = 9$, $\Theta = [0.1, 0.2, \dots, 0.9]$).

data transmission rate) $C = \log(1 + \text{SNR})$. We consider the Rayleigh fading channel model where the SNR of each channel is exponential distributed with an unknown mean. When multiple users choose the same channel, no one succeeds (collision model 2). Note that a channel that has a higher expected SNR also offers a higher expected channel capacity. It is thus equivalent to consider SNR as the reward which is exponential distributed. Specifically, we have $f(s; \theta) = 1/\theta \exp(-s/\theta)$ for $s \in \mathcal{R}^+$.

In Fig. 6, we plot the leading constant of the logarithmic order as a function of N with fixed M . We observe that in this example, eliminating pre-agreement has little impact on the performance of the TDFS policy.

C. Target Collecting in Multi-Agent Systems: Gaussian Reward Model

In this example, we consider the Gaussian reward model arisen in the application of multi-agent systems. At each time, M agents independently select one of N locations to collect targets (e.g., fishing or ore mining). The reward at each location




Fig. 6. The performance of π_F^* based on Agrawal's index policy (Exponential distributions, global horizon length $T = 5000$, $M = 2$, $\Theta = [1, 2, \dots, N]$).




Fig. 7. The convergence of the regret under π_F^* based on Lai-Robbins policy (Gaussian distributions, $\sigma = 1$, $M = 4$, $N = 9$, $\Theta = [1, 2, \dots, 9]$).

is determined by the fish size or ore quality, which has been shown in [5] and [6] to fit with log-Gaussian distribution. The reward at each location has an unknown mean but a known variance which is the same across all locations. When multiple agents choose the same location, they share the reward (collision model 1).

Note that the log-Gaussian and Gaussian reward distributions are equivalent since the arm ranks under these two distributions are the same when the arms have the same variance. We can thus focus on the Gaussian reward model, i.e., $f(s; \theta) = (2\pi\sigma^2)^{-1/2} \exp\{-(s - \theta)^2/(2\sigma^2)\}$ for $s \in \mathcal{R}$.

In Fig. 7, we plot $R_T^F / \log(T)$ as a function of global horizon T . We observe that the regret quickly converges as time goes, which implies that the policy can achieve a strong performance within a short finite period. In Fig. 8, we plot the leading constant of the logarithmic order as a function of M with fixed N . Similar to the previous example, eliminating pre-agreement has little impact on the performance of the TDFS policy.

VIII. CONCLUSION

In this paper, we have studied a decentralized formulation of the classic multi-armed bandit problem by considering multiple




Fig. 8. The performance of π_F^* based on Lai-Robbins policy (Gaussian distributions, $\sigma = 1$, global horizon length $T = 5000$, $N = 10$, $\Theta = [1, 2, \dots, 10]$).

distributed players. We have shown that the optimal system regret in the decentralized MAB grows at the same logarithmic order as that in the centralized MAB considered in the classic work by Lai and Robbins [1] and Anantharam, *et al.* [2]. A decentralized policy that achieves the optimal logarithmic order has been constructed. A lower bound on the leading constant of the logarithmic order is established for policies in the TDS class, to which the proposed policy belongs. Future work includes establishing tighter lower bound on the leading constant of the logarithmic order and investigating decentralized MAB where each arm has different means to different players.

APPENDIX A

Let \mathcal{H} denote the set of all possible values of the unknown parameter θ .

Regularity Conditions:

- C1) Existence of mean: $\mu(\theta)$ exists for any $\theta \in \mathcal{H}$.
- C2) Positive distance: $0 < I(\theta, \theta') < \infty$ whenever $\mu(\theta) < \mu(\theta')$.
- C3) Continuity of $I(\theta, \theta')$: $\forall \epsilon > 0$ and $\mu(\theta) < \mu(\theta')$, $\exists \delta > 0$ such that $|I(\theta, \theta') - I(\theta, \theta'')| < \epsilon$ whenever $\mu(\theta'') \leq \mu(\theta') + \delta$.
- C4) Denseness of \mathcal{H} : $\forall \theta$ and $\delta > 0$, $\exists \theta' \in \Theta$ such that $\mu(\theta) < \mu(\theta') < \mu(\theta) + \delta$.

Conditions on point estimate and confidence upper bound:

- C5) For any $\theta \in \mathcal{H}$, we have
 - i) $\Pr_{\theta}\{g_{t,k}(S(1), \dots, S(k)) \geq r \text{ for all } k < t\} = 1 - o(t^{-1})$ for every $r < \mu(\theta)$;

- ii) $\lim_{\epsilon \downarrow 0} (\limsup_{t \rightarrow \infty} \sum_{k=1}^{t-1} \Pr_{\theta}\{g_{t,k}(S(1), \dots, S(k)) \geq \mu(\theta') - \epsilon\} / \log t) \leq 1/I(\theta, \theta')$ whenever $\mu(\theta') > \mu(\theta)$;
- iii) $g_{t,k}$ is nondecreasing in t for every fixed $k = 1, 2, \dots$;
- iv) $g_{t,k} \geq h_k \forall t$;
- v) $\Pr_{\theta}\{\max_{\delta t \leq k \leq t} |h_k - \mu(\theta)| > \epsilon\} = o(t^{-1}) \forall \epsilon > 0, 0 < \delta < 1$.

APPENDIX B

PROOF OF LEMMA 1

Let l_t denote the leader among the arm set $\{\sigma(i), \dots, \sigma(N)\}$. Consider arm n with $\mu(\theta_n) < \mu(\theta_{\sigma(i)})$. Let $\mathcal{D}(T)$ denote the set of slots in the dominant mini-sequence up to time T .

For any $\epsilon \in (0, \mu(\theta_{\sigma(i)}) - \mu(\theta_{\sigma(i+1)}))$, let $N_1(T)$ denote the number of slots in $\mathcal{D}(T)$ at which arm n is played when the leader is the i th best arm and the difference between its point estimate and true mean does not exceed ϵ , $N_2(T)$ the number of slots in $\mathcal{D}(T)$ at which arm n is played when the leader is the i th best arm and the difference between its point estimate and true mean exceeds ϵ , and $N_3(T)$ the number of slots in $\mathcal{D}(T)$ when the leader is not the i th best arm. Recall that each arm is played once in the first N slots. We have

$$\tau_{n,T}^D \leq 1 + N_1(T) + N_2(T) + N_3(T). \quad (15)$$

Next, we show that $\mathbb{E}[N_1(T)]$, $\mathbb{E}[N_2(T)]$, and $\mathbb{E}[N_3(T)]$ are all at most in the order of $\log T$.

Consider first $\mathbb{E}[N_1(T)]$. Based on the structure of Lai-Robbins single-player policy, we have

$$\begin{aligned} N_1(T) &= |\{1 < t \leq T : t \in \{\mathcal{D}(T)\}, \mu(\theta_{l_t}) \\ &= \mu(\theta_{\sigma(i)}), |h_{\tau_{l_t,t}} - \mu(\theta_{\sigma(i)})| \\ &\leq \epsilon \text{ and arm } n \text{ is played at time } t\}| \\ &\leq 1 + |\{1 \leq j \leq T-1 : g_{s,j}(S_{n1}, \dots, S_{nj}) \\ &\geq \mu(\theta_{\sigma(i)}) - \epsilon \text{ for some } j < s \leq t\}| \\ &\leq 1 + |\{1 \leq j \leq T-1 : g_{T,j}(S_{n1}, \dots, S_{nj}) \\ &\geq \mu(\theta_{\sigma(i)}) - \epsilon\}|. \end{aligned} \quad (16)$$

Under condition C5, for any $\rho > 0$, we can choose ϵ sufficiently small such that [see the equation shown at the bottom of the page]. Thus,

$$\limsup_{T \rightarrow \infty} \frac{\mathbb{E}[N_1(T)]}{\log T} \leq \frac{1}{I(\theta_n, \theta_{\sigma(i)})}. \quad (17)$$

$$\begin{aligned} &\limsup_{T \rightarrow \infty} \frac{\mathbb{E}(|\{1 \leq j \leq T-1 : g_{T,j}(S_{n1}, \dots, S_{nj}) \geq \mu(\theta_{\sigma(i)}) - \epsilon\}|)}{\log T} \\ &= \limsup_{T \rightarrow \infty} \frac{\sum_{j=1}^{T-1} \Pr\{g_{T,j}(S_{n1}, \dots, S_{nj}) \geq \mu(\theta_{\sigma(i)}) - \epsilon\}}{\log T} \\ &\leq \frac{1 + \rho}{I(\theta_n, \theta_{\sigma(i)})}. \end{aligned}$$

Consider $\mathbb{E}[N_2(T)]$. Define $c_D(t) \triangleq |\{s : s \leq t, s, t \in \mathcal{D}(T)\}|$ as the number of slots in $\mathcal{D}(T)$ that are no larger than t . Since the number of observations obtained from l_t is at least $(c_D(t) - 1)\delta$, under condition C5, we have

$$\begin{aligned} & \Pr\{c_D(t) = s \text{ for some } s \leq t, \mu(\theta_{l_t}) \\ & \quad = \mu(\theta_{\sigma(i)}), |h_{\tau_{l_t}, t} - \mu(\theta_{\sigma(i)})| > \epsilon\} \\ & \leq \Pr\left\{\sup_{j \geq \delta(s-1)} |h_j(S_{l_t, 1}, \dots, S_{l_t, j}) - \mu(\theta_{l_t})| > \epsilon\right\} \\ & = \sum_{i=0}^{\infty} \delta^i o(s^{-1}) \\ & = o(s^{-1}). \end{aligned} \quad (18)$$

We thus have

$$\begin{aligned} \mathbb{E}[N_2(T)] & = \mathbb{E}\left(|\{1 < t \leq T : t \in \mathcal{D}(T), \mu(\theta_{l_t}) \right. \\ & \quad = \mu(\theta_{\sigma(i)}), |h_{\tau_{l_t}, t} - \mu(\theta_{\sigma(i)})| > \epsilon\}|\right) \\ & \leq \sum_{s=1}^T \Pr\{c_D(t) = s \text{ for some } s \leq t \leq T, \mu(\theta_{l_t}) \\ & \quad = \mu(\theta_{\sigma(i)}), |h_{\tau_{l_t}, t} - \mu(\theta_{\sigma(i)})| > \epsilon\} \\ & = o(\log T). \end{aligned} \quad (19)$$

Next, we show that $\mathbb{E}[N_3] = o(\log T)$.

Choose $0 < \epsilon_1 < (\mu(\theta_{\sigma(i)}) - \mu(\theta_{\sigma(i+1)}))/2$ and $c > (1 - N\delta)^{-1}$. For $r = 0, 1, \dots$, define the following events:

$$\begin{aligned} A_r & \triangleq \cap_{i \leq k \leq N} \left\{ \max_{\delta c^{r-1} \leq s} |h_s(S_{\sigma(k), 1}, \dots, S_{\sigma(k), s}) \right. \\ & \quad \left. - \mu(\theta_{\sigma(k)})| \leq \epsilon_1 \right\}, \\ B_r & \triangleq \{g_{s_m, j}(S_{\sigma(i), 1}, \dots, S_{\sigma(i), j}) \\ & \quad \geq \mu(\theta_{\sigma(i)}) - \epsilon_1 \text{ for all } 1 \leq j \\ & \quad \leq \delta m, c^{r-1} \leq m \leq c^{r+1}, \text{ and } s_m > m\}. \end{aligned}$$

By (18), we have $\Pr(\bar{A}_r) = o(c^{-r})$. Consider the following event:

$$\begin{aligned} C_r & \triangleq \{g_{m+1, j}(S_{\sigma(i), 1}, \dots, S_{\sigma(i), j}) \geq \mu(\theta_{\sigma(i)}) \\ & \quad - \epsilon_1 \text{ for all } 1 \leq j \leq \delta m, \text{ and } c^{r-1} \leq m \leq c^{r+1}\}. \end{aligned} \quad (20)$$

Under condition C5, it is easy to see that $B_r \supset C_r$. From Lemma 1 – (i) in [1], $\Pr(\bar{C}_r) = o(c^{-r})$. We thus have $\Pr(\bar{B}_r) = o(c^{-r})$.

Consider time t . Define the event $D_r \triangleq \{t \in \mathcal{D}(T), c^{r-1} \leq c_D(t) - 1 < c^{r+1}\}$. When the round-robin candidate $r_t = \sigma(i)$, we show that on the event $A_r \cap B_r \cap D_r$, $\sigma(i)$ must be played. It is sufficient to focus on the nontrivial case that $\mu(\theta_{l_t}) < \mu(\theta_{\sigma(i)})$. Since $\tau_{l_t, t} \geq (c_D(t) - 1)\delta$, on $A_r \cap D_r$, we have $h_{\tau_{l_t}, t} < \mu(\theta_{\sigma(i)}) - \epsilon_1$. We also have, on $A_r \cap B_r \cap D_r$,

$$g_{t, \tau_{\sigma(i), t}}(S_{\sigma(i), 1}, \dots, S_{\sigma(i), \tau_{\sigma(i), t}}) \geq \mu(\theta_{\sigma(i)}) - \epsilon_1. \quad (21)$$

Arm $\sigma(i)$ is thus played on $A_r \cap B_r \cap D_r$. Since $(1 - c^{-1})/N > \delta$, for any $c^r \leq s \leq c^{r+1}$, there exists an r_0 such that on $A_r \cap B_r$,

$\tau_{\sigma(i), t} \geq (1/N)(s - c^{r-1} - 2N) > \delta s$ for all $r > r_0$. It thus follows that on $A_r \cap B_r \cap D_r$, for any $c^r \leq c_D(t) - 1 < c^{r+1}$, we have $\tau_{\sigma(i), t} > (c_D(t) - 1)\delta$, and $\sigma(i)$ is thus the leader. We have, for all $r > r_0$,

$$\begin{aligned} & \Pr(l_t \neq \sigma(i), t \in \mathcal{D}(T), c^r \leq c_D(t) - 1 < c^{r+1}) \\ & \leq \Pr(\bar{A}_r \cap D_r) + \Pr(\bar{B}_r \cap D_r) = o(c^{-r}). \end{aligned} \quad (22)$$

Therefore,

$$\begin{aligned} \mathbb{E}[N_3(T)] & = \mathbb{E}\left(|\{t > 1 : t \in \mathcal{D}(T), h_{\tau_{l_t}, t} < \mu(\theta_{\sigma(i)})\}|\right) \\ & \leq \sum_{s=0}^{T-1} \Pr(c_D(t) - 1 \\ & \quad = s \text{ for some } t > s, l_t \neq \sigma(i), t \in \mathcal{D}(T)) \\ & \leq 1 + \sum_{r=0}^{\lceil \log_c T \rceil} \sum_{c^r \leq s \leq c^{r+1}} \Pr(\exists t, l_t \neq \sigma(i), t \in \mathcal{D}(T), c_D(t) - 1 = s) \\ & = 1 + \sum_{r=0}^{\lceil \log_c T \rceil} o(1) \\ & = o(\log T). \end{aligned} \quad (23)$$

From (17), (19), and (23), we arrive at Lemma 1.

APPENDIX C PROOF OF LEMMA 2

We prove by induction on i ($1 \leq i \leq M$). Consider $i = 1$. In this case, Lai–Robbins policy is applied to all arms by player 1. Let $\mathbb{E}[\hat{\tau}_{\sigma(i), T}]$ denote the expected total number of times in the first sequence that player i does not play the i th best arm up to time T . From Lemma 1, we have

$$\limsup_{T \rightarrow \infty} \frac{\mathbb{E}[\hat{\tau}_{\sigma(1), T}]}{\log T} = \sum_{n: \mu(\theta_n) < \mu(\theta_{\sigma(1)})} \limsup_{T \rightarrow \infty} \frac{\mathbb{E}[\tau_{n, T}^D]}{\log T} \leq x_1. \quad (24)$$

Since

$$\frac{(T - 2M)}{M} \leq \mathbb{E}[\hat{\tau}_{\sigma(1), T} + \hat{\tau}_{\sigma(1), T}] \leq \frac{T}{M},$$

we have

$$\limsup_{T \rightarrow \infty} \frac{T/M - \mathbb{E}[\hat{\tau}_{\sigma(1), T}]}{\log T} = \limsup_{T \rightarrow \infty} \frac{\mathbb{E}[\hat{\tau}_{\sigma(1), T}]}{\log T} \leq x_1 \quad (25)$$

which shows that (9) holds for $i = 1$. Let $\mathbb{E}[\tau_T^{D_i}]$ denote the number of slots in the dominant mini-sequence of player i up to time T . From (9), we have, for $i = 1$

$$\limsup_{T \rightarrow \infty} \frac{T/M - \mathbb{E}[\tau_T^{D_{i+1}}]}{\log T} \leq x_i. \quad (26)$$

Assume (9) and (26) hold for $i = k$. To complete the induction process, consider $i = k + 1$. Let $\mathbb{E}[\tau_{\sigma(i), T}^{D_i}]$ denote the expected number of slots that the i th best arm is played by player i in its dominant mini-sequence up to time T , and $\mathbb{E}[\tau_{\sigma(i), T}^{D_i}]$ the expected number of slots that the i th best arm is not played by

player i in its dominant mini-sequence up to time T . Based on Lemma 1, we have

$$\limsup_{T \rightarrow \infty} \frac{\mathbb{E} \left[\tau_{\sigma(i), T}^{D_i} \right]}{\log T} \leq \sum_{n: \mu(\theta_n) < \mu(\theta_{\sigma(i)})} \frac{1}{I(\theta_n, \theta_{\sigma(i)})}. \quad (27)$$

Since

$$\mathbb{E} \left[\tau_T^{D_i} \right] = \mathbb{E} \left[\tau_{\sigma(i), T}^{D_i} + \tau_{\sigma(i), T}^{D_i} \right] \leq \mathbb{E} \left[\tau_{\sigma(i), T}^{D_i} + \hat{\tau}_{\sigma(i), T} \right],$$

we have

$$\begin{aligned} \limsup_{T \rightarrow \infty} \frac{T/M - \mathbb{E} \left[\hat{\tau}_{\sigma(i), T} \right]}{\log T} \\ \leq \limsup_{T \rightarrow \infty} \frac{T/M - \mathbb{E} \left[\tau_{\sigma(i), T}^{D_i} \right]}{\log T} \\ = \limsup_{T \rightarrow \infty} \frac{T/M - \mathbb{E} \left[\tau_T^{D_i} \right] + \mathbb{E} \left[\tau_{\sigma(i), T}^{D_i} \right]}{\log T} \leq x_{k+1}. \end{aligned} \quad (28)$$

We thus proved Lemma 2 by induction.

APPENDIX D

PROOF OF THE UPPER BOUND $C(\Theta)$ IN THEOREM 2

Collision Model 1: Consider the first subsequence. Since the expected number of slots that the i th best arm is played by player i dominates that by other players, the expected total reward $\mathbb{E}[Y_i(T)]$ obtained from the i th best arm by all players up to time T satisfies

$$\limsup_{T \rightarrow \infty} \frac{T\mu(\theta_{\sigma(i)})/M - \mathbb{E}[Y_i(T)]}{\log T} \leq x_i\mu(\theta_{\sigma(i)}). \quad (29)$$

For an arm n with $\mu(\theta_n) < \mu(\theta_{\sigma(M)})$, the worst case is that player M plays arm n for $1/I(\theta_n, \theta_{\sigma(M)}) \log T - o(\log T)$ times up to time T (by Lemma 1). From (29), the expected total reward $\sum_{i=1}^N \mathbb{E}[Y_i(T)]$ obtained from all arms and players up to time T satisfies

$$\begin{aligned} \limsup_{T \rightarrow \infty} \frac{T\sum_{i=1}^M \mu(\theta_{\sigma(i)})/M - \sum_{i=1}^N \mathbb{E}[Y_i(T)]}{\log T} \\ \leq \sum_{i=1}^M x_i\mu(\theta_{\sigma(i)}) \\ - \sum_{n: \mu(\theta_n) < \mu(\theta_{\sigma(M)})} \frac{1}{I(\theta_n, \theta_{\sigma(M)})} \mu(\theta_n). \end{aligned} \quad (30)$$

The expected total reward $\sum_{j=1}^M \sum_{i=1}^N \mathbb{E}[Y_i(T)]$ obtained from all arms and players over all the M subsequences up to time T thus satisfies

$$\begin{aligned} \limsup_{T \rightarrow \infty} \frac{T\sum_{i=1}^M \mu(\theta_{\sigma(i)}) - \sum_{j=1}^M \sum_{i=1}^N \mathbb{E}[Y_i(T)]}{\log T} \\ \leq M \left(\sum_{i=1}^M x_i\mu(\theta_{\sigma(i)}) \right. \\ \left. - \sum_{n: \mu(\theta_n) < \mu(\theta_{\sigma(M)})} \frac{1}{I(\theta_n, \theta_{\sigma(M)})} \mu(\theta_n) \right). \end{aligned} \quad (31)$$

Collision Model 2: Consider the first subsequence. By Lemma 2, the expected number of slots in which the i th best arm is played by player j ($j \neq i$) that is at most $x_j \log T + o(\log T)$ up to time T . Together with (9), the expected total reward $\mathbb{E}[Y_i(T)]$ obtained from arm i by all players up to time T satisfies

$$\limsup_{T \rightarrow \infty} \frac{T\mu(\theta_{\sigma(i)})/M - \mathbb{E}[Y_i(T)]}{\log T} \leq \sum_{k=1}^M x_k\mu(\theta_{\sigma(i)}). \quad (32)$$

For an arm n with $\mu(\theta_n) < \mu(\theta_{\sigma(M)})$, the worst case is that player i plays arm n for $1/I(\theta_n, \theta_{\sigma(i)}) \log T - o(\log T)$ times up to time T . Combined with the worst case collisions, the expected reward obtained on arm n is at least $(\max\{1/I(\theta_n, \theta_M) - \sum_{i=1}^{M-1} 1/I(\theta_n, \theta_{\sigma(i)}), 0\} \mu(\theta_n)) \log T - o(\log T)$ up to time T . From (32), the expected total reward $\sum_{i=1}^N \mathbb{E}[Y_i(T)]$ obtained from all arms and players up to time T satisfies

$$\begin{aligned} \limsup_{T \rightarrow \infty} \frac{T\sum_{i=1}^M \mu(\theta_{\sigma(i)})/M - \sum_{i=1}^N \mathbb{E}[Y_i(T)]}{\log T} \\ \leq \sum_{i=1}^M \sum_{k=1}^M x_k\mu(\theta_{\sigma(i)}) - \sum_{n: \mu(\theta_n) < \mu(\theta_{\sigma(M)})} \\ \max \left\{ \frac{1}{I(\theta_n, \theta_M)} - \sum_{i=1}^{M-1} \frac{1}{I(\theta_n, \theta_{\sigma(i)})}, 0 \right\} \mu(\theta_n). \end{aligned} \quad (33)$$

The expected total reward $\sum_{j=1}^M \sum_{i=1}^N \mathbb{E}[Y_i(T)]$ obtained from all arms and players over all the M subsequences up to time T thus satisfies

$$\begin{aligned} \limsup_{T \rightarrow \infty} \frac{T\sum_{i=1}^M \mu(\theta_{\sigma(i)}) - \sum_{j=1}^M \sum_{i=1}^N \mathbb{E}[Y_i(T)]}{\log T} \\ \leq M \left(\sum_{i=1}^M \sum_{k=1}^M x_k\mu(\theta_{\sigma(i)}) - \sum_{n: \mu(\theta_n) < \mu(\theta_{\sigma(M)})} \right. \\ \left. \max \left\{ \frac{1}{I(\theta_n, \theta_M)} - \sum_{i=1}^{M-1} \frac{1}{I(\theta_n, \theta_{\sigma(i)})}, 0 \right\} \mu(\theta_n) \right). \end{aligned} \quad (34)$$

APPENDIX E

PROOF OF THEOREM 4

Based on the structure of π_F^* without pre-agreement, each player's local policy consists of disjoint rounds, each of which consists of M slots where M arms are played with a particular local offset. We say that a round is in group o ($0 \leq o \leq M-1$) if it corresponds to local offset o . We also say that a round is *normal* if the player plays the M best arms in the correct order according to its local offset, otherwise, the round is *singular*. Let t denote the global reference time. We say that slot t is *normal* if every player's round that contains slot t is normal, otherwise, slot t is *singular*.

Note that the set of slots where a reward loss incurs is a subset of the singular slots and normal slots involving collisions. To prove the logarithmic order of the system regret, it is thus sufficient to show the following: i) the expected number of singular slots is at most logarithmic with time and ii) the expected

number of normal slots involving collisions is at most logarithmic with time.

To show i), we consider an arbitrary player. Similar to the proof of Theorem 2, it can be shown that in any group o , the expected number of singular rounds is at most logarithmic with time. Summed over all M groups, the total expected number of singular rounds is at most logarithmic with time. Since that a singular round corresponds to M singular slots, the expected number of singular slots is at most logarithmic with time.

To show ii), we first show that between any two consecutive singular slots, the expected number of normal slots involving collisions is uniformly bounded by some constant (independent of the two singular slots being considered). Note that a collision occurs in a normal slot only when multiple players have the same global offset (with respect to the global time t). Since a player will randomly generate a new local offset if a collision is observed, it is easy to show that, between any two consecutive singular slots, the expected number of normal slots needed for all players to settle at different global offsets is uniformly bounded. Consequently, between any two consecutive singular slots, the expected number of normal slots involving collisions is uniformly bounded. It thus follows that the expected number of normal slots involving collisions is at the same logarithmic order as the expected number of singular slots.

The fairness of π_F^* without pre-agreement follows directly from the symmetry among players.

APPENDIX F PROOF OF LEMMA 3

For any $\delta > 0$, under condition C1–C4, we can choose a parameter λ such that

$$\mu(\theta_{\sigma(j-1)}) > \mu(\lambda) > \mu(\theta_{\sigma(j)}) > \mu(\theta_{\sigma(k)})$$

and

$$|I(\theta_{\sigma(k)}, \lambda) - I(\theta_{\sigma(k)}, \theta_{\sigma(j)})| \leq \delta I(\theta_{\sigma(k)}, \theta_{\sigma(j)}). \quad (35)$$

Consider the new parameter set Θ' where the mean of the k th best arm (say arm l) is replaced by λ . Since arm l is the j th best arm under Θ' , from (13), we have, for any $0 < b < \delta$,

$$\mathbb{E}_{\Theta'}[v(T) - \tau_{l,T}] \leq o((v(T))^b) \quad (36)$$

where $\tau_{l,T}$ is the number of times that arm l is played up to time T . There thus exists a random function $c(T)$ with $c(T) + v(T) - \tau_{l,T} \geq 0$ a.s. and $\mathbb{E}_{\Theta'}[c(T)] \geq 0$ such that⁴

$$\mathbb{E}_{\Theta'}[c(T) + v(T) - \tau_{l,T}] = o((v(T))^b). \quad (37)$$

We have

$$\begin{aligned} & \mathbb{E}_{\Theta'}[c(T) + v(T) - \tau_{l,T}] \\ &= \mathbb{E}_{\Theta'} \left[c(T) + v(T) - \tau_{l,T} \mid \tau_{l,T} < \frac{(1-\delta) \log(v(T))}{I(\theta_{\sigma(k)}, \lambda)} \right] \end{aligned}$$

⁴For example, $c(T)$ can be constructed as $c(T) = b(T) - v(T) + \tau_{l,T}$, where $b(T) = [\mathbb{E}_{\Theta'}[v(T) - \tau_{l,T}]]^+$.

$$\begin{aligned} & \times \Pr_{\Theta'} \left\{ \tau_{l,T} < \frac{(1-\delta) \log(v(T))}{I(\theta_{\sigma(k)}, \lambda)} \right\} \\ & + \mathbb{E}_{\Theta'} \left[c(T) + v(T) - \tau_{l,T} \mid \tau_{l,T} \geq \frac{(1-\delta) \log(v(T))}{I(\theta_{\sigma(k)}, \lambda)} \right] \\ & \times \Pr_{\Theta'} \left\{ \tau_{l,T} \geq \frac{(1-\delta) \log(v(T))}{I(\theta_{\sigma(k)}, \lambda)} \right\} \\ & \geq \mathbb{E}_{\Theta'} \left[c(T) + v(T) - \tau_{l,T} \mid \tau_{l,T} < \frac{(1-\delta) \log(v(T))}{I(\theta_{\sigma(k)}, \lambda)} \right] \\ & \times \Pr_{\Theta'} \left\{ \tau_{l,T} < \frac{(1-\delta) \log(v(T))}{I(\theta_{\sigma(k)}, \lambda)} \right\} \\ & \geq \left(\mathbb{E}_{\Theta'}[c(T)] + v(T) - \frac{(1-\delta) \log(v(T))}{I(\theta_{\sigma(k)}, \lambda)} \right) \\ & \times \Pr_{\Theta'} \left\{ \tau_{l,T} < \frac{(1-\delta) \log(v(T))}{I(\theta_{\sigma(k)}, \lambda)} \right\} \end{aligned} \quad (38)$$

where the first inequality is due to the fact that $c(T) + v(T) - \tau_{l,T} \geq 0$ almost surely.

Let $S_{l,1}, S_{l,2}, \dots$ denote independent observations from arm l . Define

$$L_t \triangleq \sum_{i=1}^t \log \left(\frac{f(S_{l,i}; \theta_{\sigma(k)})}{f(S_{l,i}; \lambda)} \right)$$

and event

$$C \triangleq \left\{ \tau_{l,T} < \frac{(1-\delta) \log(v(T))}{I(\theta_{\sigma(k)}, \lambda)}, L_{\tau_{l,T}} \leq (1-b) \log v(T) \right\}.$$

Following (37) and (38), we have

$$\begin{aligned} & \Pr_{\Theta'}\{C\} \\ & \leq \Pr_{\Theta'} \left\{ \tau_{l,T} < \frac{(1-\delta) \log(v(T))}{I(\theta_{\sigma(k)}, \lambda)} \right\} \\ & \leq \frac{o((v(T))^b)}{(\mathbb{E}_{\Theta'}[c(T)] + v(T) - (1-\delta) \log(v(T)) / I(\theta_{\sigma(k)}, \lambda))} \\ & \leq \frac{o((v(T))^b)}{(v(T) - (1-\delta) \log(v(T)) / I(\theta_{\sigma(k)}, \lambda))} \\ & = o((v(T))^{b-1}) \end{aligned} \quad (39)$$

where the last inequality is due to the fact that $\mathbb{E}_{\Theta'}[c(T)] \geq 0$.

We write C as the union of mutually exclusive events $C_s \triangleq \{\tau_{l,T} = s, L_s \leq (1-b) \log v(T)\}$ for each integer $s < (1-b) \log(v(T)) / I(\theta_{\sigma(k)}, \lambda)$. Note that

$$\begin{aligned} & \Pr_{\Theta'}\{C_s\} \\ &= \int_{\{\tau_{l,T}=s, L_s \leq (1-b) \log v(T)\}} d\Pr_{\Theta'} \\ &= \int_{\{\tau_{l,T}=s, L_s \leq (1-b) \log v(T)\}} \prod_{i=1}^s \frac{f(S_{l,i}; \lambda)}{f(S_{l,i}; \theta_{\sigma(k)})} d\Pr_{\Theta} \\ &\geq \exp(-(1-b) \log v(T)) \Pr_{\Theta}\{C_s\}. \end{aligned} \quad (40)$$

Based on (39) and (40), we have

$$\Pr_{\Theta}\{C\} = (\log(v(T)))^{1-b} \Pr_{\Theta'}\{C\} \rightarrow 0. \quad (41)$$

By the strong law of large numbers, $L_t/t \rightarrow I(\theta_{\sigma(k)}, \lambda) > 0$ as $t \rightarrow \infty$ a.s. under \Pr_Θ . This leads to $\max_{i \leq t} L_i/t \rightarrow I(\theta_{\sigma(k)}, \lambda) > 0$ a.s. under \Pr_Θ . Since $1 - b > 1 - \delta$, we have

$$\lim_{T \rightarrow \infty} \Pr_\Theta \left\{ L_i > (1 - b) \log v(T) \text{ for some } i < \frac{(1 - \delta) \log v(T)}{I(\theta_{\sigma(k)}, \lambda)} \right\} = 0 \quad (42)$$

which leads to

$$\lim_{T \rightarrow \infty} \Pr_\Theta \left\{ \tau_{l,T} < \frac{(1 - \delta) \log v(T)}{I(\theta_{\sigma(k)}, \lambda)}, L_{\tau_{l,T}} > (1 - b) \log v(T) \right\} = 0. \quad (43)$$

Based on (41) and (43), we have

$$\lim_{T \rightarrow \infty} \Pr_\Theta \left\{ \tau_{l,T} < \frac{(1 - \delta) \log v(T)}{I(\theta_{\sigma(k)}, \lambda)} \right\} = 0. \quad (44)$$

Based on (35), we arrive at

$$\lim_{T \rightarrow \infty} \Pr_\Theta \left\{ \tau_{l,T} < \frac{(1 - \delta) \log v(T)}{((1 + \delta) I(\theta_{\sigma(k)}, \theta_{\sigma(j)}))} \right\} = 0 \text{ for any } \delta > 0 \quad (45)$$

which leads to (14).

REFERENCES

- [1] T. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Adv. Appl. Math.*, vol. 6, no. 1, p. 4C22, 1985.
- [2] V. Anantharam, P. Varaiya, and J. Walrand, "Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays—Part I: IID rewards," *IEEE Tran. Autom. Control*, vol. 32, no. 11, p. 968C976, 1987.
- [3] Q. Zhao and B. Sadler, "A survey of dynamic spectrum access," *IEEE Signal Process. Mag.*, vol. 24, pp. 79–89, May 2007.
- [4] X. Liu, E. Chong, and N. Shroff, "Opportunistic transmission scheduling with resource-sharing constraints in wireless," *IEEE J. Sel. Areas Commun.*, vol. 19, no. 10, p. 2053, 2001.
- [5] R. Myers and P. Pepin, "The robustness of lognormal-based estimators of abundance source," *Biometrics*, vol. 46, no. 4, pp. 1185–1192, Dec. 1990.
- [6] W. E. Sharp, "A log-normal distribution of alluvial diamonds with an economic cutoff," *Econom. Geol.*, vol. 71, no. 3, pp. 648–655, May 1976.
- [7] M. A. Bean, *Probability: The Science of Uncertainty With Applications to Investments, Insurance, and Engineering*. Pacific Grove, CA: Brooks/Cole, 2009.
- [8] R. Agrawal, "Sample mean based index policies with $O(\log n)$ regret for the multi-armed bandit problem," *Adv. Appl. Probab.*, vol. 27, no. 4, pp. 1054–1078, Dec. 1995.
- [9] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Mach. Learn.*, vol. 47, pp. 235–256, 2002.
- [10] L. Lai, H. Jiang, and H. V. Poor, "Medium access in cognitive radio networks: A competitive multi-armed bandit framework," presented at the *IEEE Asilomar Conf. Signals, Systems, Computers*, Pacific Grove, CA, Oct. 2008.
- [11] A. Anandkumar, N. Michael, and A. K. Tang, "Opportunistic spectrum access with multiple players: Learning under competition," in *Proc. IEEE INFOCOM*, San Diego, CA, Mar. 2010.
- [12] K. Liu and Q. Zhao, "Decentralized multi-armed bandit with multiple distributed players," [Online]. Available: <http://arxiv.org/abs/0910.2065v1>
- [13] A. Anandkumar, N. Michael, A. Tang, and A. Swami, "Distributed learning and allocation of cognitive users with logarithmic regret," *IEEE J. Sel. Areas in Commun. on Advances in Cognitive Radio Netw. and Commun.*, Dec. 2009, submitted for publication.
- [14] Y. Gai, B. Krishnamachari, and R. Jain, "Learning multiplayer channel allocations in cognitive radio networks: A combinatorial multi-armed bandit formulation," in *IEEE Symp. New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2010.



Keqin Liu received the B.S. degree in automation from Southeast University, China, in 2005 and the M.S. and Ph.D. degrees in electrical engineering from the University of California at Davis, in 2008 and 2010, respectively.

He is currently working as a Postdoctoral Researcher at the Department of Electrical Engineering, University of California at Davis. His research interests are stochastic optimization in dynamic systems, distributed control in multi-agent systems, and signal processing in wireless networks.



Qing Zhao received the Ph.D. degree in electrical engineering from Cornell University, Ithaca, NY, in 2001.

In August 2004, she joined the Department of Electrical and Computer Engineering at the University of California at Davis (UC Davis), where she is currently an Associate Professor. Her research interests are in the general area of dynamic systems and communication networks.

Dr. Zhao received the 2000 Young Author Best Paper Award from the IEEE Signal Processing Society and the 2008 Outstanding Junior Faculty Award from the UC Davis College of Engineering.