

Learning of Uncontrolled Restless Bandits with Logarithmic Strong Regret

Cem Tekin, *Member, IEEE*, Mingyan Liu, *Senior Member, IEEE*

Abstract

In this paper we consider the problem of learning the optimal dynamic policy for uncontrolled restless bandit problems. In an uncontrolled restless bandit problem, there is a finite set of arms, each of which when played yields a non-negative reward. There is a player who sequentially selects one of the arms at each time step. The goal of the player is to maximize its undiscounted reward over a time horizon T . The reward process of each arm is a finite state Markov chain, whose transition probabilities are unknown to the player. State transitions of each arm is independent of the player's actions, thus "uncontrolled". We propose a learning algorithm with near-logarithmic regret uniformly over time with respect to the optimal (dynamic) finite horizon policy, referred to as strong regret, to contrast with commonly studied notion of weak regret which is with respect to the optimal (static) single-action policy. We also show that when an upper bound on a function of the system parameters is known, our learning algorithm achieves logarithmic regret. Our results extend the literature on optimal adaptive learning of Markov Decision Processes (MDPs) to Partially Observed Markov Decision Processes (POMDPs). Finally, we provide numerical results on a variation of our proposed learning algorithm and compare its performance and running time with other bandit algorithms.

Index Terms

Online learning, restless bandits, POMDPs, regret, exploration-exploitation tradeoff

C. Tekin is with the Electrical Engineering Department, University of California, Los Angeles, CA, USA, cmtkn@ucla.edu. M. Liu is with the Electrical Engineering and Computer Science Department, University of Michigan, Ann Arbor, MI 48105, USA, mingyan@eecs.umich.edu. Part of this work was done when C. Tekin was a graduate student in the Electrical Engineering and Computer Science Department, University of Michigan. A preliminary version of this work appeared in Allerton 2011. The work is partially supported by the NSF under grants CIF-0910765 and CNS-1217689, and by the ARO under Grant W911NF-11-1-0532.

I. INTRODUCTION

In an *uncontrolled restless bandit problem* (URBP) there is a set of arms indexed by $1, 2, \dots, K$, whose state process is discrete and follows a discrete time Markov rule independent of each other. There is a user/player who chooses one arm at each of the discrete time steps, gets a reward dependent on the state of the arm, and observes the current state of the selected arm. The control action, i.e., the arm selection, does not affect the state transition, therefore the underlying system dynamics is uncontrolled. However, judicious arm selections allow the player to obtain high instantaneous reward (exploitation) and decrease the uncertainty about the current state of the system (exploration), and the key in designing an optimal policy lies in a good balance between exploration and exploitation.

If the structure of the system, i.e., the state transition probabilities and the rewards of the arms are known, then the optimal policy for an infinite horizon problem can be found by using dynamic programming. In the case of infinite horizon with discounted reward, stationary optimal policies can be found by using contraction properties of the dynamic programming operator. For the infinite horizon average reward case, stationary optimal policies can be found under some assumptions on the transition probabilities [1], [2].

In this paper, rather than finding the optimal policy given the structure of the system (referred to as the optimization problem), we consider the *learning* version of the problem, where we assume that initially the player has no knowledge on the transition probabilities of the arms. This type of learning problem arise in many applications. Examples include sequential channel selection in a multi-channel wireless system where a user initially has no information on the channel statistics, and target tracking where initially the statistics of the target's movement is unknown. Our goal is to design learning algorithms with the fastest convergence rate, i.e., the minimum regret, where the *regret* of a learning policy at time t is defined as the difference between the total undiscounted reward of the optimal dynamic policy for the finite t -horizon undiscounted reward problem given full statistics of the system model, and that of the learning policy up to time t . It should be noted that this is a form of *strong regret*, as the comparison benchmark is the optimal dynamic policy, the best causal policy that can be achieved given full statistics of the system. By contrast, a much more commonly used performance criterion is the *weak regret*, which is the difference between a learning policy and the best *single-action policy*, a static policy that always plays the same arm. Also, to determine the best single-action policy one does not need to know the full statistics of the system, but only the average reward of each arm. For simplicity throughout the paper, the term *regret* refers to the *strong regret* defined above.

In this paper, we show that when the transition probability between any two states of the same arm is always positive, and when the player knows an upper bound on a function of the system parameters, which we will define later, an algorithm with logarithmic regret uniform in time for the finite horizon undiscounted reward problem exists. If such a bound is not known, we show that near-logarithmic regret can be achieved. We would like to note that the iid bandit problem where the rewards for arm k are drawn from a distribution f_k with finite support (such as the Bernoulli bandit problem) is a special case of the URBP. Since it is proven in [3] that iid bandits have a logarithmic lower bound on the regret, this bound also holds for URBP.

To the best of our knowledge, this is the first attempt to extend optimal adaptive learning in Markov decision processes (MDPs) to partially observable Markov decision processes (POMDPs), of which the URBP is a sub-class, where there are an uncountably infinite number of information states. A parallel work [?] considered the idea of state aggregation to solve URBP, and proved a $O(\sqrt{T})$ lower bound on the regret. This bound does not contradict our $O(\log T)$ bound since it is derived for an adversarially selected problem instance. In contrast, the problem instance (including the states, rewards and transition probability matrices of the arms) is fixed in our case, and the constant in the regret that multiplies the logarithmic time order depends on these parameters. Therefore our $O(\log T)$ bound is instance dependent. Similarly, it is known from [?] that even in the iid bandit problem, for any time horizon T there exists a set of reward distributions such that no algorithm can achieve a regret bound better than $O(\sqrt{T})$, while logarithmic regret upper bounds dependent on instance-specific constants have been proved by many existing works (see, e.g., [6], [5]) in the iid setting.

The remainder of this paper is organized as follows. Related work is given in Section II. In Section III, we present the problem formulation, notations and some preliminaries, including an equivalent countable representation of the information state. Then in Section IV we present a learning algorithm, followed by a finite partition of the information states in Section V, which is used to bound the strong regret of the learning algorithm. We analyze the regret of this learning algorithm in Section VI and prove that it increases logarithmically in time. We then introduce two variants of the learning algorithm aimed at relaxing certain assumptions, in Sections VII and VIII, respectively. Extensive numerical results on the performance of our online learning algorithms and comparison with other online learning algorithms in prior works is given in Section IX. Section X concludes the paper.

II. RELATED WORK

Related work on multi-armed bandit problems started with the seminal paper by Lai and Robbins [3], where asymptotically optimal adaptive policies for arms with iid reward processes (referred to as the “iid problem” for simplicity below) were constructed. These are index policies, and it was shown that they achieve an asymptotically optimal $O(\log t)$ regret, for single parameterized bandit problems, meaning that the regret is optimal both in terms of the logarithmic order in time t (referred to as order optimality), and optimal among all algorithms with logarithmic order in time (referred to as optimality in the constant), asymptotically. Later, Agrawal [4] considered the same iid problem and provided sample mean based index policies which are easier to compute, order optimal but not optimal in terms of the constant in general. Multiple simultaneous plays was considered in Anantharam et al [5], and an asymptotically optimal policy was proposed. All the above work assumed parametrized distributions for the reward process of the arms. The problem class we consider in this paper, i.e., URBP, includes parameterized iid bandits with finite states such as the Bernoulli bandits. Therefore, there exists instances of URBP for which no online learning algorithm can achieve better than logarithmic regret. However, in general proving lower bounds for bandit problems is a challenging task and is out of the scope of this paper. Rather than proving lower bounds on the strong regret, in this paper we focus on proving upper bounds on the strong regret, and evaluating the real time performance of our learning algorithms numerically. In contrast to the work on parameterized bandits, Auer et al [6] proposed sample mean based index policies for the iid problem with logarithmic regret when reward processes have a bounded support. Their upper bound holds uniformly over time rather than asymptotically but this bound is not asymptotically optimal.

There have been efforts to extend this type of learning from iid bandits to Markovian bandits. Markovian bandits are further divided into two groups: *rested* bandits whereby the state of an arm can only change when it is played or activated, and *restless* bandits whereby the state of an arm changes according to different Markovian rules depending on whether it is played or not. Optimization version of the rested bandits, when the player knows the transition probabilities, for a discounted reward criterion was solved by Gittins and Jones [7], while the restless bandits was shown to be intractable by Papadimitriou and Tsitsiklis [8] even in the optimization version. Nevertheless, heuristics, approximations and exact solutions under different/stronger assumptions on the arms have been studied by many, see e.g., [9], [10], [11]. The learning version of the rested bandits was considered by Anantharam et. al. [12], and an asymptotically optimal algorithm was proposed. Note that when the arms are driven by iid processes, the optimal policy (given all statistics) is a static policy that always plays the arm with the highest expected reward. Thus

in this case weak regret coincides with strong regret, and the problem is greatly simplified in that arm selection is based on a function of the sample mean of the sequence of rewards from the arms, but not on the sequence itself. The same observation holds when the arms evolve according to a rested Markovian process, since the optimal policy for average reward is also a static policy when the time horizon is sufficiently large.

Because of the aforementioned difficulties in efficiently computing the optimal dynamic policy for the restless bandits even when the transition probabilities are known, the performance criteria used in designing computationally efficient learning algorithms for restless bandits is typically the weak regret, by comparing to the best static, or single-action policy, one that always plays the arm with the highest expected reward. In particular, following the approach by Auer et al [6], we in [13], [14] provided policies with uniformly logarithmic weak regret bounds with respect to the best-single arm policy for both restless and rested bandit problems and extended the results to single-player multiple-play and decentralized multi-player models in [15]. We achieved this by merging blocks of observed states from the same arm to form a continuous sample path of the underlying Markov process. In a parallel work by Liu et. al. [16], a similar result is obtained through deterministic sequencing of exploration and exploitation, in which the player explores or exploits an arm in blocks whose lengths increase geometrically over time. Also related are decentralized multi-player versions of the iid problem under different collision models, see e.g., [17], [18], [19], [20].

There has been relatively less work in designing learning algorithms with stronger regret notions than the weak regret, with the exception of [21], in which a bandit problem with identical two-state arms was considered, whereby the optimal policy belongs to a finite set given to the player, and [22], in which we considered a special case of the restless bandits called the *feedback* bandits, a more general version of the problem studied in [21] and proposed a computationally efficient learning algorithm that is approximately optimal compared to the optimal dynamic policy, based on an optimization algorithm proposed by Guha et. al. [10].

Outside the bandit literature, there has been a lot of work in adaptive learning of Markov Decision Processes (MDP) with finite state and action spaces, where the goal is to learn the optimal policy with the smallest possible strong regret. Burnetas and Katehakis [23] proposed index policies with asymptotically logarithmic regret, where the indices are the inflations of right-hand-sides of the estimated average reward optimality equations based on Kullback Leibler (KL) divergence, and showed that these are asymptotically optimal both in terms of the order and the constant. However, they assumed that the support of the transition probabilities are known. Tewari and Bartlett [24] proposed a learning algorithm that uses l_1

distance instead of KL divergence with the same order of regret but a larger constant; their proof is simpler than that found in [23] and does not require the support of the transition probabilities to be known. Auer and Ortner [25] proposed another algorithm with logarithmic regret and reduced computation for the MDP problem, which solves the average reward optimality equations only when a confidence interval is halved. In all the above work the MDPs are assumed to be irreducible.

The rested bandit problem with Markovian reward processes may be viewed as a special case of the MDP problem, and therefore the above literature applies to the rested bandit learning and can be used to obtain strong regret results. By contrast, the restless bandit problem is a special case of the partially observed MDP (or POMDP) problems, which has an uncountably infinite state space and the irreducibility condition may not hold in general. The goal of this paper is to develop learning algorithms that achieve logarithmic strong regret for the restless bandit problem.

It is also worth noting the difference between regret learning, the approach taken in this paper, and Q-learning. The learning approach we take in this paper is model-based, in which the player estimates the transition probabilities and exploits the estimates to learn how to play. By contrast, Q-learning is a model-free approach that estimates Q-values for state-action pairs (rather than transition probabilities). However, the convergence guarantees of Q-learning algorithms are weaker than our regret bounds. In general, Q-learning algorithms do not have sublinear regret guarantees, and convergence in terms of average reward only takes place when all state-action pairs are observed infinitely many times. See [26], [27], [28] for examples of Q-learning in finite MDPs and [29], [30] for Q-learning in POMDPs. Specifically in [28], a Q-learning algorithm for the rested bandit problem is proposed, in which estimates of the Gittins indices for the arms are obtained via Q-learning. In [30], a POMDP problem with *landmark states* in which each state results in a unique observation is studied, and under standard stochastic approximation conditions asymptotic convergence of Q-functions to the optimal Q-values was proved.

III. PROBLEM FORMULATION AND PRELIMINARIES

A. Problem Formulation and Notations

Consider K mutually independent uncontrolled restless Markovian arms, indexed by the set $\mathcal{K} = \{1, 2, \dots, K\}$ whose states evolve in discrete time steps $t = 1, 2, \dots$ according to a finite-state Markov chain with unknown transition probabilities.

Let S^k be the state space of arm k . For simplicity of presentation, we assume that for state $x \in S^k$, $r_x^k = x$, i.e., the state of an arm also represents its reward under that state. This is without loss of generality as long as one of the following is true: either the state is perfectly observed when played, or

that the reward is perfectly observed when received which uniquely identifies a state for a given arm (i.e., no two states have the same reward). It follows that the state space of the system is the Cartesian product of the state spaces of individual arms, denoted by $\mathcal{S} = S^1 \times \dots \times S^K$. Let p_{ij}^k denote the transition probability from state i to state j of arm k . The transition probability matrix of arm k is denoted by P^k , whose ij th element is p_{ij}^k . The set of transition probability matrices is denoted by $\mathbf{P} = (P^1, \dots, P^K)$. We assume that P^k s are such that each arm is ergodic. This implies that, for each arm there exists a unique stationary distribution which is given by $\pi^k = (\pi_x^k)_{x \in S^k}$. At each time step, the state of the system is a K -dimensional vector of states of arms which is given by $\mathbf{x} = (x^1, \dots, x^K) \in \mathcal{S}$.

The following notation will be frequently used throughout the paper. Let e_x^k represent the unit vector with dimension $|S^k|$, whose x th element is 1, and all other elements are 0. $\mathbb{N} = \{1, 2, \dots\}$ denotes the set of natural numbers, $\mathbb{Z}_+ = \{0, 1, \dots\}$ the set of non-negative integers, $(\mathbf{v} \bullet \mathbf{w})$ the standard inner product of vectors \mathbf{v} and \mathbf{w} , $\|\mathbf{v}\|_1$ and $\|\mathbf{v}\|_\infty$ respectively the l_1 and l_∞ norms of vector \mathbf{v} , and $\|P\|_1$ the induced maximum row sum norm of matrix P . For a vector \mathbf{v} , (\mathbf{v}_{-u}, v') denotes the vector whose u th element is v' , while all other elements are the same as in \mathbf{v} . For a vector of matrices \mathbf{P} , (\mathbf{P}_{-u}, P') denotes the vector of matrices whose u th matrix is P' , while all other matrices are the same as in \mathbf{P} . The transpose of a vector \mathbf{v} or matrix P is denoted by \mathbf{v}^T or P^T , respectively. In addition, the following quantities frequently appear in this paper:

- $\beta = \sum_{t=1}^{\infty} 1/t^2$, $\pi_{\min}^k = \min_{x \in S^k} \pi_x^k$;
- $\pi_{\min} = \min_{k \in \mathcal{K}} \pi_{\min}^k$;
- $r_{\max} = \max_{x \in S^k, k \in \mathcal{K}} r_x^k$;
- $S_{\max} = \max_{k \in \mathcal{K}} |S^k|$.

There is a player who selects one of the K arms at each time step t , and gets a bounded reward depending on the state of the selected arm at time t . Without loss of generality, we assume that the state rewards are non-negative. Let $r^k(t)$ be the random variable which denotes the reward from arm k at time t . The objective of the player is to maximize the undiscounted sum of the rewards over any finite horizon $T > 0$. However, the player does not know the set of transition probability matrices \mathbf{P} . In addition, at any time step t the player can only observe the state of the arm it selects but not the states of the other arms. Intuitively, in order to maximize its reward, the player needs to explore/sample the arms to estimate their transition probabilities and to reduce the uncertainty about the current state $\mathbf{x} \in \mathcal{S}$ of the system, while it also needs to exploit the information it has acquired about the system to select arms that yield high rewards. The exploration and exploitation need to be carefully balanced to yield the

maximum reward for the player. In a more general sense, the player is learning to play optimally in an uncontrolled POMDP.

We denote the set of all possible stochastic matrices with $|S^k|$ rows and $|S^k|$ columns by Ξ^k , and let $\Xi = (\Xi^1, \Xi^2, \dots, \Xi^K)$. Since \mathbf{P} is unknown to the player, at time t the player has an estimate of \mathbf{P} , denoted by $\hat{\mathbf{P}}_t \in \Xi$. For two vectors of transition probability matrices \mathbf{P} and $\tilde{\mathbf{P}}$, the distance between them is defined as $\|\mathbf{P} - \tilde{\mathbf{P}}\|_1 := \sum_{k=1}^K \|P^k - \tilde{P}^k\|_1$. Let X_t^k be the random variable representing the state of arm k at time t . Then, the random vector $\mathbf{X}_t = (X_t^1, X_t^2, \dots, X_t^K)$ represents the state of the system at time t .

The action space U of the player is equal to \mathcal{K} since it chooses an arm in \mathcal{K} at each time step, and the observation space Y of the player is equal to $\cup_{k=1}^K S^k$, since it observes the state of the arm it selects at each time step. Since the player can distinguish different arms, for simplicity we will assume $S^k \cap S^l = \emptyset$ for $k \neq l$, so that these states may be labeled distinctly. Let $u_t \in U$ be the arm selected by the player at time t , and $y_t \in Y$ be the state/reward observed by the player at time t . The history of the player at time t consists of all the actions and observations of the player by time t , which is denoted by $\mathbf{z}^t = (u_1, y_1, u_2, y_2, \dots, u_t, y_t)$. Let H^t denote the set of histories at time t . An algorithm $\alpha = (\alpha(1), \alpha(2), \dots)$ for the player, is a sequence of mappings from the set of histories to actions, i.e., $\alpha(t) : H^t \rightarrow U$. Since the history depends on the stochastic evolution of the arms, let U_t and Y_t be the random variables representing the action and the observation at time t , respectively. Let $Q_{\mathbf{P}}(y|u)$ be the sub-stochastic transition probability matrix such that

$$(Q_{\mathbf{P}}(y|u))_{xx'} = P_{\mathbf{P}}(\mathbf{X}_t = \mathbf{x}', Y_t = y | \mathbf{X}_{t-1} = \mathbf{x}, U_t = u),$$

where $P_{\mathbf{P}}(\cdot)$ denotes the conditional probability with respect to distribution \mathbf{P} . For URBP, $Q_{\mathbf{P}}(y|u)$ is the zero matrix for $y \notin S^u$, and for $y \in S^u$, only nonzero entries of $Q_{\mathbf{P}}(y|u)$ are the ones for which $x^u = y$.

Let Γ be the set of admissible policies, i.e., policies γ' for which $\gamma'(t) : H^t \rightarrow U$. Note that the set of admissible policies include the set of optimal policies which are computed by dynamic programming based on \mathbf{P} . Let ψ_0 be the initial belief of the player, which is a probability distribution over \mathcal{S} . Since we assume that the player knows nothing about the state of the system initially, ψ_0 can be taken as the uniform distribution over \mathcal{S} .

Let $E_{\psi, \gamma}^{\mathbf{P}}[\cdot]$ denote the expectation taken with respect to an algorithm or policy γ , initial state ψ , and the set of transition probability matrices \mathbf{P} . The performance of an algorithm α can be measured by

its strong regret, whose value at time t is the difference between performance of the algorithm and performance of the optimal policy by time t . It is given by

$$R^\alpha(T) = \sup_{\gamma' \in \Gamma} \left(E_{\psi_0, \gamma'}^{\mathbf{P}} \left[\sum_{t=1}^T r^{\gamma'(t)}(t) \right] \right) - E_{\psi_0, \alpha}^{\mathbf{P}} \left[\sum_{t=1}^T r^{\alpha(t)}(t) \right]. \quad (1)$$

It is easy to see that the time average reward of any algorithm with sublinear regret, i.e., regret $O(T^\rho)$, $\rho < 1$, converges to the time average reward of the optimal policy. For any algorithm with sublinear regret, its regret is a measure of its convergence rate to the average reward. In Section IV, we will give an algorithm whose regret grows logarithmically in time, which is the best possible rate of convergence.

B. Solutions to the Average Reward Optimality Equation

As mentioned earlier, if the transition probability matrices of the arms are known by the player, then the URBP becomes an optimization problem (POMDP) rather than a learning problem. In this section we discuss the solution approach to this optimization problem. This approach is then used in subsequent sections by the player in the learning context using *estimated* transition probability matrices.

A POMDP problem is often presented using the belief space (or information state), i.e., the set of probability distributions over the state space. For the URBP with the set of transition probability matrices \mathbf{P} , the belief space is given by

$$\Psi := \left\{ \psi : \psi^T \in \mathbb{R}^{|\mathcal{S}|}, \psi_{\mathbf{x}} \geq 0, \forall \mathbf{x} \in \mathcal{S}, \sum_{\mathbf{x} \in \mathcal{S}} \psi_{\mathbf{x}} = 1 \right\},$$

which is the unit simplex in $\mathbb{R}^{|\mathcal{S}|}$. Let ψ_t denote the belief of the player at time t . Then the probability that the player observes y given it selects arm u when the belief is ψ is given by

$$V_{\mathbf{P}}(\psi, y, u) := \psi Q_{\mathbf{P}}(y|u) \mathbf{1},$$

where $\mathbf{1}$ is the $|\mathcal{S}|$ dimensional column vector of 1s. Given arm u is chosen under belief state ψ and y is observed, the next belief state is

$$T_{\mathbf{P}}(\psi, y, u) := \frac{\psi Q_{\mathbf{P}}(y|u)}{V_{\mathbf{P}}(\psi, y, u)}.$$

The average reward optimality equation (AROE) is

$$g + h(\psi) = \max_{u \in \mathcal{U}} \left\{ \bar{r}(\psi, u) + \sum_{y \in \mathcal{S}^u} V_{\mathbf{P}}(\psi, y, u) h(T_{\mathbf{P}}(\psi, y, u)) \right\}, \quad (2)$$

where g is a constant and h is a function from $\Psi \rightarrow \mathbb{R}$,

$$\bar{r}(\psi, u) = (\psi \bullet r(u)) = \sum_{x^u \in S^u} x^u \phi_{u,x^u}(\psi)$$

is the expected reward of action u under belief ψ , $\phi_{u,x^u}(\psi)$ is the probability that arm u is in state x^u given belief ψ , $r(u) = (r(\mathbf{x}, u))_{\mathbf{x} \in S}$ and $r(\mathbf{x}, u) = x^u$ is the reward when arm u is chosen in state \mathbf{x} .

Assumption 1: $p_{ij}^k > 0, \forall k \in \mathcal{K}, i, j \in S^k$.

When Assumption 1 holds, the existence of a bounded, convex continuous solution to (2) is guaranteed. The set of Markov chains for which Assumption 1 holds is a subset of the class of aperiodic Markov chains. From any periodic or aperiodic Markov chain, we can obtain a Markov chain which belongs to this class by adding a small uniform perturbation to the state transition probabilities.

Let V denote the space of bounded real-valued functions on Ψ . Next, we define the undiscounted dynamic programming operator $F : V \rightarrow V$. Let $v \in V$, we have

$$(Fv)(\psi) = \max_{u \in U} \left\{ \bar{r}(\psi, u) + \sum_{y \in S^u} V_{\mathbf{P}}(\psi, y, u) v(T_{\mathbf{P}}(\psi, y, u)) \right\}. \quad (3)$$

In the following lemma, we give some of the properties of the solutions to the average reward optimality equation and the dynamic programming operator defined above.

Lemma 1: Let $h_+ = h - \inf_{\psi \in \Psi} (h(\psi))$, $h_- = h - \sup_{\psi \in \Psi} (h(\psi))$ and

$$h_{T,\mathbf{P}}(\psi) = \sup_{\gamma \in \Gamma} \left(E_{\psi,\gamma}^{\mathbf{P}} \left[\sum_{t=1}^T r^\gamma(t) \right] \right).$$

Given that Assumption 1 is true, the following holds:

S-1 Consider a sequence of functions v_0, v_1, v_2, \dots in V such that $v_0 = 0$, and $v_l = Fv_{l-1}$, $l = 1, 2, \dots$

This sequence converges uniformly to a convex continuous function v^* for which $Fv^* = v^* + g$ where g is a finite constant. In terms of (2), this result means that there exists a finite constant $g_{\mathbf{P}}$ and a bounded convex continuous function $h_{\mathbf{P}} : \Psi \rightarrow \mathbb{R}$ which is a solution to (2).

S-2 $h_{\mathbf{P}-}(\psi) \leq h_{T,\mathbf{P}}(\psi) - Tg_{\mathbf{P}} \leq h_{\mathbf{P}+}(\psi)$, $\forall \psi \in \Psi$.

S-3 $h_{T,\mathbf{P}}(\psi) = Tg_{\mathbf{P}} + h_{\mathbf{P}}(\psi) + O(1)$ as $T \rightarrow \infty$.

Proof: Sufficient conditions for the existence of a bounded convex continuous solution to the AROE are investigated in [1]. According to Theorem 4 of [1], if reachability and detectability conditions are satisfied then S-1 holds. Below, we directly prove that reachability condition in [1] is satisfied. To prove that detectability condition is satisfied, we show another condition, i.e., subrectangular substochastic

matrices, holds which implies the detectability condition.

We note that $P(\mathbf{X}_{t+1} = \mathbf{x}' | \mathbf{X}_t = \mathbf{x}) > 0, \forall \mathbf{x}, \mathbf{x}' \in \mathcal{S}$ since by Assumption 1, $p_{ij}^k > 0 \forall i, j \in \mathcal{S}^k, \forall k \in \mathcal{K}$.

Condition 1: (Reachability) There is a $\rho < 1$ and an integer ξ such that for all $\mathbf{x} \in \mathcal{S}$

$$\sup_{\gamma \in \Gamma} \max_{0 \leq t \leq \xi} P(\mathbf{X}_t = \mathbf{x} | \psi_0) \geq 1 - \rho, \quad \forall \psi_0 \in \Psi.$$

Set $\rho = 1 - \min_{\mathbf{x}, \mathbf{x}'} P(\mathbf{X}_{t+1} = \mathbf{x}' | \mathbf{X}_t = \mathbf{x}), \xi = 1$. Since the system is uncontrolled, state transitions are independent of the arm selected by the player. Therefore,

$$\begin{aligned} \sup_{\gamma \in \Gamma} P(\mathbf{X}_1 = \mathbf{x} | \psi_0) &= P(\mathbf{X}_1 = \mathbf{x} | \psi_0) \\ &\geq \min_{\mathbf{x}, \mathbf{x}'} P(\mathbf{X}_{t+1} = \mathbf{x}' | \mathbf{X}_t = \mathbf{x}) = 1 - \rho. \end{aligned}$$

Condition 2: (Subrectangular matrices) For any substochastic matrix $Q(y|u), y \in Y, u \in U$, and for any $i, i', j, j' \in \mathcal{S}$,

$$(Q(y|u))_{ij} > 0 \text{ and } (Q(y|u))_{i'j'} > 0 \Rightarrow (Q(y|u))_{ij'} > 0 \text{ and } (Q(y|u))_{i'j} > 0.$$

$Q(y|u)$ is subrectangular for $y \notin S^u$ since it is the zero matrix. For $y \in S^u$ all entries of $Q(y|u)$ is positive since $P(\mathbf{X}_{t+1} = \mathbf{x}' | \mathbf{X}_t = \mathbf{x}) > 0, \forall \mathbf{x}, \mathbf{x}' \in \mathcal{S}$.

S-2 holds by Lemma 1 in [1], and S-3 is a consequence of S-2 and the boundedness property in S-1. ■

C. Countable Representation of the Information State

The belief space is uncountable. Since the problem we consider is a learning problem, it is natural to assume that the player does not have an initial belief about the state of the system. However, in a learning context there is no loss of generality in adopting an initial belief formed by playing each arm at least once. Assume that the initial K steps are such that the player selects arm k at the k th step. Then the POMDP for the player can be written as a countable-state MDP. In this case a more succinct way of representing the information state at time t is given by

$$(\mathbf{s}_t, \boldsymbol{\tau}_t) = ((s_t^1, s_t^2 \dots, s_t^K), (\tau_t^1, \tau_t^2 \dots, \tau_t^K)),$$

where s_t^k and τ_t^k are the last observed state of arm k and how long ago (from t) the last observation of arm k was made, respectively. Note that the countable state MDP obtained this way is a subset of

the POMDP for the bandit problem in which the player can only be in one of the countably many points in the belief space Ψ at any time step t . Our approach is to exploit the continuity property of the AROE to bound the regret of the player. In order to do this we need to work with both methods of state representation. We thus make a distinction between ψ_t , which is a probability distribution over the state space, and (s_t, τ_t) which is a sufficient statistic for the player to calculate ψ_t when \mathbf{P} is given. Subsequently we will call $\psi \in \Psi$, the *belief* or *belief state*, and (s, τ) the *information state*¹.

The contribution of the initial K steps to the regret is at most Kr_{\max} , which we will subsequently ignore in our analysis. We will only analyze the time steps after this initialization, and set $t = 0$ upon the completion of the initialization phase. The initial information state of the player can be written as (s_0, τ_0) . Let \mathcal{C} be the set of all possible information states that the player can be in. Since the player selects a single arm at each time step, at any time $\tau^k = 1$ for the last selected arm k .

The player can compute its belief state $\psi_t \in \Psi$ by using its transition probability estimates $\hat{\mathbf{P}}_t$ together with the information state (s_t, τ_t) . We let $\psi_{\mathbf{P}}(s_t, \tau_t)$ be the belief that corresponds to information state (s_t, τ_t) when the set of transition probability matrices is \mathbf{P} . The player knows the information state exactly, but it only has an estimate of the belief that corresponds to the information state, because it does not know the transition probabilities. The true belief computed with the knowledge of exact transition probabilities and information state at time t is denoted by ψ_t , while the estimated belief computed with estimated transition probabilities and information state at time t is denoted by $\hat{\psi}_t$.

When the belief is ψ and the set of transition probability matrices is \mathbf{P} , the set of optimal actions which are the maximizers of (2) is denoted by $O(\psi; \mathbf{P})$. When the information state is (s_t, τ_t) , and the set of transition probability matrices is \mathbf{P} , we denote the set of optimal actions by $O((s, \tau); \mathbf{P}) := O(\psi_{\mathbf{P}}((s, \tau)); \mathbf{P})$.

IV. AVERAGE REWARD WITH ESTIMATED PROBABILITIES (AREP) ALGORITHM

In this section we propose the algorithm *Average Reward with Estimated Probabilities* (AREP) given in Figure 1, as a learning algorithm for the player. AREP consists of exploration and exploitation phases. In the exploration phase the player plays each arm for a certain amount of time to form estimates of the transition probabilities, while in the exploitation phase the player selects an arm according to the optimal policy based on the estimated transition probabilities. At each time step, the player decides if it is an exploration phase or an exploitation phase based on the accuracy of the transition probability estimates.

¹We note that in the POMDP literature these two terms are generally used interchangeably.

Average Reward with Estimated Probabilities (AREP)

```

1: Initialize:  $f(t)$  given for  $t \in \{1, 2, \dots\}$ ,  $t = 1$ ,  $N^k = 0$ ,  $N_{i,j}^k = 0$ ,  $C_i^k = 0$ ,  $\forall k \in \mathcal{K}$ ,  $i, j \in S^k$ . Play each arm
   once to set the initial information state  $(\mathbf{s}, \boldsymbol{\tau})_0$ . Pick  $\alpha(0)$  randomly.
2: while  $t \geq 1$  do
3:    $\bar{p}_{ij}^k = (I(N_{i,j}^k = 0) + N_{i,j}^k) / (|S^k|I(C_i^k = 0) + C_i^k)$ 
4:    $\hat{p}_{ij}^k = (\bar{p}_{ij}^k) / (\sum_{l \in S^k} \bar{p}_{il}^k)$ 
5:    $W = \{k \in \mathcal{K} : \text{there exists } i \in S^k \text{ such that } C_i^k \leq f(t)\}$ .
6:   if  $W \neq \emptyset$  then
7:     EXPLORE
8:     if  $\alpha(t-1) \in W$  then
9:        $\alpha(t) = \alpha(t-1)$ 
10:    else
11:      select  $\alpha(t) \in W$  arbitrarily
12:    end if
13:  else
14:    EXPLOIT
15:    solve  $\hat{g}_t + \hat{h}_t(\psi) = \max_{u \in U} \{\bar{r}(\psi, u) + \sum_{y \in S^u} V(\psi, y, u) \hat{h}_t(T_{\hat{P}_t}(\psi, y, u))\}$ ,  $\forall \psi \in \Psi$ .
16:    Let  $\hat{\psi}_t$  be the estimate of the belief at time  $t$  based on  $(\mathbf{s}_t, \boldsymbol{\tau}_t)$  and  $\hat{P}_t$ .
17:    compute the indices of all actions at  $\hat{\psi}_t$ :
18:     $\forall u \in U$ ,  $\mathcal{I}_t(\hat{\psi}_t, u) = \bar{r}(\hat{\psi}_t, u) + \sum_{y \in S^u} V(\hat{\psi}_t, y, u) \hat{h}_t(T_{\hat{P}_t}(\hat{\psi}_t, y, u))$ .
19:    Let  $u^*$  be the arm with the highest index (arbitrarily select one if there is more than one such arm).
20:     $\alpha(t) = u^*$ .
21:  end if
22:  Receive reward  $r^{\alpha(t)}(t)$ , i.e., state of  $\alpha(t)$  at  $t$ 
23:  Compute  $(\mathbf{s}_{t+1}, \boldsymbol{\tau}_{t+1})$ 
24:  if  $\alpha(t-1) = \alpha(t)$  then
25:    for  $i, j \in S^{\alpha(t)}$  do
26:      if State  $j$  is observed at  $t$ , state  $i$  is observed at  $t-1$  then
27:         $N_{i,j}^{\alpha(t)} = N_{i,j}^{\alpha(t)} + 1$ ,  $C_i^{\alpha(t)} = C_i^{\alpha(t)} + 1$ .
28:      end if
29:    end for
30:  end if
31:   $t := t + 1$ 
32: end while

```

Fig. 1. pseudocode for the Average Reward with Estimated Probabilities (AREP) algorithm

Let $N^k(t)$ be the number of times arm k is selected by time t , $N_{i,j}^k(t)$ be the number of times a transition from state i to state j of arm k is observed by the player by time t , and $C_i^k(t)$ be the number of times a transition from state i of arm k to any state of arm k is observed by time t . Clearly,

$$C_i^k(t) = \sum_{j \in S^k} N_{i,j}^k(t).$$

Let $f(t)$ be a non-negative, increasing function which sets a condition on the accuracy of the estimates. If $C_i^k(t) < f(t)$ for some $k \in \mathcal{K}$, $i \in S^k$, the player explores at time t . Otherwise, the player exploits at time t . In an exploration step, in order to update the estimate of p_{ij}^k , $j \in S^k$, the player keeps playing arm k until state i is observed, and then plays arm k one more time to observe the state following i . Note that in between the player can update other estimates depending on what states are observed. Then

the player forms the following sample mean estimates of the transition probabilities:

$$\bar{p}_{ij,t}^k := \frac{N_{i,j}^k(t)}{C_i^k(t)}, i, j \in S^k,$$

followed by the normalization

$$\hat{p}_{ij,t}^k := \frac{\bar{p}_{ij,t}^k}{\sum_{l \in S^k} \bar{p}_{il,t}^k}.$$

If AREP is in the exploitation phase at time t , then the player first computes $\hat{\psi}_t$, the estimated belief at time t , using the set of estimated transition probability matrices \hat{P}_t . Then, it solves the AROE using \hat{P}_t , to which the solution is given by \hat{g}_t and \hat{h}_t . For now we will ignore complexity issues and assume the player can compute the solution at every time step. More is discussed in the Conclusion. This solution is used to compute the indices (given on line 18 of AREP) as

$$\mathcal{I}_t(\hat{\psi}_t, u) = \bar{r}(\hat{\psi}_t, u) + \sum_{y \in S^u} V(\hat{\psi}_t, y, u) \hat{h}_t(T_{\hat{P}_t}(\hat{\psi}_t, y, u)),$$

for each action $u \in U$ at estimated belief $\hat{\psi}_t$. $\mathcal{I}_t(\hat{\psi}_t, u)$ represents the advantage of choosing action u starting from information state $\hat{\psi}_t$, i.e, the sum of gain and bias. After computing the indices for each action, the player selects the action with the highest index. In case of a tie, one of the actions with the highest index is randomly selected. Note that it is possible to update the state transition probabilities even in the exploitation phase given that the arms selected at times $t - 1$ and t are the same. Thus $C_i^k(t)$ may also increase in an exploitation phase, and the number of explorations may be smaller than the number of explorations needed in the worst case, where the estimates are only updated during exploration steps.

In subsequent sections we will bound the strong regret of AREP by bounding the number of times a suboptimal arm selection is made at any information state. Since there are infinitely many information states, our approach to bounding the sum of the number of suboptimal plays is to introduce a finite partition of the space of information states. We do this in the next section. For the remainder of the paper we will denote AREP by α .

V. FINITE PARTITIONS OF THE INFORMATION STATE

Note that even if the player knows the optimal policy as a function of the belief state for any time horizon T , it may not be able to play optimally because it does not know the exact belief ψ_t at time t . One way to ensure that the player plays optimally in this case is to show that there exists an $\epsilon > 0$ such that if $\|\psi_t - \hat{\psi}_t\|_1 < \epsilon$, the optimal actions in $\hat{\psi}_t$ belong to a subset of the set of optimal actions

in ψ_t . This is indeed the case, and we prove it by exploiting the continuity of the solution to (2) under Assumption 1.

We start by introducing a finite partition of the space of information states \mathcal{C} .

Definition 1: Let $\tau_{\text{tr}} > 0$ be an integer which denotes a threshold in time lag. This threshold is used to group all information states of an arm which has not been played for more than this threshold as a single group. Consider a vector $\mathbf{i} = (i^1, \dots, i^K)$ such that either $i^k = \tau_{\text{tr}}$ or $i^k = (s_i^k, \tau_i^k)$, $\tau_i^k < \tau_{\text{tr}}$, $s_i^k \in S^k$. For a finite τ_{tr} there are only a finite number of such vectors. Each vector defines a set of information states; it contains either a single information state or infinitely many information states of the arms for which $i^k = \tau_{\text{tr}}$. Together these vectors form a finite partition of \mathcal{C} , and we will call \mathbf{i} a *partition vector*. Let $\mathcal{G}_{\tau_{\text{tr}}}$ denote the partition formed by τ_{tr} , and let $\mathcal{M}(\mathbf{i}) := \{k : i^k = \tau_{\text{tr}}\}$ be the set of arms that are played at least τ_{tr} time steps ago, while $\overline{\mathcal{M}}(\mathbf{i}) := \mathcal{K} - \mathcal{M}(\mathbf{i})$. Vector \mathbf{i} represents the following set in the partition $\mathcal{G}_{\tau_{\text{tr}}}$:

$$G_{\mathbf{i}} = \{(\mathbf{s}, \boldsymbol{\tau}) \in \mathcal{C} : s^k = s_i^k, \tau^k = \tau_i^k, \forall k \in \overline{\mathcal{M}}(\mathbf{i}), s^k \in S^k, \tau^k \geq \tau_{\text{tr}}, \forall k \in \mathcal{M}(\mathbf{i})\}. \quad (4)$$

Let $A(\tau_{\text{tr}})$ be the number of sets in partition $\mathcal{G}_{\tau_{\text{tr}}}$. Re-index the sets in $\mathcal{G}_{\tau_{\text{tr}}}$ as $G_1, G_2, \dots, G_{A(\tau_{\text{tr}})}$. For a set $G_l \in \mathcal{G}_{\tau_{\text{tr}}}$, given a set of transition probability matrices \mathbf{P} , we define its center as follows. If G_l only contains a single information state, then the belief corresponding to that information state is the center of G_l . If G_l contains infinitely many information states, then the center belief of G_l is the belief in which all arms for which $i^k = \tau_{\text{tr}}$ are in their stationary distribution based on \mathbf{P} . In both cases, the center belief of G_l is denoted by $\psi^*(G_l; \mathbf{P})$. Let \mathbf{s}_l be the information state corresponding to the center belief. Although $\psi^*(G_l; \mathbf{P})$ depends on \mathbf{P} , \mathbf{s}_l does not depend on \mathbf{P} . Since each arm is ergodic, when we map a set G_l with infinitely many information states to the belief space using $\psi_{\mathbf{P}}$, for any $\delta > 0$, only a finite number of information states in G_l will lie outside the radius- δ ball around the center belief.

Let $O^*(G_l; \mathbf{P})$ be the set of optimal actions at the center belief. Note that as τ_{tr} increases, the number of sets with infinitely many elements increases, as does the number of sets with a single information state. The points in the belief space corresponding to these sets are shown in Figure 2. Below is an example of of such a finite partition of \mathcal{C} with $\tau_{\text{tr}} = 3$ when $K = 2$.

Example 1: Let $K = 2$, $S^1 = \{0, 2\}$, $S^2 = \{1\}$ and $\tau_{\text{tr}} = 3$. For convenience we will rewrite




Fig. 2. Partition of \mathcal{C} on Ψ based on \mathbf{P} and τ_{tr} . G_l is a set with a single information state and $G_{l'}$ is a set with infinitely many information states.

$(\mathbf{s}, \boldsymbol{\tau}) = ((s^1, \tau^1), (s^2, \tau^2))$. Then the partition formed by τ_{tr} , i.e., $\mathcal{G}_{\tau_{tr}}$ contains the following sets:

$$\begin{aligned}
 G_1 &= \{((0, 1), (1, 2))\}, & G_2 &= \{((2, 1), (1, 2))\}, \\
 G_3 &= \{((0, 2), (1, 1))\}, & G_4 &= \{((2, 2), (1, 1))\}, \\
 G_5 &= \{((0, 1), (1, 3)), ((0, 1), (1, 4)), \dots\}, \\
 G_6 &= \{((2, 1), (1, 3)), ((2, 1), (1, 4)), \dots\}, \\
 G_7 &= \{((0, 3), (1, 1)), ((2, 3), (1, 1)), ((0, 4), (1, 1)), ((2, 4), (1, 1)), \dots\}
 \end{aligned}$$

Next we define extensions of the sets G_l on the belief space. For a set $B \in \Psi$ let $B(\epsilon)$ be the ϵ -extension of that set, i.e.,

$$B(\epsilon) = \{\psi \in \Psi : \psi \in B \text{ or } d_1(\psi, B) < \epsilon\},$$

where $d_1(\psi, B)$ is the minimum l_1 distance between ψ and any element of B . The ϵ -extension of $G_l \in \mathcal{G}_{\tau_{tr}}$ corresponding to \mathbf{P} is the ϵ -extension of the convex-hull of the points $\psi_{\mathbf{P}}(\mathbf{s}, \boldsymbol{\tau})$ such that $(\mathbf{s}, \boldsymbol{\tau}) \in G_l$. Let $J_{l,\epsilon}$ denote the ϵ -extension of G_l . Examples of $J_{l,\epsilon}$ on the belief space are given in Figure V.

Let the diameter of a set B be the maximum distance between any two elements of that set. We note that when τ_{tr} increases, the diameter of the convex-hull of the points of an infinite-set in $\mathcal{G}_{\tau_{tr}}$ decreases. The following lemma shows that when τ_{tr} is chosen large enough, there exists $\epsilon > 0$ such for all $G_l \in \mathcal{G}_{\tau_{tr}}$, we have non-overlapping ϵ -extensions in which only a subset of the actions in $O^*(G_l; \mathbf{P})$ is optimal.




Fig. 3. ϵ -extensions of the sets in \mathcal{G}_{τ_r} on the belief space.

Lemma 2: For any \mathbf{P} , $\exists \tau_r > 0$ and $\epsilon > 0$ such that for all $G_l \in \mathcal{G}_{\tau_r}$, its ϵ -extension $J_{l,\epsilon}$ has the following properties:

- i For any $\psi \in J_{l,\epsilon}$, $O(\psi; \mathbf{P}) \subset O^*(G_l; \mathbf{P})$.
- ii For $l \neq l'$, $J_{l,\epsilon} \cap J_{l',\epsilon} = \emptyset$.

Proof: For $G_l \in \mathcal{G}_{\tau_r}$ consider its center $\psi^*(G_l; \mathbf{P})$. For any $\psi \in \Psi$ the suboptimality gap is defined as

$$\Delta(\psi, \mathbf{P}) = \max_{u \in U} \left\{ \bar{r}(\psi, u) + \sum_{y \in S^u} V_{\mathbf{P}}(\psi, y, u) h(T_{\mathbf{P}}(\psi, y, u)) \right\} - \max_{u \in U - O(\psi; \mathbf{P})} \left\{ \bar{r}(\psi, u) + \sum_{y \in S^u} V_{\mathbf{P}}(\psi, y, u) h(T_{\mathbf{P}}(\psi, y, u)) \right\}. \quad (5)$$

Since r, h, V and T are continuous in ψ , we can find an $\epsilon > 0$ such that for any $\psi \in B_{2\epsilon}(\psi^*(G_l; \mathbf{P}))$ and for all $u \in U$,

$$\left| \bar{r}(\psi^*(G_l; \mathbf{P}), u) + \sum_{y \in S^u} V_{\mathbf{P}}(\psi^*(G_l; \mathbf{P}), y, u) h(T_{\mathbf{P}}(\psi^*(G_l; \mathbf{P}), y, u)) - \bar{r}(\psi, u) + \sum_{y \in S^u} V_{\mathbf{P}}(\psi, y, u) h(T_{\mathbf{P}}(\psi, y, u)) \right| < \Delta(\psi^*(G_l; \mathbf{P}), \mathbf{P})/2, \quad (6)$$

and $B_{2\epsilon}(\psi^*(G_l; \mathbf{P})) \cap B_{2\epsilon}(\psi^*(G_{l'}; \mathbf{P})) = \emptyset$ for $l \neq l'$. Therefore, any action u which is not in $O^*(G_l; \mathbf{P})$ cannot be optimal for any $\psi \in B_{2\epsilon}(\psi^*(G_l; \mathbf{P}))$. Since the diameter of the convex-hull of the sets that

contains infinitely many information states decreases with τ_{tr} , there exists $\tau_{\text{tr}} > 0$ such that for any $G_l \in \mathcal{G}_{\tau_{\text{tr}}}$, the diameter of the convex-hull $J_{l,0}$ is less than ϵ . Let τ_{tr} be the smallest integer such that this holds. Then, the ϵ -extension of the convex hull $J_{l,\epsilon}$ is included in the ball $B_{2\epsilon}(\psi^*(G_l; \mathbf{P}))$ for all $G_l \in \mathcal{G}_{\tau_{\text{tr}}}$. This concludes the proof. ■

Remark 1: According to Lemma 2, although we can find an ϵ -extension in which a subset of $O^*(G_l; \mathbf{P})$ is optimal for any $\psi, \psi' \in J_{l,\epsilon}$, the set of optimal actions for ψ may be different from the set of optimal actions for ψ' . Note that the player's estimated belief $\hat{\psi}_t$ is different from the true belief ψ_t . If no matter how close $\hat{\psi}_t$ is to ψ_t , their respective sets of optimal actions are different, then the player can make a suboptimal decision even if it knows the optimal policy. Thus in this case the performance loss of the player, which can be bounded by the number of suboptimal decisions, may grow linearly over time. This turns out to be a major challenge. In this paper, we present two different approaches that lead to performance loss (regret) growing logarithmically in time. The first approach is based on an assumption about the structure of the optimal policy, while the second approach is to construct an algorithm that will almost always choose near-optimal actions, whose sub-optimality can be controlled by a function of the time horizon T . We shall take the first approach below and the second approach in Section VIII.

Assumption 2: There exists $\tau_{\text{tr}} \in \mathbb{N}$ such that for any $G_l \in \mathcal{G}_{\tau_{\text{tr}}}$, there exists $\epsilon > 0$ such that the same subset of $O^*(G_l; \mathbf{P})$ is optimal for any $\psi \in J_{l,\epsilon} - \psi^*(G_l; \mathbf{P})$.

When this assumption is correct, if ψ_t and $\hat{\psi}_t$ are sufficiently close to each other, then the player will always chose an optimal arm. Assume that this assumption is false. Consider the *stationary* information states for which $\tau^k = \infty$ for some arm k . Then for any $\tau_{\text{tr}} > 0$, there exists a set $G_l \in \mathcal{G}_{\tau_{\text{tr}}}$ and a sequence of information states $(s, \tau)_n$, $n = 1, 2, \dots$, such that $\psi_{\mathbf{P}}((s, \tau)_n)$ converges to $\psi^*(G_l; \mathbf{P})$ but there exists infinitely many n 's for which $O((s, \tau)_n; \mathbf{P}) \neq O((s, \tau)_{n+1}; \mathbf{P})$.

For simplicity of analysis, we focus on the following version of Assumption 2, although our results in Section VI will also hold when Assumption 2 is true.

Assumption 3: There exists $\tau_{\text{tr}} \in \mathbb{N}$ such that for any $G_l \in \mathcal{G}_{\tau_{\text{tr}}}$, a single action is optimal for $\psi^*(G_l; \mathbf{P})$.

In the next lemma, we show that when the set of transition probability matrices, i.e., \mathbf{P} , is drawn from a continuous distribution with a bounded density function (which is unknown to the player), before the play begins, then Assumption 3 will hold with probability one. In other words, the set of the set of transition probability matrices for which Assumption 3 does not hold is a measure zero subset of the set Ξ .

Lemma 3: Let $f_{\Xi}(\cdot)$ be the density function of the distribution from which \mathbf{P} is drawn as an instance of the bandit problem. Let $\sup_{\mathbf{P} \in \Xi} f_{\Xi}(\mathbf{P}) \leq f_{\max} < \infty$. Then, Assumption 3 holds with probability

1. In other words, when w is a realization of transition probability matrices, and \mathcal{A} be the event that Assumption 3 holds, then $P(w \in \mathcal{A}) = 1$.

Proof: If Assumption 3 does not hold for some $\mathbf{P} \in \Xi$, this means that there exists $G_l \in \mathcal{G}_{\tau_r}$ such that for some distinct arms u and u' we have

$$\bar{r}_{\mathbf{P}}(s_l, u) + \sum_{y \in S^u} V_{\mathbf{P}}(s_l, y, u) h_{\mathbf{P}}(T_{\mathbf{P}}(s_l, y, u)) = \bar{r}(s_l, u') + \sum_{y \in S^{u'}} V_{\mathbf{P}}(s_l, y, u') h_{\mathbf{P}}(T_{\mathbf{P}}(s_l, y, u')), \quad (7)$$

where $\bar{r}_{\mathbf{P}}(s_l, u)$, $V_{\mathbf{P}}(s_l, y, u)$ and $h_{\mathbf{P}}(T_{\mathbf{P}}(s_l, y, u))$ are equivalents of the expressions given in Section III-B for the case when we derive the optimal policy with respect to the countable information state instead of the belief state, where s_l is the center information state of G_l .

In order to have $P(w \in \mathcal{A}) < 1$, the integral of $f_{\Xi}(\cdot)$ over the subset of Ξ for which Assumption 3 does not hold should be positive. This means that there should be at least one \mathbf{P} for which Assumption 3 does not hold for the ϵ neighborhood of \mathbf{P} , for some $\epsilon > 0$. Next, we will show that for any $\epsilon > 0$, there exists \mathbf{P}' such that $\|\mathbf{P}' - \mathbf{P}\| \leq \epsilon$ and Assumption 3 holds for \mathbf{P}' . Let $\epsilon_{p,\min} = \min_{k \in \mathcal{K}, i, j \in S^k} p_{i,j}^k$ and $\epsilon_s = \min\{\epsilon/(2S_{\max}), \epsilon_{p,\min}/2\}$. For an arm k let \bar{x}_k be the state with the highest reward and x_k be state with the lowest reward. We form \mathbf{P}' from \mathbf{P} as follows: For all arms other than u and u' , state transition probabilities of \mathbf{P}' is the same as \mathbf{P} . For arm u , let $(p')_{j,\bar{x}_u}^u = (p)_{j,\bar{x}_u}^u - \epsilon_s$ and $(p')_{j,x_u}^u = (p)_{j,x_u}^u + \epsilon_s$ for all $j \in S^u$. For arm u' let $(p')_{j,\bar{x}_{u'}}^{u'} = (p)_{j,\bar{x}_{u'}}^{u'} - \epsilon_s$ and $(p')_{j,x_{u'}}^{u'} = (p)_{j,x_{u'}}^{u'} + \epsilon_s$ for all $j \in S^{u'}$. By construction we have $\|\mathbf{P}' - \mathbf{P}\| \leq \epsilon$.

Due to the special assignment of probabilities in \mathbf{P}' , for any information state (s, τ) , we have $\bar{r}_{\mathbf{P}'}((s, \tau), u) > \bar{r}_{\mathbf{P}}((s, \tau), u)$ and $\bar{r}_{\mathbf{P}'}((s, \tau), u') < \bar{r}_{\mathbf{P}}((s, \tau), u')$. Since arms evolve independently of each other, the continuation value of choosing arm u in information state (s, τ) under \mathbf{P}' is greater than or equal to the continuation value of choosing arm u under \mathbf{P} , while the continuation value of choosing arm u' in information state (s, τ) under \mathbf{P}' is less than or equal to the continuation value of choosing arm u' under \mathbf{P} . These and (7) together implies that

$$\bar{r}_{\mathbf{P}'}(s_l, u) + \sum_{y \in S^u} V_{\mathbf{P}'}(s_l, y, u) h_{\mathbf{P}'}(T_{\mathbf{P}'}(s_l, y, u)) > \bar{r}(s_l, u') + \sum_{y \in S^{u'}} V_{\mathbf{P}'}(s_l, y, u') h_{\mathbf{P}'}(T_{\mathbf{P}'}(s_l, y, u')).$$

Note that for \mathbf{P}' , although the tie between arms u and u' is broken in favor of arm u for center belief s_l , there can be another center belief $s_{l'}$ of $G_{l'} \in \mathcal{G}_{\tau_r}$ for which two different arms a and a' became optimal in \mathbf{P}' , while only one arm is optimal for $s_{l'}$ in \mathbf{P} . If such a thing happens, then we can define a new transition probability matrix \mathbf{P}'' from \mathbf{P} by subtracting and adding $\epsilon_s/2$ similar to the construction of \mathbf{P}' . This will both break the ties between u , u' and a , a' in favor of u and a . Since the number of

arms and the number of center beliefs in \mathcal{G}_{τ_r} is finite, after repeating this procedure for a finite number of times we will find a $\tilde{\mathbf{P}}$ for which none of the center beliefs have more than one optimal arm such that $\|\tilde{\mathbf{P}} - \mathbf{P}\| \leq \epsilon$. ■

Lemma 3 implies that even though there might exist some \mathbf{P} for which Assumption 3 does not hold, for a perturbation of \mathbf{P} Assumption 3 will hold. Assumption 3 may not hold for some symmetric setting such as the case when there are identical arms. We illustrate this in the next example.

Example 2: Consider 3 arms, with $S^k = \{0, 1\}$ and $p_{01}^k := p_{01}$, $p_{10}^k := p_{10}$ and $p_{01} + p_{10} < 1$ for $k = \{1, 2, 3\}$. Under this setting, it is shown in [11] that the following myopic policy is optimal: At each time step, select the arm with the highest probability of being in state 1. Consider the center belief $(\pi^1, (p_{00}, p_{01}), \pi^3)$ which occurs right after arm 2 is played and state 0 is observed. Since the arms are symmetric we have $\pi^1 = \pi^3$. Since the myopic policy is optimal it would be optimal to select either of arm 1 and 3 in the above center belief. Therefore Assumption 3 does not hold for this case.

Remark 2: Although Assumption 3 does not hold in the example, note however the center belief $(\pi^1, (p_{00}, p_{01}), \pi^3)$ can only be reached under the following condition: (1) if arms 1 and 3 start with the same initial belief and are never played, or (2) arms 1 and 3 have not been played for infinitely long. This is because only one arm can be played at a time, which results in asynchronous update of the belief states once an arm is played. Therefore in practice this center belief will never be reached in finite time. Therefore, in reality AREP never computes the values of its indices at this center belief. But AREP might be required to compute the value of its indices near this center belief.

Corollary 1: Let $\tau_r \in \mathbb{N}$ be the minimum integer such that Assumption 3 holds. Then, there exists $\bar{\epsilon} > 0$, depending on τ_r , such that for all $\epsilon \leq \bar{\epsilon}$, and any $\psi \in \mathcal{J}_{l, \epsilon}$, a single action is optimal.

Proof: This result follows from Assumption 3 and Lemma 2. ■

Remark 3: Although we do not know of a way to check if Assumption 3 holds given a set of transition probability matrices \mathbf{P} , we conjecture that it holds for a large set of \mathbf{P} s for the following reason. This is because the player's selection does not affect state transitions of the arms; it only affects the player's information state. Moreover, each arm evolves independently from each other. If \mathbf{P} is arbitrarily selected from Ξ , and the state rewards r_x^k , $x \in S^k$ are arbitrarily selected from $[0, r_{\max}]$, then at any information state $(s, \tau) \in \mathcal{C}$, the probability that the reward distribution of two arms are the same will be zero. We therefore claim that Assumption 3 holds with probability one if the arm rewards and \mathbf{P} are chosen from the uniform distribution on $\Psi \times [0, r_{\max}]$. In other words, the set of arm rewards and transition probabilities for which Assumption 3 does not hold is a measure zero subset of $\Psi \times [0, r_{\max}]$.

VI. ANALYSIS OF THE REGRET OF AREP

In this section we show that when \mathbf{P} is such that Assumptions 1 and 3 hold and when using AREP with $f(t) = L \log t$ with L sufficiently large (i.e., $L \geq C(\mathbf{P})$, a constant dependent on \mathbf{P}), the regret due to explorations is logarithmic in time, while the regret due to all other terms are finite, independent of t . Note that since the player does not know \mathbf{P} , it cannot know how large it should chose L . For simplicity we assume that the player starts with an L that is large enough without knowing $C(\mathbf{P})$. We also prove a near-logarithmic regret result in Section VII when the player sets $f(t) = L(t) \log t$, where $L(t)$ is a positive increasing function over time such that $\lim_{t \rightarrow \infty} L(t) = \infty$.

In what follows, we first provide an upper bound on the regret as the summation of a number of components. We then proceed to bound these individual terms separately.

A. An Upper Bound on the Regret

For any admissible policy α , the regret with respect to the optimal T horizon policy is given in (1), which we restate below:

$$\sup_{\gamma \in \Gamma} \left(E_{\psi_0, \gamma}^{\mathbf{P}} \left[\sum_{t=1}^T r^{\gamma(t)}(t) \right] \right) - E_{\psi_0, \alpha}^{\mathbf{P}} \left[\sum_{t=1}^T r^{\alpha(t)}(t) \right].$$

We first derive the regret with respect to the optimal policy as a function of the number of suboptimal plays. Before proceeding, we introduce expressions to compactly represent the RHS of the AROE. Let

$$\begin{aligned} \mathcal{L}(\psi, u, h, \mathbf{P}) &:= \bar{r}(\psi, u) + (V(\psi, \cdot, u) \bullet h(T_{\mathbf{P}}(\psi, \cdot, u))) \\ \mathcal{L}^*(\psi, \mathbf{P}) &:= \max_{u \in U} \mathcal{L}(\psi, u, h_{\mathbf{P}}, \mathbf{P}) \\ \Delta(\psi, u; \mathbf{P}) &:= \mathcal{L}^*(\psi, \mathbf{P}) - \mathcal{L}(\psi, u, h_{\mathbf{P}}, \mathbf{P}), \end{aligned} \tag{8}$$

where the last one denotes the degree of suboptimality of action u at information state ψ when the set of transition probability matrices is \mathbf{P} .

From Proposition 1 in [23], we have for all $\gamma \in \Gamma$

$$R_{(\psi_0; \mathbf{P})}^{\gamma}(T) = \sum_{t=1}^T E_{\psi_0, \gamma}^{\mathbf{P}}[\Delta(\psi_t, U_t; \mathbf{P})] + \bar{C}_{\mathbf{P}}, \tag{9}$$

for some constant $\bar{C}_{\mathbf{P}}$ dependent on \mathbf{P} , and U_t is the random variable denoting the arm selected by the player at time t which depends on the policy used. We have used the subscript $(\psi_0; \mathbf{P})$ to denote the dependence of regret on the initial belief and the transition probabilities. We assume that initially

all the arms are sampled once thus the initial belief is $\psi_0 = \psi_{\mathbf{P}}((s_0, \tau_0))$. For the true set of transition probability matrices \mathbf{P} , let τ_{tr} and ϵ be such that Corollary 1 is true. Specifically, let τ_{tr} be the minimum over all possible values so that Corollary 1 is true, and ϵ be the maximum over all possible values given τ_{tr} so that Corollary 1 is true. Denote the ϵ -extension of the set $G_l \in \mathcal{G}_{\tau_{\text{tr}}}$ by $J_{l,\epsilon}$. Note that at any t , the belief $\psi_t \in J_{l,\epsilon}$ for some l . For simplicity of notation, when ϵ is clear from the context we will re-write $J_{l,\epsilon}$ as J_l . Let

$$\bar{\Delta}(J_l, u; \mathbf{P}) := \sup_{\psi \in J_l} \Delta(\psi, u; \mathbf{P}).$$

Note that if $U_t \in O(\psi_t; \mathbf{P})$ then $\Delta(\psi_t, U_t; \mathbf{P}) = 0$; otherwise $U_t \notin O(\psi_t; \mathbf{P})$, and then $\Delta(\psi_t, U_t; \mathbf{P}) \leq \bar{\Delta}(J_l, U_t; \mathbf{P})$ with probability one. Let

$$N_T(J_l, u) := \sum_{t=1}^T I(\psi_t \in J_l, U_t = u).$$

Lemma 4: For any admissible policy γ ,

$$R_{(\psi_0; \mathbf{P})}^{\gamma}(T) \leq \sum_{l=1}^{A(\tau_{\text{tr}})} \sum_{u \notin O(J_l; \mathbf{P})} E_{\psi_0, \gamma}^{\mathbf{P}}[N_T(J_l, u)] \bar{\Delta}(J_l, u; \mathbf{P}) + \bar{C}_{\mathbf{P}}.$$

Proof:

$$\begin{aligned} R_{(\psi_0; \mathbf{P})}^{\gamma}(T) &\leq \sum_{t=1}^T E_{\psi_0, \gamma}^{\mathbf{P}} \left[\sum_{l=1}^A \sum_{u \notin O(J_l; \mathbf{P})} I(\psi_t \in J_l, U_t = u) \bar{\Delta}(J_l, u; \mathbf{P}) \right] + \bar{C}_{\mathbf{P}} \\ &= \sum_{l=1}^A \sum_{u \notin O(J_l; \mathbf{P})} E_{\psi_0, \gamma}^{\mathbf{P}} \left[\sum_{t=1}^T I(\psi_t \in J_l, U_t = u) \right] \bar{\Delta}(J_l, u; \mathbf{P}) + \bar{C}_{\mathbf{P}} \\ &= \sum_{l=1}^A \sum_{u \notin O(J_l; \mathbf{P})} E_{\psi_0, \gamma}^{\mathbf{P}}[N_T(J_l, u)] \bar{\Delta}(J_l, u; \mathbf{P}) + \bar{C}_{\mathbf{P}}. \end{aligned}$$

■

Now consider AREP, which is denoted by α . We will upper bound $N_T(J_l, u)$ for suboptimal actions u by a sum of expressions which we will then bound individually. Let \mathcal{E}_t be the event that AREP is in an exploitation step at time t and $\mathcal{F}_t(\epsilon) := \left\{ \left\| \hat{h}_t - h_{\mathbf{P}} \right\|_{\infty} \leq \epsilon \right\}$. For an event \mathcal{F} , denote its complement by \mathcal{F}^c . Consider the following random variables which count the number of times a certain event has

happened by time T (the dependence on T is dropped from the notation for convenience).

$$\begin{aligned}
D_{1,1}(\epsilon, J_l, u) &:= \sum_{t=1}^T I\left(\hat{\psi}_t \in J_l, U_t = u, \mathcal{E}_t, \mathcal{F}_t(\epsilon)\right), \\
D_{1,2}(\epsilon) &:= \sum_{t=1}^T I(\mathcal{E}_t, \mathcal{F}_t^c(\epsilon)), \\
D_1(\epsilon, J_l, u) &:= D_{1,1}(\epsilon, J_l, u) + D_{1,2}(\epsilon), \\
D_{2,1}(\epsilon) &:= \sum_{t=1}^T I(\|\psi_t - \hat{\psi}_t\|_1 > \epsilon, \mathcal{E}_t), \\
D_{2,2}(\epsilon, J_l) &:= \sum_{t=1}^T I(\|\psi_t - \hat{\psi}_t\|_1 \leq \epsilon, \hat{\psi}_t \notin J_l, \psi_t \in J_l, \mathcal{E}_t), \\
D_2(\epsilon, J_l) &:= D_{2,1}(\epsilon) + D_{2,2}(\epsilon, J_l).
\end{aligned}$$

Lemma 5: For any \mathbf{P} satisfying Assumption 2, we have

$$E_{\psi_0, \gamma}^{\mathbf{P}}[N_T(J_l, u)] \leq E_{\psi_0, \gamma}^{\mathbf{P}}[D_1(\epsilon, J_l, u)] + E_{\psi_0, \gamma}^{\mathbf{P}}[D_2(\epsilon, J_l)] + E_{\psi_0, \gamma}^{\mathbf{P}}\left[\sum_{t=1}^T I(\mathcal{E}_t^c)\right]. \quad (10)$$

Proof:

$$\begin{aligned}
N_T(J_l, u) &= \sum_{t=1}^T (I(\psi_t \in J_l, U_t = u, \mathcal{E}_t) + I(\psi_t \in J_l, U_t = u, \mathcal{E}_t^c)) \\
&\leq \sum_{t=1}^T I(\psi_t \in J_l, \hat{\psi}_t \in J_l, U_t = u, \mathcal{E}_t) + \sum_{t=1}^T I(\psi_t \in J_l, \hat{\psi}_t \notin J_l, U_t = u, \mathcal{E}_t) \\
&\quad + \sum_{t=1}^T I(\mathcal{E}_t^c) \\
&\leq \sum_{t=1}^T I(\hat{\psi}_t \in J_l, U_t = u, \mathcal{E}_t) + \sum_{t=1}^T I(\psi_t \in J_l, \hat{\psi}_t \notin J_l, \mathcal{E}_t) + \sum_{t=1}^T I(\mathcal{E}_t^c) \\
&\leq D_{1,1}(\epsilon, J_l, u) + D_{1,2}(\epsilon) + D_{2,1}(\epsilon) + D_{2,2}(\epsilon, J_l) \\
&\quad + \sum_{t=1}^T I(\mathcal{E}_t^c).
\end{aligned}$$

The result follows from taking the expectation on both sides. ■

B. Bounding the Expected Number of Explorations

The following lemma bounds the number of explorations by time T .

Lemma 6:

$$E_{\psi_0, \alpha}^{\mathbf{P}} \left[\sum_{t=1}^T I(\mathcal{E}_t^c) \right] \leq \left(\sum_{k=1}^K |S^k| \right) L \log T (1 + T_{\max}), \quad (11)$$

where $T_{\max} = \max_{k \in \mathcal{K}, i, j \in S^k} E[T_{ij}^k] + 1$ and T_{ij}^k is the hitting time of state j of arm k starting from state i of arm k . Since all arms are ergodic $E[T_{ij}^k]$ is finite for all $k \in \mathcal{K}, i, j \in S^k$.

Proof: Assume that state i of arm k is *under-sampled*, i.e., $C_i^k(t) < L \log t$. Since arms are ergodic, if the player keeps playing arm k , the expected number of time steps until a transition out of state i is observed is at most $(1 + T_{\max})$. If by time T , transitions out of state i of arm k is observed at least $L \log T$ times, for all states i of all arms k , then the player will not explore at time T . Therefore there can be at most $\sum_{k=1}^K \sum_{i \in S^k} L \log T$ such transitions by time T that take place in an exploration step. In the worst-case each of these transitions takes $(1 + T_{\max})$ expected time steps. ■

C. Bounding $E_{\psi_0, \alpha}^{\mathbf{P}}[D_1(\epsilon, J_l, u)]$ for a suboptimal action $u \notin O(J_l; \mathbf{P})$

We begin with the following lemma, based on the Chernoff-Hoeffding bound, that shows that the probability that an estimated transition probability is significantly different from the true transition probability given AREP is in an exploitation phase is very small.

Lemma 7: For any $\epsilon > 0$, for a player using AREP with constant $L \geq C_{\mathbf{P}}(\epsilon)$, we have

$$P \left(|\hat{p}_{ij,t}^k - p_{ij}^k| > \epsilon, \mathcal{E}_t \right) := P \left(\{|\hat{p}_{ij,t}^k - p_{ij}^k| > \epsilon\} \cap \mathcal{E}_t \right) \leq \frac{2S_{\max} + 2}{t^2},$$

for all $t > 0$, $i, j \in S^k$, $k \in \mathcal{K}$, where $C_{\mathbf{P}}(\epsilon)$ is a constant that depends on \mathbf{P} and ϵ .

Proof: See Appendix B. ■

We next bound $E_{\psi_0, \alpha}^{\mathbf{P}}[D_{1,1}(\epsilon, J_l, u)]$ for any suboptimal u . Let

$$\underline{\Delta}(J_l; \mathbf{P}) := \min_{\psi \in J_l, u \notin O(J_l; \mathbf{P})} \Delta(\psi, u; \mathbf{P}).$$

By Corollary 1, $\underline{\Delta}(J_l; \mathbf{P}) > 0$ for all $l = 1, \dots, A(\tau_{\text{tr}})$. Let

$$\underline{\Delta} := \min_{l=1, \dots, A(\tau_{\text{tr}})} \underline{\Delta}(J_l; \mathbf{P}).$$

In the following lemma we show that when the transition probability estimates are sufficiently accurate and the estimated solution to the AROE is sufficiently close to the true solution, a suboptimal action cannot be chosen by the player.

Lemma 8: Let $\delta_e > 0$ be the greatest real number such that

$$\|\hat{\mathbf{P}}_t - \mathbf{P}\|_1 < \delta_e \Rightarrow \left| \mathcal{L}(\psi, u, h_{\mathbf{P}}, \mathbf{P}) - \mathcal{L}(\psi, u, h_{\mathbf{P}}, \hat{\mathbf{P}}_t) \right| \leq \underline{\Delta}/4,$$

for all $\psi \in \Psi$. Such δ_e exists because $T_{\mathbf{P}}(\psi, y, u)$ is continuous in \mathbf{P} , and $h_{\mathbf{P}}(\psi)$ is continuous in ψ . Then for a player using AREP with $L \geq C_{\mathbf{P}}(\delta_e/(KS_{\max}^2))$, for any suboptimal action $u \notin O(J_l; \mathbf{P})$, we have

$$E_{\psi_0, \alpha}^{\mathbf{P}}[D_{1,1}(\epsilon, J_l, u)] \leq 2KS_{\max}^2(S_{\max} + 1)\beta,$$

for $\epsilon < \underline{\Delta}/4$, where $\beta = \sum_{t=1}^{\infty} 1/t^2$ and $C_{\mathbf{P}}(\cdot)$ is the constant given in Lemma 7.

Proof: See Appendix C. ■

Next we bound $E_{\psi_0, \alpha}^{\mathbf{P}}[D_{1,2}(\epsilon)]$. To do this we introduce the following lemma which implies that when the estimated transition probabilities get close to the true values, the difference between the solutions to the AROE based on the estimated and true values diminishes.

Lemma 9: For any $\epsilon > 0$, there exists $\varsigma > 0$ depending on ϵ such that if $\|P^k - \hat{P}^k\|_1 < \varsigma, \forall k \in \mathcal{K}$ then $\|h_{\mathbf{P}} - h_{\hat{\mathbf{P}}}\|_{\infty} < \epsilon$.

Proof: See Appendix D. ■

The following lemma bounds $E_{\psi_0, \alpha}^{\mathbf{P}}[D_{1,2}(\epsilon)]$.

Lemma 10: For any $\epsilon > 0$, let $\varsigma > 0$ be such that Lemma 9 holds. Then for a player using AREP with $L \geq C_{\mathbf{P}}(\varsigma/S_{\max}^2)$, we have

$$E_{\psi_0, \alpha}^{\mathbf{P}}[D_{1,2}(\epsilon)] \leq 2KS_{\max}^2(S_{\max} + 1)\beta, \quad (12)$$

where $C_{\mathbf{P}}(\cdot)$ is the constant given in Lemma 7.

Proof: See Appendix E. ■

D. Bounding $E_{\psi_0, \alpha}^{\mathbf{P}}[D_2(\epsilon, J_l)]$

Lemma 11: For a player using AREP with exploration constant $L \geq C_{\mathbf{P}}(\epsilon/(KS_{\max}^2|S^1| \dots |S^K|C_1(\mathbf{P})))$, we have

$$E_{\psi_0, \alpha}^{\mathbf{P}}[D_{2,1}(\epsilon)] \leq 2KS_{\max}^2(S_{\max} + 1)\beta, \quad (13)$$

where $C_{\mathbf{P}}(\cdot)$ is the constant given in Lemma 7, $C_1(\mathbf{P}) = \max_{k \in \mathcal{K}} C_1(P^k, \infty)$ and $C_1(P^k, t)$ is a constant that can be found in Lemma 13 in Appendix A.

Proof: See Appendix F. ■

Next we will bound $E_{\psi_0, \alpha}^{\mathbf{P}}[D_{2,2}(\epsilon, J_l)]$.

Lemma 12: Let τ_{tr} be such that Assumption 3 holds. Then for $\epsilon < \bar{\epsilon}/2$, where $\bar{\epsilon}$ is given in Corollary 1, $E_{\psi_0, \alpha}^{\mathbf{P}}[D_{2,2}(\epsilon, J_l)] = 0, l = 1, \dots, A(\tau_{\text{tr}})$.

Proof: By Corollary 1, any $\psi_t \in J_l$ is at least $\bar{\epsilon}$ away from the boundary of J_l . Thus given $\hat{\psi}_t$ is at most ϵ away from ψ_t , it is at least $\bar{\epsilon}/2$ away from the boundary of J_l . ■

E. Logarithmic regret upper bound

Theorem 1: Assume that Assumptions 1 and 3 are true. Let τ_{tr} be the minimum threshold, and $\epsilon = \bar{\epsilon}$ be the number given in Corollary 1. Under these assumptions, for a player using AREP with L sufficiently large (depending on \mathbf{P} and ϵ), for any arm (action) $u \in U$ which is suboptimal for the belief vectors in J_l , we have

$$E_{\psi_0, \alpha}^{\mathbf{P}}[N_T(J_l, u)] \leq \left(\sum_{k=1}^K |S^k| \right) L \log T(1 + T_{\max}) + 6K S_{\max}^2 (S_{\max} + 1) \beta ,$$

for some $\delta > 0$ depending on L . Therefore,

$$\begin{aligned} R_{\psi_0; \mathbf{P}}^{\alpha}(T) &\leq \left(\left(\sum_{k=1}^K |S^k| \right) L \log T(1 + T_{\max}) + 6K S_{\max}^2 (S_{\max} + 1) \beta \right) \\ &\quad \times \sum_{l=1}^{A(\tau_{\text{tr}})} \sum_{u \notin O(J_l; \mathbf{P})} \bar{\Delta}(J_l, u; \mathbf{P}) + \bar{C}_{\mathbf{P}}. \end{aligned}$$

When the arm rewards are in $[0, 1]$, strong regret at time T given as $R_{\psi_0; \mathbf{P}}^{\alpha}(T)$ can also be upper bounded by

$$\left(\left(\sum_{k=1}^K |S^k| \right) L \log T(1 + T_{\max}) + 6K S_{\max}^2 (S_{\max} + 1) \beta \right) (KA(\tau_{\text{tr}})) + \bar{C}_{\mathbf{P}} .$$

Proof: The result follows from combining Lemmas 4, 6, 10 11 and 12. ■

Remark 4: Our regret bound depends on $A(\tau_{\text{tr}})$. However, the player does not need to know the value of τ_{tr} for which Corollary 1 is true. It only needs to choose L large enough so that the number of exploration steps is sufficient to ensure a bounded number of errors in exploitation steps.

VII. AREP WITH AN ADAPTIVE EXPLORATION FUNCTION

In this section, we consider an adaptive exploration function for AREP, by which the player can achieve near-logarithmic regret without knowing how large it should chose the exploration constant L , which

depends on \mathbf{P} . First note that the analysis in Section VI holds when AREP is run with a sufficiently large exploration constant $L \geq C(\mathbf{P})$ such that in each exploitation step the estimated transition probabilities $\hat{\mathbf{P}}$ is close enough to \mathbf{P} to guarantee that all regret terms in (10) is finite except that due to explorations. Practically, since the player does not know the transition probabilities initially, it may be unreasonable to assume that it can check if $L \geq C(\mathbf{P})$. One possible solution is to assume that the player knows a compact set $\tilde{\Xi} \subset \Xi$, the set of transition probability matrices where \mathbf{P} lies in. If this is the case, then it can compute $\tilde{C} = \max_{\tilde{\mathbf{P}} \in \tilde{\Xi}} C(\tilde{\mathbf{P}})$, and choose $L > \tilde{C}$.

In this section, we present another exploration function for AREP such that the player can achieve near-logarithmic regret even without knowing $C(\mathbf{P})$ or \tilde{C} . Let $f(t) = L(t) \log t$ where $L(t)$ is an increasing function such that $L(1) = 1$ and $\lim_{t \rightarrow \infty} L(t) = \infty$. The intuition behind this exploration function is that after some time T_0 , $L(t)$ will be large enough so that the estimated transition probabilities are sufficiently accurate, and the regret due to incorrect calculations becomes a constant independent of time.

Theorem 2: When \mathbf{P} is such that Assumptions 1 and 3 hold, if the player uses AREP with $f(t) = L(t) \log t$, for some increasing $L(t)$ such that $L(1) = 1$ and $\lim_{t \rightarrow \infty} L(t) = \infty$, then there exists $\tau_{\text{tr}}(\mathbf{P}) > 0$, $T_0(L, \mathbf{P}) > 0$ such that the strong regret is upper bounded by

$$\begin{aligned} R_{\psi_0; \mathbf{P}}^\alpha(T) &\leq r_{\max} \left(T_0(L, \mathbf{P}) + \left(\sum_{k=1}^K |S^k| \right) L(T) \log T (1 + T_{\max}) \right. \\ &\quad \left. + 6K S_{\max}^2 (S_{\max} + 1) \beta \left(\sum_{l=1}^{A(\tau_{\text{tr}})} \sum_{u \notin O(J_l; \mathbf{P})} \bar{\Delta}(J_l, u; \mathbf{P}) \right) \right) \\ &\leq r_{\max} \left(T_0(L, \mathbf{P}) + \left(\sum_{k=1}^K |S^k| \right) L(T) \log T (1 + T_{\max}) \right. \\ &\quad \left. + 6K S_{\max}^2 (S_{\max} + 1) \beta (\tau_{\text{tr}})^M \left(\sum_{k=1}^K |S^k| \right) \max_{l \in \{1, \dots, A(\tau_{\text{tr}})\}} \bar{\Delta}(J_l, u; \mathbf{P}) \right) + \bar{C}_{\mathbf{P}} . \end{aligned}$$

Proof: The regret up to $T_0(L, \mathbf{P})$ can be at most $r_{\max} T_0(L, \mathbf{P})$. After $T_0(L, \mathbf{P})$, since $L(t) \geq C(\mathbf{P})$, transition probabilities at exploitation steps sufficiently accurate so that all regret terms in (10) except the regret due to explorations is finite. Since time t is an exploration step whenever $C_i^k(t) < L(t) \log t$, the regret due to explorations is at most $r_{\max} \left(\sum_{k=1}^K |S^k| \right) L(T) \log T (1 + T_{\max})$. ■

Remark 5: There is a tradeoff between choosing a rapidly increasing L or a slowly increasing L . The regret of AREP up to time $T_0(L, \mathbf{P})$ is linear. Since $T_0(L, \mathbf{P})$ is a decreasing function in L , a rapidly increasing L will have better performance when the time horizon is small. However, in terms of asymptotic performance as $T \rightarrow \infty$, L should be a slowly diverging sequence. For example if $L = \log(\log t)$, then

the asymptotic regret will be $O(\log(\log t) \log t)$.

VIII. AREP WITH FINITE PARTITIONS

In this section we present a modified version of AREP and prove that it can achieve logarithmic regret without Assumption 2 or 3 if the player knows the time horizon T . We call this variant AREP with finite partitions (AREP-FP).

Basically, AREP-FP takes as input the threshold or mixing time τ_{tr} and then forms the $\mathcal{G}_{\tau_{\text{tr}}}$ partition of the set of information states \mathcal{C} . At each exploitation step (time t) AREP-FP solves the estimated AROE based on the transition probability estimate $\hat{\mathbf{P}}_t$. If the information state $(\mathbf{s}_t, \boldsymbol{\tau}_t) \in G_l$, the player arbitrarily picks an arm in $O^*(G_l; \hat{\mathbf{P}}_t)$, instead of picking an arm in $O(\psi_{\hat{\mathbf{P}}_t}((\mathbf{s}_t, \boldsymbol{\tau}_t)); \hat{\mathbf{P}}_t)$. If the arm selected by the player is indeed in $O(\psi_{\mathbf{P}}((\mathbf{s}_t, \boldsymbol{\tau}_t)); \mathbf{P})$, then it ends up playing optimally at that time step. Else if the selected arm is in $O^*(G_l; \mathbf{P})$ but not in $O(\psi_{\mathbf{P}}((\mathbf{s}_t, \boldsymbol{\tau}_t)); \mathbf{P})$ such that $(\mathbf{s}_t, \boldsymbol{\tau}_t) \in G_l$, then it plays near-optimally. Finally, it plays suboptimally if the selected arm is neither in $O(\psi_{\mathbf{P}}((\mathbf{s}_t, \boldsymbol{\tau}_t)); \mathbf{P})$ nor in $O^*(G_l; \mathbf{P})$. By Lemma 2 we know that when τ_{tr} is chosen sufficiently large, for any $G_l \in \mathcal{G}_{\tau_{\text{tr}}}$, and $(\mathbf{s}, \boldsymbol{\tau}) \in G_l$, $O(\psi_{\mathbf{P}}((\mathbf{s}, \boldsymbol{\tau})); \mathbf{P})$ is a subset of $O^*(G_l; \mathbf{P})$. Since the solution to the AROE is a continuous function, by choosing a sufficiently large τ_{tr} we can control the regret due to near-optimal actions. The regret due to suboptimal actions can be bounded in the same way as in Theorem 1. The following theorem gives a logarithmic upper bound on the regret of AREP-FP.

Theorem 3: When the true set of transition probabilities \mathbf{P} is such that Assumption 1 is true, for a player using AREP-FP with exploration constant L , and threshold τ_{tr} sufficiently large such that for any $(\mathbf{s}, \boldsymbol{\tau}) \in G_l$, $G_l \in \mathcal{G}_{\tau_{\text{tr}}}$, we have $|h_{\mathbf{P}}(\psi_{\mathbf{P}}((\mathbf{s}, \boldsymbol{\tau}))) - h_{\mathbf{P}}(\psi^*(G_l; \mathbf{P}))| < C/2T$, where $C > 0$ is a constant and T is the time horizon, the regret of AREP-FP is upper bounded by

$$C + (L \log T (1 + T_{\max}) + 6K S_{\max}^2 \beta (S_{\max} + 1)) \times \sum_{l=1}^{A(\tau_{\text{tr}})} \sum_{u \notin O(J_l; \mathbf{P})} \bar{\Delta}(J_l, u; \mathbf{P}) + \bar{C}_{\mathbf{P}},$$

for some $\delta > 0$ which depends on L and τ_{tr} .

Proof: The regret at time T is upper bounded by Lemma 4. Consider any t which is an exploitation step. Let l be such that $(\mathbf{s}_t, \boldsymbol{\tau}_t) \in G_l$. If the selected arm $\alpha(t) \in O(\psi_{\mathbf{P}}((\mathbf{s}_t, \boldsymbol{\tau}_t)); \mathbf{P})$, then an optimal decision is made at t , so the contribution to regret in time step t is zero. Next, we consider the case when $\alpha(t) \notin O(\psi_{\mathbf{P}}((\mathbf{s}_t, \boldsymbol{\tau}_t)); \mathbf{P})$. In this case there are two possibilities: either $\alpha(t) \in O^*(G_l; \mathbf{P})$ or not. We know that when $O^*(G_l; \hat{\mathbf{P}}_t) \subset O^*(G_l; \mathbf{P})$ we have $\alpha(t) \in O^*(G_l; \mathbf{P})$. Since $|h_{\mathbf{P}}(\psi_{\mathbf{P}}((\mathbf{s}, \boldsymbol{\tau}))) -$

$h_{\mathbf{P}}(\psi^*(G_l; \mathbf{P}))| < C/2T$ for all $(\mathbf{s}, \boldsymbol{\tau}) \in G_l$, we have by (8),

$$\Delta(\psi_t, \alpha(t); \mathbf{P}) = \mathcal{L}^*(\psi_t, \mathbf{P}) - \mathcal{L}(\psi_t, \alpha(t), h_{\mathbf{P}}, \mathbf{P}) \leq C/T. \quad (14)$$

Therefore, contribution of a near-optimal action to regret is at most C/T .

Finally, consider the case when $\alpha(t) \notin O^*(G_l; \mathbf{P})$. This implies that either the estimated belief $\hat{\psi}_t$ is not close enough to ψ_t or the estimated solution to the AROE, i.e., \hat{h}_t , is not close enough to $h_{\mathbf{P}}$. Due to the non-vanishing suboptimality gap at any belief vector $\psi^*(G_l; \mathbf{P})$, and since decisions of AREP-FP is only based on belief vectors corresponding to $(\mathbf{s}, \boldsymbol{\tau}) \in \mathcal{C}$, the regret due to suboptimal actions can be bounded by Theorem 1. We get the regret bound by combining all these results. \blacksquare

Note that the regret bound in Theorem 3 depends on τ_{tr} which further depends on T : τ_{tr} is chosen so that for every G_l in the partition created by τ_{tr} , the function $h_{\mathbf{P}}$ varies by at most $C/2T$. Clearly since $h_{\mathbf{P}}$ is a continuous function, the variation of $h_{\mathbf{P}}$ over G_l , i.e., the difference between the maximum and minimum values of $h_{\mathbf{P}}$ over G_l decreases with the diameter of G_l on the belief space. Note that there is a term in regret that depends linearly on the number of sets $A(\tau_{\text{tr}})$ in the partition generated by τ_{tr} , and $A(\tau_{\text{tr}})$ increases proportional to $(\tau_{\text{tr}})^K$. This tradeoff is not taken into account in Theorem 3. For example, if $(\tau_{\text{tr}})^K \geq T$ then the regret bound in Theorem 3 is of no use. Another approach is to jointly optimize the regret due to suboptimal and near-optimal actions by balancing the number of sets $A(\tau_{\text{tr}})$ and the variation of $h_{\mathbf{P}}$ on sets in $\mathcal{G}_{\tau_{\text{tr}}}$. For example, given $0 < \theta \leq 1$, we can find a $\tau_{\text{tr}}(\theta)$ such that for any $(\mathbf{s}, \boldsymbol{\tau}) \in G_l$, $G_l \in \mathcal{G}_{\tau_{\text{tr}}(\theta)}$, and $C > 0$, we have

$$|h_{\mathbf{P}}(\psi_{\mathbf{P}}((\mathbf{s}, \boldsymbol{\tau}))) - h_{\mathbf{P}}(\psi^*(G_l; \mathbf{P}))| < \frac{C}{2T^\theta}.$$

Then, the regret due to near-optimal decisions will be proportional to $CT^{1-\theta}$, and the regret due to suboptimal decision will be proportional to $(\tau_{\text{tr}}(\theta))^K$. Let $C = \sup_{\psi \in \Psi} h_{\mathbf{P}}(\psi) - \inf_{\psi \in \Psi} h_{\mathbf{P}}(\psi)$. Since $T^{1-\theta}$ is decreasing in θ and $\tau_{\text{tr}}(\theta)$ is increasing in θ , there exists $\theta \in [0, 1]$, such that $\theta = \arg \min_{\theta' \in [0, 1]} |T^{1-\theta'} - (\tau_{\text{tr}}(\theta'))^K|$, $|T^{1-\theta} - (\tau_{\text{tr}}(\theta))^K| \leq (\tau_{\text{tr}}(\theta) + 1)^K - (\tau_{\text{tr}}(\theta))^K$. If the optimal value of θ is in $(0, 1)$, then given θ , the player can balance the tradeoff, and achieve sublinear regret proportional to $T^{1-\theta}$. However, since the player does not know \mathbf{P} initially, it may not know the optimal value of θ . Online learning algorithms for the player to estimate the optimal value of θ is a future research direction.

IX. NUMERICAL EXAMPLES AND COMPARISON WITH EXISTING ONLINE LEARNING ALGORITHMS

In this section we compare the performance of AREP with existing online learning algorithms including ABA proposed in [22] and the online version of the myopic policy proposed in [11]. Although AREP is

computationally inefficient, we can solve the AROE approximately by using belief state quantization and relative value iteration. We modify AREP such that instead of solving the AROE at each exploitation step and computing a new estimated optimal policy, it only solves the AROE when the number of observations of transitions out of every state of every arm has increased by C_{inc} percent compared to the last time the AROE is solved. For example, if the minimum number of observations out of any state of any arm was N_{old} the last time AROE is solved, it will be solved again with the new estimated transition probabilities when the minimum number of observations out of any state of any arm exceeds $(1 + C_{\text{inc}}/100)N_{\text{old}}$. A similar idea is used in [31], in which the optimal policy is recomputed whenever the confidence interval for transition probabilities is halved.

We call time steps t in which the AROE is solved by AREP as the *computation steps*. The AROE is solved in the following way. A finite information-state Markov Decision Problem (MDP) is formed based on the estimated transition probabilities \hat{P}_t , at a computation step t . The state space of this finite MDP is

$$S^{\text{fin}} := \left\{ (s, \tau) : s^k \in S^k, \tau^k \in \{1, \dots, \tau_{\text{tr}}\}, \tau^k = 1 \text{ for only one } k \in \mathcal{K} \right\},$$

and the transition probability matrix is exactly the same as the transition probability matrix of the original information state MDP, except when arm k is in state (s_k, τ_m) and not played, its state remains (s_k, τ_m) instead of evolving to $(s_k, \tau_m + 1)$. Then the optimal average reward policy for this MDP is computed using relative value iteration. This is in some way similar to AREP-FP but there is no guarantee that this finite state approximation will yield sublinear regret. We call AREP used with this approximation method AREP with finite approximation (AREP-FA).

In this section we will show that AREP-FA achieves results that are much better than that of prior work in online learning, thus even though AREP and AREP-FP may be computationally inefficient, they are practically implementable using approximation algorithms.

Our numerical experiment considers the following setting: $M = 2$ (two arms) with state spaces $S^1 = S^2 = \{0, 1\}$, and rewards $r_0^1 = r_0^2 = 0$, $r_1^1 = r_1^2 = 1$. We consider three different sets of transition probability matrices C1, C2 and C3, given in Table I. In C1 arms are identical and bursty, i.e., $p_{01}^k + p_{11}^k < 1$, in C2 the first arm is bursty, while the second arm is not bursty, and in C3, neither arm is bursty. We consider the following types of algorithms.

When the transition probabilities are known, the myopic policy [11] chooses the arm with the highest one-step expected reward and is shown to be optimal in terms of the average reward when the arms

have only two states and are identical in terms of their state rewards and transition probabilities, and they are bursty. We define the online version of the myopic policy as follows. Similar to AREP, the online myopic policy keeps counters $N_{i,j}^k(t)$ and $C_i^k(t)$, $i, j \in S^k$, $k \in \mathcal{K}$, which are used to form the transition probability estimates defined in the same way as in Section IV. Whenever there is an arm k and state $i \in S^k$ such that $C_i^k(t) \leq L \log t$, the player explores arm k . Otherwise when $C_i^k(t) > L \log t$ for all arms and states, the player chooses the arm with the highest one-step expected reward based on the information state (s_t, τ_t) and $\hat{\mathbf{P}}_t$.

ABA is an online learning algorithm proposed in [22], which is a threshold variant of the approximation algorithm proposed in [10] to solve two-state non-identical, bursty restless bandit problems. This algorithm gives an index policy which computes indices for each arm separately based only on the transition probability estimates of that arm and chooses an arm to play by comparing their indices. Given $\epsilon > 0$ as an input parameter to the algorithm, it is guaranteed to achieve at least $1/(2 + \epsilon)$ of the average expected reward of the optimal policy for two-state, bursty arms. Unless otherwise stated we assume that the exploration constant $L = 30$ for AREP-FA, ABA and the online myopic policy.

We compute the total reward of AREP-FA, ABA, online myopic policy and the myopic policy with known transition probabilities for $T = 20000$. Note that the myopic policy only exploits since it knows the transition probabilities and is thus used as a benchmark. We average our results over 100 runs of the algorithms. Average rewards of these algorithms for different transition probabilities are shown in Table II, when AREP-FA is run with $C_{\text{inc}} = 5$ and ABA is run with $\epsilon = 0.02$.

We see that for C1, AREP-FA, ABA and online myopic perform roughly equally well and their total reward is very close to the myopic policy which is optimal for this case. It can be verified that Assumption 3 does not hold in this case, since arms are identical. The information state which violates this assumption is $(s, s)(\infty, \infty)$, $s \in \{0, 1\}$. This information state is the center information state of a $G_l \in G_{\tau_r}$. However, the information state $(s, s)(\infty, \infty)$, $s \in \{0, 1\}$ that violates Assumption 3 will never be reached by AREP-FA, thus this assumption will not be violated during the runtime of AREP-FA. Therefore, although theoretical regret bounds requires this assumption to hold, in this example, AREP-FA performs equally well with the online myopic policy. The difference between the total rewards of the myopic policy and online learning algorithms is due to the exploration steps which are necessary to learn the unknown transition probabilities.

For C2, since the arms are not identical the myopic policy is suboptimal. In addition, arm 2 is not bursty since $p_{01}^2 + p_{10}^2 = 1$. However, it is almost bursty since a slight decrease in p_{01}^2 or p_{10}^2 will make it bursty. We see that AREP-FA and ABA perform nearly equally well, while the online myopic policy has

	C1	C2	C3
p_{01}^1, p_{10}^1	0.2, 0.2	0.05, 0.05	0.9, 0.9
p_{01}^2, p_{10}^2	0.2, 0.2	0.5, 0.5	0.8, 0.5

TABLE I

STATE TRANSITION PROBABILITIES OF THE ARMS. TWO BURSTY ARMS IN C1, ONE BURSTY ARM IN C2, NO BURSTY ARM IN C3.

	AREP-FA ($C_{\text{inc}} = 5$)	ABA	Online myopic	Myopic
C1	12833	12784	12809	13008
C2	13262	13220	10935	10233
C3	14424	10076	14408	14690

TABLE II

TOTAL REWARDS OF AREP-FA ($C_{\text{inc}} = 5$), ABA, ONLINE MYOPIC POLICY AND MYOPIC POLICY (WITH KNOWN TRANSITION PROBABILITIES) FOR THREE DIFFERENT CASES C1-C3.

a total reward 18% less than AREP-FA and the myopic policy has a total reward 23% less than AREP-FA. From the results for C1 and C2 we also see that although ABA is proved to be an approximation algorithm, its actual performance is very close to AREP-FA. It works well for C2 because arm 2 is almost bursty.

For C3, we see that the online myopic policy performs almost as well as AREP-FA while ABA has a total reward 30% less than AREP-FA. It is expected that ABA should perform poorly since both arms are not bursty, but it came as a surprise that the online myopic policy performs so well. The myopic policy achieves the highest total reward for this case since it knows the transition probabilities and does not need to explore. We would like to note that for two state arms, bursty and identical arms assumption is a sufficient condition for the optimality of the myopic policy but it is not a necessary condition. There may be other cases in which the myopic policy can also be optimal, maybe including C3. For example, in [?], the authors derive sufficient conditions on the optimality of the myopic policy in both finite and infinite horizon discounted cases, both for bursty and non-bursty arms. However, the discussion on the conditions that guarantees optimality of the myopic policy is out of the scope of this paper. The results for C1-C3 shows that AREP-FA can potentially significantly outperform the other online learning algorithms.

In Table III we compare the time-averaged rewards of AREP-FA with the average reward of the best arm for C1, C2 and C3. Since online learning algorithms proposed in [15] and [32] have sublinear regret with respect to the best arm, and since they learn to play the best arm, their performances will be very poor compared to the performance of AREP-FA.

The number of time steps in which AROE is re-solved by AREP-FA by time T and the total reward of AREP-FA as a function of C_{inc} is given in Table IV for C2. Although in theory exploration and exploitation should continue indefinitely or until the final time T , in order to have sublinear regret with

	C1	C2	C3
AREP-FA ($C_{\text{inc}} = 5$)	0.642	0.663	0.721
The best arm	0.5	0.5	0.615

TABLE III
COMPARISON OF THE AVERAGE REWARD OF AREP-FA WITH THE AVERAGE EXPECTED REWARD OF THE BEST ARM FOR CASES C1-C3.

C_{inc}	5	10	50
Total reward	13262	13229	13264
Average number of computations of AROE	8.63	4.51	1.21

TABLE IV
THE TOTAL REWARD OF AREP-FA AND THE AVERAGE NUMBER OF COMPUTATIONS OF AROE AS A FUNCTION OF C_{inc} FOR C2.

respect to the optimal allocation, from Table IV we see that solving the AROE only when the number of observations that are used to estimate the transition probabilities multiply by some fixed amount works equally well in practice. We see that the performance does not improve as the number of computations increase. The reason for this can be that the structure of the optimal policy for the information states that are visited by AREP-FA in run time does not change when $\hat{\mathbf{P}}_t$ is slightly different from \mathbf{P} . By having sufficient exploration steps we guarantee that probability of $\hat{\mathbf{P}}_t$ being very different from \mathbf{P} is very small each time the AROE is solved.

The average CPU times required to run AREP-FA, ABA and online myopic policy in MATLAB are given in Table V for different values of C_{inc} for C1. The running time of ABA is the highest since it computes the approximately optimal policy at each exploitation step. From this table we see that although AREP is computationally intractable, it can be practically implemented in much the same way as ABA and the online myopic policy using approximation methods.

Finally, we give the performance of AREP-FA as a function of the exploration constant L . Recall that for logarithmic regret bound to hold, L should be chosen large enough for AREP. If no bound on L is known, L should be chosen increasingly over time such that our bounds will hold after some time. However, this increases regret since the number of explorations increases over time. The average reward of AREP-FA as a function of L is given in Table VI for C2. We see that choosing L very small has a much less negative impact than choosing L very large. This is due to the fact that the transition

	AREP-FA ($C_{\text{inc}} = 5$)	— (10)	— (50)	ABA	Online myopic policy
Avg run time (sec)	2.01	1.26	0.67	17.2	0.39

TABLE V
AVERAGE RUN TIME OF AREP-FA AS A FUNCTION OF C_{inc} , AND THE AVERAGE RUN TIMES OF ABA AND THE ONLINE MYOPIC POLICY IN MATLAB FOR C1.

Exploration constant L	0.3	3	30	300
Total reward ($C_{\text{inc}} = 5$)	12264	12919	13262	11463

TABLE VI
THE TOTAL REWARD OF AREP-FA AS A FUNCTION OF L FOR C2.

probabilities can still be learned over exploitation steps, and the loss of reward due to the suboptimal decisions made on exploitation steps as a result of the poorly estimated transition probabilities is not much greater than the loss incurred due to explorations. On the contrary, when L is very large, e.g., $L = 300$, there are 13341 exploration steps on average, thus losses from explorations is much more than the gains in exploitation steps.

X. CONCLUSION

We showed for an uncontrolled restless bandit problem there exist index policies with logarithmic regret uniformly in time with respect to the optimal finite horizon expected total reward policy. This result assumes that the player is able to solve the average reward optimization problem sufficiently often, which can be computationally costly. In practice, the player can use value iteration and belief state space discretization to obtain an approximate solution. Furthermore, the player can simply choose to use an explicit/structured policy (e.g., a greedy or one-step lookahead policy) that does not require much computational effort. In some instances such policies may be optimal or suboptimal with performance guarantees. Provided that the approximate solution holds similar continuity properties, then using AREP (substituting the part solving the AROE with an approximation or an explicit policy) results in a learning algorithm with logarithmic regret w.r.t the chosen substitute. We showed via numerical results that the real-time performance of the learning algorithms proposed in this paper is much better than the previous online learning algorithms, and our algorithms can be efficiently implemented using approximation methods.

APPENDIX A

RESULTS REGARDING DEVIATIONS OF ESTIMATED TRANSITION PROBABILITIES

Certain results from the large deviation theory are frequently used, e.g., to relate the accuracy of the player's transition probability estimates to its probability of deviating from the optimal action. We begin with the definition of a uniformly ergodic Markov chain.

Definition 2: [33] A Markov chain $X = \{X_t, t \in \mathbb{Z}_+\}$ on a measurable space $(\mathcal{S}, \mathcal{B})$, with transition kernel $P(x, \mathcal{G})$ is uniformly ergodic if there exists constants $\rho < 1, C < \infty$ such that for all $x \in \mathcal{S}$,

$$\|e_x P^t - \pi\| \leq C \rho^t, t \in \mathbb{Z}_+, \quad (15)$$

where e_x is the unit vector indicating that the initial state is x , and the total variation norm is used. For finite and countable vectors this corresponds to l_1 norm, and the induced matrix norm corresponds to the maximum absolute row sum norm. Clearly, for a finite state Markov chain uniform ergodicity is equivalent to ergodicity. The next is a bound on a perturbation to a uniformly ergodic Markov chain.

Lemma 13: ([33] Theorem 3.1.) Let $X = \{X_t, t \in \mathbb{Z}_+\}$ be a uniformly ergodic Markov chain for which (15) holds. Let $\hat{X} = \{\hat{X}_t, t \in \mathbb{Z}_+\}$ be the perturbed chain with transition kernel \hat{P} . Given the two chains have the same initial distribution, let $\psi_t, \hat{\psi}_t$ be the distribution of X, \hat{X} at time t , respectively. Then,

$$\|\psi_t - \hat{\psi}_t\| \leq C_1(P, t) \|\hat{P} - P\|, \quad (16)$$

where $C_1(P, t) = \left(\hat{t} + C \frac{\rho^{\hat{t}} - \rho^t}{1 - \rho}\right)$ and $\hat{t} = \lceil \log_\rho C^{-1} \rceil$.

Next, the Chernoff-Hoeffding bound is frequently used in our proofs that bounds the difference between the sample mean and the expected reward on distributions with bounded support.

Lemma 14: (Chernoff-Hoeffding Bound) Let X_1, \dots, X_T be random variables with common range $[0, 1]$, such that $E[X_t | X_{t-1}, \dots, X_1] = \mu$ for $t = 2, 3, \dots, T$. Let $S_T = X_1 + \dots + X_T$. Then for all $\epsilon \geq 0$

$$P(|S_T - T\mu| \geq \epsilon) \leq 2e^{-\frac{2\epsilon^2}{T}}.$$

The following lemma is used to relate the estimate estimated belief state of the player with the true belief state; it gives an upper bound on the difference between the product of two equal-sized sets of numbers in the unit interval, in terms of the sum of the absolute values of the pairwise differences between the numbers taken from each set.

Lemma 15: for $\rho_k, \rho'_k \in [0, 1]$ we have

$$|\rho_1 \dots \rho_K - \rho'_1 \dots \rho'_K| \leq \sum_{k=1}^K |\rho_k - \rho'_k|. \quad (17)$$

Proof: First consider $|\rho_1 \rho_2 - \rho'_1 \rho'_2|$ where $\rho_1, \rho_2, \rho'_1, \rho'_2 \in [0, 1]$. Let $\epsilon = \rho'_2 - \rho_2$. Then

$$|\rho_1 \rho_2 - \rho'_1 \rho'_2| = |\rho_1 \rho_2 - \rho'_1(\rho_2 + \epsilon)| = |\rho_2(\rho_1 - \rho'_1) - \rho'_1 \epsilon| \leq \rho_2 |\rho_1 - \rho'_1| + \rho'_1 |\epsilon|.$$

But we have

$$|\rho_1 - \rho'_1| + |\rho_2 - \rho'_2| = |\rho_1 - \rho'_1| + |\epsilon| \geq \rho_2 |\rho_1 - \rho'_1| + \rho'_1 |\epsilon|.$$

Thus

$$|\rho_1\rho_2 - \rho'_1\rho'_2| \leq |\rho_1 - \rho'_1| + |\rho_2 - \rho'_2|.$$

We now use induction. Clearly (17) holds for $K = 1$. Assume it holds for some $K > 1$. Then

$$|\rho_1 \cdots \rho_{K+1} - \rho'_1 \cdots \rho'_{K+1}| \leq |\rho_1 \cdots \rho_K - \rho'_1 \cdots \rho'_K| + |\rho_{K+1} - \rho'_{K+1}| \leq \sum_{k=1}^K |\rho_k - \rho'_k|.$$

■

The next lemma bounds the distance between the true and the estimated (but not normalized) transition probabilities for a player using AREP.

Lemma 16: For any $\epsilon > 0$, for a player using AREP with constant $L \geq 3/(2\epsilon^2)$,

$$P\left(|\bar{p}_{ij,t}^k - p_{ij}^k| > \epsilon, \mathcal{E}_t\right) := P\left(\{|\bar{p}_{ij,t}^k - p_{ij}^k| > \epsilon\} \cap \mathcal{E}_t\right) \leq \frac{2}{t^2},$$

for all $t > 0$, $i, j \in S^k$, $k \in \mathcal{K}$.

Proof: Let $t(l)$ be the time at which $C_{t(l)}^k(i) = l$. We have,

$$\bar{p}_{ij,t}^k = \frac{N_t^k(i, j)}{C_t^k(i)} = \frac{\sum_{l=1}^{C_t^k(i)} I(X_{t(l)-1}^k = i, X_{t(l)}^k = j)}{C_t^k(i)}.$$

Note that $I(X_{t(l)-1}^k = i, X_{t(l)}^k = j)$, $l = 1, 2, \dots, C_t^k(i)$ are IID random variables with mean p_{ij}^k . Then

$$\begin{aligned} P\left(|\bar{p}_{ij,t}^k - p_{ij}^k| > \epsilon, \mathcal{E}_t\right) &= P\left(\left|\frac{\sum_{l=1}^{C_t^k(i)} I(X_{t(l)-1}^k = i, X_{t(l)}^k = j)}{C_t^k(i)} - p_{ij}^k\right| \geq \epsilon, \mathcal{E}_t\right) \\ &= \sum_{b=1}^t P\left(\left|\sum_{l=1}^{C_t^k(i)} I(X_{t(l)-1}^k = i, X_{t(l)}^k = j) - C_t^k(i)p_{ij}^k\right| \geq C_t^k(i)\epsilon, C_t^k(i) = b, \mathcal{E}_t\right) \\ &\leq \sum_{b=1}^t 2e^{-\frac{2(L \log t \epsilon)^2}{L \log t}} = 2 \sum_{b=1}^t e^{-2L \log t (\epsilon)^2} = 2 \sum_{b=1}^t \frac{1}{t^{2L\epsilon^2}} = \frac{1}{t^{2L\epsilon^2-1}} \leq \frac{2}{t^2}, \end{aligned} \quad (18)$$

where we used Lemma 14 and the fact that $C_t^k(i) \geq L \log t$ with probability one in the event \mathcal{E}_t . ■

APPENDIX B
PROOF OF LEMMA 7

From symmetry we have

$$\begin{aligned} P\left(|\hat{p}_{ij,t}^k - p_{ij}^k| > \epsilon, \mathcal{E}_t\right) &= P\left(\hat{p}_{ij,t}^k - p_{ij}^k > \epsilon, \mathcal{E}_t\right) + P\left(\hat{p}_{ij,t}^k - p_{ij}^k < -\epsilon, \mathcal{E}_t\right) \\ &= 2P\left(\hat{p}_{ij,t}^k - p_{ij}^k > \epsilon, \mathcal{E}_t\right). \end{aligned} \quad (19)$$

Then

$$\begin{aligned} P\left(\hat{p}_{ij,t}^k - p_{ij}^k > \epsilon, \mathcal{E}_t\right) &= P\left(\frac{\bar{p}_{ij}^k}{\sum_{l \in S^k} \bar{p}_{il}^k} - p_{ij}^k > \epsilon, \mathcal{E}_t\right) \\ &= P\left(\frac{\bar{p}_{ij}^k}{\sum_{l \in S^k} \bar{p}_{il}^k} - p_{ij}^k > \epsilon, \left|\sum_{l \in S^k} \bar{p}_{il}^k - 1\right| < \delta, \mathcal{E}_t\right) \\ &\quad + P\left(\frac{\bar{p}_{ij}^k}{\sum_{l \in S^k} \bar{p}_{il}^k} - p_{ij}^k > \epsilon, \left|\sum_{l \in S^k} \bar{p}_{il}^k - 1\right| \geq \delta, \mathcal{E}_t\right) \\ &\leq P\left(\frac{\bar{p}_{ij}^k}{1 - \delta} - p_{ij}^k > \epsilon, \mathcal{E}_t\right) + P\left(\left|\sum_{l \in S^k} \bar{p}_{il}^k - 1\right| \geq \delta, \mathcal{E}_t\right). \end{aligned} \quad (20)$$

We have

$$P\left(\frac{\bar{p}_{ij}^k}{1 - \delta} - p_{ij}^k > \epsilon, \mathcal{E}_t\right) \leq P\left(\bar{p}_{ij}^k - p_{ij}^k > \epsilon(1 - \delta) - \delta, \mathcal{E}_t\right). \quad (21)$$

Note that $\epsilon(1 - \delta) - \delta$ is decreasing in δ . Choose δ small enough such that $\epsilon(1 - \delta) - \delta > \epsilon/2$. Then

$$P\left(\bar{p}_{ij}^k - p_{ij}^k > \epsilon(1 - \delta) - \delta, \mathcal{E}_t\right) \leq P\left(\bar{p}_{ij}^k - p_{ij}^k > \frac{\epsilon}{2}, \mathcal{E}_t\right) \leq \frac{2}{t^2}, \quad (22)$$

for $L \geq 6/(\epsilon^2)$ using Lemma 14. We also have

$$P\left(\left|\sum_{l \in S^k} \bar{p}_{il}^k - 1\right| \geq \delta, \mathcal{E}_t\right) \leq P\left(\sum_{l \in S^k} \left|\bar{p}_{il}^k - p_{il}^k\right| \geq \delta, \mathcal{E}_t\right).$$

Consider the events $\mathcal{A}_{k,i} = \{|\bar{p}_{il}^k - p_{il}^k| < \delta/|S^k|, \forall k \in \mathcal{K}\}$ and $\mathcal{B}_{k,i} = \{\sum_{l \in S^k} |\bar{p}_{il}^k - p_{il}^k| < \delta\}$. If $\omega \in \mathcal{A}_{k,i}$, then $\omega \in \mathcal{B}_{k,i}$. Thus, $\mathcal{A}_{k,i} \subset \mathcal{B}_{k,i}$, $\mathcal{B}_{k,i}^c \subset \mathcal{A}_{k,i}^c$. Then

$$\begin{aligned}
P\left(\sum_{l \in S^k} |\bar{p}_{il}^k - p_{il}^k| \geq \delta, \mathcal{E}_t\right) &= P(\mathcal{B}_{k,i}^c, \mathcal{E}_t) \\
&\leq P(\mathcal{A}_{k,i}^c, \mathcal{E}_t) \\
&= P\left(\bigcup_{l \in S^k} \{|\bar{p}_{il}^k - p_{il}^k| \geq \delta/|S^k|\}, \mathcal{E}_t\right) \\
&\leq \sum_{l \in S^k} P\left(|\bar{p}_{il}^k - p_{il}^k| \geq \delta/S_{\max}, \mathcal{E}_t\right) \\
&\leq \frac{2S_{\max}}{t^2}, \tag{23}
\end{aligned}$$

for $L \geq S_{\max}^2/(2\delta^2)$ by Lemma 14. Combining (19), (22) and (23) we get

$$P\left(|\hat{p}_{ij,t}^k - p_{ij}^k| > \epsilon, \mathcal{E}_t\right) \leq \frac{2S_{\max} + 2}{t^2},$$

for $L \geq \max\{6/(\epsilon^2), S_{\max}^2/(2\delta^2)\} = C_{\mathcal{P}}(\epsilon)$.

APPENDIX C

PROOF OF LEMMA 8

When the estimated belief is in J_l , for any suboptimal action u , we have

$$\mathcal{L}^*(\psi_t, \mathbf{P}) - \mathcal{L}(\psi_t, u, h_{\mathbf{P}}, \mathbf{P}) \geq \underline{\Delta}. \tag{24}$$

Let $\epsilon < \underline{\Delta}/4$. When $\mathcal{F}_t(\epsilon)$ occurs, we have

$$\left|\mathcal{I}_t(\hat{\psi}_t, u) - \mathcal{L}(\hat{\psi}_t, u, h_{\mathbf{P}}, \hat{\mathbf{P}}_t)\right| \leq \epsilon, \tag{25}$$

for all $u \in U$. Since $T_{\mathbf{P}}(\psi, y, u)$ is continuous in \mathbf{P} , and $h_{\mathbf{P}}(\psi)$ is continuous in ψ , there exists $\delta_e > 0$ such that $\|\hat{\mathbf{P}}_t - \mathbf{P}\|_1 < \delta_e$ implies that

$$\left|\mathcal{L}(\hat{\psi}_t, u, h_{\mathbf{P}}, \mathbf{P}) - \mathcal{L}(\hat{\psi}_t, u, h_{\mathbf{P}}, \hat{\mathbf{P}}_t)\right| \leq \underline{\Delta}/4, \tag{26}$$

for all $u \in U$. Let $u^* \in O(J_l; \mathbf{P})$. Using (24), (25) and (26), we have

$$\begin{aligned}
\mathcal{I}(\hat{\psi}_t, u^*) &\geq \mathcal{L}(\hat{\psi}_t, u^*, h_{\mathbf{P}}, \hat{\mathbf{P}}_t) - \epsilon \\
&\geq \mathcal{L}(\hat{\psi}_t, u^*, h_{\mathbf{P}}, \mathbf{P}) - \epsilon - \underline{\Delta}/4 \\
&= \mathcal{L}^*(\hat{\psi}_t, \mathbf{P}) - \epsilon - \underline{\Delta}/4 \\
&\geq \mathcal{L}(\hat{\psi}_t, u, h_{\mathbf{P}}, \mathbf{P}) + 3\underline{\Delta}/4 - \epsilon \\
&\geq \mathcal{L}(\hat{\psi}_t, u, h_{\mathbf{P}}, \hat{\mathbf{P}}_t) + \underline{\Delta}/2 - \epsilon \\
&\geq \mathcal{I}(\hat{\psi}_t, u) + \underline{\Delta}/2 - 2\epsilon \\
&> \mathcal{I}(\hat{\psi}_t, u) .
\end{aligned}$$

Therefore, we have

$$\left\{ \hat{\psi}_t \in J_l, U_t = u, \|\hat{\mathbf{P}}_t - \mathbf{P}\|_1 < \delta_e, \mathcal{E}_t, \mathcal{F}_t \right\} = \emptyset . \quad (27)$$

Recall that for any $u \notin O(J_l; \mathbf{P})$,

$$\begin{aligned}
E_{\psi_0, \alpha}^{\mathbf{P}}[D_{1,1}(T, \epsilon, J_l, u)] &= \sum_{t=1}^T P\left(\hat{\psi}_t \in J_l, U_t = u, \mathcal{E}_t, \mathcal{F}_t\right) \\
&= \sum_{t=1}^T P\left(\hat{\psi}_t \in J_l, U_t = u, \|\hat{\mathbf{P}}_t - \mathbf{P}\|_1 < \delta_e, \mathcal{E}_t, \mathcal{F}_t\right) \\
&\quad + \sum_{t=1}^T P\left(\hat{\psi}_t \in J_l, U_t = u, \|\hat{\mathbf{P}}_t - \mathbf{P}\|_1 \geq \delta_e, \mathcal{E}_t, \mathcal{F}_t\right) \\
&\leq \sum_{t=1}^T P\left(\|\hat{\mathbf{P}}_t - \mathbf{P}\|_1 \geq \delta_e, \mathcal{E}_t\right), \quad (28)
\end{aligned}$$

where (28) follows from (27). Therefore for any $u \notin O(J_l; \mathbf{P})$,

$$\begin{aligned}
E_{\psi_0, \alpha}^{\mathbf{P}}[D_{1,1}(T, \epsilon, J_l, u)] &\leq \sum_{t=1}^T P\left(\|\hat{\mathbf{P}}_t - \mathbf{P}\|_1 \geq \delta_e, \mathcal{E}_t\right) \\
&\leq \sum_{t=1}^T P\left(\left\{|\hat{p}_{ij,t}^k - p_{ij}^k| \geq \frac{\delta_e}{K S_{\max}^2}, \text{ for some } k \in \mathcal{K}, i, j \in S^k\right\}, \mathcal{E}_t\right) \\
&\leq \sum_{t=1}^T \sum_{k=1}^K \sum_{(i,j) \in S^k \times S^k} P\left(|\hat{p}_{ij,t}^k - p_{ij}^k| \geq \frac{\delta_e}{K S_{\max}^2}, \mathcal{E}_t\right) \\
&\leq 2K S_{\max}^2 (S_{\max} + 1)\beta,
\end{aligned}$$

for $L \geq C_{\mathbf{P}}(\delta_e/(KS_{\max}^2))$, where the last inequality follows from Lemma 7.

APPENDIX D

PROOF OF LEMMA 9

Since $h_{\tilde{\mathbf{P}}}$ is continuous in ψ by Lemma 1 for any $\tilde{\mathbf{P}}$ such that Assumption 1 holds, and since $\bar{r}(\psi), V_{\tilde{\mathbf{P}}}, T_{\tilde{\mathbf{P}}}$ are continuous in $\tilde{\mathbf{P}}$, we have for any $\psi \in \Psi$:

$$\begin{aligned} g_{\hat{\mathbf{P}}} + h_{\hat{\mathbf{P}}}(\psi) &= \arg \max_{u \in U} \left\{ \bar{r}(\psi, u) + \sum_{y \in S^u} V_{\hat{\mathbf{P}}}(\psi, y, u) h_{\hat{\mathbf{P}}}(T_{\hat{\mathbf{P}}}(\psi, y, u)) \right\} \\ &= \arg \max_{u \in U} \left\{ \bar{r}(\psi, u) + \sum_{y \in S^u} V_{\mathbf{P}}(\psi, y, u) h_{\hat{\mathbf{P}}}(T_{\mathbf{P}}(\psi, y, u)) + q(\mathbf{P}, \hat{\mathbf{P}}, \psi, u) \right\}, \end{aligned} \quad (29)$$

for some function q such that $\lim_{\hat{\mathbf{P}} \rightarrow \mathbf{P}} q(\mathbf{P}, \hat{\mathbf{P}}, \psi, u) = 0, \quad \forall \psi \in \Psi, u \in U$. Let $\bar{r}(\mathbf{P}, \hat{\mathbf{P}}, \psi, u) = \bar{r}(\psi, u) + q(\mathbf{P}, \hat{\mathbf{P}}, \psi, u)$. We can write (29) as

$$g_{\hat{\mathbf{P}}} + h_{\hat{\mathbf{P}}}(\psi) = \arg \max_{u \in U} \left\{ \bar{r}(\mathbf{P}, \hat{\mathbf{P}}, \psi, u) + \sum_{y \in S^u} V_{\mathbf{P}}(\psi, y, u) h_{\hat{\mathbf{P}}}(T_{\mathbf{P}}(\psi, y, u)) \right\}. \quad (30)$$

Note that (30) is the AROE for a system with set of transition probability matrices \mathbf{P} , and perturbed rewards $\bar{r}(\mathbf{P}, \hat{\mathbf{P}}, \psi, u)$. Since $\lim_{\hat{\mathbf{P}} \rightarrow \mathbf{P}} \bar{r}(\mathbf{P}, \hat{\mathbf{P}}, \psi, u) = \bar{r}(\psi, u), \quad \forall \psi \in \Psi, u \in U$, we expect $h_{\hat{\mathbf{P}}}$ to converge to $h_{\mathbf{P}}$. Below we prove that this is true. Let $F_{\hat{\mathbf{P}}}$ denote the dynamic programming operator defined in (3), with transition probabilities \mathbf{P} and rewards $\bar{r}(\mathbf{P}, \hat{\mathbf{P}}, \psi, u)$. Then, by S-1 of Lemma (1), there exists a sequence of functions $v_{0, \hat{\mathbf{P}}}, v_{1, \hat{\mathbf{P}}}, v_{2, \hat{\mathbf{P}}}, \dots$ such that $v_{0, \hat{\mathbf{P}}} = 0, v_{l, \hat{\mathbf{P}}} = F_{\hat{\mathbf{P}}} v_{l-1, \hat{\mathbf{P}}}$ and another sequence of functions $v_{0, \mathbf{P}}, v_{1, \mathbf{P}}, v_{2, \mathbf{P}}, \dots$ such that $v_{0, \mathbf{P}} = 0, v_{l, \mathbf{P}} = F_{\mathbf{P}} v_{l-1, \mathbf{P}}$, for which

$$\lim_{l \rightarrow \infty} v_{l, \hat{\mathbf{P}}} = h_{\hat{\mathbf{P}}}, \quad (31)$$

$$\lim_{l \rightarrow \infty} v_{l, \mathbf{P}} = h_{\mathbf{P}}, \quad (32)$$

uniformly in ψ . Next we prove that for any $l \in \{1, 2, \dots\}$, $\lim_{\hat{\mathbf{P}} \rightarrow \mathbf{P}} v_{l, \hat{\mathbf{P}}} = v_{l, \mathbf{P}}$ uniformly in ψ . Let

$$q_{\max}(\mathbf{P}, \hat{\mathbf{P}}) = \sup_{u \in U, \psi \in \Psi} |q(\mathbf{P}, \hat{\mathbf{P}}, \psi, u)|.$$

By Equation 2.27 in [1], we have

$$\begin{aligned}
\sup_{\psi \in \Psi} \left\{ |v_{l, \hat{\mathbf{P}}}(\psi) - v_{l, \mathbf{P}}(\psi)| \right\} &= \sup_{\psi \in \Psi} \left\{ |F^{l-1}v_{1, \hat{\mathbf{P}}}(\psi) - F^{l-1}v_{1, \mathbf{P}}(\psi)| \right\} \\
&\leq \sup_{\psi \in \Psi} \left\{ |v_{1, \hat{\mathbf{P}}}(\psi) - v_{1, \mathbf{P}}(\psi)| \right\} \\
&\leq 2q_{\max}(\mathbf{P}, \hat{\mathbf{P}}),
\end{aligned} \tag{33}$$

where the last inequality follows from $v_{0, \mathbf{P}} = 0$, $v_{0, \hat{\mathbf{P}}} = 0$, and

$$\begin{aligned}
v_{1, \mathbf{P}}(\psi) &= \max_{u \in U} \{ \bar{r}(\psi, u) \} \\
v_{1, \hat{\mathbf{P}}}(\psi) &= \max_{u \in U} \left\{ \bar{r}(\psi, u) + q(\mathbf{P}, \hat{\mathbf{P}}, \psi, u) \right\}.
\end{aligned}$$

Consider a sequence $\{\hat{\mathbf{P}}_n\}_{n=1}^{\infty}$ which converges to \mathbf{P} . Since $\lim_{n \rightarrow \infty} q_{\max}(\mathbf{P}, \hat{\mathbf{P}}_n) = 0$, for any $\epsilon > 0$, there exists N_0 such that for all $n > N_0$ we have $q_{\max}(\mathbf{P}, \hat{\mathbf{P}}_n) < \epsilon/2$, which implies by (33) that $\sup_{\psi \in \Psi} \left\{ |v_{l, \hat{\mathbf{P}}_n}(\psi) - v_{l, \mathbf{P}}(\psi)| \right\} < \epsilon$, for all $\psi \in \Psi$. Therefore, for any $l \in \{1, 2, \dots\}$, we have

$$\lim_{\hat{\mathbf{P}} \rightarrow \mathbf{P}} v_{l, \hat{\mathbf{P}}} = v_{l, \mathbf{P}}, \tag{34}$$

uniformly in ψ . Using (31) and (32), for any $\epsilon > 0$ and any $n \in \{1, 2, \dots\}$, there exists $N_1(n)$ such that for any $l > N_1(n)$ and $\psi \in \Psi$, we have

$$\begin{aligned}
|v_{l, \hat{\mathbf{P}}_n}(\psi) - h_{\hat{\mathbf{P}}_n}(\psi)| &< \epsilon/3, \\
|v_{l, \mathbf{P}}(\psi) - h_{\mathbf{P}}(\psi)| &< \epsilon/3.
\end{aligned}$$

Similarly using (34), for any $\epsilon > 0$, there exists N_0 such that for all $n > N_0$ and $\psi \in \Psi$, we have

$$|v_{l, \hat{\mathbf{P}}_n}(\psi) - v_{l, \mathbf{P}}(\psi)| \leq \epsilon/3. \tag{35}$$

These imply that for any $\epsilon > 0$, there exists $N_2 \geq N_0$ and such that for all $n > N_2$, such that for all $\psi \in \Psi$, we have

$$|h_{\mathbf{P}}(\psi) - h_{\hat{\mathbf{P}}_n}(\psi)| \leq |h_{\mathbf{P}}(\psi) - v_{l, \mathbf{P}}(\psi)| + |v_{l, \hat{\mathbf{P}}_n}(\psi) - v_{l, \mathbf{P}}(\psi)| + |v_{l, \hat{\mathbf{P}}_n}(\psi) - h_{\hat{\mathbf{P}}_n}(\psi)| < \epsilon,$$

since there exists some $l > N_1(n)$ such that (35) holds. Therefore, for any $\epsilon > 0$ there exists some $\eta > 0$ such that $|\mathbf{P} - \hat{\mathbf{P}}| < \eta$ implies $|h_{\mathbf{P}}(\psi) - h_{\hat{\mathbf{P}}}(\psi)|_{\infty} \leq \epsilon$.

APPENDIX E
PROOF OF LEMMA 10

We have by Lemma 9,

$$\left\{ \left\| P^k - \hat{P}_t^k \right\|_1 < \varsigma, \forall k \in \mathcal{K} \right\} \subset \{ \|h_{\mathbf{P}} - h_t\|_\infty < \epsilon \},$$

which implies

$$\left\{ \left\| P^k - \hat{P}_t^k \right\|_1 \geq \varsigma, \text{ for some } k \in \mathcal{K} \right\} \supset \{ \|h_{\mathbf{P}} - h_t\|_\infty \geq \epsilon \}.$$

Then

$$\begin{aligned} E_{\psi_0, \alpha}^{\mathbf{P}}[D_{1,2}(T, \epsilon)] &= E_{\psi_0, \alpha}^{\mathbf{P}} \left[\sum_{t=1}^T I(\mathcal{E}_t, \mathcal{F}_t^c(\epsilon)) \right] \\ &\leq \sum_{t=1}^T P \left(\left\| P^k - \hat{P}_t^k \right\|_1 \geq \varsigma, \text{ for some } k \in \mathcal{K}, \mathcal{E}_t \right) \\ &\leq \sum_{k=1}^K \sum_{(i,j) \in S^k \times S^k} \sum_{t=1}^T P \left(|p_{ij}^k - \hat{p}_{ij,t}^k| > \frac{\varsigma}{S_{\max}^2}, \mathcal{E}_t \right) \\ &\leq 2K S_{\max}^2 (S_{\max} + 1) \beta. \end{aligned}$$

APPENDIX F
PROOF OF LEMMA 11

Consider $t > 0$

$$\begin{aligned} |(\hat{\psi}_t)_{\mathbf{x}} - (\psi_t)_{\mathbf{x}}| &= \left| \prod_{k=1}^K \left((\hat{P}_t^k)^{\tau_k} e_{s^k}^k \right)_{x^k} - \prod_{k=1}^K \left((P^k)^{\tau_k} e_{s^k}^k \right)_{x^k} \right| \\ &\leq \sum_{k=1}^K \left| \left((\hat{P}_t^k)^{\tau_k} e_{s^k}^k \right)_{x^k} - \left((P^k)^{\tau_k} e_{s^k}^k \right)_{x^k} \right| \\ &\leq \sum_{k=1}^K \left\| (\hat{P}_t^k)^{\tau_k} e_{s^k}^k - (P^k)^{\tau_k} e_{s^k}^k \right\|_1 \\ &\leq C_1(\mathbf{P}) \sum_{k=1}^K \left\| \hat{P}_t^k - P^k \right\|_1, \end{aligned} \tag{36}$$

where last inequality follows from Lemma 13. By (36)

$$\left\| \hat{\psi}_t - \psi_t \right\|_1 \leq |S^1| \dots |S^K| C_1(\mathbf{P}) \sum_{k=1}^K \left\| \hat{P}_t^k - P^k \right\|_1.$$

Thus we have

$$\begin{aligned}
& P\left(\left\|\hat{\psi}_t - \psi_t\right\|_1 > \epsilon, \mathcal{E}_t\right) \\
& \leq P\left(\sum_{k=1}^K \left\|\hat{P}_t^k - P^k\right\|_1 > \epsilon/(|S^1| \dots |S^K| C_1(\mathbf{P})), \mathcal{E}_t\right) \\
& \leq \sum_{k=1}^K P\left(\left\|\hat{P}_t^k - P^k\right\|_1 > \epsilon/(K|S^1| \dots |S^K| C_1(\mathbf{P})), \mathcal{E}_t\right) \\
& \leq \sum_{k=1}^K \sum_{(i,j) \in S^k \times S^k} P\left(\left|\hat{p}_{ij,t}^k - p_{ij}^k\right| > \frac{\epsilon}{(K S_{\max}^2 |S^1| \dots |S^K| C_1(\mathbf{P}))}, \mathcal{E}_t\right) \\
& \leq 2K S_{\max}^2 \frac{S_{\max} + 1}{t^2},
\end{aligned}$$

where last inequality follows from Lemma 7 since $L \geq C_{\mathbf{P}}(\epsilon/(K S_{\max}^2 |S^1| \dots |S^K| C_1(\mathbf{P})))$. Then,

$$E_{\psi_0, \alpha}^{\mathbf{P}}[D_{2,1}(T, \epsilon)] = \sum_{t=1}^T P_{\psi_0, \alpha}\left(\left\|\psi_t - \hat{\psi}_t\right\|_1 > \epsilon, \mathcal{E}_t\right) \leq 2K S_{\max}^2 (S_{\max} + 1)\beta.$$

REFERENCES

- [1] L. K. Platzman, "Optimal infinite-horizon undiscounted control of finite probabilistic systems," *SIAM J. Control Optim.*, vol. 18, pp. 362–380, 1980.
- [2] S. P. Hsu, D. M. Chuang, and A. Arapostathis, "On the existence of stationary optimal policies for partially observed mdps under the long-run average cost criterion," *Systems and Control Letters*, vol. 55, pp. 165–173, 2006.
- [3] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in Applied Mathematics*, vol. 6, pp. 4–22, 1985.
- [4] R. Agrawal, "Sample mean based index policies with $O(\log(n))$ regret for the multi-armed bandit problem," *Advances in Applied Probability*, vol. 27, no. 4, pp. 1054–1078, December 1995.
- [5] V. Anantharam, P. Varaiya, and J. Walrand, "Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays-Part I: IID rewards," *IEEE Trans. Automat. Contr.*, pp. 968–975, November 1987.
- [6] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine Learning*, vol. 47, pp. 235–256, 2002.
- [7] J. Gittins and D. Jones, "A dynamic allocation index for sequential design of experiments," *Progress in Statistics, Euro. Meet. Statist.*, vol. 1, pp. 241–266, 1972.
- [8] C. Papadimitriou and J. Tsitsiklis, "The complexity of optimal queuing network control," *Mathematics of Operations Research*, vol. 24, no. 2, pp. 293–305, May 1999.
- [9] P. Whittle, "Restless bandits: Activity allocation in a changing world," *Journal of applied probability*, pp. 287–298, 1988.
- [10] S. Guha, K. Munagala, and P. Shi, "Approximation algorithms for restless bandit problems," *Journal of the ACM (JACM)*, vol. 58, no. 1, p. 3, 2010.
- [11] S. Ahmad, M. Liu, T. Javidi, Q. Zhao, and B. Krishnamachari, "Optimality of myopic sensing in multichannel opportunistic access," *Information Theory, IEEE Transactions on*, vol. 55, no. 9, pp. 4040–4050, 2009.

- [12] V. Anantharam, P. Varaiya, and J. Walrand, "Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays-Part II: Markovian rewards," *IEEE Trans. Automat. Contr.*, pp. 977–982, November 1987.
- [13] C. Tekin and M. Liu, "Online algorithms for the multi-armed bandit problem with markovian rewards," in *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*. IEEE, 2010, pp. 1675–1682.
- [14] —, "Online learning in opportunistic spectrum access: A restless bandit approach," in *Proc. of the 30th Annual IEEE International Conference on Computer Communications (INFOCOM)*, April 2011, pp. 2462–2470.
- [15] —, "Online learning of rested and restless bandits," *IEEE Trans. on Information Theory*, vol. 58, no. 8, pp. 5588–5611, August 2012.
- [16] H. Liu, K. Liu, and Q. Zhao, "Learning in a changing world: Non-bayesian restless multi-armed bandit," *Technical Report, UC Davis*, October 2010.
- [17] C. Tekin and M. Liu, "Performance and convergence of multi-user online learning," in *Proc. of the 2nd International Conference on Game Theory for Networks (GAMENETS)*, April 2011.
- [18] —, "Online learning in decentralized multi-user spectrum access with synchronized explorations," in *submitted to IEEE MILCOM 2012*, 2012.
- [19] K. Liu and Q. Zhao, "Distributed learning in multi-armed bandit with multiple players," *Signal Processing, IEEE Transactions on*, vol. 58, no. 11, pp. 5667–5681, 2010.
- [20] A. Anandkumar, N. Michael, A. Tang, and A. Swami, "Distributed algorithms for learning and cognitive medium access with logarithmic regret," *Selected Areas in Communications, IEEE Journal on*, vol. 29, no. 4, pp. 731–745, 2011.
- [21] W. Dai, Y. Gai, B. Krishnamachari, and Q. Zhao, "The non-bayesian restless multi-armed bandit: A case of near-logarithmic regret," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*. IEEE, 2011, pp. 2940–2943.
- [22] C. Tekin and M. Liu, "Approximately optimal adaptive learning in opportunistic spectrum access," in *Proc. of the 31st Annual IEEE International Conference on Computer Communications (INFOCOM)*, March 2012.
- [23] A. Burnetas and M. Katehakis, "Optimal adaptive policies for markov decision processes," *Mathematics of Operations Research*, pp. 222–255, 1997.
- [24] A. Tewari and P. Bartlett, "Optimistic linear programming gives logarithmic regret for irreducible mdps," *Advances in Neural Information Processing Systems*, vol. 20, pp. 1505–1512, 2008.
- [25] P. Auer, T. Jaksch, and R. Ortner, "Near-optimal regret bounds for reinforcement learning," 2009.
- [26] C. Watkins and P. Dayan, "Q-learning," *Machine learning*, vol. 8, no. 3, pp. 279–292, 1992.
- [27] L. Kaelbling, M. Littman, and A. Moore, "Reinforcement learning: A survey," *Journal of Artificial Intelligence Research*, vol. 4, pp. 237–285, 1996.
- [28] M. Du, "Q-learning for bandit problems," in *Machine Learning: Proceedings of the Twelfth International Conference on Machine Learning*. Citeseer, 1995, pp. 209–217.
- [29] H. Kimura, K. Miyazaki, and S. Kobayashi, "Reinforcement learning in POMDPs with function approximation," in *MACHINE LEARNING-INTERNATIONAL WORKSHOP THEN CONFERENCE-*. MORGAN KAUFMANN PUBLISHERS, INC., 1997, pp. 152–160.
- [30] M. James and S. Singh, "SarsaLandmark: an algorithm for learning in POMDPs with landmarks," in *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2009, pp. 585–591.

- [31] P. Ortner, "Logarithmic online regret bounds for undiscounted reinforcement learning," in *Advances in Neural Information Processing Systems 19: Proceedings of the 2006 Conference*, vol. 19. The MIT Press, 2007, p. 49.
- [32] W. Dai, Y. Gai, B. Krishnamachari, and Q. Zhao, "The non-bayesian restless multi-armed bandit: A case of near-logarithmic regret," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*. IEEE, 2011, pp. 2940–2943.
- [33] A. Y. Mitrophanov, "Sensitivity and convergence of uniformly ergodic markov chains," *J. Appl. Prob.*, vol. 42, pp. 1003–1014, 2005.