

Information-Theoretic Regret Bounds for Bandits with Fixed Expert Advice

Khaled Eldowa*, Nicolò Cesa-Bianchi*, Alberto Maria Metelli†, Marcello Restelli†

*Università degli Studi di Milano, Milan, Italy

†Politecnico di Milano, Milan, Italy

{khaled.eldowa, nicolo.cesa-bianchi}@unimi.it, {albertomaria.metelli, marcello.restelli}@polimi.it

Abstract—We investigate the problem of bandits with expert advice when the experts are fixed and known distributions over the actions. Improving on previous analyses, we show that the regret in this setting is controlled by information-theoretic quantities that measure the similarity between experts. In some natural special cases, this allows us to obtain the first regret bound for EXP4 that can get arbitrarily close to zero if the experts are similar enough. While for a different algorithm, we provide another bound that describes the similarity between the experts in terms of the KL-divergence, and we show that this bound can be smaller than the one of EXP4 in some cases. Additionally, we provide lower bounds for certain classes of experts showing that the algorithms we analyzed are nearly optimal in some cases.

I. INTRODUCTION

Bandits with expert advice (see, e.g., [1]) is a well-known variant of the non-stochastic bandits problem in which, at the beginning of each round, N experts each make a recommendation to the learner in the form of a distribution over the K available actions. The algorithm EXP4 [2] solves this problem with a regret against the best expert bounded by $\sqrt{2TK \log N}$, where T is the horizon. When $N \gg K$, this bound shows the ability of EXP4 to leverage the structure of the problem, as opposed to running a bandit algorithm over the N experts achieving a bound of \sqrt{TN} . An almost matching lower bound of order $\sqrt{TK \log N / \log K}$ was proved in [3] (for deterministic experts). In this work, we study a variant of bandits with expert advice in which the distributions recommended by the experts are *fixed* and *known*. In the following, we will use the term *policies* to denote these fixed experts. Our goal is to determine the best possible dependence of the regret on the structure of the policy set Θ irrespective of the assigned sequence of losses.

This problem is closely related to linear bandits [4] (with finite decision sets), where the structure of the decision/policy set can be provably leveraged. Our problem can also be viewed as a non-stochastic version of *bandits with mediator feedback* [5], where the learner’s access to actions is mediated by the fixed policy set. When losses are stochastic rather than being adversarial, regret bounds were proved in [5] that scale with the largest pair-wise exponentiated 2-Rényi divergence¹ between the policies in the context of policy-based reinforcement learning. Comparable bounds were also proved

in [6] in the setting of contextual bandits. In our setting, where losses are adversarially generated, the best known bound is $\sqrt{2TS(\Theta) \log N}$ from [7], where $S(\Theta) \leq \min\{K, N\}$ is a notion describing the similarity between policies, see Section III-A for its definition and a new information-theoretic interpretation. Since $S(\Theta) \geq 1$ for all Θ , this bound cannot get arbitrarily small no matter how similar the policies are, and becomes vacuous when the policies are identical.

Our first contribution (Theorem 1) is a new regret bound for EXP4 of the form $\sqrt{2TS^*(\Theta) \log N}$, where $S^*(\Theta)$ is a new index of similarity between policies that is never larger than $S(\Theta)$ and reduces to twice the total variation distance when $N = |\Theta| = 2$. In particular, we show that $S^*(\Theta)$ can indeed become arbitrarily small, depending on the policy set. Note that such guarantees cannot be obtained only as a consequence of the reduced range of the losses caused by the similarity of the policies, see [8]. Additionally, we show in Theorem 2 an algorithm whose regret is bounded by $\sqrt{2TKD^*(\Theta)}$, where $D^*(\Theta)$ is a notion of the “width” of Θ in terms of the KL-divergence that reduces to the information radius when Θ is symmetric. This bound is never worse than the EXP4 bound $\sqrt{2TK \log N}$. Moreover, we construct sets Θ where $KD^*(\Theta) < S^*(\Theta) \log N$. Finally, we prove lower bounds for a number of policy set structures and contrast them with the upper bounds we derived. We illustrate, in particular, some examples where the bounds are nearly matching.

II. PROBLEM FORMULATION

We consider a non-stochastic multi-armed bandits problem with a finite action set $\mathcal{A} = [K]$ containing K actions, and a (fixed) policy set $\Theta \in \Delta_{K-1}$ consisting of N probability distributions over the actions. Here Δ_{K-1} denotes the probability simplex in \mathbb{R}^K and, for a policy $\theta \in \Theta$ and $j \in [K]$, $\theta(j)$ is the probability with which policy θ picks action j . We additionally assume that each arm is in the support of at least one policy. With a time horizon of T rounds, an instance of the problem is characterized by an unknown sequence of loss vectors $(\ell_t)_{t=1}^T$, where $\ell_t(j) \in [0, 1]$, for $j \in [K]$, denotes the loss assigned to action j at round $t \in [T]$. A decision maker interacts with the environment as follows: at each round t , the decision maker selects a policy $\theta_t \in \Theta$; an action $A_t \in [K]$ is then sampled from θ_t ; the decision maker subsequently suffers the loss $\ell_t(A_t)$ and observes the pair $(A_t, \ell_t(A_t))$. With a slight abuse of notation we denote by $\ell_t(\theta)$ the expected

¹This divergence is related to the chi-squared divergence. Note that pair-wise, these divergences can be infinite in non-trivial cases, see Example 2.

Algorithm 1 EXP4 With Fixed Expert Advice

-
- 1: **Input:** K, Θ, η
 - 2: **Initialize:** $\forall \theta \in \Theta, \hat{\ell}_0(\theta) = 0$
 - 3: **for** $t = 1, \dots, T$ **do**
 - 4: Draw $\theta_t \sim P_t$, where $P_t(\theta) = \frac{\exp(-\eta \sum_{s=0}^{t-1} \hat{\ell}_s(\theta))}{\sum_{\xi \in \Theta} \exp(-\eta \sum_{s=0}^{t-1} \hat{\ell}_s(\xi))}$
 - 5: Draw $A_t \sim \theta_t$, and observe loss $\ell_t(A_t)$
 - 6: $\forall \theta \in \Theta$, set $\hat{\ell}_t(\theta) = \frac{\theta(A_t)}{\sum_{\xi \in \Theta} P_t(\xi) \xi(A_t)} \ell_t(A_t)$
-

loss (at round t) given that policy θ was selected; that is, $\ell_t(\theta) = \sum_{j=1}^K \theta(j) \ell_t(j)$. The objective is to minimize the regret, which we define as follows:

$$R_T = \mathbb{E} \sum_{t=1}^T \ell_t(\theta_t) - \sum_{t=1}^T \ell_t(\theta^*),$$

where $\theta^* \in \arg \min_{\theta \in \Theta} \sum_{t=1}^T \ell_t(\theta)$ and the expectation is over the internal randomization of the player.

III. EXP4 REGRET ANALYSIS

It is possible to show² that the EXP4 algorithm with a suitable tuning of the learning rate satisfies a similar regret bound to the one proven for the algorithm developed in [7]. In our setting, this bound is $\sqrt{2T\mathcal{S}(\Theta) \log N}$, where $\mathcal{S}(\Theta) = \sum_{j=1}^K \max_{\theta \in \Theta} \theta(j)$ is a notion of similarity for the policy set. It is easy to see that $\mathcal{S}(\Theta) \leq \min\{K, N\}$, and thus, it can be interpreted as the effective “number” of policies.

A. An Information-Theoretic Interpretation of $\mathcal{S}(\Theta)$

An alternative characterization of the policy set similarity can be derived by observing that $\mathcal{S}(\Theta) = 1 + \mathcal{TV}(\Theta)$, where for any ordering of the policies $(\theta_i)_{i=1}^N$ we define:

$$\mathcal{TV}(\Theta) = \sum_{i=2}^N \sum_{j: \theta_i(j) > \theta_{[i-1]}(j)} (\theta_i(j) - \theta_{[i-1]}(j)), \quad (1)$$

where $\theta_{[i-1]}(j) = \max_{\theta \in \{\theta_1, \dots, \theta_{i-1}\}} \theta(j)$, and the second sum is the upper variation of the signed measure $\theta_i - \theta_{[i-1]}$. $\mathcal{TV}(\Theta)$ can be seen as a generalization of the total variation distance to describe the overall divergence of the policy set. Indeed, it is easy to see that when $\Theta = \{\theta_1, \theta_2\}$, $\mathcal{TV}(\Theta)$ reduces to the total variation between the two policies:

$$D_{\text{TV}}(\theta_1, \theta_2) = \sum_{j: \theta_2(j) > \theta_1(j)} (\theta_2(j) - \theta_1(j)).$$

Moreover, an upper bound on $\mathcal{TV}(\Theta)$ can be derived by noting that for any $\tau \in \Delta_{K-1}$ we have:

$$\mathcal{TV}(\Theta) = \mathcal{S}(\Theta) - 1 = \sum_{\theta \in \Theta} \sum_{j \in B(\theta)} (\theta(j) - \tau(j)),$$

where $(B(\theta))_{\theta \in \Theta}$ is any partition of $[K]$ such that for $j \in B(\theta)$ we have $\theta \in \arg \max_{\theta' \in \Theta} \theta'(j)$. Then, it follows that:

$$\mathcal{TV}(\Theta) \leq \min_{\tau \in \Delta_{K-1}} \sum_{\theta \in \Theta} D_{\text{TV}}(\theta, \tau). \quad (2)$$

²See also Theorem 18.3 in [1].

It should be noted that this bound can be loose. For instance, if the policy set is partitioned into clusters of similar policies, then it is not hard to see that the right-hand side of (2) will be wasteful compared to (1). It is also noteworthy that quantities related to $\mathcal{S}(\Theta)$ are used when studying the minimax risk in statistical estimation problems. In particular, Theorem II.1 in [9] can be used to derive upper bounds similar to (2) in terms of any f -divergence, though often in implicit form.

B. An Improved Bound

Nevertheless, the stated regret bound scales with $\mathcal{S}(\Theta)$, not with $\mathcal{TV}(\Theta)$. The main limitation is that $\mathcal{S}(\Theta) \geq 1$, no matter how close $\mathcal{TV}(\Theta)$ is to zero. Thus, the bound $\sqrt{2T\mathcal{S}(\Theta) \log N}$ is never smaller than $\sqrt{2T \log N}$ regardless of the structure. One might wonder if this is necessary. The following theorem provides the first regret bound for EXP4 that can get arbitrarily close to zero if the policies are similar enough. Our bound depends on the key quantity:

$$\mathcal{S}^*(\Theta) = \sum_{j=1}^K \left(\max_{\theta \in \Theta} \theta(j) - \min_{\theta' \in \Theta} \theta'(j) \right),$$

that is easily seen to satisfy $\mathcal{TV}(\Theta) \leq \mathcal{S}^*(\Theta) \leq \mathcal{S}(\Theta)$.

Theorem 1. *Algorithm 1 run with $\eta = \sqrt{\frac{2 \log N}{T\mathcal{S}^*(\Theta)}}$ satisfies $R_T \leq \sqrt{2T\mathcal{S}^*(\Theta) \log N}$.*

Proof. For a policy θ , we define a modified version of the loss at time t as $\zeta_t(\theta) = \sum_{j=1}^K (\theta(j) - q(j)) \ell_t(j)$, where $q(j) = \min_{\theta' \in \Theta} \theta'(j)$. Notice that for any two policies θ_i and θ_j , $\zeta_t(\theta_i) - \zeta_t(\theta_j) = \ell_t(\theta_i) - \ell_t(\theta_j)$. An estimate for this modified loss can be constructed in a standard manner: $\hat{\zeta}_t(\theta) = (\theta(A_t) - q(A_t)) \frac{\ell_t(A_t)}{\psi_t(A_t)}$, where $\psi_t(j) = \sum_{\theta' \in \Theta} P_t(\theta') \theta'(A_t)$ is the probability of playing arm j at time t given P_t . Let \mathbb{E}_{t-1} be the expectation conditioned on the events up to round $t-1$, it can be verified³ that (for any θ) $\mathbb{E}_{t-1} \hat{\zeta}_t(\theta) = \zeta_t(\theta)$ and $\mathbb{E}_{t-1} \zeta_t(\theta_t) = \mathbb{E}_{t-1} \sum_{\theta \in \Theta} P_t(\theta) \hat{\zeta}_t(\theta)$. Hence, we have that

$$\begin{aligned} R_T &= \mathbb{E} \left[\sum_t \zeta_t(\theta_t) - \sum_t \zeta_t(\theta^*) \right] \\ &= \mathbb{E} \left[\sum_t \sum_{\theta \in \Theta} P_t(\theta) \hat{\zeta}_t(\theta) - \sum_t \zeta_t(\theta^*) \right]. \quad (3) \end{aligned}$$

Notice that: $\sum_{s=1}^{t-1} \hat{\ell}_s(\theta) = \sum_{s=1}^{t-1} \hat{\zeta}_s(\theta) + \sum_{s=1}^{t-1} q(A_s) \frac{\ell_s(A_s)}{\psi_s(A_s)}$. And since the second sum does not depend on θ , we can rewrite⁴ P_t as follows:

$$P_t(\theta) = \frac{\exp(-\eta \sum_{s=0}^{t-1} \hat{\zeta}_s(\theta))}{\sum_{\theta' \in \Theta} \exp(-\eta \sum_{s=0}^{t-1} \hat{\zeta}_s(\theta'))}.$$

One can then bound (3) using a standard manipulation (see, for example, Theorem 1.5 in [11]) to yield that:⁵

$$R_T \leq \frac{\eta}{2} \mathbb{E} \sum_t \sum_{\theta \in \Theta} P_t(\theta) \hat{\zeta}_t^2(\theta) + \frac{\log N}{\eta}. \quad (4)$$

³See the proof of Theorem 4.2 in [10] for similar arguments.

⁴Like $\hat{\ell}_0(\theta)$, we initialize $\hat{\zeta}_0(\theta)$ as zero.

⁵This requires that $\hat{\zeta}_t(\theta)$ is non-negative, which is indeed the case.

Define $s(j) = \max_{\theta \in \Theta} \theta(j)$, we now bound the second moment term at step t :

$$\begin{aligned} & \mathbb{E}_{t-1} \sum_{\theta \in \Theta} P_t(\theta) \hat{\zeta}_t^2(\theta) \\ &= \mathbb{E}_{t-1} \sum_{\theta \in \Theta} P_t(\theta) (\theta(A_t) - q(A_t))^2 \frac{\ell_t^2(A_t)}{\psi_t^2(A_t)} \\ &\leq \mathbb{E}_{t-1} (s(A_t) - q(A_t)) \frac{\ell_t^2(A_t)}{\psi_t(A_t)} \frac{\sum_{\theta \in \Theta} P_t(\theta) (\theta(A_t) - q(A_t))}{\sum_{\theta \in \Theta} P_t(\theta) \theta(A_t)} \\ &\leq \mathbb{E}_{t-1} (s(A_t) - q(A_t)) \frac{\ell_t^2(A_t)}{\psi_t(A_t)} \\ &= \sum_j (s(j) - q(j)) \ell_t^2(j) \leq \mathcal{S}^*(\Theta). \end{aligned}$$

Combining this with (4) and the specified choice of η concludes the proof. \square

In general, $\mathcal{S}^*(\Theta)$ is not guaranteed to be strictly smaller $\mathcal{S}(\Theta)$; they can be equal in some cases regardless of how small $\mathcal{TV}(\Theta)$ is. In other cases, however, there can be an improvement. To see this, note that for any $\tau \in \Delta_{K-1}$,

$$\mathcal{S}^*(\Theta) = \sum_{\theta \in \Theta} \sum_{j \in B(\theta)} (\theta(j) - \tau(j)) + \sum_{j \in B'(\theta)} (\tau(j) - \theta(j)),$$

where $(B'(\theta))_{\theta \in \Theta}$ is a partition of $[K]$ such that for every $j \in B'(\theta)$ we have $\theta \in \arg \min_{\theta' \in \Theta} \theta'(j)$. Then, analogously to (2), we have that:

$$\mathcal{S}^*(\Theta) \leq 2 \min_{\tau \in \Delta_{K-1}} \sum_{\theta \in \Theta} D_{\text{TV}}(\theta, \tau).$$

Like (2), this bound can be loose, but it serves to indicate that if $\min_{\tau \in \Delta_{K-1}} \sum_{\theta \in \Theta} D_{\text{TV}}(\theta, \tau)$ is small (it can get arbitrarily so), $\mathcal{S}^*(\Theta)$ is guaranteed to be of at most the same order.

C. Examples

In the following, we compare the quantities $\mathcal{S}(\Theta)$, $\mathcal{S}^*(\Theta)$, and $\mathcal{TV}(\Theta)$ for a selection of policy set structures.

Example 1 (Two Policies). As we have seen before, for the two policies case, i.e., $\Theta = \{\theta_1, \theta_2\}$, $\mathcal{S}(\Theta) = 1 + D_{\text{TV}}(\theta_1, \theta_2)$. Whereas $\mathcal{S}^*(\Theta) = \sum_{j=1}^K |\theta_1(j) - \theta_2(j)| = 2D_{\text{TV}}(\theta_1, \theta_2)$.

The next two examples concern the case in which each policy is a uniform distribution over a support of $M \leq K$ arms. In this scenario, we get that⁶ $\mathcal{S}(\Theta) = \frac{K}{M}$, while $\mathcal{S}^*(\Theta)$ depends on the number of arms common to all policies.

Example 2 (Radially Symmetric Uniform Policies). Consider a structure where the intersection of the supports of any pair of policies is the same.⁷ Let $V \leq M$ be the number of arms common to all policies, we have that $\mathcal{S}(\Theta) = \frac{N(M-V)+V}{M}$, that is bounded from below by 1, while $\mathcal{S}^*(\Theta) = \frac{N}{M} \frac{M-V}{M}$ and $\mathcal{TV}(\Theta) = (N-1) \frac{M-V}{M}$ are not.⁸

Example 3 (Failure of \mathcal{S}^*). On the other hand, if $K = M+1$, and the policy set contains all possible M -supported uniform

policies, then $\mathcal{S}(\Theta) = \frac{M+1}{M}$ which approaches 1 as M increases, thus $\mathcal{TV}(\Theta) = \frac{1}{M}$ approaches 0. However, $\mathcal{S}^*(\Theta)$ is always equal to $\mathcal{S}(\Theta)$ since for each arm j , $\min_{\theta' \in \Theta} \theta'(j) = 0$.

Example 4 (ϵ -Uniform Policies). If $N = K$ and each policy θ is associated (one-to-one) with an arm a_θ so that, for an arm j , $\theta(j) = \frac{1-\epsilon}{K} + \epsilon \mathbb{I}\{j = a_\theta\}$, where $0 \leq \epsilon \leq 1$. Then, $\mathcal{S}(\Theta) = \epsilon K + 1 - \epsilon$, while $\mathcal{S}^*(\Theta) = \epsilon K$ and $\mathcal{TV}(\Theta) = \epsilon(K-1)$.

IV. AN ALTERNATIVE APPROACH

Since we can randomize our policy choice at each round, we can interpret our setting as a bandits problem in which the player has to randomize over the actions choosing a distribution from (and also competing with) the convex hull $\text{co}(\Theta)$ of the available policies. A simple approach, outlined in Algorithm 2, is to adapt the Online Stochastic Mirror Descent (OSMD) interpretation of EXP3 [2] to our setting. The main distinction is that we need to project onto $\text{co}(\Theta)$ at each step. Denote by $D(P, Q)$ the KL-divergence between distributions P and Q , and for $\tau \in \Delta_{K-1}$, define $D(\Theta || \tau) = \max_{\theta \in \Theta} D(\theta, \tau)$. The following regret bound for Algorithm 2 uses a notion of the ‘‘width’’ of Θ in terms of the KL-divergence defined by:⁹

$$D^*(\Theta) = \min_{\tau \in \Delta_{K-1}} D(\Theta || \tau).$$

Theorem 2. Algorithm 2 run with

$$\tau^* \in \arg \min_{\tau \in \Delta_{K-1}} D(\Theta || \tau) \quad \text{and} \quad \eta = \sqrt{\frac{2D^*(\Theta)}{TK}}$$

satisfies $R_T \leq \sqrt{2TD^*(\Theta)K}$.

Proof. Noting that $\hat{\ell}_t$ is unbiased given x_t , we get that $R_T = \mathbb{E} \sum_t \langle x_t - \theta^*, \hat{\ell}_t \rangle$. This expression is the regret of OMD on the estimated losses with a decision set $\text{co}(\Theta)$ and the negative entropy function $\psi(x) = \sum_j x(j) \log x(j)$ as the regularizer. Using Lemma 6.14 in [13],¹⁰ we get that:

$$R_T \leq \frac{D(\theta^*, \tau^*)}{\eta} + \frac{\eta}{2} \mathbb{E} \sum_t \sum_j x_t(j) \hat{\ell}_t^2(j), \quad (5)$$

Similar to the last step of the proof of Theorem 1, we can show that $\mathbb{E} \sum_j x_t(j) \hat{\ell}_t^2(j) \leq K$. The proof concludes by bounding $D(\theta^*, \tau^*)$ with $D^*(\Theta)$ and plugging in the value of η . \square

It can be shown that $D^*(\Theta) \leq \log N$. Moreover, if the policy set is symmetric, in the sense that the KL-divergence between any policy and the uniform mixture is the same, then $D^*(\Theta)$ (attained at the uniform mixture) coincides with the Jensen-Shannon divergence [14] (or the information radius) of the policy set.

Algorithm 1 can also be seen as an instance of OSMD with the negative entropy regularizer. The main difference is that the decision space is the entire probability simplex over the policies. Hence, the regularization in Algorithm 1

⁶Recall that we assume that an arm is in the support of at least one policy.

⁷This means that any arm is either in the support of all policies or exclusively in the support of a single one.

⁸Note that $\frac{M-V}{M}$ is the total variation distance between any two policies.

⁹Note that the minimum value can only be attained in $\text{co}(\Theta)$; see Theorem 11.6.1 in [12].

¹⁰See also, still in [13], the discussion leading up to Theorem 10.2 concerning the negative entropy regularizer.

Algorithm 2 OSMD on the Convex Hull of Policies

-
- 1: **Input:** $K, \Theta, \eta, \tau \in \text{co}(\Theta) : \tau(j) > 0 \forall j \in [K]$
 - 2: **Initialize:** $x_1 = \tau$
 - 3: **for** $t = 1, \dots, T$ **do**
 - 4: Pick distribution P_t on Θ such that $\sum_{\theta \in \Theta} P_t(\theta)\theta = x_t$
 - 5: Draw $\theta_t \sim P_t$, then $A_t \sim \theta_t$, and observe loss $\ell_t(A_t)$
 - 6: $\forall j \in [K]$, set $\hat{\ell}_t(j) = \frac{\mathbb{I}\{j=A_t\}}{x_t(j)} \ell_t(A_t)$
 - 7: Update $x_{t+1} = \arg \min_{x \in \text{co}(\Theta)} \eta \langle x, \hat{\ell}_t \rangle + D(x, x_t)$
-

favors exploring uniformly over the policies, whereas it favors exploring uniformly over actions in Algorithm 2. Analysis-wise (compare in particular (4) and (5)), Algorithm 1 takes advantage of the similarity between policies to reduce the variance of their loss estimates (compared to K), whereas Algorithm 2 takes advantage of the (possibly) limited size of $\text{co}(\Theta)$ to reduce the bias term (compared to $\log N$).

Consider the ϵ -uniform structure of Example 4. $D^*(\Theta)$ in this case is given by

$$\frac{K-1}{K}(1-\epsilon)\log(1-\epsilon) + \frac{1+\epsilon(K-1)}{K}\log(1+\epsilon(K-1)).$$

If we now compare $KD^*(\Theta)$ and $\epsilon K \log K$, we see that both are equal when $\epsilon \in \{0, 1\}$. However, the former is strictly convex for $\epsilon \in (0, 1)$, while the latter is linear. Thus, in this case, the bound of Theorem 2 is better than that of Theorem 1. However, if we consider the structure of Example 2, we have that $D^*(\Theta) = \frac{M-V}{M} \log N$. Thus, the bound of Theorem 2 becomes $\sqrt{2T \frac{M-V}{M} K \log N}$ which is worse than the $\sqrt{2T \frac{M-V}{M} N \log N}$ bound of Theorem 1 since $K \geq N$.

V. LOWER BOUNDS

In this section, we prove minimax lower bounds for specific classes of policy sets and contrast them with the regret bounds discussed thus far. More precisely, with a fixed policy set Θ , we prove lower bounds on $\inf_{\pi} \sup_{(\ell_t)_{t=1}^T} R_T$, where π is the player's strategy. To this end, we will consider a class of stochastic environments, each identified by the vector $\mu \in [0, 1]^K$ such that, for $j \in [K]$ and any t , $\ell_t(j)$ is drawn from a Bernoulli distribution with mean $\mu(j)$. For any $t \leq T$, let $H_t = (\theta_1, A_1, \ell_1(A_1), \dots, \theta_t, A_t, \ell_t(A_t))$ be the interaction history up to round T . Each environment μ (together with strategy π) induces a probability distribution P_{μ} on H_T . Define:

$$R_T(\mu) = \max_{\theta^* \in \Theta} \mathbb{E}_{\mu} \sum_{t=1}^T \sum_{j=1}^K (\theta_t(j) - \theta^*(j)) \mu(j),$$

where the subscript in \mathbb{E}_{μ} emphasizes the dependence on P_{μ} . In the following, we will prove lower bounds on $\sup_{\mu} R_T(\mu)$ that hold for any algorithm. We will rely on the following lemma, which is an immediate extension of a standard result (see Lemma 15.1 in [1]) to our setting.

Lemma 3. Fix a strategy π , a policy set Θ , and a horizon T ; and let μ and μ' be two environments. Then,

$$D(P_{\mu}, P_{\mu'}) = \sum_{\theta \in \Theta} N_{\mu}(\theta; T) \sum_j \theta(j) d(\mu(j), \mu'(j)),$$

where $N_{\mu}(\theta; T) = \mathbb{E}_{\mu} \sum_{t=1}^T \mathbb{I}\{\theta_t = \theta\}$, and $d(a, b)$ is the KL-divergence between two Bernoulli distributions with means a and b .

A. Radially Symmetric Uniform Policies

The first lower bound concerns the case discussed in Example 2. The construction of the lower bound (which is an adaptation of the standard approach in [2]) leverages the fact that the policies are uniform and equidistant.

Theorem 4. If Θ conforms to the structure of Example 2 such that $M > V$. Then for any algorithm and $T \geq \frac{NM}{4 \log(4/3)(M-V)}$, there exists a sequence of losses such that $R_T \geq \frac{1}{18} \sqrt{N \frac{M-V}{M} T}$.

Since $\mathcal{S}^*(\Theta) = N \frac{M-V}{M}$, it follows that the bound of Theorem 1 is optimal in this case, up to a logarithmic factor.

Proof. We denote by $U(\theta) \subseteq [K]$ the support of θ and $C = \bigcap_{\theta \in \Theta} U(\theta)$. We consider N environments $\{\mu_{\theta}\}_{\theta \in \Theta}$ such that for μ_{θ} and arm j , $\mu_{\theta}(j) = \frac{1}{2} - \Delta \mathbb{I}\{j \in U(\theta) \setminus C\}$, where $0 < \Delta < \frac{1}{2}$ is to be tuned later. Let μ_0 be an environment such that $\mu_0(j) = \frac{1}{2}$ for any arm j . Note that θ is the optimal policy in μ_{θ} , and for $\theta' \in \Theta \setminus \{\theta\}$, we have that:

$$\sum_{j=1}^K (\theta'(j) - \theta(j)) \mu(j) = \Delta \sum_{j \in U(\theta) \setminus C} \left(\frac{1}{M} - 0 \right) = \Delta \frac{M-V}{M}.$$

And thus,

$$\begin{aligned} R_T(\mu_{\theta}) &= \Delta \frac{M-V}{M} (T - N_{\mu_{\theta}}(\theta; T)) \\ &\geq \Delta \frac{M-V}{M} \left(T - N_{\mu_0}(\theta; T) - T \sqrt{\frac{1}{2} D(P_{\mu_0}, P_{\mu_{\theta}})} \right), \end{aligned}$$

where the inequality follows from the fact that¹¹ $N_{\mu_{\theta}}(\theta; T) - N_{\mu_0}(\theta; T) \leq TD_{\text{TV}}(P_{\mu_0}, P_{\mu_{\theta}})$ and by Pinsker's inequality. Starting from Lemma 3, we have that:

$$\begin{aligned} D(P_{\mu_0}, P_{\mu_{\theta}}) &= \sum_{\theta' \in \Theta} N_{\mu_0}(\theta'; T) \sum_{j=1}^K \theta'(j) d(\mu_0(j), \mu_{\theta}(j)) \\ &= N_{\mu_0}(\theta; T) \sum_{j \in U(\theta) \setminus C} \frac{1}{M} d\left(\frac{1}{2}, \frac{1}{2} - \Delta\right) \\ &= \frac{M-V}{M} N_{\mu_0}(\theta; T) \left(-\frac{1}{2} \log(1 - 4\Delta^2) \right) \\ &\leq \frac{M-V}{M} N_{\mu_0}(\theta; T) c \Delta^2, \end{aligned}$$

¹¹See Exercise 14.4 in [1] for a general version of this inequality.

where the second equality holds since the only arms whose means have changed between the two environments lie exclusively in the support of θ , and the inequality holds for $\Delta \leq \frac{1}{4}$ with $c = 8 \log \frac{4}{3}$. Hence:

$$\begin{aligned} \sup_{\mu} R_T(\mu) &\geq \frac{1}{N} \sum_{\theta \in \Theta} R_T(\mu_{\theta}) \\ &\geq \Delta \frac{M-V}{M} \left(T - \frac{T}{N} - T \sqrt{\frac{1}{2} \frac{M-V}{M} \frac{T}{N} c \Delta^2} \right) \\ &\geq \Delta \frac{M-V}{M} T \left(\frac{1}{2} - \Delta \sqrt{\frac{1}{2} \frac{M-V}{M} \frac{T}{N} c} \right), \end{aligned}$$

where the second inequality holds by the concavity of the square root, and the third since $N \geq 2$. The theorem then follows by setting $\Delta = \frac{1}{4} \sqrt{\frac{2MN}{c(M-V)T}}$ and verifying that the condition on T ensures that $\Delta \leq \frac{1}{4}$. \square

B. ϵ -Uniform Policies

The next bound concerns the ϵ -Uniform structure of Example 4. While all policies are still equidistant, this case does not enjoy the peculiar discrete structure of Example 2.

Theorem 5. *If Θ conforms to the ϵ -uniform structure of Example 4. Then for $T \geq \frac{K}{4 \log(4/3)}$ and any algorithm, there exists a sequence of losses such that $R_T \geq \frac{1}{18} \epsilon \sqrt{KT}$.*

The proof, see Appendix A of the extended version [15], is similar to that of Theorem 4 apart from the fact that all policies contribute to the KL-divergence between any two environments (see Lemma 3) since all policies have full support. Notice that the bound is of order $\sqrt{\epsilon^2 KT}$ instead of $\sqrt{\epsilon KT} = \sqrt{S^*(\Theta)T}$, which would have nearly matched the bound of Theorem 1. This was expected since we have shown that Algorithm 2 enjoys a better regret bound in this case. It is interesting to see if a matching lower bound could be proved.

C. The Two Policies Case

For the two policies case, we can prove a lower bound of order $\sqrt{H^2(\theta_1, \theta_2)T}$ as asserted by the following theorem, where $H^2(\theta_1, \theta_2) = \frac{1}{2} \sum_j (\sqrt{\theta_1(j)} - \sqrt{\theta_2(j)})^2$ is the squared Hellinger distance. Relative to the total variation, we have that in general: $\frac{1}{2} D_{TV}^2(\theta_1, \theta_2) \leq H^2(\theta_1, \theta_2) \leq D_{TV}(\theta_1, \theta_2)$.

Theorem 6. *Assume that $\Theta = \{\theta_1, \theta_2\}$ and $H^2(\theta_1, \theta_2) > 0$. Then for any algorithm and $T \geq \frac{1}{8 \log(4/3) H^2(\theta_1, \theta_2)}$, there exists a sequence of losses such that $R_T \geq \frac{1}{13} \sqrt{H^2(\theta_1, \theta_2)T}$.*

The proof, see [15, Appendix B], follows a similar scheme as before, so we only sketch the main distinctions. For an arm j , we define $z_1(j) = \frac{\sqrt{\theta_1(j)} - \sqrt{\theta_2(j)}}{\sqrt{\theta_1(j) + \sqrt{\theta_2(j)}}$, and $z_2(j) = -z_1(j)$. We use two environments μ_1 and μ_2 , where $\mu_1(j) = 1/2 - \Delta z_1(j)$, with $\mu_2(j)$ defined analogously. Subsequently, focusing on μ_1 and θ_1 , we can bound $R_T(\mu_1)$ from below by

$$2\Delta H^2(\theta_1, \theta_2) \left(T - N_{\mu_0}(\theta; T) - T \sqrt{\frac{1}{2} D(P_{\mu_0}, P_{\mu_1})} \right),$$

whereas we can show that $D(P_{\mu_0}, P_{\mu_1}) \leq 2c\Delta^2 H^2(\theta_1, \theta_2)T$. To see the latter, it suffices to start from Lemma 3 and to use that $d(\mu_0(j), \mu_1(j)) \leq c\Delta^2 z_1(j)^2$ for sufficiently small Δ , and that $\sum_j \theta_1(j) z_1(j)^2 \leq 2H^2(\theta_1, \theta_2)$.

The squared Hellinger distance can be related to other measures of divergence. For instance, it is of the same order as the Jensen-Shannon divergence and the triangular discrimination [16], [17]. Thus, the bound of Theorem 6 can be stated, up to small constants, in terms of these measures as well.

D. A Matching Lower Bound for a Class of Policy Sets

Lastly, we provide a lower bound that almost matches both Theorems 1 and 2 for a certain class of policy sets that we discuss shortly. This bound is analogous to the $\sqrt{KT \log N / \log K}$ lower bound proved in [3]. However, unlike [3], we rely on *fixed* sets of *stochastic* policies.

Theorem 7. *For any integer $q \geq 2$, there exists a problem structure where $K \geq 2$, $S^*(\Theta) = q$, and $S^*(\Theta) \log N = KD^*(\Theta)$; such that any algorithm, for sufficiently large T , suffers $\Omega\left(\sqrt{S^*(\Theta)T \frac{\log N}{\log S^*(\Theta)}}\right)$ regret.*

In the type of structure referred to in the theorem, the arms are divided into M sections (where $\frac{K}{M} = q$) and each policy is a uniform distributions supported over M arms such that its support contains an arm from each section. When the policy set contains all such policies ($(\frac{K}{M})^M$ in total), this problem becomes equivalent to playing M bandit problems simultaneously with the choice of policy at each round dictating an arm choice at each game. The distinction is that only the loss of one such arm is observed, while the player aims to minimize the average regret of the M games. This type of structure (albeit with a different type of feedback) is commonly used to prove lower bounds for combinatorial bandits, see [18] for example. An adaptation of the proof of Theorem 5 in [18] to our case (see [15, Appendix C]) leads to a lower bound of $\Omega(\sqrt{KT})$, from which the theorem follows by using that $S^*(\Theta) = S(\Theta) = \frac{K}{M}$ and that $N = (\frac{K}{M})^M$. For what concerns the bound of Theorem 2, we have that $D^*(\Theta) = \log(\frac{K}{M})$. Thus, Theorems 1 and 2 provide the same bound since $S^*(\Theta) \log N = KD^*(\Theta)$.

VI. CONCLUSION

We analyzed two algorithms providing regret bounds that depend on information-theoretic quantities describing the policy set. We proved lower bounds for certain classes of policy sets highlighting instances where our regret bounds are nearly matched. Nevertheless, it remains to be seen if better guarantees can be proved in cases like Example 3 where $S^*(\Theta) \geq 1$ even if $\mathcal{TV}(\Theta)$ can be smaller. It is also interesting to see what are the optimal rates for cases like Example 4 where the bound of Theorem 1 is suboptimal, as we learned in Section IV, even if $S^*(\Theta)$ is of the same order as $\mathcal{TV}(\Theta)$. Another direction is to investigate the possibility of adapting to the policy set structure if the distributions are not known beforehand.

REFERENCES

- [1] T. Lattimore and C. Szepesvári, *Bandit algorithms*. Cambridge University Press, 2020.
- [2] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, “The non-stochastic multiarmed bandit problem,” *SIAM Journal on Computing*, vol. 32, no. 1, pp. 48–77, 2002.
- [3] Y. Seldin and G. Lugosi, “A lower bound for multi-armed bandits with expert advice,” in *EWRL*, 2016.
- [4] S. Bubeck, N. Cesa-Bianchi, and S. M. Kakade, “Towards minimax policies for online linear optimization with bandit feedback,” in *Conference on Learning Theory*. JMLR Workshop and Conference Proceedings, 2012, pp. 41–1.
- [5] A. M. Metelli, M. Papini, P. D’Oro, and M. Restelli, “Policy optimization as online learning with mediator feedback,” in *AAAI*, vol. 35, 2021, pp. 8958–8966.
- [6] R. Sen, K. Shanmugam, and S. Shakkottai, “Contextual bandits with stochastic experts,” in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2018, pp. 852–861.
- [7] H. B. McMahan and M. J. Streeter, “Tighter bounds for multi-armed bandits with expert advice,” in *COLT*, 2009.
- [8] S. Gerchinovitz and T. Lattimore, “Refined lower bounds for adversarial bandits,” *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [9] A. Guntuboyina, “Lower bounds for the minimax risk using f-divergences, and applications,” *IEEE Trans. Inf. Theor.*, vol. 57, no. 4, p. 2386–2399, apr 2011.
- [10] S. Bubeck, N. Cesa-Bianchi *et al.*, “Regret analysis of stochastic and nonstochastic multi-armed bandit problems,” *Foundations and Trends® in Machine Learning*, vol. 5, no. 1, pp. 1–122, 2012.
- [11] E. Hazan, “Introduction to online convex optimization,” *Foundations and Trends® in Optimization*, vol. 2, no. 3-4, pp. 157–325, 2016.
- [12] T. M. Cover and J. A. Thomas, *Elements of information theory (2. ed.)*. Wiley, 2006.
- [13] F. Orabona, “A modern introduction to online learning,” 2019. [Online]. Available: <https://arxiv.org/abs/1912.13213>
- [14] J. Lin, “Divergence measures based on the shannon entropy,” *IEEE Transactions on Information Theory*, vol. 37, no. 1, pp. 145–151, 1991.
- [15] K. Eldowa, N. Cesa-Bianchi, A. M. Metelli, and M. Restelli, “Information-theoretic regret bounds for bandits with fixed expert advice,” 2023. [Online]. Available: <https://arxiv.org/abs/2303.08102>
- [16] F. Topsoe, “Some inequalities for information divergence and related measures of discrimination,” *IEEE Transactions on information theory*, vol. 46, no. 4, pp. 1602–1609, 2000.
- [17] G. Lu and B. Li, “A class of new metrics based on triangular discrimination,” *Information*, vol. 6, no. 3, pp. 361–374, 2015.
- [18] J.-Y. Audibert, S. Bubeck, and G. Lugosi, “Regret in Online Combinatorial Optimization,” arXiv, Tech. Rep. arXiv:1204.4710, 2013.