



HAL
open science

Combinatorial Bandits Revisited

Richard Combes, Sadegh Talebi, Alexandre Proutière, Marc Lelarge

► **To cite this version:**

Richard Combes, Sadegh Talebi, Alexandre Proutière, Marc Lelarge. Combinatorial Bandits Revisited. NIPS 2015 - Twenty-ninth Conference on Neural Information Processing Systems, Dec 2015, Montreal, Canada. hal-01257796

HAL Id: hal-01257796

<https://hal.science/hal-01257796>

Submitted on 9 Apr 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Combinatorial Bandits Revisited

Richard Combes* M. Sadegh Talebi† Alexandre Proutiere† Marc Lelarge‡

* Centrale-Supelec, L2S, Gif-sur-Yvette, FRANCE

† Department of Automatic Control, KTH, Stockholm, SWEDEN

‡ INRIA & ENS, Paris, FRANCE

richard.combes@supelec.fr, {mstms,alepro}@kth.se, marc.lelarge@ens.fr

Abstract

This paper investigates stochastic and adversarial combinatorial multi-armed bandit problems. In the stochastic setting under semi-bandit feedback, we derive a problem-specific regret lower bound, and discuss its scaling with the dimension of the decision space. We propose ESCB, an algorithm that efficiently exploits the structure of the problem and provide a finite-time analysis of its regret. ESCB has better performance guarantees than existing algorithms, and significantly outperforms these algorithms in practice. In the adversarial setting under bandit feedback, we propose COMBEXP, an algorithm with the same regret scaling as state-of-the-art algorithms, but with lower computational complexity for some combinatorial problems.

1 Introduction

Multi-Armed Bandits (MAB) problems [1] constitute the most fundamental sequential decision problems with an exploration vs. exploitation trade-off. In such problems, the decision maker selects an arm in each round, and observes a realization of the corresponding unknown reward distribution. Each decision is based on past decisions and observed rewards. The objective is to maximize the expected cumulative reward over some time horizon by balancing exploitation (arms with higher observed rewards should be selected often) and exploration (all arms should be explored to learn their average rewards). Equivalently, the performance of a decision rule or algorithm can be measured through its expected regret, defined as the gap between the expected reward achieved by the algorithm and that achieved by an oracle algorithm always selecting the best arm. MAB problems have found applications in many fields, including sequential clinical trials, communication systems, economics, see e.g. [2, 3].

In this paper, we investigate generic combinatorial MAB problems with linear rewards, as introduced in [4]. In each round $n \geq 1$, a decision maker selects an arm M from a finite set $\mathcal{M} \subset \{0, 1\}^d$ and receives a reward $M^\top X(n) = \sum_{i=1}^d M_i X_i(n)$. The reward vector $X(n) \in \mathbb{R}_+^d$ is unknown. We focus here on the case where all arms consist of the same number m of *basic* actions in the sense that $\|M\|_1 = m, \forall M \in \mathcal{M}$. After selecting an arm M in round n , the decision maker receives some feedback. We consider both (i) semi-bandit feedback under which after round n , for all $i \in \{1, \dots, d\}$, the component $X_i(n)$ of the reward vector is revealed if and only if $M_i = 1$; (ii) bandit feedback under which only the reward $M^\top X(n)$ is revealed. Based on the feedback received up to round $n - 1$, the decision maker selects an arm for the next round n , and her objective is to maximize her cumulative reward over a given time horizon consisting of T rounds. The challenge in these problems resides in the very large number of arms, i.e., in its combinatorial structure: the size of \mathcal{M} could well grow as d^m . Fortunately, one may hope to exploit the problem structure to speed up the exploration of sub-optimal arms.

We consider two instances of combinatorial bandit problems, depending on how the sequence of reward vectors is generated. We first analyze the case of stochastic rewards, where for all

Algorithm	LLR [9]	CUCB [10]	CUCB [11]	ESCB (Theorem 5)
Regret	$\mathcal{O}\left(\frac{m^3 d \Delta_{\max}}{\Delta_{\min}^2} \log(T)\right)$	$\mathcal{O}\left(\frac{m^2 d}{\Delta_{\min}} \log(T)\right)$	$\mathcal{O}\left(\frac{m d}{\Delta_{\min}} \log(T)\right)$	$\mathcal{O}\left(\frac{\sqrt{m d}}{\Delta_{\min}} \log(T)\right)$

Table 1: Regret upper bounds for stochastic combinatorial optimization under semi-bandit feedback.

$i \in \{1, \dots, d\}$, $(X_i(n))_{n \geq 1}$ are i.i.d. with Bernoulli distribution of unknown mean. The reward sequences are also independent across i . We then address the problem in the adversarial setting where the sequence of vectors $X(n)$ is arbitrary and selected by an *adversary* at the beginning of the experiment. In the stochastic setting, we provide sequential arm selection algorithms whose performance exceeds that of existing algorithms, whereas in the adversarial setting, we devise simple algorithms whose regret have the same scaling as that of state-of-the-art algorithms, but with lower computational complexity.

2 Contribution and Related Work

2.1 Stochastic combinatorial bandits under semi-bandit feedback

Contribution. (a) We derive an asymptotic (as the time horizon T grows large) regret lower bound satisfied by any algorithm (Theorem 1). This lower bound is *problem-specific* and *tight*: there exists an algorithm that attains the bound on all problem instances, although the algorithm might be computationally expensive. To our knowledge, such lower bounds have not been proposed in the case of stochastic combinatorial bandits. The dependency in m and d of the lower bound is unfortunately not explicit. We further provide a simplified lower bound (Theorem 2) and derive its scaling in (m, d) in specific examples.

(b) We propose ESCB (Efficient Sampling for Combinatorial Bandits), an algorithm whose regret scales at most as $\mathcal{O}(\sqrt{m d} \Delta_{\min}^{-1} \log(T))$ (Theorem 5), where Δ_{\min} denotes the expected reward difference between the best and the second-best arm. ESCB assigns an index to each arm. The index of given arm can be interpreted as performing likelihood tests with vanishing risk on its average reward. Our indexes are the natural extension of KL-UCB indexes defined for unstructured bandits [5]. Numerical experiments for some specific combinatorial problems are presented in the supplementary material, and show that ESCB significantly outperforms existing algorithms.

Related work. Previous contributions on stochastic combinatorial bandits focused on specific combinatorial structures, e.g. m -sets [6], matroids [7], or permutations [8]. Generic combinatorial problems were investigated in [9, 10, 11, 12]. The proposed algorithms, LLR and CUCB are variants of the UCB algorithm, and their performance guarantees are presented in Table 1. Our algorithms improve over LLR and CUCB by a multiplicative factor of \sqrt{m} .

2.2 Adversarial combinatorial problems under bandit feedback

Contribution. We present algorithm COMBEXP, whose regret is $\mathcal{O}\left(\sqrt{m^3 T (d + m^{1/2} \lambda^{-1}) \log \mu_{\min}^{-1}}\right)$, where $\mu_{\min} = \min_{i \in [d]} \frac{1}{m|\mathcal{M}|} \sum_{M \in \mathcal{M}} M_i$ and λ is the smallest nonzero eigenvalue of the matrix $\mathbb{E}[MM^\top]$ when M is uniformly distributed over \mathcal{M} (Theorem 6). For most problems of interest $m(d\lambda)^{-1} = \mathcal{O}(1)$ [4] and $\mu_{\min}^{-1} = \mathcal{O}(\text{poly}(d))$, so that COMBEXP has $\mathcal{O}(\sqrt{m^3 d T \log(d/m)})$ regret. A known regret lower bound is $\Omega(m\sqrt{dT})$ [13], so the regret gap between COMBEXP and this lower bound scales at most as $m^{1/2}$ up to a logarithmic factor.

Related work. Adversarial combinatorial bandits have been extensively investigated recently, see [13] and references therein. Some papers consider specific instances of these problems, e.g., shortest-path routing [14], m -sets [15], and permutations [16]. For generic combinatorial problems, known regret lower bounds scale as $\Omega(\sqrt{m d T})$ and $\Omega(m\sqrt{dT})$ (if $d \geq 2m$) in the case of semi-bandit and bandit feedback, respectively [13]. In the case of semi-bandit feedback, [13] proposes

Algorithm	Regret
Lower Bound [13]	$\Omega\left(m\sqrt{dT}\right)$, if $d \geq 2m$
COMBAND [4]	$\mathcal{O}\left(\sqrt{m^3 d T \log \frac{d}{m} \left(1 + \frac{2m}{d\lambda}\right)}\right)$
EXP2 WITH JOHN'S EXPLORATION [18]	$\mathcal{O}\left(\sqrt{m^3 d T \log \frac{d}{m}}\right)$
COMBEXP (Theorem 6)	$\mathcal{O}\left(\sqrt{m^3 d T \left(1 + \frac{m^{1/2}}{d\lambda}\right) \log \mu_{\min}^{-1}}\right)$

Table 2: Regret of various algorithms for adversarial combinatorial bandits with bandit feedback. Note that for most combinatorial classes of interests, $m(d\lambda)^{-1} = \mathcal{O}(1)$ and $\mu_{\min}^{-1} = \mathcal{O}(\text{poly}(d))$.

OSMD, an algorithm whose regret upper bound matches the lower bound. [17] presents an algorithm with $\mathcal{O}(m\sqrt{dL_T^* \log(d/m)})$ regret where L_T^* is the total reward of the best arm after T rounds.

For problems with bandit feedback, [4] proposes COMBAND and derives a regret upper bound which depends on the structure of action set \mathcal{M} . For most problems of interest, the regret under COMBAND is upper-bounded by $\mathcal{O}(\sqrt{m^3 d T \log(d/m)})$. [18] addresses generic linear optimization with bandit feedback and the proposed algorithm, referred to as EXP2 WITH JOHN'S EXPLORATION, has a regret scaling at most as $\mathcal{O}(\sqrt{m^3 d T \log(d/m)})$ in the case of combinatorial structure. As we show next, for many combinatorial structures of interest (e.g. m -sets, matchings, spanning trees), COMBEXP yields the same regret as COMBAND and EXP2 WITH JOHN'S EXPLORATION, with lower computational complexity for a large class of problems. Table 2 summarises known regret bounds.

Example 1: m -sets. \mathcal{M} is the set of all d -dimensional binary vectors with m non-zero coordinates. We have $\mu_{\min} = \frac{m}{d}$ and $\lambda = \frac{m(d-m)}{d(d-1)}$ (refer to the supplementary material for details). Hence when $m = o(d)$, the regret upper bound of COMBEXP becomes $\mathcal{O}(\sqrt{m^3 d T \log(d/m)})$, which is the same as that of COMBAND and EXP2 WITH JOHN'S EXPLORATION.

Example 2: matchings. The set of arms \mathcal{M} is the set of perfect matchings in $\mathcal{K}_{m,m}$. $d = m^2$ and $|\mathcal{M}| = m!$. We have $\mu_{\min} = \frac{1}{m}$, and $\lambda = \frac{1}{m-1}$. Hence the regret upper bound of COMBEXP is $\mathcal{O}(\sqrt{m^5 T \log(m)})$, the same as for COMBAND and EXP2 WITH JOHN'S EXPLORATION.

Example 3: spanning trees. \mathcal{M} is the set of spanning trees in the complete graph \mathcal{K}_N . In this case, $d = \binom{N}{2}$, $m = N - 1$, and by Cayley's formula \mathcal{M} has N^{N-2} arms. $\log \mu_{\min}^{-1} \leq 2N$ for $N \geq 2$ and $\frac{m}{d\lambda} < 7$ when $N \geq 6$. The regret upper bound of COMBAND and EXP2 WITH JOHN'S EXPLORATION becomes $\mathcal{O}(\sqrt{N^5 T \log(N)})$. As for COMBEXP, we get the same regret upper bound $\mathcal{O}(\sqrt{N^5 T \log(N)})$.

3 Models and Objectives

We consider MAB problems where each arm M is a subset of m basic actions taken from $[d] = \{1, \dots, d\}$. For $i \in [d]$, $X_i(n)$ denotes the reward of basic action i in round n . In the stochastic setting, for each i , the sequence of rewards $(X_i(n))_{n \geq 1}$ is i.i.d. with Bernoulli distribution with mean θ_i . Rewards are assumed to be independent across actions. We denote by $\theta = (\theta_1, \dots, \theta_d)^\top \in \Theta = [0, 1]^d$ the vector of unknown expected rewards of the various basic actions. In the adversarial setting, the reward vector $X(n) = (X_1(n), \dots, X_d(n))^\top \in [0, 1]^d$ is arbitrary, and the sequence $(X(n), n \geq 1)$ is decided (but unknown) at the beginning of the experiment.

The set of arms \mathcal{M} is an arbitrary subset of $\{0, 1\}^d$, such that each of its elements M has m basic actions. Arm M is identified with a binary column vector $(M_1, \dots, M_d)^\top$, and we have $\|M\|_1 = m$, $\forall M \in \mathcal{M}$. At the beginning of each round n , a policy π , selects an arm $M^\pi(n) \in \mathcal{M}$ based on the arms chosen in previous rounds and their observed rewards. The reward of arm $M^\pi(n)$ selected in round n is $\sum_{i \in [d]} M_i^\pi(n) X_i(n) = M^\pi(n)^\top X(n)$.

We consider both semi-bandit and bandit feedbacks. Under semi-bandit feedback and policy π , at the end of round n , the outcome of basic actions $X_i(n)$ for all $i \in M^\pi(n)$ are revealed to the decision maker, whereas under bandit feedback, $M^\pi(n)^\top X(n)$ only can be observed.

Let Π be the set of all feasible policies. The objective is to identify a policy in Π maximizing the cumulative expected reward over a finite time horizon T . The expectation is here taken with respect to possible randomness in the rewards (in the stochastic setting) and the possible randomization in the policy. Equivalently, we aim at designing a policy that minimizes regret, where the regret of policy $\pi \in \Pi$ is defined by:

$$R^\pi(T) = \max_{M \in \mathcal{M}} \mathbb{E} \left[\sum_{n=1}^T M^\top X(n) \right] - \mathbb{E} \left[\sum_{n=1}^T M^\pi(n)^\top X(n) \right].$$

Finally, for the stochastic setting, we denote by $\mu_M(\theta) = M^\top \theta$ the expected reward of arm M , and let $M^*(\theta) \in \mathcal{M}$, or M^* for short, be any arm with maximum expected reward: $M^*(\theta) \in \arg \max_{M \in \mathcal{M}} \mu_M(\theta)$. In what follows, to simplify the presentation, we assume that the optimal M^* is unique. We further define: $\mu^*(\theta) = M^{*\top} \theta$, $\Delta_{\min} = \min_{M \neq M^*} \Delta_M$ where $\Delta_M = \mu^*(\theta) - \mu_M(\theta)$, and $\Delta_{\max} = \max_M (\mu^*(\theta) - \mu_M(\theta))$.

4 Stochastic Combinatorial Bandits under Semi-bandit Feedback

4.1 Regret Lower Bound

Given θ , define the set of parameters that cannot be distinguished from θ when selecting action $M^*(\theta)$, and for which arm $M^*(\theta)$ is suboptimal:

$$B(\theta) = \{\lambda \in \Theta : M_i^*(\theta)(\theta_i - \lambda_i) = 0, \forall i, \mu^*(\lambda) > \mu^*(\theta)\}.$$

We define $\mathcal{X} = (\mathbb{R}^+)^{|\mathcal{M}|}$ and $\text{kl}(u, v)$ the Kullback-Leibler divergence between Bernoulli distributions of respective means u and v , i.e., $\text{kl}(u, v) = u \log(u/v) + (1-u) \log((1-u)/(1-v))$. Finally, for $(\theta, \lambda) \in \Theta^2$, we define the vector $\text{kl}(\theta, \lambda) = (\text{kl}(\theta_i, \lambda_i))_{i \in [d]}$.

We derive a regret lower bound valid for any *uniformly good* algorithm. An algorithm π is uniformly good iff $R^\pi(T) = o(T^\alpha)$ for all $\alpha > 0$ and all parameters $\theta \in \Theta$. The proof of this result relies on a general result on controlled Markov chains [19].

Theorem 1 For all $\theta \in \Theta$, for any uniformly good policy $\pi \in \Pi$, $\liminf_{T \rightarrow \infty} \frac{R^\pi(T)}{\log(T)} \geq c(\theta)$, where $c(\theta)$ is the optimal value of the optimization problem:

$$\inf_{x \in \mathcal{X}} \sum_{M \in \mathcal{M}} x_M (M^*(\theta) - M)^\top \theta \quad \text{s.t.} \quad \left(\sum_{M \in \mathcal{M}} x_M M \right)^\top \text{kl}(\theta, \lambda) \geq 1, \forall \lambda \in B(\theta). \quad (1)$$

Observe first that optimization problem (3) is a semi-infinite linear program which can be solved for any fixed θ , but its optimal value is difficult to compute explicitly. Determining how $c(\theta)$ scales as a function of the problem dimensions d and m is not obvious. Also note that (3) has the following interpretation: assume that (3) has a unique solution x^* . Then any uniformly good algorithm must select action M at least $x_M^* \log(T)$ times over the T first rounds. From [19], we know that there exists an algorithm which is asymptotically optimal, so that its regret matches the lower bound of Theorem 1. However this algorithm suffers from two problems: it is computationally infeasible for large problems since it involves solving (3) T times, furthermore the algorithm has no finite time performance guarantees, and numerical experiments suggests that its finite time performance on typical problems is rather poor. Further remark that if \mathcal{M} is the set of singletons (classical bandit), Theorem 1 reduces to the Lai-Robbins bound [20] and if \mathcal{M} is the set of m -sets (bandit with multiple plays), Theorem 1 reduces to the lower bound derived in [6]. Finally, Theorem 1 can be generalized in a straightforward manner for when rewards belong to a one-parameter exponential family of distributions (e.g., Gaussian, Exponential, Gamma etc.) by replacing kl by the appropriate divergence measure.

A Simplified Lower Bound We now study how the regret $c(\theta)$ scales as a function of the problem dimensions d and m . To this aim, we present a simplified regret lower bound. Given θ , we say that a set $\mathcal{H} \subset \mathcal{M} \setminus M^*$ has property $P(\theta)$ iff, for all $(M, M') \in \mathcal{H}^2$, $M \neq M'$ we have $M_i M'_i (1 - M_i^*(\theta)) = 0$ for all i . We may now state Theorem 2.

Theorem 2 *Let \mathcal{H} be a maximal (inclusion-wise) subset of \mathcal{M} with property $P(\theta)$. Define $\beta(\theta) = \min_{M \neq M^*} \frac{\Delta_M}{|\mathcal{M} \setminus M^*|}$. Then:*

$$c(\theta) \geq \sum_{M \in \mathcal{H}} \frac{\beta(\theta)}{\max_{i \in M \setminus M^*} \text{kl}\left(\theta_i, \frac{1}{|\mathcal{M} \setminus M^*|} \sum_{j \in M^* \setminus M} \theta_j\right)}.$$

Corollary 1 *Let $\theta \in [a, 1]^d$ for some constant $a > 0$ and \mathcal{M} be such that each arm $M \in \mathcal{M}$, $M \neq M^*$ has at most k suboptimal basic actions. Then $c(\theta) = \Omega(|\mathcal{H}|/k)$.*

Theorem 2 provides an explicit regret lower bound. Corollary 1 states that $c(\theta)$ scales at least with the size of \mathcal{H} . For most combinatorial sets, $|\mathcal{H}|$ is proportional to $d - m$ (see supplementary material for some examples), which implies that in these cases, one cannot obtain a regret smaller than $\mathcal{O}((d - m)\Delta_{\min}^{-1} \log(T))$. This result is intuitive since $d - m$ is the number of parameters not observed when selecting the optimal arm. The algorithms proposed below have a regret of $\mathcal{O}(d\sqrt{m}\Delta_{\min}^{-1} \log(T))$, which is acceptable since typically, \sqrt{m} is much smaller than d .

4.2 Algorithms

Next we present ESCB, an algorithm for stochastic combinatorial bandits that relies on arm indexes as in UCB1 [21] and KL-UCB [5]. We derive finite-time regret upper bounds for ESCB that hold even if we assume that $\|M\|_1 \leq m$, $\forall M \in \mathcal{M}$, instead of $\|M\|_1 = m$, so that arms may have different numbers of basic actions.

4.2.1 Indexes

ESCB relies on arm indexes. In general, an index of arm M in round n , say $b_M(n)$, should be defined so that $b_M(n) \geq M^\top \theta$ with high probability. Then as for UCB1 and KL-UCB, applying the principle of optimism against uncertainty, a natural way to devise algorithms based on indexes is to select in each round the arm with the highest index. Under a given algorithm, at time n , we define $t_i(n) = \sum_{s=1}^n M_i(s)$ the number of times basic action i has been sampled. The empirical mean reward of action i is then defined as $\hat{\theta}_i(n) = (1/t_i(n)) \sum_{s=1}^n X_i(s) M_i(s)$ if $t_i(n) > 0$ and $\hat{\theta}_i(n) = 0$ otherwise. We define the corresponding vectors $t(n) = (t_i(n))_{i \in [d]}$ and $\hat{\theta}(n) = (\hat{\theta}_i(n))_{i \in [d]}$.

The indexes we propose are functions of the round n and of $\hat{\theta}(n)$. Our first index for arm M , referred to as $b_M(n, \hat{\theta}(n))$ or $b_M(n)$ for short, is an extension of KL-UCB index. Let $f(n) = \log(n) + 4m \log(\log(n))$. $b_M(n, \hat{\theta}(n))$ is the optimal value of the following optimization problem:

$$\max_{q \in \Theta} M^\top q \quad \text{s.t.} \quad (Mt(n))^\top \text{kl}(\hat{\theta}(n), q) \leq f(n), \quad (2)$$

where we use the convention that for $v, u \in \mathbb{R}^d$, $vu = (v_i u_i)_{i \in [d]}$. As we show later, $b_M(n)$ may be computed efficiently using a line search procedure similar to that used to determine KL-UCB index.

Our second index $c_M(n, \hat{\theta}(n))$ or $c_M(n)$ for short is a generalization of the UCB1 and UCB-tuned indexes:

$$c_M(n) = M^\top \hat{\theta}(n) + \sqrt{\frac{f(n)}{2} \left(\sum_{i=1}^d \frac{M_i}{t_i(n)} \right)}$$

Note that, in the classical bandit problems with independent arms, i.e., when $m = 1$, $b_M(n)$ reduces to the KL-UCB index (which yields an asymptotically optimal algorithm) and $c_M(n)$ reduces to the UCB-tuned index. The next theorem provides generic properties of our indexes. An important consequence of these properties is that the expected number of times where $b_{M^*}(n, \hat{\theta}(n))$ or $c_{M^*}(n, \hat{\theta}(n))$ underestimate $\mu^*(\theta)$ is *finite*, as stated in the corollary below.

Theorem 3 (i) For all $n \geq 1$, $M \in \mathcal{M}$ and $\tau \in [0, 1]^d$, we have $b_M(n, \tau) \leq c_M(n, \tau)$.
(ii) There exists $C_m > 0$ depending on m only such that, for all $M \in \mathcal{M}$ and $n \geq 2$:

$$\mathbb{P}[b_M(n, \hat{\theta}(n)) \leq M^\top \theta] \leq C_m n^{-1} (\log(n))^{-2}.$$

Corollary 2 $\sum_{n \geq 1} \mathbb{P}[b_{M^*}(n, \hat{\theta}(n)) \leq \mu^*] \leq 1 + C_m \sum_{n \geq 2} n^{-1} (\log(n))^{-2} < \infty$.

Statement (i) in the above theorem is obtained combining Pinsker and Cauchy-Schwarz inequalities. The proof of statement (ii) is based on a concentration inequality on sums of empirical KL divergences proven in [22]. It enables to control the fluctuations of multivariate empirical distributions for exponential families. It should also be observed that indexes $b_M(n)$ and $c_M(n)$ can be extended in a straightforward manner to the case of continuous linear bandit problems, where the set of arms is the unit sphere and one wants to maximize the dot product between the arm and an unknown vector. $b_M(n)$ can also be extended to the case where reward distributions are not Bernoulli but lie in an exponential family (e.g. Gaussian, Exponential, Gamma, etc.), replacing kl by a suitably chosen divergence measure. A close look at $c_M(n)$ reveals that the indexes proposed in [10], [11], and [9] are too conservative to be optimal in our setting: there the ‘‘confidence bonus’’ $\sum_{i=1}^d \frac{M_i}{t_i(n)}$ was replaced by (at least) $m \sum_{i=1}^d \frac{M_i}{t_i(n)}$. Note that [10], [11] assume that the various basic actions are arbitrarily correlated, while we assume independence among basic actions. When independence does not hold, [11] provides a problem instance where the regret is at least $\mathcal{O}(\frac{md}{\Delta_{\min}} \log(T))$. This does not contradict our regret upper bound (scaling as $\mathcal{O}(\frac{d\sqrt{m}}{\Delta_{\min}} \log(T))$), since we have added the independence assumption.

4.2.2 Index computation

While the index $c_M(n)$ is explicit, $b_M(n)$ is defined as the solution to an optimization problem. We show that it may be computed by a simple line search. For $\lambda \geq 0$, $w \in [0, 1]$ and $v \in \mathbb{N}$, define:

$$g(\lambda, w, v) = \left(1 - \lambda v + \sqrt{(1 - \lambda v)^2 + 4wv\lambda}\right) / 2.$$

Fix n , M , $\hat{\theta}(n)$ and $t(n)$. Define $I = \{i : M_i = 1, \hat{\theta}_i(n) \neq 1\}$, and for $\lambda > 0$, define:

$$F(\lambda) = \sum_{i \in I} t_i(n) \text{kl}(\hat{\theta}_i(n), g(\lambda, \hat{\theta}_i(n), t_i(n))).$$

Theorem 4 If $I = \emptyset$, $b_M(n) = \|M\|_1$. Otherwise: (i) $\lambda \mapsto F(\lambda)$ is strictly increasing, and $F(\mathbb{R}^+) = \mathbb{R}^+$. (ii) Define λ^* as the unique solution to $F(\lambda) = f(n)$. Then $b_M(n) = \|M\|_1 - |I| + \sum_{i \in I} g(\lambda^*, \hat{\theta}_i(n), t_i(n))$.

Theorem 4 shows that $b_M(n)$ can be computed using a line search procedure such as bisection, as this computation amounts to solving the nonlinear equation $F(\lambda) = f(n)$, where F is strictly increasing. The proof of Theorem 4 follows from KKT conditions and the convexity of KL divergence.

4.2.3 The ESCB Algorithm

The pseudo-code of ESCB is presented in Algorithm 1. We consider two variants of the algorithm based on the choice of the index $\xi_M(n)$: ESCB-1 when $\xi_M(n) = b_M(n)$ and ESCB-2 if $\xi_M(n) = c_M(n)$. In practice, ESCB-1 outperforms ESCB-2. Introducing ESCB-2 is however instrumental in the regret analysis of ESCB-1 (in view of Theorem 3 (i)). The following theorem provides a finite time analysis of our ESCB algorithms. The proof of this theorem borrows some ideas from the proof of [11, Theorem 3].

Theorem 5 The regret under algorithms $\pi \in \{\text{ESCB-1}, \text{ESCB-2}\}$ satisfies for all $T \geq 1$:

$$R^\pi(T) \leq 16d\sqrt{m}\Delta_{\min}^{-1}f(T) + 4dm^3\Delta_{\min}^{-2} + C'_m,$$

where $C'_m \geq 0$ does not depend on θ , d and T . As a consequence $R^\pi(T) = \mathcal{O}(d\sqrt{m}\Delta_{\min}^{-1} \log(T))$ when $T \rightarrow \infty$.

Algorithm 1 ESCB

for $n \geq 1$ **do**
 Select arm $M(n) \in \arg \max_{M \in \mathcal{M}} \xi_M(n)$.
 Observe the rewards, and update $t_i(n)$ and $\hat{\theta}_i(n), \forall i \in M(n)$.
end for

Algorithm 2 COMBEXP

Initialization: Set $q_0 = \mu^0, \gamma = \frac{\sqrt{m \log \mu_{\min}^{-1}}}{\sqrt{m \log \mu_{\min}^{-1}} + \sqrt{C(Cm^2d+m)T}}$ and $\eta = \gamma C$, with $C = \frac{\lambda}{m^{3/2}}$.
for $n \geq 1$ **do**
 Mixing: Let $q'_{n-1} = (1 - \gamma)q_{n-1} + \gamma\mu^0$.
 Decomposition: Select a distribution p_{n-1} over \mathcal{M} such that $\sum_M p_{n-1}(M)M = mq'_{n-1}$.
 Sampling: Select a random arm $M(n)$ with distribution p_{n-1} and incur a reward $Y_n = \sum_i X_i(n)M_i(n)$.
 Estimation: Let $\Sigma_{n-1} = \mathbb{E}[MM^\top]$, where M has law p_{n-1} . Set $\tilde{X}(n) = Y_n \Sigma_{n-1}^+ M(n)$, where Σ_{n-1}^+ is the pseudo-inverse of Σ_{n-1} .
 Update: Set $\tilde{q}_n(i) \propto q_{n-1}(i) \exp(\eta \tilde{X}_i(n)), \forall i \in [d]$.
 Projection: Set q_n to be the projection of \tilde{q}_n onto the set \mathcal{P} using the KL divergence.
end for

ESCB with time horizon T has a complexity of $\mathcal{O}(|\mathcal{M}|T)$ as neither b_M nor c_M can be written as $M^\top y$ for some vector $y \in \mathbb{R}^d$. Assuming that the offline (static) combinatorial problem is solvable in $\mathcal{O}(V(\mathcal{M}))$ time, the complexity of CUCB algorithm in [10] and [11] after T rounds is $\mathcal{O}(V(\mathcal{M})T)$. Thus, if the offline problem is efficiently implementable, i.e., $V(\mathcal{M}) = \mathcal{O}(\text{poly}(d))$, CUCB is efficient, whereas ESCB is not since $|\mathcal{M}|$ may have exponentially many elements. In §2.5 of the supplement, we provide an extension of ESCB called EPOCH-ESCB, that attains almost the same regret as ESCB while enjoying much better computational complexity.

5 Adversarial Combinatorial Bandits under Bandit Feedback

We now consider adversarial combinatorial bandits with bandit feedback. We start with the following observation:

$$\max_{M \in \mathcal{M}} M^\top X = \max_{\mu \in Co(\mathcal{M})} \mu^\top X,$$

with $Co(\mathcal{M})$ the convex hull of \mathcal{M} . We embed \mathcal{M} in the d -dimensional simplex by dividing its elements by m . Let \mathcal{P} be this scaled version of $Co(\mathcal{M})$.

Inspired by OSMD [13, 18], we propose the COMBEXP algorithm, where the KL divergence is the Bregman divergence used to project onto \mathcal{P} . Projection using the KL divergence is addressed in [23]. We denote the KL divergence between distributions q and p in \mathcal{P} by $\text{KL}(p, q) = \sum_{i \in [d]} p(i) \log \frac{p(i)}{q(i)}$. The projection of distribution q onto a closed convex set Ξ of distributions is $p^* = \arg \min_{p \in \Xi} \text{KL}(p, q)$.

Let $\underline{\lambda}$ be the smallest nonzero eigenvalue of $\mathbb{E}[MM^\top]$, where M is uniformly distributed over \mathcal{M} . We define the exploration-inducing distribution $\mu^0 \in \mathcal{P}$: $\mu_i^0 = \frac{1}{m|\mathcal{M}|} \sum_{M \in \mathcal{M}} M_i, \forall i \in [d]$, and let $\mu_{\min} = \min_i m\mu_i^0$. μ^0 is the distribution over basic actions $[d]$ induced by the uniform distribution over \mathcal{M} . The pseudo-code for COMBEXP is shown in Algorithm 2. The KL projection in COMBEXP ensures that $mq_{n-1} \in Co(\mathcal{M})$. There exists λ , a distribution over \mathcal{M} such that $mq_{n-1} = \sum_M \lambda(M)M$. This guarantees that the system of linear equations in the *decomposition step* is consistent. We propose to perform the *projection step* (the KL projection of \tilde{q} onto \mathcal{P}) using interior-point methods [24]. We provide a simpler method in §3.4 of the supplement. The *decomposition step* can be efficiently implemented using the algorithm of [25]. The following theorem provides a regret upper bound for COMBEXP.

Theorem 6 For all $T \geq 1$: $R^{\text{COMBEXP}}(T) \leq 2\sqrt{m^3 T \left(d + \frac{m^{1/2}}{\lambda}\right) \log \mu_{\min}^{-1}} + \frac{m^{5/2}}{\lambda} \log \mu_{\min}^{-1}$.

For most classes of \mathcal{M} , we have $\mu_{\min}^{-1} = \mathcal{O}(\text{poly}(d))$ and $m(d\lambda)^{-1} = \mathcal{O}(1)$ [4]. For these classes, COMBEXP has a regret of $\mathcal{O}(\sqrt{m^3 d T \log(d/m)})$, which is a factor $\sqrt{m \log(d/m)}$ off the lower bound (see Table 2).

It might not be possible to compute the projection step exactly, and this step can be solved up to accuracy ϵ_n in round n . Namely we find q_n such that $\text{KL}(q_n, \tilde{q}_n) - \min_{p \in \Xi} \text{KL}(p, \tilde{q}_n) \leq \epsilon_n$. Proposition 1 shows that for $\epsilon_n = \mathcal{O}(n^{-2} \log^{-3}(n))$, the approximate projection gives the same regret as when the projection is computed exactly. Theorem 7 gives the computational complexity of COMBEXP with approximate projection. When $Co(\mathcal{M})$ is described by polynomially (in d) many linear equalities/inequalities, COMBEXP is efficiently implementable and its running time scales (almost) linearly in T . Proposition 1 and Theorem 7 easily extend to other OSMD-type algorithms and thus might be of independent interest.

Proposition 1 *If the projection step of COMBEXP is solved up to accuracy $\epsilon_n = \mathcal{O}(n^{-2} \log^{-3}(n))$, we have:*

$$R^{\text{COMBEXP}}(T) \leq 2\sqrt{2m^3 T \left(d + \frac{m^{1/2}}{\lambda}\right) \log \mu_{\min}^{-1}} + \frac{2m^{5/2}}{\lambda} \log \mu_{\min}^{-1}.$$

Theorem 7 *Assume that $Co(\mathcal{M})$ is defined by c linear equalities and s linear inequalities. If the projection step is solved up to accuracy $\epsilon_n = \mathcal{O}(n^{-2} \log^{-3}(n))$, then COMBEXP has time complexity $\mathcal{O}(T[\sqrt{s}(c+d)^3 \log(T) + d^4])$.*

The time complexity of COMBEXP can be reduced by exploiting the structure of \mathcal{M} (See [24, page 545]). In particular, if inequality constraints describing $Co(\mathcal{M})$ are box constraints, the time complexity of COMBEXP is $\mathcal{O}(T[c^2 \sqrt{s}(c+d) \log(T) + d^4])$.

The computational complexity of COMBEXP is determined by the structure of $Co(\mathcal{M})$ and COMBEXP has $\mathcal{O}(T \log(T))$ time complexity due to the efficiency of interior-point methods. In contrast, the computational complexity of COMBAND depends on the complexity of sampling from \mathcal{M} . COMBAND may have a time complexity that is super-linear in T (see [16, page 217]). For instance, consider the matching problem described in Section 2. We have $c = 2m$ equality constraints and $s = m^2$ box constraints, so that the time complexity of COMBEXP is: $\mathcal{O}(m^3 T \log(T))$. It is noted that using [26, Algorithm 1], the cost of decomposition in this case is $\mathcal{O}(m^4)$. On the other hand, COMBBAND has a time complexity of $\mathcal{O}(m^{10} F(T))$, with F a super-linear function, as it requires to approximate a permanent, requiring $\mathcal{O}(m^{10})$ operations per round. Thus, COMBEXP has much lower complexity than COMBAND and achieves the same regret.

6 Conclusion

We have investigated stochastic and adversarial combinatorial bandits. For stochastic combinatorial bandits with semi-bandit feedback, we have provided a tight, problem-dependent regret lower bound that, in most cases, scales at least as $\mathcal{O}((d-m)\Delta_{\min}^{-1} \log(T))$. We proposed ESCB, an algorithm with $\mathcal{O}(d\sqrt{m}\Delta_{\min}^{-1} \log(T))$ regret. We plan to reduce the gap between this regret guarantee and the regret lower bound, as well as investigate the performance of EPOCH-ESCB. For adversarial combinatorial bandits with bandit feedback, we proposed the COMBEXP algorithm. There is a gap between the regret of COMBEXP and the known regret lower bound in this setting, and we plan to reduce it as much as possible.

Acknowledgments

A. Proutiere's research is supported by the ERC FSA grant, and the SSF ICT-Psi project.

References

- [1] Herbert Robbins. Some aspects of the sequential design of experiments. In *Herbert Robbins Selected Papers*, pages 169–177. Springer, 1985.
- [2] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–222, 2012.
- [3] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*, volume 1. Cambridge University Press Cambridge, 2006.
- [4] Nicolò Cesa-Bianchi and Gábor Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422, 2012.
- [5] Aurélien Garivier and Olivier Cappé. The KL-UCB algorithm for bounded stochastic bandits and beyond. In *Proc. of COLT*, 2011.
- [6] Venkatachalam Anantharam, Pravin Varaiya, and Jean Walrand. Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays-part i: iid rewards. *Automatic Control, IEEE Transactions on*, 32(11):968–976, 1987.
- [7] Branislav Kveton, Zheng Wen, Azin Ashkan, Hoda Eydgahi, and Brian Eriksson. Matroid bandits: Fast combinatorial optimization with learning. In *Proc. of UAI*, 2014.
- [8] Yi Gai, Bhaskar Krishnamachari, and Rahul Jain. Learning multiuser channel allocations in cognitive radio networks: A combinatorial multi-armed bandit formulation. In *Proc. of IEEE DySpan*, 2010.
- [9] Yi Gai, Bhaskar Krishnamachari, and Rahul Jain. Combinatorial network optimization with unknown variables: Multi-armed bandits with linear rewards and individual observations. *IEEE/ACM Trans. on Networking*, 20(5):1466–1478, 2012.
- [10] Wei Chen, Yajun Wang, and Yang Yuan. Combinatorial multi-armed bandit: General framework and applications. In *Proc. of ICML*, 2013.
- [11] Branislav Kveton, Zheng Wen, Azin Ashkan, and Csaba Szepesvari. Tight regret bounds for stochastic combinatorial semi-bandits. In *Proc. of AISTATS*, 2015.
- [12] Zheng Wen, Azin Ashkan, Hoda Eydgahi, and Branislav Kveton. Efficient learning in large-scale combinatorial semi-bandits. In *Proc. of ICML*, 2015.
- [13] Jean-Yves Audibert, Sébastien Bubeck, and Gábor Lugosi. Regret in online combinatorial optimization. *Mathematics of Operations Research*, 39(1):31–45, 2013.
- [14] András György, Tamás Linder, Gábor Lugosi, and György Ottucsák. The on-line shortest path problem under partial monitoring. *Journal of Machine Learning Research*, 8(10), 2007.
- [15] Satyen Kale, Lev Reyzin, and Robert Schapire. Non-stochastic bandit slate problems. *Advances in Neural Information Processing Systems*, pages 1054–1062, 2010.
- [16] Nir Ailon, Kohei Hatano, and Eiji Takimoto. Bandit online optimization over the permutahedron. In *Algorithmic Learning Theory*, pages 215–229. Springer, 2014.
- [17] Gergely Neu. First-order regret bounds for combinatorial semi-bandits. In *Proc. of COLT*, 2015.
- [18] Sébastien Bubeck, Nicolò Cesa-Bianchi, and Sham M. Kakade. Towards minimax policies for online linear optimization with bandit feedback. *Proc. of COLT*, 2012.
- [19] Todd L. Graves and Tze Leung Lai. Asymptotically efficient adaptive choice of control laws in controlled markov chains. *SIAM J. Control and Optimization*, 35(3):715–743, 1997.
- [20] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- [21] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002.
- [22] Stefan Magureanu, Richard Combes, and Alexandre Proutiere. Lipschitz bandits: Regret lower bounds and optimal algorithms. *Proc. of COLT*, 2014.
- [23] I. Csiszár and P.C. Shields. *Information theory and statistics: A tutorial*. Now Publishers Inc, 2004.
- [24] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [25] H. D. Sherali. A constructive proof of the representation theorem for polyhedral sets based on fundamental definitions. *American Journal of Mathematical and Management Sciences*, 7(3-4):253–270, 1987.
- [26] David P. Helmbold and Manfred K. Warmuth. Learning permutations with exponential weights. *Journal of Machine Learning Research*, 10:1705–1736, 2009.
- [27] Richard Combes and Alexandre Proutiere. Unimodal bandits: Regret lower bounds and optimal algorithms. *arXiv:1405.5096*, 2014.

- [28] Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- [29] J.W. Moon and L. Moser. On cliques in graphs. *Israel Journal of Mathematics*, 3:23–28, 1965.
- [30] Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2003.

Supplementary Materials and Proofs

A Stochastic Combinatorial Bandits: Regret Lower Bounds

A.1 Proof of Theorem 1

To derive regret lower bounds, we apply the techniques used by Graves and Lai [19] to investigate efficient adaptive decision rules in controlled Markov chains. First we give an overview of their general framework.

Consider a controlled Markov chain $(X_n)_{n \geq 0}$ on a finite state space \mathcal{S} with a control set U . The transition probabilities given control $u \in U$ are parameterized by θ taking values in a compact metric space Θ : the probability to move from state x to state y given the control u and the parameter θ is $p(x, y; u, \theta)$. The parameter θ is not known. The decision maker is provided with a finite set of stationary control laws $G = \{g_1, \dots, g_K\}$, where each control law g_j is a mapping from \mathcal{S} to U : when control law g_j is applied in state x , the applied control is $u = g_j(x)$. It is assumed that if the decision maker always selects the same control law g , the Markov chain is then irreducible with stationary distribution π_θ^g . Now the reward obtained when applying control u in state x is denoted by $r(x, u)$, so that the expected reward achieved under control law g is: $\mu_\theta(g) = \sum_x r(x, g(x))\pi_\theta^g(x)$. There is an optimal control law given θ whose expected reward is denoted by $\mu_\theta^* = \max_{g \in G} \mu_\theta(g)$. Now the objective of the decision maker is to sequentially select control laws so as to maximize the expected reward up to a given time horizon T . As for MAB problems, the performance of a decision scheme can be quantified through the notion of regret which compares the expected reward to that obtained by always applying the optimal control law.

Proof. The parameter θ takes values in $[0, 1]^d$. The Markov chain has values in $\mathcal{S} = \{0, 1\}^d$. The set of controls corresponds to the set of feasible actions \mathcal{M} , and the set of control laws is also \mathcal{M} . These laws are constant, in the sense that the control applied by control law $M \in \mathcal{M}$ does not depend on the state of the Markov chain, and corresponds to selecting action M . The transition probabilities are given as follows: for all $x, y \in \mathcal{S}$,

$$p(x, y; M, \theta) = p(y; M, \theta) = \prod_{i \in [d]} p_i(y_i; M, \theta),$$

where for all $i \in [d]$, if $M_i = 0$, $p_i(0; M, \theta) = 1$, and if $M_i = 1$, $p_i(y_i; M, \theta) = \theta_i^{y_i} (1 - \theta_i)^{1 - y_i}$. Finally, the reward $r(y, M)$ is defined by $r(y, M) = M^\top y$. Note that the state space of the Markov chain is here finite, and so, we do not need to impose any cost associated with switching control laws (see the discussion on page 718 in [19]).

We can now apply Theorem 1 in [19]. Note that the KL number under action M is

$$\text{kl}^M(\theta, \lambda) = \sum_{i \in [d]} M_i \text{kl}(\theta_i, \lambda_i).$$

From [19, Theorem 1], we conclude that for any uniformly good rule π ,

$$\liminf_{T \rightarrow \infty} \frac{R^\pi(T)}{\log(T)} \geq c(\theta),$$

where $c(\theta)$ is the optimal value of the following optimization problem:

$$\inf_{x_M \geq 0, M \in \mathcal{M}} \sum_{M \neq M^*} x_M (\mu^* - \mu_M(\theta)), \quad (3)$$

$$\text{s.t.} \quad \inf_{\lambda \in B(\theta)} \sum_{Q \neq M^*} x_Q \text{kl}^Q(\theta, \lambda) \geq 1. \quad (4)$$

The result is obtained by observing that $B(\theta) = \bigcup_{M \neq M^*} B_M(\theta)$, where

$$B_M(\theta) = \{\lambda \in \Theta : M_i^*(\theta)(\theta_i - \lambda_i) = 0, \forall i, \mu^*(\theta) < \mu_M(\lambda)\}.$$

□

A.2 Proof of Theorem 2

The proof proceeds in three steps. In the subsequent analysis, given the optimization problem \mathbf{P} , we use $\text{val}(\mathbf{P})$ to denote its optimal value.

Step 1. In this step, first we introduce an equivalent formulation for problem (3) above by simplifying its constraints. We show that constraint (4) is equivalent to:

$$\inf_{\lambda \in B_M(\theta)} \sum_{i \in M \setminus M^*} \text{kl}(\theta_i, \lambda_i) \sum_{Q \in \mathcal{M}} Q_i x_Q \geq 1, \quad \forall M \neq M^*.$$

Observe that:

$$\sum_{Q \neq M^*} x_Q \text{kl}^Q(\theta, \lambda) = \sum_{Q \neq M^*} x_Q \sum_{i \in [d]} Q_i \text{kl}(\theta_i, \lambda_i) = \sum_{i \in [d]} \text{kl}(\theta_i, \lambda_i) \sum_{Q \neq M^*} Q_i x_Q.$$

Fix $M \neq M^*$. In view of the definition of $B_M(\theta)$, we can find $\lambda \in B_M(\theta)$ such that $\lambda_i = \theta_i, \forall i \in ([d] \setminus M) \cup M^*$. Thus, for the r.h.s. of the M -th constraint in (4), we get:

$$\begin{aligned} \inf_{\lambda \in B_M(\theta)} \sum_{Q \neq M^*} x_Q \text{kl}^Q(\theta, \lambda) &= \inf_{\lambda \in B_M(\theta)} \sum_{i \in [d]} \text{kl}(\theta_i, \lambda_i) \sum_{Q \neq M^*} Q_i x_Q \\ &= \inf_{\lambda \in B_M(\theta)} \sum_{i \in M \setminus M^*} \text{kl}(\theta_i, \lambda_i) \sum_Q Q_i x_Q, \end{aligned}$$

and therefore problem (3) can be equivalently written as:

$$c(\theta) = \inf_{x_M \geq 0, M \in \mathcal{M}} \sum_{M \neq M^*} x_M (\mu^* - \mu_M(\theta)), \quad (5)$$

$$\text{s.t.} \quad \inf_{\lambda \in B_M(\theta)} \sum_{i \in M \setminus M^*} \text{kl}(\theta_i, \lambda_i) \sum_Q Q_i x_Q \geq 1, \quad \forall M \neq M^*. \quad (6)$$

Next, we formulate an LP whose value gives a lower bound for $c(\theta)$. Define $\hat{\lambda}(M) = (\hat{\lambda}_i(M), i \in [d])$ with

$$\hat{\lambda}_i(M) = \begin{cases} \frac{1}{|M \setminus M^*|} \sum_{j \in M^* \setminus M} \theta_j & \text{if } i \in M \setminus M^*, \\ \theta_i & \text{otherwise.} \end{cases}$$

Clearly $\hat{\lambda}(M) \in B_M(\theta)$, and therefore:

$$\inf_{\lambda \in B_M(\theta)} \sum_{i \in M \setminus M^*} \text{kl}(\theta_i, \lambda_i) \sum_Q Q_i x_Q \leq \sum_{i \in M \setminus M^*} \text{kl}(\theta_i, \hat{\lambda}_i(M)) \sum_Q Q_i x_Q,$$

Then, we can write:

$$c(\theta) \geq \inf_{x \geq 0} \sum_{M \neq M^*} \Delta_M x_M \quad (7)$$

$$\text{s.t.} \quad \sum_{i \in M \setminus M^*} \text{kl}(\theta_i, \hat{\lambda}_i(M)) \sum_Q Q_i x_Q \geq 1, \quad \forall M \neq M^*. \quad (8)$$

For any $M \neq M^*$ introduce: $g_M = \max_{i \in M \setminus M^*} \text{kl}(\theta_i, \hat{\lambda}_i(M))$. Now we form $\mathbf{P1}$ as follows:

$$\mathbf{P1}: \quad \inf_{x \geq 0} \sum_{M \neq M^*} \Delta_M x_M \quad (9)$$

$$\text{s.t.} \quad \sum_{i \in M \setminus M^*} \sum_Q Q_i x_Q \geq \frac{1}{g_M}, \quad \forall M \neq M^*. \quad (10)$$

Observe that $c(\theta) \geq \text{val}(\mathbf{P1})$ since the feasible set of problem (7) is contained in that of $\mathbf{P1}$.

Step 2. In this step, we formulate an LP to give a lower bound for $\text{val}(\mathbf{P1})$. To this end, for any suboptimal basic action $i \in [d]$, we define $z_i = \sum_M M_i x_M$. Further, we let $z = [z_i, i \in [d]]$. Next, we represent the objective of $\mathbf{P1}$ in terms of z , and give a lower bound for it as follows:

$$\begin{aligned}
\sum_{M \neq M^*} \Delta_M x_M &= \sum_{M \neq M^*} x_M \sum_{i \in M \setminus M^*} \frac{\Delta_M}{|M \setminus M^*|} \\
&= \sum_{M \neq M^*} x_M \sum_{i \in [d] \setminus M^*} \frac{\Delta_M}{|M \setminus M^*|} M_i \\
&\geq \min_{M \neq M^*} \frac{\Delta_M}{|M \setminus M^*|} \cdot \sum_{i \in [d] \setminus M^*} \sum_{M' \neq M^*} M'_i x_{M'} \\
&= \min_{M \neq M^*} \frac{\Delta_M}{|M \setminus M^*|} \cdot \sum_{i \in [d] \setminus M^*} z_i \\
&= \beta(\theta) \sum_{i \in [d] \setminus M^*} z_i.
\end{aligned}$$

Then, defining

$$\begin{aligned}
\mathbf{P2}: \quad &\inf_{z \geq 0} \beta(\theta) \sum_{i \in [d] \setminus M^*} z_i \\
\text{s.t.} \quad &\sum_{i \in M \setminus M^*} z_i \geq \frac{1}{g_M}, \quad \forall M \neq M^*,
\end{aligned}$$

yields: $\text{val}(\mathbf{P1}) \geq \text{val}(\mathbf{P2})$.

Step 3. Introduce set \mathcal{H} satisfying property $P(\theta)$ as stated in Section 4. Now define

$$\mathcal{Z} = \left\{ z \in \mathbb{R}_+^d : \sum_{i \in M \setminus M^*} z_i \geq \frac{1}{g_M}, \quad \forall M \in \mathcal{H} \right\},$$

and

$$\mathbf{P3}: \quad \inf_{z \in \mathcal{Z}} \beta(\theta) \sum_{i \in [d] \setminus M^*} z_i.$$

Observe that $\text{val}(\mathbf{P2}) \geq \text{val}(\mathbf{P3})$ since the feasible set of $\mathbf{P2}$ is contained in \mathcal{Z} . The definition of \mathcal{H} implies that $\sum_{i \in [d] \setminus M^*} z_i = \sum_{M \in \mathcal{H}} \sum_{i \in M \setminus M^*} z_i$. It then follows that

$$\begin{aligned}
\text{val}(\mathbf{P3}) &= \sum_{M \in \mathcal{H}} \frac{\beta(\theta)}{g_M} \\
&\geq \sum_{M \in \mathcal{H}} \frac{\beta(\theta)}{\max_{i \in M \setminus M^*} \text{kl}(\theta_i, \hat{\lambda}_i(M))} \\
&= \sum_{M \in \mathcal{H}} \frac{\beta(\theta)}{\max_{i \in M \setminus M^*} \text{kl}\left(\theta_i, \frac{1}{|M \setminus M^*|} \sum_{j \in M^* \setminus M} \theta_j\right)}.
\end{aligned}$$

The proof is completed by observing that: $c(\theta) \geq \text{val}(\mathbf{P1}) \geq \text{val}(\mathbf{P2}) \geq \text{val}(\mathbf{P3})$. □




Figure 1: Matchings in $\mathcal{K}_{4,4}$: (a) The optimal matching M^* , (b)-(g) Elements of \mathcal{H} .

A.3 Proof of Corollary 1

Fix $M \neq M^*$. For any $i \in M \setminus M^*$, we have:

$$\begin{aligned}
\text{kl}\left(\theta_i, \frac{1}{|M \setminus M^*|} \sum_{j \in M^* \setminus M} \theta_j\right) &\leq \frac{1}{|M \setminus M^*|} \sum_{j \in M^* \setminus M} \text{kl}(\theta_i, \theta_j) \quad (\text{By convexity of } \text{kl}(\cdot, \cdot)) \\
&\leq \frac{1}{|M \setminus M^*|} \sum_{j \in M^* \setminus M} \frac{(\theta_i - \theta_j)^2}{\theta_j(1 - \theta_j)} \\
&\leq \frac{1}{|M \setminus M^*|} \sum_{j \in M^* \setminus M} \frac{(1 - \theta_j)^2}{\theta_j(1 - \theta_j)} \\
&\leq \frac{1}{|M \setminus M^*|} \sum_{j \in M^* \setminus M} \left(\frac{1}{\theta_j} - 1\right) \\
&\leq \frac{1}{\min_{j \in M^* \setminus M} \theta_j} - 1 \\
&\leq \frac{1}{a} - 1,
\end{aligned}$$

where the second inequality follows from the inequality $\text{kl}(p, q) \leq \frac{(p-q)^2}{q(1-q)}$ for all $(p, q) \in [0, 1]^2$. Moreover, we have that

$$\beta(\theta) = \min_{M \neq M^*} \frac{\Delta^M}{|M \setminus M^*|} \geq \frac{\Delta_{\min}}{\max_M |M \setminus M^*|} = \frac{\Delta_{\min}}{k}.$$

Applying Theorem 2, we get:

$$c(\theta) \geq \sum_{M \in \mathcal{H}} \frac{\beta(\theta)}{\max_{i \in M \setminus M^*} \text{kl}\left(\theta_i, \frac{1}{|M \setminus M^*|} \sum_{j \in M^* \setminus M} \theta_j\right)} \geq \frac{\Delta_{\min}^a}{k(1-a)} |\mathcal{H}|,$$

which gives the required lower bound and completes the proof. \square

A.4 Examples of Scaling of the Lower Bound

A.4.1 Matchings

In the first example, we assume that \mathcal{M} is the set of perfect matchings in the complete bipartite graph $\mathcal{K}_{m,m}$, with $|\mathcal{M}| = m!$ and $d = m^2$. A maximal subset \mathcal{H} of \mathcal{M} satisfying property $P(\theta)$ can be constructed by adding all matchings that differ from the optimal matching by only two edges, see Figure 1 for illustration in the case of $m = 4$. Here $|\mathcal{H}| = \binom{m}{2}$ and thus, $|\mathcal{H}|$ scales as $m^2 = d$.

A.4.2 Spanning trees

Consider the problem of finding the minimum spanning tree in a complete graph \mathcal{K}_N . This corresponds to letting \mathcal{M} be the set of all spanning trees in \mathcal{K}_N , where $|\mathcal{M}| = N^{N-2}$ (Cayley's formula). In this case, we have $d = \binom{N}{2} = \frac{N(N-1)}{2}$, which is the number of edges of \mathcal{K}_N , and $m = N - 1$.




Figure 2: Spanning trees in \mathcal{K}_5 : (a) The optimal spanning tree M^* , (b)-(g) Elements of \mathcal{H} .




Figure 3: Routing in a grid: (a) Grid topology with source (red) and destination (blue) nodes, (b) Optimal path M^* , (c)-(e) Elements of \mathcal{H} .

A maximal subset \mathcal{H} of \mathcal{M} satisfying property $P(\theta)$ can be constructed by composing all spanning trees that differ from the optimal tree by one edge only, see Figure 2. In this case, \mathcal{H} has $d - m = \frac{(N-1)(N-2)}{2}$ elements.

A.4.3 Routing in a grid

Now we give an example, in which $|\mathcal{H}|$ is not scaling as $\Omega(d)$. Consider routing in an N -by- N directed grid, whose topology is shown in Figure 3(a) where the source (resp. destination) node is shown in red (resp. blue). Here \mathcal{M} is the set of all $\binom{2N-2}{N-1}$ paths with $m = 2(N - 1)$ edges. We further have $d = 2N(N - 1)$. In this example, elements of any maximal set \mathcal{H} satisfying $P(\theta)$ do not cover all basic actions. For instance, for the grid shown in Figure 3(a), the two edges incident to the right lower corner do not appear in any arm in \mathcal{H} . It can be easily verified that in this case, $|\mathcal{H}|$ scales as N rather than $N^2 = d$.

A.5 Lower Bound Example

Here we provide an example, motivated by [11], to investigate the tightness of the regret bounds of our algorithms. Consider the topology shown in Figure 4, where there are $\frac{d}{m}$ paths, each consisting

of m links. Let parameter θ be defined such that

$$\theta_i = \begin{cases} 0.5 & \text{if } i \text{ belongs to the first path} \\ 0.5 - \delta & \text{otherwise.} \end{cases}$$

The first path is the optimal path and for any $M \neq 1$ we have: $\Delta_M = \Delta = m\delta$. Since various paths are independent, this problem reduces to a classical MAB problem with $\frac{d}{m}$ arms. It is observed that the total reward of each path is the sum of m independent Bernoulli random variables with the same parameter. Hence, it is distributed according to a binomial distribution. It then follows that




Figure 4: Lower bound example

$$\begin{aligned} \liminf_{T \rightarrow \infty} \frac{R(T)}{\log(T)} &\geq \sum_{M \neq M^*} \frac{\Delta_M}{\text{KL}(\text{Bin}(m, 0.5 - \delta), \text{Bin}(m, 0.5))} \\ &= \left(\frac{d}{m} - 1 \right) \cdot \frac{\Delta}{m \text{kl}(0.5 - \delta, 0.5)} \\ &\geq \frac{(d - m)\Delta}{4m^2\delta^2} \\ &= \frac{d - m}{4\Delta}, \end{aligned}$$

where the first equality follows from the fact that the KL divergence between two Binomial distributions with respective parameters (m, u) and (m, v) is $m \text{kl}(u, v)$, and where the last step is due to inequality $\text{kl}(x, y) \leq \frac{(x-y)^2}{y(1-y)}$ for all $x, y \in (0, 1)$.

B Stochastic Combinatorial Bandits: Regret Analysis of ESCB

We use the convention that for $v, u \in \mathbb{R}^d$, $vu = (v_i u_i)_{i \in [d]}$.

B.1 A concentration inequality

We first recall Lemma 1, a concentration inequality derived in [22, Theorem 2].

Lemma 1 *There exists a number $C_m > 0$ depending only on m such that, for all M and all $n \geq 2$:*

$$\mathbb{P}[(Mt(n))^\top \text{kl}(\hat{\theta}(n), \theta) \geq f(n)] \leq C_m n^{-1} (\log(n))^{-2}.$$

B.2 Proof of Theorem 3

First statement:

Consider $q \in \Theta$, and apply the Cauchy-Schwartz inequality:

$$M^\top (q - \hat{\theta}(n)) = \sum_{i=1}^d \sqrt{t_i(n)} (q_i - \hat{\theta}_i(n)) \frac{M_i}{\sqrt{t_i(n)}} \leq \sqrt{\sum_{i=1}^d M_i t_i(n) (q_i - \hat{\theta}_i(n))^2} \sqrt{\sum_{i=1}^d \frac{M_i}{t_i(n)}}$$

By Pinsker's inequality, for all $(p, q) \in [0, 1]^2$ we have $2(p - q)^2 \leq \text{kl}(p, q)$ so that:

$$M^\top (q - \hat{\theta}(n)) \leq \sqrt{\frac{(Mt(n))^\top \text{kl}(\hat{\theta}(n), q)}{2}} \sqrt{\sum_{i=1}^d \frac{M_i}{t_i(n)}}$$

Hence, $(Mt(n))^\top \text{kl}(\hat{\theta}(n), q) \leq f(n)$ implies:

$$M^\top q = M^\top \hat{\theta}(n) + M^\top (q - \hat{\theta}(n)) \leq M^\top \hat{\theta}(n) + \sqrt{\frac{f(n)}{2} \sum_{i=1}^d \frac{M_i}{t_i(n)}} = c_M(n).$$

so that, by definition of $b_M(n)$, we have $b_M(n) \leq c_M(n)$.

Second statement:

If $(Mt(n))^\top \text{kl}(\hat{\theta}(n), \theta) \leq f(n)$ then, by definition of $b_M(n)$ we have $b_M(n) \geq M^\top \theta$. Therefore, using Lemma 1, there exists C_m such that for all $n \geq 2$ we have:

$$\mathbb{P}[b_M(n) < M^\top \theta] \leq \mathbb{P}[(Mt(n))^\top \text{kl}(\hat{\theta}(n), \theta) \geq f(n)] \leq C_m n^{-1} (\log(n))^{-2},$$

which concludes the proof.

B.3 Proof of Theorem 4

We recall the following facts about the KL divergence kl , for all $p \in [0, 1]$:

- (i) $q \mapsto \text{kl}(p, q)$ is strictly convex on $[0, 1]$ and attains its minimum at p , with $\text{kl}(p, p) = 0$.
- (ii) Its derivative with respect to the second parameter $q \mapsto \text{kl}'(p, q) = \frac{q-p}{q(1-q)}$ is strictly increasing on $(p, 1)$.
- (iii) For $p < 1$, we have $\text{kl}(p, q) \xrightarrow{q \rightarrow 1^-} \infty$ and $\text{kl}'(p, q) \xrightarrow{q \rightarrow 1^-} \infty$.

Consider M and n fixed throughout the proof. Define $I = \{i \in M : \hat{\theta}_i(n) \neq 1\}$. Consider $q^* \in \Theta$ the optimal solution of optimization problem:

$$\begin{aligned} \max_{q \in \Theta} M^\top q \\ \text{s.t. } (Mt(n))^\top \text{kl}(\hat{\theta}(n), q) \leq f(n). \end{aligned}$$

so that $b_M(n) = M^\top q^*$. Consider $i \notin M$, then $M^\top q$ does not depend on q_i and from (i) we get $q_i = \hat{\theta}_i(n)$. Now consider $i \in M$. From (i) we get that $1 \geq q_i^* \geq \hat{\theta}_i(n)$. Hence $q_i^* = 1$ if $\hat{\theta}_i(n) = 1$. If I is empty, then $q_i^* = 1$ for all $i \in M$, so that $b_M(n) = \|M\|_1$.

Consider the case where $I \neq \emptyset$. From (iii) and the fact that $t(n)^\top \text{kl}(\hat{\theta}(n), q^*) < \infty$ we get $\hat{\theta}_i(n) \leq q_i^* < 1$. From the Karush-Kuhn-Tucker (KKT) conditions, there exists $\lambda^* > 0$ such that for all $i \in I$:

$$1 = \lambda^* t_i(n) \text{kl}'(\hat{\theta}_i(n), q_i^*).$$

For $\lambda > 0$ define $\hat{\theta}_i(n) \leq \bar{q}_i(\lambda) < 1$ a solution to the equation:

$$1 = \lambda t_i(n) \text{kl}'(\hat{\theta}_i(n), \bar{q}_i(\lambda)).$$

From (i) we have that $\lambda \mapsto \bar{q}_i(\lambda)$ is uniquely defined, is strictly decreasing and $\hat{\theta}_i(n) < \bar{q}_i(\lambda) < 1$.

From (iii) we get that $\bar{q}_i(\mathbb{R}^+) = [\hat{\theta}_i(n), 1]$. Define the function:

$$F(\lambda) = \sum_{i \in I} t_i(n) \text{kl}(\hat{\theta}(n), \bar{q}_i(\lambda)).$$

From the reasoning below, F is well defined, strictly increasing and $F(\mathbb{R}^+) = \mathbb{R}^+$. Therefore, λ^* is the unique solution to $F(\lambda^*) = f(n)$, and $q_i^* = \bar{q}_i(\lambda^*)$. Furthermore, replacing kl' by its expression we obtain the quadratic equation:

$$\bar{q}_i(\lambda)^2 + \bar{q}_i(\lambda)(\lambda t_i(n) - 1) - \lambda t_i(n) \hat{\theta}_i(n) = 0.$$

Solving for $\bar{q}_i(\lambda)$, we obtain that $\bar{q}_i(\lambda) = g(\lambda, \hat{\theta}_i(n), t_i(n))$, which concludes the proof. \square

B.4 Proof of Theorem 5

To prove Theorem 5, we borrow some ideas from proof of [11, Theorem 3].

For any $n \in \mathbb{N}$, $s \in \mathbb{R}^d$, and $M \in \mathcal{M}$ define $h_{n,s,M} = \sqrt{\frac{f(n)}{2} \sum_{i=1}^d \frac{M_i}{s_i}}$, and introduce the following events:

$$\begin{aligned} G_n &= \{(M^*t(n))^\top \text{kl}(\hat{\theta}(n), \theta) > f(n)\}, \\ H_{i,n} &= \{M_i(n) = 1, |\hat{\theta}_i(n) - \theta_i| \geq m^{-1} \Delta_{\min}/2\}, \quad H_n = \cup_{i=1}^d H_{i,n}, \\ F_n &= \{\Delta_{M(n)} \leq 2h_{T,t(n),M(n)}\}. \end{aligned}$$

Then the regret can be bounded as:

$$\begin{aligned} R^\pi(T) &= \mathbb{E}\left[\sum_{n=1}^T \Delta_{M(n)}\right] \leq \mathbb{E}\left[\sum_{n=1}^T \Delta_{M(n)}(\mathbb{1}\{G_n\} + \mathbb{1}\{H_n\})\right] + \mathbb{E}\left[\sum_{n=1}^T \Delta_{M(n)} \mathbb{1}\{\overline{G_n}, \overline{H_n}\}\right] \\ &\leq m \mathbb{E}\left[\sum_{n=1}^T (\mathbb{1}\{G_n\} + \mathbb{1}\{H_n\})\right] + \mathbb{E}\left[\sum_{n=1}^T \Delta_{M(n)} \mathbb{1}\{\overline{G_n}, \overline{H_n}\}\right], \end{aligned}$$

since $\Delta_{M(n)} \leq m$.

Next we show that for any n such that $M(n) \neq M^*$, it holds that $\overline{G_n} \cup \overline{H_n} \subset F_n$. Recall that $c_M(n) \geq b_M(n)$ for any M and n (Theorem 3). Moreover, if $\overline{G_n}$ holds, we have $(M^*t(n))^\top \text{kl}(\hat{\theta}(n), \theta) \leq f(n)$, which by definition of b_M implies: $b_{M^*}(n) \geq M^{*\top} \theta$. Hence we have:

$$\begin{aligned} \mathbb{1}\{\overline{G_n}, \overline{H_n}, M(n) \neq M^*\} &= \mathbb{1}\{\overline{G_n}, \overline{H_n}, \xi_{M(n)}(n) \geq \xi_{M^*}(n)\} \\ &\leq \mathbb{1}\{\overline{H_n}, c_{M(n)}(n) \geq M^{*\top} \theta\} \\ &= \mathbb{1}\{\overline{H_n}, M(n)^\top \hat{\theta}(n) + h_{n,t(n),M(n)} \geq M^{*\top} \theta\} \\ &\leq \mathbb{1}\{M(n)^\top \theta + \Delta_{M(n)}/2 + h_{n,t(n),M(n)} \geq M^{*\top} \theta\} \\ &= \mathbb{1}\{2h_{n,t(n),M(n)} \geq \Delta_{M(n)}\} \\ &\leq \mathbb{1}\{2h_{T,t(n),M(n)} \geq \Delta_{M(n)}\} \\ &= \mathbb{1}\{F_n\}, \end{aligned}$$

where the second inequality follows from the fact that event $\overline{G_n}$ implies: $M(n)^\top \hat{\theta}(n) \leq M(n)^\top \theta + \Delta_{\min}/2 \leq M(n)^\top \theta + \Delta_{M(n)}/2$.

Hence, the regret is upper bounded by:

$$R^\pi(T) \leq m \mathbb{E}\left[\sum_{n=1}^T \mathbb{1}\{G_n\}\right] + m \mathbb{E}\left[\sum_{n=1}^T \mathbb{1}\{H_n\}\right] + \mathbb{E}\left[\sum_{n=1}^T \Delta_{M(n)} \mathbb{1}\{F_n\}\right].$$

We will prove the following inequalities: (i) $\mathbb{E}[\sum_{n=1}^T \mathbb{1}\{G_n\}] \leq m^{-1} C'_m$, with $C'_m \geq 0$ independent of θ , d , and T , (ii) $\mathbb{E}[\sum_{n=1}^T \mathbb{1}\{H_n\}] \leq 4dm^2 \Delta_{\min}^{-2}$, and (iii) $\mathbb{E}[\sum_{n=1}^T \Delta_{M(n)} \mathbb{1}\{F_n\}] \leq 16d\sqrt{m} \Delta_{\min}^{-1} f(T)$.

Hence as announced:

$$R^\pi(T) \leq 16d\sqrt{m} \Delta_{\min}^{-1} f(T) + 4dm^3 \Delta_{\min}^{-2} + C'_m.$$

Inequality (i): An application of Lemma 1 gives

$$\begin{aligned} \mathbb{E}\left[\sum_{n=1}^T \mathbb{1}\{G_n\}\right] &= \sum_{n=1}^T \mathbb{P}[(M^*t(n))^\top \text{kl}(\hat{\theta}(n), \theta) > f(n)] \\ &\leq 1 + \sum_{n \geq 2} C_m n^{-1} (\log(n))^{-2} \equiv m^{-1} C'_m < \infty. \end{aligned}$$

Inequality (ii): Fix i and n . Define $s = \sum_{n'=1}^n \mathbb{1}\{H_{n',i}\}$. Observe that $H_{n',i}$ implies $M_i(n') = 1$, hence $t_i(n) \geq s$. Therefore, applying [27, Lemma B.1], we have that $\sum_{n=1}^T \mathbb{P}[H_{n,i}] \leq 4m^2 \Delta_{\min}^{-2}$. Using the union bound: $\sum_{n=1}^T \mathbb{P}[H_n] \leq 4dm^2 \Delta_{\min}^{-2}$.

Inequality (iii): Let $\ell > 0$. For any n introduce the following events:

$$\begin{aligned} S_n &= \{i \in M(n) : t_i(n) \leq 4mf(T)\Delta_{M(n)}^{-2}\}, \\ A_n &= \{|S_n| \geq \ell\}, \\ B_n &= \{|S_n| < \ell, [\exists i \in M(n) : t_i(n) \leq 4\ell f(T)\Delta_{M(n)}^{-2}]\}. \end{aligned}$$

We claim that for any n such that $M(n) \neq M^*$, we have $F_n \subset (A_n \cup B_n)$. To prove this, we show that when F_n holds and $M(n) \neq M^*$, the event $\overline{A_n \cup B_n}$ cannot happen. Let n be a time instant such that $M(n) \neq M^*$ and F_n holds, and assume that $\overline{A_n \cup B_n} = \{|S_n| < \ell, [\forall i \in M(n) : t_i(n) > 4\ell f(T)\Delta_{M(n)}^{-2}]\}$ happens. Then F_n implies:

$$\begin{aligned} \Delta_{M(n)} &\leq 2h_{T,t(n),M(n)} = 2\sqrt{\frac{f(T)}{2}} \sqrt{\sum_{i \in [d] \setminus S_n} \frac{M_i(n)}{t_i(n)} + \sum_{i \in S_n} \frac{M_i(n)}{t_i(n)}} \\ &< 2\sqrt{\frac{f(T)}{2}} \sqrt{m \frac{\Delta_{M(n)}^2}{4mf(T)} + |S_n| \frac{\Delta_{M(n)}^2}{4\ell f(T)}} < \Delta_{M(n)}, \end{aligned} \quad (11)$$

where the last inequality uses the observation that $\overline{A_n \cup B_n}$ implies $|S_n| < \ell$. Clearly, (11) is a contradiction. Thus $F_n \subset (A_n \cup B_n)$ and consequently:

$$\sum_{n=1}^T \Delta_{M(n)} \mathbb{1}\{F_n\} \leq \sum_{n=1}^T \Delta_{M(n)} \mathbb{1}\{A_n\} + \sum_{n=1}^T \Delta_{M(n)} \mathbb{1}\{B_n\}. \quad (12)$$

To further bound the r.h.s. of the above, we introduce the following events for any i :

$$\begin{aligned} A_{i,n} &= A_n \cap \{i \in M(n), t_i(n) \leq 4mf(T)\Delta_{M(n)}^{-2}\}, \\ B_{i,n} &= B_n \cap \{i \in M(n), t_i(n) \leq 4\ell f(T)\Delta_{M(n)}^{-2}\}. \end{aligned}$$

It is noted that:

$$\sum_{i \in [d]} \mathbb{1}\{A_{i,n}\} = \mathbb{1}\{A_n\} \sum_{i \in [d]} \mathbb{1}\{i \in S_n\} = |S_n| \mathbb{1}\{A_n\} \geq \ell \mathbb{1}\{A_n\},$$

and hence: $\mathbb{1}\{A_n\} \leq \frac{1}{\ell} \sum_{i \in [d]} \mathbb{1}\{A_{i,n}\}$. Moreover $\mathbb{1}\{B_n\} \leq \sum_{i \in [d]} \mathbb{1}\{B_{i,n}\}$. Let each basic action i belong to K_i suboptimal arms, ordered based on their gaps as: $\Delta^{i,1} \geq \dots \geq \Delta^{i,K_i} > 0$. Also define $\Delta^{i,0} = \infty$. Plugging the above inequalities into (12), we have

$$\begin{aligned} \sum_{n=1}^T \Delta_{M(n)} \mathbb{1}\{F_n\} &\leq \sum_{n=1}^T \sum_{i=1}^d \frac{\Delta_{M(n)}}{\ell} \mathbb{1}\{A_{i,n}\} + \sum_{n=1}^T \sum_{i=1}^d \Delta_{M(n)} \mathbb{1}\{B_{i,n}\} \\ &= \sum_{n=1}^T \sum_{i=1}^d \frac{\Delta_{M(n)}}{\ell} \mathbb{1}\{A_{i,n}, M(n) \neq M^*\} + \sum_{n=1}^T \sum_{i=1}^d \Delta_{M(n)} \mathbb{1}\{B_{i,n}, M(n) \neq M^*\} \\ &\leq \sum_{n=1}^T \sum_{i=1}^d \sum_{k \in [K_i]} \frac{\Delta^{i,k}}{\ell} \mathbb{1}\{A_{i,n}, M(n) = k\} + \sum_{n=1}^T \sum_{i=1}^d \sum_{k \in [K_i]} \Delta^{i,k} \mathbb{1}\{B_{i,n}, M(n) = k\} \\ &\leq \sum_{i=1}^d \sum_{n=1}^T \sum_{k \in [K_i]} \frac{\Delta^{i,k}}{\ell} \mathbb{1}\{i \in M(n), t_i(n) \leq 4mf(T)(\Delta^{i,k})^{-2}, M(n) = k\} \\ &\quad + \sum_{i=1}^d \sum_{n=1}^T \sum_{k \in [K_i]} \Delta^{i,k} \mathbb{1}\{i \in M(n), t_i(n) \leq 4\ell f(T)(\Delta^{i,k})^{-2}, M(n) = k\} \\ &\leq \frac{8df(T)}{\Delta_{\min}} \left(\frac{m}{\ell} + \ell \right), \end{aligned}$$

where the last inequality follows from Lemma 2, which is proven next. The proof is completed by setting $\ell = \sqrt{m}$. \square

Lemma 2 *Let $C > 0$ be a constant independent of n . Then for any i such that $K_i \geq 1$:*

$$\sum_{n=1}^T \sum_{k=1}^{K_i} \mathbb{1}\{i \in M(n), t_i(n) \leq C(\Delta^{i,k})^{-2}, M(n) = k\} \Delta^{i,k} \leq \frac{2C}{\Delta_{\min}}.$$

Proof. We have:

$$\begin{aligned} & \sum_{n=1}^T \sum_{k=1}^{K_i} \mathbb{1}\{i \in M(n), t_i(n) \leq C(\Delta^{i,k})^{-2}, M(n) = k\} \Delta^{i,k} \\ &= \sum_{n=1}^T \sum_{k=1}^{K_i} \sum_{j=1}^k \mathbb{1}\{i \in M(n), t_i(n) \in (C(\Delta^{i,j-1})^{-2}, C(\Delta^{i,j})^{-2}], M(n) = k\} \Delta^{i,k} \\ &\leq \sum_{n=1}^T \sum_{k=1}^{K_i} \sum_{j=1}^k \mathbb{1}\{i \in M(n), t_i(n) \in (C(\Delta^{i,j-1})^{-2}, C(\Delta^{i,j})^{-2}], M(n) = k\} \Delta^{i,j} \\ &\leq \sum_{n=1}^T \sum_{k=1}^{K_i} \sum_{j=1}^{K_i} \mathbb{1}\{i \in M(n), t_i(n) \in (C(\Delta^{i,j-1})^{-2}, C(\Delta^{i,j})^{-2}], M(n) = k\} \Delta^{i,j} \\ &\leq \sum_{n=1}^T \sum_{j=1}^{K_i} \mathbb{1}\{i \in M(n), t_i(n) \in (C(\Delta^{i,j-1})^{-2}, C(\Delta^{i,j})^{-2}], M(n) \neq M^*\} \Delta^{i,j} \\ &\leq \frac{C}{\Delta^{i,1}} + \sum_{j=2}^{K_i} C((\Delta^{i,j})^{-2} - (\Delta^{i,j-1})^{-2}) \Delta^{i,j} \\ &\leq \frac{C}{\Delta^{i,1}} + \int_{\Delta^{i,K_i}}^{\Delta^{i,2}} Cx^{-2} dx \leq \frac{2C}{\Delta^{i,K_i}} \leq \frac{2C}{\Delta_{\min}}, \end{aligned}$$

which completes the proof. \square

B.5 EPOCH-ESCB: An algorithm with lower computational complexity

ESCB with time horizon T has a complexity of $\mathcal{O}(|\mathcal{M}|T)$ as neither b_M nor c_M can be written as $M^\top y$ for some vector $y \in \mathbb{R}^d$. Since \mathcal{M} typically has exponentially many elements, we deduce that ESCB is not computationally efficient. Assuming that the offline (static) combinatorial problem is solvable in $\mathcal{O}(V(\mathcal{M}))$ time, the complexity of CUCB algorithm in [10] and [11] after T rounds is $\mathcal{O}(V(\mathcal{M})T)$. Thus, if the offline problem is efficiently implementable, i.e., $V(\mathcal{M}) = \mathcal{O}(\text{poly}(d))$, CUCB is efficient, whereas ESCB is not. We next propose an extension to ESCB, called EPOCH-ESCB, that attains almost the same regret as ESCB while enjoying much better computational complexity.

EPOCH-ESCB algorithm in epochs of varying lengths. Epoch k comprises rounds $\{N_k, \dots, N_{k+1} - 1\}$, where N_{k+1} (and thus the length of the k -th epoch) is determined at time $n = N_k$. The algorithm simply consists in playing the arm with the maximal index at the beginning of every epoch, and playing the current leader (i.e., the arm with the highest empirical average reward) in the rest of rounds. If the leader is the arm with the maximal index, the length of epoch k will be set twice as long as the previous epoch $k - 1$, i.e., $N_{k+1} = N_k + 2(N_k - N_{k-1})$. Otherwise, it will be set to 1. In contrast to ESCB, EPOCH-ESCB computes the maximal index infrequently, and more precisely (almost) at an exponentially decreasing rate. Thus, one might expect that after T rounds, the maximal index will be computed $\mathcal{O}(\log(T))$ times. The pseudo-code of EPOCH-ESCB is presented in Algorithm 3.

We assess the performance of EPOCH-ESCB through numerical experiments in the next subsection, and leave the analysis of its regret as a future work. These experiments corroborate our conjecture that the complexity of EPOCH-ESCB after T rounds will be $\mathcal{O}(V(\mathcal{M})T + \log(T)|\mathcal{M}|)$. Compared to CUCB, the complexity is penalized by $|\mathcal{M}| \log(T)$, which may become dominated by the term $V(\mathcal{M})T$ as T grows large.

Algorithm 3 EPOCH-ESCB

Initialization: Set $k = 1$ and $N_0 = N_1 = 1$.

for $n \geq 1$ **do**

 Compute $L(n) \in \arg \max_{M \in \mathcal{M}} M^\top \hat{\theta}(n)$.

if $n = N_k$ **then**

 Select arm $M(n) \in \arg \max_{M \in \mathcal{M}} \xi_M(n)$.

if $M(n) = L(n)$ **then**

 Set $N_{k+1} = N_k + 2(N_k - N_{k-1})$.

else

 Set $N_{k+1} = N_k + 1$.

end if

 Increment k .

else

 Select arm $M(n) = L(n)$.

end if

 Observe the rewards, and update $t_i(n)$ and $\hat{\theta}_i(n), \forall i \in M(n)$.

end for




Figure 5: Regret of various algorithms for matchings with $a = 0.7$ and $b = 0.5$.

B.6 Numerical Experiments

In this section, we compare the performance of ESCB against existing algorithms through numerical experiments for some classes of \mathcal{M} . When implementing ESCB we replace $f(n)$ by $\log(n)$, ignoring the term proportional to $\log(\log(n))$, as is done when implementing KL-UCB in practice.

B.6.1 Experiment 1: Matching

In our first experiment, we consider the matching problem with $N_1 = N_2 = 5$, which corresponds to $d = 5^2 = 25$ and $m = 5$. We also set θ such that $\theta_i = a$ if $i \in M^*$, and $\theta_i = b$ otherwise, with $0 < b < a < 1$. In this case the lower bound becomes $c(\theta) = \frac{m(m-1)(a-b)}{2\text{kl}(b,a)}$.

Figure 5(a)-(b) depicts the regret of various algorithms for the case of $a = 0.7$ and $b = 0.5$. The curves in Figure 5(a) are shown with a 95% confidence interval. We observe that ESCB-1 has the lowest regret. Moreover, ESCB-2 significantly outperforms CUCB and LLR, and is close to ESCB-1. Moreover, we observe that the regret of EPOCH-ESCB attains is quite close to that of ESCB-2.

Figures 6(a)-(b) presents the regret of various algorithms for the case of $a = 0.95$ and $b = 0.3$. The difference compared to the former case is that ESCB-1 significantly outperforms ESCB-2. The reason is that in the former case, mean rewards of the most of the basic actions were close to $1/2$, for which the performance of UCB-type algorithms are closer to their KL-divergence based counterparts. On the other hand, when mean rewards are not close to $1/2$, there exists a significant performance gap between ESCB-1 and ESCB-2. Comparing the results with the ‘lower bound’ curve, we highlight that ESCB-1 gives close-to-optimal performance in both cases. Furthermore, similar to previous experiment, EPOCH-ESCB attains a regret whose curve is almost indistinguishable from that of ESCB-2.




Figure 6: Regret of various algorithms for matchings with $a = 0.95$ and $b = 0.3$.




Figure 7: Number of epochs in EPOCH-ESCB vs. time for Experiment 1 and 2 (%95 confidence interval).

The number of epochs in EPOCH-ESCB vs. time for the two examples is displayed in Figure 7(a)-(b), where the curves are shown with 95% confidence intervals. We observe that in both cases, the number of epochs grows at a rate proportional to $\log(n)/n$ at round n . Since the number of epochs is equal to the number of times the algorithm computes indexes, these curves suggest that index computation after n rounds requires a number of operations that scales as $|\mathcal{M}| \log(n)$.

B.6.2 Experiment 2: Spanning Trees

In the second experiment, we consider spanning trees problem described in Section A.4.2 for the case of $N = 5$. In this case, we have $d = \binom{5}{2} = 10$, $m = 4$, and $|\mathcal{M}| = 5^3 = 125$.

Figure 8 portrays the regret of various algorithms with 95% confidence intervals, with $\Delta_{\min} = 0.54$. Our algorithms significantly outperform CUCB and LLR.




Figure 8: Regret of various algorithms for spanning trees with $N = 5$ and $\Delta_{\min} = 0.54$.

C Proofs for Adversarial Combinatorial Bandits

C.1 Proof of Theorem 6

We first prove a simple result:

Lemma 3 For all $x \in \mathbb{R}^d$, we have $\Sigma_{n-1}^+ \Sigma_{n-1} x = \bar{x}$, where \bar{x} is the orthogonal projection of x onto $\text{span}(\mathcal{M})$, the linear space spanned by \mathcal{M} .

Proof: Note that for all $y \in \mathbb{R}^d$, if $\Sigma_{n-1} y = 0$, then we have

$$y^\top \Sigma_{n-1} y = \mathbb{E} [y^\top M M^\top y] = \mathbb{E} [(y^\top M)^2] = 0, \quad (13)$$

where M has law p_{n-1} such that $\sum_M M_i p_{n-1}(M) = q'_{n-1}(i)$, $\forall i \in [d]$ and $q'_{n-1} = (1-\gamma)q_{n-1} + \gamma\mu^0$. By definition of μ^0 , each $M \in \mathcal{M}$ has a positive probability. Hence, by (13), $y^\top M = 0$ for all $M \in \mathcal{M}$. In particular, we see that the linear application Σ_{n-1} restricted to $\text{span}(\mathcal{M})$ is invertible and is zero on $\text{span}(\mathcal{M})^\perp$, hence we have $\Sigma_{n-1}^+ \Sigma_{n-1} x = \bar{x}$. \square

Lemma 4 We have for any $\eta \leq \frac{\gamma\lambda}{m^{3/2}}$ and any $q \in \mathcal{P}$,

$$\sum_{n=1}^T q^\top \tilde{X}(n) - \sum_{n=1}^T q_{n-1}^\top \tilde{X}(n) \leq \frac{\eta}{2} \sum_{n=1}^T q_{n-1}^\top \tilde{X}^2(n) + \frac{\text{KL}(q, q_0)}{\eta},$$

where $\tilde{X}^2(n)$ is the vector that is the coordinate-wise square of $\tilde{X}(n)$.

Proof: We have

$$\text{KL}(q, \tilde{q}_n) - \text{KL}(q, q_{n-1}) = \sum_{i \in [d]} q(i) \log \frac{q_{n-1}(i)}{\tilde{q}_n(i)} = -\eta \sum_{i \in [d]} q(i) \tilde{X}_i(n) + \log Z_n,$$

with

$$\begin{aligned} \log Z_n &= \log \sum_{i \in [d]} q_{n-1}(i) \exp(\eta \tilde{X}_i(n)) \\ &\leq \log \sum_{i \in [d]} q_{n-1}(i) (1 + \eta \tilde{X}_i(n) + \eta^2 \tilde{X}_i^2(n)) \end{aligned} \quad (14)$$

$$\leq \eta q_{n-1}^\top \tilde{X}(n) + \eta^2 q_{n-1}^\top \tilde{X}^2(n), \quad (15)$$

where we used $\exp(z) \leq 1 + z + z^2$ for all $|z| \leq 1$ in (14) and $\log(1 + z) \leq z$ for all $z > -1$ in (15). Later we verify the condition for the former inequality.

Hence we have

$$\text{KL}(q, \tilde{q}_n) - \text{KL}(q, q_{n-1}) \leq \eta q_{n-1}^\top \tilde{X}(n) - \eta q^\top \tilde{X}(n) + \eta^2 q_{n-1}^\top \tilde{X}^2(n).$$

Generalized Pythagorean inequality (see Theorem 3.1 in [23]) gives

$$\text{KL}(q, q_n) + \text{KL}(q_n, \tilde{q}_n) \leq \text{KL}(q, \tilde{q}_n).$$

Since $\text{KL}(q_n, \tilde{q}_n) \geq 0$, we get

$$\text{KL}(q, q_n) - \text{KL}(q, q_{n-1}) \leq \eta q_{n-1}^\top \tilde{X}(n) - \eta q^\top \tilde{X}(n) + \eta^2 q_{n-1}^\top \tilde{X}^2(n).$$

Finally, summing over n gives

$$\sum_{n=1}^T (q^\top \tilde{X}(n) - q_{n-1}^\top \tilde{X}(n)) \leq \eta \sum_{n=1}^T q_{n-1}^\top \tilde{X}^2(n) + \frac{\text{KL}(q, q_0)}{\eta}.$$

To satisfy the condition for the inequality (14), i.e., $\eta|\tilde{X}_i(n)| \leq 1$, $\forall i \in [d]$, we find the upper bound for $\max_{i \in [d]} |\tilde{X}_i(n)|$ as follows:

$$\begin{aligned}
\max_{i \in [d]} |\tilde{X}_i(n)| &\leq \|\tilde{X}(n)\|_2 \\
&= \|\Sigma_{n-1}^+ M(n) Y_n\|_2 \\
&\leq m \|\Sigma_{n-1}^+ M(n)\|_2 \\
&\leq m \sqrt{M(n)^\top \Sigma_{n-1}^+ \Sigma_{n-1}^+ M(n)} \\
&\leq m \|M(n)\|_2 \sqrt{\lambda_{\max}(\Sigma_{n-1}^+ \Sigma_{n-1}^+)} \\
&= m^{3/2} \sqrt{\lambda_{\max}(\Sigma_{n-1}^+ \Sigma_{n-1}^+)} \\
&= m^{3/2} \lambda_{\max}(\Sigma_{n-1}^+) \\
&= \frac{m^{3/2}}{\lambda_{\min}(\Sigma_{n-1})},
\end{aligned}$$

where $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ respectively denote the maximum and the minimum nonzero eigenvalue of matrix A . Note that μ^0 induces uniform distribution over \mathcal{M} . Thus by $q'_{n-1} = (1-\gamma)q_{n-1} + \gamma\mu^0$ we see that p_{n-1} is a mixture of uniform distribution and the distribution induced by q_{n-1} . Note that, we have:

$$\lambda_{\min}(\Sigma_{n-1}) = \min_{\|x\|_2=1, x \in \text{span}(\mathcal{M})} x^\top \Sigma_{n-1} x.$$

Moreover, we have

$$x^\top \Sigma_{n-1} x = \mathbb{E}[x^\top M(n) M(n)^\top x] = \mathbb{E}[(M(n)^\top x)^2] \geq \gamma \mathbb{E}[(M^\top x)^2],$$

where in the last inequality M has law μ^0 . By definition, we have for any $x \in \text{span}(\mathcal{M})$ with $\|x\|_2 = 1$,

$$\mathbb{E}[(M^\top x)^2] \geq \lambda,$$

so that in the end, we get $\lambda_{\min}(\Sigma_{n-1}) \geq \gamma\lambda$, and hence $\eta|\tilde{X}_i(n)| \leq \frac{\eta m^{3/2}}{\gamma\lambda}$, $\forall i \in [d]$. Finally, we choose $\eta \leq \frac{\gamma\lambda}{m^{3/2}}$ to satisfy the condition for the inequality we used in (14). \square

We have

$$\mathbb{E}_n [\tilde{X}(n)] = \mathbb{E}_n [Y_n \Sigma_{n-1}^+ M(n)] = \mathbb{E}_n [\Sigma_{n-1}^+ M(n) M(n)^\top X(n)] = \Sigma_{n-1}^+ \Sigma_{n-1} X(n) = \overline{X(n)},$$

where the last equality follows from Lemma 3 and $\overline{X(n)}$ is the orthogonal projection of $X(n)$ onto $\text{span}(\mathcal{M})$. In particular, for any $mq' \in Co(\mathcal{M})$, we have

$$\mathbb{E}_n [mq'^\top \tilde{X}(n)] = mq'^\top \overline{X(n)} = mq'^\top X(n).$$

Moreover, we have:

$$\begin{aligned}
\mathbb{E}_n [q_{n-1}^\top \tilde{X}^2(n)] &= \sum_{i \in [d]} q_{n-1}(i) \mathbb{E}_n [\tilde{X}_i^2(n)] \\
&= \sum_{i \in [d]} \frac{q'_{n-1}(i) - \gamma\mu^0(i)}{1-\gamma} \mathbb{E}_n [\tilde{X}_i^2(n)] \\
&\leq \frac{1}{m(1-\gamma)} \sum_{i \in [d]} mq'_{n-1}(i) \mathbb{E}_n [\tilde{X}_i^2(n)] \\
&= \frac{1}{m(1-\gamma)} \mathbb{E}_n \left[\sum_{i \in [d]} \tilde{M}_i(n) \tilde{X}_i^2(n) \right],
\end{aligned}$$

where $\tilde{M}(n)$ is a random arm with the same law as $M(n)$ and independent of $M(n)$. Note that $\tilde{M}_i^2(n) = \tilde{M}_i(n)$, so that we have

$$\begin{aligned} \mathbb{E}_n \left[\sum_{i \in [d]} \tilde{M}_i(n) \tilde{X}_i^2(n) \right] &= \mathbb{E}_n \left[X(n)^\top M(n) M(n)^\top \Sigma_{n-1}^+ \tilde{M}(n) \tilde{M}(n)^\top \Sigma_{n-1}^+ M(n) M(n)^\top X(n) \right] \\ &\leq m^2 \mathbb{E}_n [M(n)^\top \Sigma_{n-1}^+ M(n)], \end{aligned}$$

where we used the bound $M(n)^\top X(n) \leq m$. By [4, Lemma 15], $\mathbb{E}_n [M(n)^\top \Sigma_{n-1}^+ M(n)] \leq d$, so that we have:

$$\mathbb{E}_n \left[q_{n-1}^\top \tilde{X}^2(n) \right] \leq \frac{md}{1-\gamma}.$$

Observe that

$$\begin{aligned} \mathbb{E}_n \left[q^{*\top} \tilde{X}(n) - q_{n-1}^\top \tilde{X}(n) \right] &= \mathbb{E}_n \left[q^{*\top} \tilde{X}(n) - (1-\gamma)q_{n-1}^\top \tilde{X}(n) - \gamma\mu^{0\top} \tilde{X}(n) \right] \\ &= \mathbb{E}_n \left[q^{*\top} \tilde{X}(n) - q_{n-1}^\top \tilde{X}(n) \right] + \gamma q_{n-1}^\top X(n) - \gamma\mu^{0\top} X(n) \\ &\leq \mathbb{E}_n \left[q^{*\top} \tilde{X}(n) - q_{n-1}^\top \tilde{X}(n) \right] + \gamma q_{n-1}^\top X(n) \\ &\leq \mathbb{E}_n \left[q^{*\top} \tilde{X}(n) - q_{n-1}^\top \tilde{X}(n) \right] + \gamma. \end{aligned}$$

Using Lemma 4 and the above bounds, we get with mq^* the optimal arm, i.e. $q^*(i) = \frac{1}{m}$ iff $M_i^* = 1$,

$$\begin{aligned} R^{\text{COMBEXP}}(T) &= \mathbb{E} \left[\sum_{n=1}^T mq^{*\top} \tilde{X}(n) - \sum_{n=1}^T mq_{n-1}^\top \tilde{X}(n) \right] \\ &\leq \mathbb{E} \left[\sum_{n=1}^T mq^{*\top} \tilde{X}(n) - \sum_{n=1}^T mq_{n-1}^\top \tilde{X}(n) \right] + m\gamma T \\ &\leq \frac{\eta m^2 d T}{1-\gamma} + \frac{m \log \mu_{\min}^{-1}}{\eta} + m\gamma T, \end{aligned}$$

since

$$\text{KL}(q^*, q_0) = -\frac{1}{m} \sum_{i \in M^*} \log m\mu_i^0 \leq \log \mu_{\min}^{-1}.$$

Choosing $\eta = \gamma C$ with $C = \frac{\lambda}{m^{3/2}}$ gives

$$\begin{aligned} R^{\text{COMBEXP}}(T) &\leq \frac{\gamma C m^2 d T}{1-\gamma} + \frac{m \log \mu_{\min}^{-1}}{\gamma C} + m\gamma T \\ &= \frac{C m^2 d + m - m\gamma}{1-\gamma} \gamma T + \frac{m \log \mu_{\min}^{-1}}{\gamma C} \\ &\leq \frac{(C m^2 d + m) \gamma T}{1-\gamma} + \frac{m \log \mu_{\min}^{-1}}{\gamma C}. \end{aligned}$$

The proof is completed by setting $\gamma = \frac{\sqrt{m \log \mu_{\min}^{-1}}}{\sqrt{m \log \mu_{\min}^{-1}} + \sqrt{C(C m^2 d + m) T}}$. \square

C.2 Proof of Proposition 1

We first provide a simple result:

Lemma 5 *The KL-divergence $z \mapsto \text{KL}(z, q)$ is 1-strongly convex with respect to the $\|\cdot\|_1$ norm.*

Proof. To prove the lemma, it suffices to show that for any $x, y \in \mathcal{P}$:

$$(\nabla \text{KL}(x, q) - \nabla \text{KL}(y, q))^\top (x - y) \geq \|x - y\|_1^2.$$

We have

$$\begin{aligned} (\nabla \text{KL}(x, q) - \nabla \text{KL}(y, q))^\top (x - y) &= \sum_{i \in [d]} \left(1 + \log \frac{x(i)}{q(i)} - 1 - \log \frac{y(i)}{q(i)}\right) (x(i) - y(i)) \\ &= \sum_{i \in [d]} (1 + \log x(i) - 1 - \log y(i)) (x(i) - y(i)) \\ &= \left(\nabla \sum_{i \in [d]} x(i) \log x(i) - \nabla \sum_{i \in [d]} y(i) \log y(i) \right)^\top (x - y) \\ &\geq \|x - y\|_1^2, \end{aligned}$$

where the last inequality follows from strong convexity of the entropy function $z \mapsto \sum_{i \in [d]} z_i \log z_i$ with respect to the $\|\cdot\|_1$ norm [28, Proposition 5.1]. \square

Recall that $u_n = \arg \min_{p \in \mathcal{P}} \text{KL}(p, \tilde{q}_n)$ and that q_n is an ϵ_n -optimal solution for the projection step, that is

$$\text{KL}(u_n, \tilde{q}_n) \geq \text{KL}(q_n, \tilde{q}_n) - \epsilon_n.$$

By Lemma 5, we have

$$\text{KL}(q_n, \tilde{q}_n) - \text{KL}(u_n, \tilde{q}_n) \geq (q_n - u_n)^\top \nabla \text{KL}(u_n, \tilde{q}_n) + \frac{1}{2} \|q_n - u_n\|_1^2 \geq \frac{1}{2} \|q_n - u_n\|_1^2,$$

where we used $(q_n - u_n)^\top \nabla \text{KL}(u_n, \tilde{q}_n) \geq 0$ due to first-order optimality condition for u_n . Hence $\text{KL}(q_n, \tilde{q}_n) - \text{KL}(u_n, \tilde{q}_n) \leq \epsilon_n$ implies that $\|q_n - u_n\|_\infty \leq \|q_n - u_n\|_1 \leq \sqrt{2\epsilon_n}$.

Consider q^* , the distribution over \mathcal{P} for the optimal arm, i.e. $q^*(i) = \frac{1}{m}$ iff $M_i^* = 1$. Recall that from proof of Lemma 4, for $q = q^*$ we have

$$\text{KL}(q^*, \tilde{q}_n) - \text{KL}(q^*, q_{n-1}) \leq \eta q_{n-1}^\top \tilde{X}(n) - \eta q^{*\top} \tilde{X}(n) + \eta^2 q_{n-1}^\top \tilde{X}^2(n). \quad (16)$$

Generalized Pythagorean Inequality (see Theorem 3.1 in [23]) gives

$$\text{KL}(q^*, \tilde{q}_n) \geq \text{KL}(q^*, u_n) + \text{KL}(u_n, \tilde{q}_n). \quad (17)$$

Let $\underline{q}_n = \min_i q_n(i)$. Observe that

$$\begin{aligned} \text{KL}(q^*, u_n) &= \sum_{i \in [d]} q^*(i) \log \frac{q^*(i)}{u_n(i)} = -\frac{1}{m} \sum_{i \in M^*} \log m u_n(i) \\ &\geq -\frac{1}{m} \sum_{i \in M^*} \log m (q_n(i) + \sqrt{2\epsilon_n}) \geq -\frac{1}{m} \sum_{i \in M^*} \left(\log m q_n(i) + \frac{\sqrt{2\epsilon_n}}{\underline{q}_n} \right) \\ &\geq -\frac{\sqrt{2\epsilon_n}}{\underline{q}_n} - \frac{1}{m} \sum_{i \in M^*} \log m q_n(i) = -\frac{\sqrt{2\epsilon_n}}{\underline{q}_n} + \text{KL}(q^*, q_n), \end{aligned}$$

Plugging this into (17), we get

$$\text{KL}(q^*, \tilde{q}_n) \geq \text{KL}(q^*, q_n) - \frac{\sqrt{2\epsilon_n}}{\underline{q}_n} + \text{KL}(u_n, \tilde{q}_n) \geq \text{KL}(q^*, q_n) - \frac{\sqrt{2\epsilon_n}}{\underline{q}_n}.$$

Putting this together with (16) yields

$$\text{KL}(q^*, q_n) - \text{KL}(q^*, q_{n-1}) \leq \eta q_{n-1}^\top \tilde{X}(n) - \eta q^{*\top} \tilde{X}(n) + \eta^2 q_{n-1}^\top \tilde{X}^2(n) + \frac{\sqrt{2\epsilon_n}}{\underline{q}_n}.$$

Finally, summing over n gives

$$\sum_{n=1}^T \left(q^{*\top} \tilde{X}(n) - q_{n-1}^\top \tilde{X}(n) \right) \leq \eta \sum_{n=1}^T q_{n-1}^\top \tilde{X}^2(n) + \frac{\text{KL}(q^*, q_0)}{\eta} + \frac{1}{\eta} \sum_{n=1}^T \frac{\sqrt{2\epsilon_n}}{\underline{q}_n}.$$

Defining

$$\epsilon_n = \frac{\left(\underline{q}_n \log \mu_{\min}^{-1}\right)^2}{32n^2 \log^3(n)}, \quad \forall n \geq 1,$$

and recalling that $\text{KL}(q^*, q_0) \leq \log \mu_{\min}^{-1}$, we get

$$\begin{aligned} \sum_{n=1}^T \left(q^{*\top} \tilde{X}(n) - q_{n-1}^\top \tilde{X}(n) \right) &\leq \eta \sum_{n=1}^T q_{n-1}^\top \tilde{X}^2(n) + \frac{\log \mu_{\min}^{-1}}{\eta} + \frac{\log \mu_{\min}^{-1}}{\eta} \sum_{n=1}^T \sqrt{\frac{2}{32n^2 \log^3(n+1)}} \\ &\leq \eta \sum_{n=1}^T q_{n-1}^\top \tilde{X}^2(n) + \frac{2 \log \mu_{\min}^{-1}}{\eta}, \end{aligned}$$

where we used the fact $\sum_{n \geq 1} n^{-1} (\log(n+1))^{-3/2} \leq 4$. We remark that by the properties of KL divergence and since $q'_{n-1} \geq \gamma \mu^0 > 0$, we have $\underline{q}_n > 0$ at every round n , so that $\epsilon_n > 0$ at every round n .

Using the above result and following the same lines as in the proof of Theorem 6, we have

$$R^{\text{COMBEXP}}(T) \leq \frac{\eta m^2 d T}{1 - \gamma} + \frac{2m \log \mu_{\min}^{-1}}{\eta} + m \gamma T.$$

Choosing $\eta = \gamma C$ with $C = \frac{\lambda}{m^{3/2}}$ gives

$$R^{\text{COMBEXP}}(T) \leq \frac{(Cm^2 d + m) \gamma T}{1 - \gamma} + \frac{2m \log \mu_{\min}^{-1}}{\gamma C}.$$

The proof is completed by setting $\gamma = \frac{\sqrt{2m \log \mu_{\min}^{-1}}}{\sqrt{2m \log \mu_{\min}^{-1}} + \sqrt{C(Cm^2 d + m)T}}$. \square

C.3 Proof of Theorem 7

We calculate the time complexity of the various steps of COMBEXP at round $n \geq 1$.

- (i) Mixing: This step requires $\mathcal{O}(d)$ time.
- (ii) Decomposition: Using the algorithm of [25], the vector mq'_{n-1} may be represented as a convex combination of at most $d+1$ arms in $\mathcal{O}(d^4)$ time, so that p_{n-1} may have at most $d+1$ non-zero elements (observe that the existence of such a representation follows from Carathéodory Theorem).
- (iii) Sampling: This step takes $\mathcal{O}(d)$ time since p_{n-1} has at most $d+1$ non-zero elements.
- (iv) Estimation: The construction of matrix Σ_{n-1} is done in time $\mathcal{O}(d^2)$ since p_n has at most $d+1$ non-zero elements and MM^\top is formed in $\mathcal{O}(d)$ time. Computing the pseudo-inverse of Σ_{n-1} costs $\mathcal{O}(d^3)$.
- (v) Update: This step requires $\mathcal{O}(d)$ time.
- (vi) Projection: The projection step is equivalent to solving a convex program up to accuracy $\epsilon_n = \mathcal{O}(n^{-2} \log^{-3}(n))$. We use the Interior-Point Method (Barrier method). The total number of Newton iterations to achieve accuracy ϵ_n is $\mathcal{O}(\sqrt{s} \log(s/\epsilon_n))$ [24, Ch. 11]. Moreover, the cost of each iteration is $\mathcal{O}((d+c)^3)$ [24, Ch. 10], so that the total cost of this step becomes $\mathcal{O}(\sqrt{s}(c+d)^3 \log(s/\epsilon_n))$. Plugging $\epsilon_n = \mathcal{O}(n^{-2} \log^{-3}(n))$ and noting that $\mathcal{O}(\sum_{n=1}^T \log(s/\epsilon_n)) = \mathcal{O}(T \log(T))$, the cost of this step is $\mathcal{O}(\sqrt{s}(c+d)^3 T \log(T))$.

Hence the total time complexity after T rounds is $\mathcal{O}(T[\sqrt{s}(c+d)^3 \log(T) + d^4])$, which completes the proof. \square

C.4 Implementation: The Case of Graph Coloring

In this subsection, we present an iterative algorithm for the projection step of COMBEXP, for the graph coloring problem described next.

Consider a graph $G = (V, E)$ consisting of m nodes indexed by $i \in [m]$. Each node can use one of the $c \geq m$ available colors indexed by $j \in [c]$. A feasible coloring is represented by a matrix $M \in \{0, 1\}^{m \times c}$, where $M_{ij} = 1$ if and only if node i is assigned color j . Coloring M is feasible if (i) for all i , node i uses at most one color, i.e., $\sum_{j \in [c]} M_{ij} \in \{0, 1\}$; (ii) neighboring nodes are assigned different colors, i.e., for all $i, i' \in [m]$, $(i, i') \in E$ implies for all $j \in [c]$, $M_{ij}M_{i'j} = 0$. In the following we denote by $\mathcal{K} = \{\mathcal{K}_\ell, \ell \in [k]\}$ the set of maximal cliques of the graph G . We also introduce $K_{\ell i} \in \{0, 1\}$ such that $K_{\ell i} = 1$ if and only if node i belongs to the maximal clique \mathcal{K}_ℓ .

There is a specific case where our algorithm can be efficiently implementable: when the convex hull $Co(\mathcal{M})$ can be captured by polynomial in m many constraints. Note that this cannot be ensured unless restrictive assumptions are made on the graph G since there are up to $3^{m/3}$ maximal cliques in a graph with m vertices [29]. There are families of graphs in which the number of cliques is polynomially bounded. These families include chordal graphs, complete graphs, triangle-free graphs, interval graphs, and planar graphs. Note however, that a limited number of cliques does not ensure a priori that $Co(\mathcal{M})$ can be captured by a limited number of constraints. To the best of our knowledge, this problem is open and only particular cases have been solved as for the stable set polytope (corresponding to the case $c = 2$, $X_{i1} = 1$ and $X_{i2} = 0$ with our notation) [30].

For the coloring problem described above we have

$$Co(\mathcal{M}) = Co\{\forall i, \sum_{j \in [c]} M_{ij} \leq 1, \quad \forall \ell, j, \sum_{i \in [m]} K_{\ell i} M_{ij} \leq 1\}. \quad (18)$$

Note that in the special case where G is the complete graph, such a representation becomes

$$Co(\mathcal{M}) = Co\{\sum_{j \in [c]} M_{ij} \leq 1, \quad \forall i, \sum_{i \in [m]} M_{ij} \leq 1, \quad \forall j\}.$$

We now give an algorithm for the projection a distribution p onto \mathcal{P} using KL divergence. Since \mathcal{P} is a scaled version of $Co(\mathcal{M})$, we give an algorithm for the projection of mp onto $Co(\mathcal{M})$ given by (18).

Set $\lambda_i(0) = \mu_j(0) = 0$ for all i, j and then define for $t \geq 0$,

$$\forall i \in [m], \lambda_i(t+1) = \log\left(\sum_j mp_{ij} e^{-\mu_j(t)}\right) \quad (19)$$

$$\forall j \in [c], \mu_j(t+1) = \max_\ell \log\left(\sum_i K_{i\ell} mp_{ij} e^{-\lambda_i(t+1)}\right). \quad (20)$$

We can show that

Proposition 2 *Let $p_{ij}^* = \lim_{t \rightarrow \infty} p_{ij} e^{-\lambda_i(t) - \mu_j(t)}$. Then mp^* is the projection of mp onto $Co(\mathcal{M})$ using the KL divergence.*

Although this algorithm is shown to converge, we must stress that the step (20) might be expensive as the number of distinct values of ℓ might be exponential in m . When G is a complete graph, this step is easy and our algorithm reduces to Sinkhorn's algorithm (see [26] for a discussion).

Proof: First note that the definition of projection can be extended to non-negative vectors thanks to the relation

$$KL(p^*, q) = \min_{p \in \Xi} KL(p, q).$$

More precisely, given an alphabet A and a vector $q \in \mathbb{R}_+^A$, we have for any probability vector $p \in \mathbb{R}_+^A$

$$\sum_{a \in A} p(a) \log \frac{p(a)}{q(a)} \geq \sum_a p(a) \log \frac{\sum_a p(a)}{\sum_a q(a)} = \log \frac{1}{\|q\|_1},$$

thanks to the log-sum inequality. Hence we see that $p^*(a) = \frac{q(a)}{\|q\|_1}$ is the projection of q onto the simplex of \mathbb{R}_+^A .

Now define $\mathcal{A}_i = \text{Co}\{M_{ij}, \sum_j M_{ij} \leq 1\}$ and $\mathcal{B}_{\ell_j} = \text{Co}\{M_{ij}, \sum_i K_{\ell_i} M_{ij} \leq 1\}$. Hence $\bigcap_i \mathcal{A}_i \cap \bigcap_{\ell_j} \mathcal{B}_{\ell_j} = \text{Co}(\mathcal{M})$. By the argument described above, iteration (19) (resp. (20)) corresponds to the projection onto \mathcal{A}_i (resp. $\bigcap_{\ell} \mathcal{B}_{\ell_j}$) and the proposition follows from Theorem 5.1 in [23]. \square

C.5 Examples

In this subsection, we compare the performance of COMBEXP against state-of-the-art algorithms (refer to Table 2 for the summary of regret of various algorithms).

C.5.1 m -sets

In this case, \mathcal{M} is the set of all d -dimensional binary vectors with m ones. We have

$$\mu_{\min} = \min_i \frac{1}{\binom{d}{m}} \sum_M M_i = \frac{\binom{d-1}{m-1}}{\binom{d}{m}} = \frac{m}{d}.$$

Moreover, according to [4, Proposition 12], we have $\underline{\lambda} = \frac{m(d-m)}{d(d-1)}$. When $m = o(d)$, the regret of COMBEXP becomes $O(\sqrt{m^3 d T \log(d/m)})$, namely it has the same performance as COMBAND and EXP2 WITH JOHN'S EXPLORATION.

C.5.2 Matching

Let \mathcal{M} be the set of perfect matchings in $\mathcal{K}_{m,m}$, where we have $d = m^2$ and $|\mathcal{M}| = m!$. We have

$$\mu_{\min} = \min_i \frac{1}{m!} \sum_M M_i = \frac{(m-1)!}{m!} = \frac{1}{m},$$

Furthermore, from [4, Proposition 4] we have that $\underline{\lambda} = \frac{1}{m-1}$, thus giving $R^{\text{COMBEXP}}(T) = O(\sqrt{m^5 T \log(m)})$, which is the same as the regret of COMBAND and EXP2 WITH JOHN'S EXPLORATION in this case.

C.5.3 Spanning Trees

In our next example, we assume that \mathcal{M} is the set of spanning trees in the complete graph \mathcal{K}_N . In this case, we have $d = \binom{N}{2}$, $m = N - 1$, and by Cayley's formula \mathcal{M} has N^{N-2} elements. Observe that

$$\mu_{\min} = \min_i \frac{1}{N^{N-2}} \sum_M M_i = \frac{(N-1)^{N-3}}{N^{N-2}},$$

which gives for $N \geq 2$

$$\begin{aligned} \log \mu_{\min}^{-1} &= \log \left(\frac{N^{N-2}}{(N-1)^{N-3}} \right) \\ &= (N-3) \log \left(\frac{N}{N-1} \right) + \log N \\ &\leq (N-3) \log 2 + \log(N) \leq 2N. \end{aligned}$$

From [4, Corollary 7], we also get $\underline{\lambda} \geq \frac{1}{N} - \frac{17}{4N^2}$. For $N \geq 6$, the regret of COMBAND takes the form $O(\sqrt{N^5 T \log(N)})$ since $\frac{m}{d\lambda} < 7$ when $N \geq 6$. Further, EXP2 WITH JOHN'S EXPLORATION attains the same regret. On the other hand, we get

$$R^{\text{COMBEXP}}(T) = O(\sqrt{N^5 T \log(N)}), \quad N \geq 6,$$

and therefore it gives the same regret as COMBAND and EXP2 WITH JOHN'S EXPLORATION.

C.5.4 Cut sets

Consider the case where \mathcal{M} is the set of balanced cuts of the complete graph \mathcal{K}_{2N} , where a balanced cut is defined as the set of edges between a set of N vertices and its complement. It is easy to verify that $d = \binom{2N}{2}$ and $m = N^2$. Moreover, \mathcal{M} has $\binom{2N}{N}$ balanced cuts and hence

$$\mu_{\min} = \min_i \frac{1}{\binom{2N}{N}} \sum_M M_i = \frac{\binom{2N-2}{N-1}}{\binom{2N}{N}} = \frac{N}{4N-2},$$

Moreover, by [4, Proposition 9], we have

$$\underline{\lambda} = \frac{1}{4} + \frac{8N-7}{4(2N-1)(2N-3)}, \quad N \geq 2,$$

and consequently, the regret of COMBEXP becomes $O(N^4\sqrt{T})$ for $N \geq 2$, which is the same as that of COMBAND and EXP2 WITH JOHN'S EXPLORATION.