

Tight Lower Bounds for Combinatorial Multi-Armed Bandits

Nadav Merlis

Faculty of Electrical Engineering, Technion, Israel Institute of Technology.

MERLIS@CAMPUS.TECHNION.AC.IL

Shie Mannor

Faculty of Electrical Engineering, Technion, Israel Institute of Technology.

Nvidia Research.

SHIE@EE.TECHNION.AC.IL

Editors: Jacob Abernethy and Shivani Agarwal

Abstract

The Combinatorial Multi-Armed Bandit problem is a sequential decision-making problem in which an agent selects a set of arms at each round, observes feedback for each of these arms and aims to maximize a known reward function of the arms it chose. While previous work proved regret upper bounds in this setting for general reward functions, only a few works provided matching lower bounds, all for specific reward functions. In this work, we prove regret lower bounds for combinatorial bandits that hold under mild assumptions for all smooth reward functions. We derive both problem-dependent and problem-independent bounds and show that the recently proposed Gini-weighted smoothness parameter (Merlis and Mannor, 2019) also determines the lower bounds for monotone reward functions. Notably, this implies that our lower bounds are tight up to log-factors.

Keywords: Combinatorial Multi-Armed Bandits, Lower Bounds, Gini-Weighted Smoothness

1. Introduction

Combinatorial Multi-Armed Bandits (CMABs) are a well-known extension of Multi-Armed Bandits (MABs) (Robbins, 1952), where instead of choosing a single arm at each round, the agent selects a set of arms. It then observes noisy feedback for each arm in this set (‘semi-bandit feedback’) and aims to maximize a known reward function of the selected arms and their parameters. More specifically, it aims to minimize its regret, which is the expected cumulative difference between the reward of the best action and the reward of the agent’s actions. The applications of this framework are numerous and vary between reward functions; the most common one is the linear reward function (Kveton et al., 2015c), which can be applied for problems such as spectrum allocation, shortest paths, routing problems and more (Gai et al., 2012). Another common application is the Probabilistic Maximum Coverage (PMC) problem (Merlis and Mannor, 2019), which is closely related to problems such as influence maximization and ranked recommendations.

Due to its usefulness, many previous works analyze regret upper bounds for different variants of this setting. While some works focus on specific reward functions, others derive bounds that hold for general reward functions. In these cases, the bounds usually depend on some measure of smoothness of the reward, for example, its global Lipschitz constant, or its Gini-weighted smoothness. The latter is a more refined smoothness criterion, recently suggested in (Merlis and Mannor, 2019), that takes into account the interaction between the local gradients of the reward and concentration properties of the arms. On the other hand, there are almost no works on matching lower bounds; to the best of our knowledge, all existing lower bounds for CMABs were derived for specific reward functions – either

the linear one or the PMC problem. Notably, there is no characterization of lower bounds for general reward functions, and it is unclear whether existing upper bounds are tight.

The gain from general lower bounds is threefold: (i) When the bounds are loose, understanding which quantities affect the lower bounds allows devising tighter algorithms; (ii) When the bounds are tight, the instances on which the bounds were derived can help to determine under which additional assumptions the lower bounds do not hold. Such assumptions might allow us to derive improved upper bounds; (iii) When we can control some parameters of the problem, e.g., the number of arms in an action, their effect on the lower bound can help us tune them for each application.

In this work, we derive problem-dependent (**Theorem 1**) and problem-independent (**Theorem 2**) lower bounds that hold for general reward functions under mild assumptions. The problem-dependent bound shows that for any ‘good’ bandit strategy, there exists a CMAB instance such that the asymptotic regret must be larger than a certain logarithmic rate. The problem-independent bound shows that for any strategy and any large enough horizon T , there exists a horizon-dependent instance with a \sqrt{T} regret. To derive these bounds, we define a family of action sets for CMAB problems, which we call *I-disjoint*. There, a subset of arms I appear in all actions and independent of other actions, while the rest of the arms appear in a single action. We then prove that for *I-disjoint* problems, both bounds depend on a new modified Gini-smoothness measure; specifically, they reproduce existing lower bounds for both the linear reward function and the PMC problem. If the reward function is also monotone, as in most practical applications, we derive an additional bound that depends on the Gini-smoothness of the reward and matches the upper bound of (Merlis and Mannor, 2019) up to logarithmic factors (**Proposition 4**). Thus, our results demonstrate that without any additional assumptions, the bounds are tight for almost any reward function.

2. Related Work

The general framework of combinatorial bandits with semi-bandit feedback was first presented in (Chen et al., 2013). Since then, it has had many extensions, e.g., for the case of probabilistically-triggered arms, where the set of arms in an action might be random (Chen et al., 2016a; Wang and Chen, 2017), and for reward functions that depend on the arm distribution (Chen et al., 2016b). Moreover, many previous works focus on specific instances of this problem, e.g., linear reward functions (Kveton et al., 2015c; Combes et al., 2015; Degenne and Perchet, 2016), cascading bandits (Kveton et al., 2015a,b) and more. Recently, Merlis and Mannor (2019) presented BC-UCB, a Bernstein-based UCB algorithm with regret bounds that depend on a new smoothness measure, which they call the *Gini-weighted smoothness*. Specifically, they show that by combining the reward nonlinearity with the local behavior of the confidence intervals, the dependency of previous regret bounds in the maximal action size can be removed. In this work, we show that for monotone reward functions, the Gini-smoothness also characterizes the lower bounds for CMAB problems and therefore prove that this upper bound is tight. In addition, while all previously stated papers assume that the reward function is monotone, a few papers also support non-monotone reward functions (Wang and Chen, 2018; Hüyük and Tekin, 2019). We also present lower bounds for this scenario.

Although there has been extensive work on regret upper bounds for CMABs, there are almost no results on lower bounds for this setting. Kveton et al. (2015c) derived lower bounds for the linear reward function with general arm distributions, and when arms are also independent, lower bounds can be found in (Degenne and Perchet, 2016; Combes et al., 2015). Also, Kveton et al. (2015a) derived lower bounds for cascading bandits and Merlis and Mannor (2019) derived bounds for the

Table 1: Upper (UB) and lower (LB) bounds of different CMAB problems for arbitrary action sets. Dep./Ind. are problem-dependent and problem-independent bounds, and the notations follow Section 3. γ_∞ is the global Lipschitz constant of a reward function, and for the Gini-smoothness γ_g , it holds that $\gamma_g \leq \sqrt{K}\gamma_\infty$ (Merlis and Mannor, 2019). Δ_{\min} is the minimal gap.

CMAB problem	Type	Previous UB	Previous LB	Theorem 1 or 2	Proposition 4
General reward functions	Dep.	$\mathcal{O}\left(\frac{\gamma_\infty^2 m K \ln T}{\Delta_{\min}}\right)^\dagger$	None	$\Omega\left(\frac{\max_{\mu, I} \gamma_g^2(\mu; I) m \ln T}{\Delta_{\min} K_I}\right)$	NA
	Ind.	None	None	$\Omega\left(\max_{\mu, I} \sqrt{\frac{\gamma_g^2(\mu; I) m T}{K_I}}\right)$	NA
Monotone reward functions	Dep.	$\mathcal{O}\left(\frac{\gamma_g^2 m \ln^2 K \ln T}{\Delta_{\min}}\right)^\ddagger$	None	$\Omega\left(\frac{\max_{\mu, I} \gamma_g^2(\mu; I) m \ln T}{\Delta_{\min} K_I}\right)$	$\tilde{\Omega}\left(\frac{\gamma_g^2 m \ln T}{\Delta_{\min}}\right)$
	Ind.	$\mathcal{O}\left(\gamma_g \ln K \sqrt{mT}\right)^\ddagger$	None	$\Omega\left(\max_{\mu, I} \sqrt{\frac{\gamma_g^2(\mu; I) m T}{K_I}}\right)$	$\tilde{\Omega}\left(\gamma_g \sqrt{mT}\right)$
Linear reward function $r(S; \mu) = \sum_{i \in S} \mu_i$	Dep.	$\mathcal{O}\left(\frac{mK \ln T}{\Delta_{\min}}\right)^\S$	$\Omega\left(\frac{mK \ln T}{\Delta_{\min}}\right)^\S$	$\Omega\left(\frac{mK \ln T}{\Delta_{\min}}\right)$	$\Omega\left(\frac{mK \ln T}{(\ln K) \Delta_{\min}}\right)$
	Ind.	$\mathcal{O}\left(\sqrt{mKT}\right)^\S$	$\Omega\left(\sqrt{mKT}\right)^\S$	$\Omega\left(\sqrt{mKT}\right)$	$\Omega\left(\sqrt{\frac{mKT}{\ln K}}\right)$
PMC problem $r(S; \mu) = \sum_{i=1}^M \left(1 - \prod_{j \in S} (1 - \mu_{ij})\right)$	Dep.	$\mathcal{O}\left(\frac{mM \ln^2 K \ln T}{\Delta_{\min}}\right)^\ddagger$	$\Omega\left(\frac{mM \ln T}{\Delta_{\min}}\right)^\ddagger$	$\Omega\left(\frac{mM \ln T}{\Delta_{\min}}\right)$	$\Omega\left(\frac{mM \ln T}{(\ln K)^2 \Delta_{\min}}\right)$
	Ind.	$\mathcal{O}\left(\ln K \sqrt{mMT}\right)^\ddagger$	$\Omega\left(\sqrt{mMT}\right)^\ddagger$	$\Omega\left(\sqrt{mMT}\right)$	$\Omega\left(\sqrt{\frac{mMT}{\ln K}}\right)$

† (Wang and Chen, 2018), requires independent arms. ‡ (Merlis and Mannor, 2019) § (Kveton et al., 2015c)

PMC problem. Nevertheless, and to the best of our knowledge, there are no lower bounds for general reward functions. A comparison of our bounds to previous related bounds can be found in Table 1.

In contrast to the CMAB problem, the lower bounds for MABs are well characterized. In their seminal work, Lai and Robbins (1985) presented the first general problem-dependent lower bound for MABs, which was later extended by Burnetas and Katehakis (1996). In terms of problem-independent bounds, Auer et al. (2002) derived an $\Omega(\sqrt{KT})$ lower bound for K -armed bandit problems with time horizon T , whose constants were later improved by Cesa-Bianchi and Lugosi (2006). Also, Mannor and Tsitsiklis (2004) proved problem-independent lower bounds with both linear and logarithmic regimes. Recently, Garivier et al. (2018) presented a general tool that allows deriving various lower bounds for MABs. We adapt this tool for the CMAB problem to derive our new regret bounds.

3. Preliminaries and Notations

We start with some notations. Let $[n] = \{1, \dots, n\}$, and for any vector $x \in \mathbb{R}^n$ and set $I \subset [n]$, denote by x_I , a sub-vector of x that contains only elements from I . We denote the Kullback-Leibler (KL) divergence between two distributions $\underline{\nu}, \underline{\nu}'$ by $D_{\text{KL}}(\underline{\nu}, \underline{\nu}')$, and the KL divergence between two Bernoulli random variables with expectations p, q by $\text{kl}(p, q)$. For any vector $x \in \mathbb{R}^n$, let x^s be a permutation such that $x_1^s \leq \dots \leq x_n^s$, and define the increasing permutation of vector $x \in \mathbb{R}^n$ w.r.t. a set I as $p^{x, I} = [x_{I^c}^s, x_I^s] \in \mathbb{R}^n$; namely, the beginning of the vector $p^{x, I}$ contains a sorted permutation of the elements of x in $I^c = [n]/I$, and its end contains the elements of x in I . Finally, for any set I of bounded size $|I| \leq K$, we denote by $K_I = K - |I|$ the size of the complementary set w.r.t. K .

We work under the combinatorial multi-armed bandit setting with semi-bandit feedback. Denote the number of arms ('base arms') by m , and let $\mathcal{S} \subset 2^{[m]}$ be the set of possible actions ('super arms'), that is, the set that contains all valid combinations of base arms that the agent can choose. The number of base arms in each action $S \in \mathcal{S}$ is bounded by $|S| \leq K$, and w.l.o.g., assume that $|S| = K$. At the beginning of each round t , the arms generate an observation vector $X^{(t)} =$

$(X_1^{(t)}, \dots, X_m^{(t)}) \in [0, 1]^m$, sampled from a fixed distribution independently of other rounds. Then, the agent chooses an action $S_t \in \mathcal{S}$ and observes feedback $X_{S_t}^{(t)} \triangleq \{(i, X_i^{(t)})\}, \forall i \in S_t$. Denote the means of base arms by $\mathbb{E}[X^{(t)}] = \boldsymbol{\mu} = (\mu_1, \dots, \mu_m)$. The goal of the agent is to maximize a known reward function $r(S; \boldsymbol{\mu})$, without knowing $\boldsymbol{\mu}$. Specifically, the agent aims to minimize its regret $R(T) = \sum_{t=1}^T (r(S^*; \boldsymbol{\mu}) - r(S_t; \boldsymbol{\mu})) \triangleq \sum_{t=1}^T \Delta_{S_t}$, where $S^* \in \arg \max_{S \in \mathcal{S}} r(S; \boldsymbol{\mu})$ is an optimal action¹ and $\Delta_{S_t} = r(S^*; \boldsymbol{\mu}) - r(S_t; \boldsymbol{\mu})$ is the suboptimality gap of S_t . To prove the lower bounds, we require a mild assumption on the reward function, which we call *index invariance*:

Definition 1 A reward function $r(S; \boldsymbol{\mu}) : \mathcal{S} \times [0, 1]^m \rightarrow \mathbb{R}$ is called differentiable if for any $S \in \mathcal{S}$, it is differentiable in $\boldsymbol{\mu} \in [0, 1]^m$.

Definition 2 A differentiable reward function $r(S; \boldsymbol{\mu}) : \mathcal{S} \times [0, 1]^m \rightarrow \mathbb{R}$ is called smooth index invariant if for any $S \in \mathcal{S}$, it only depends on the arms in S , i.e., $r(S; \boldsymbol{\mu}) = r(\boldsymbol{\mu}_S)$.

When the function is index invariant, and with a slight abuse of notations, we also write $r(\boldsymbol{\mu})$, with $\boldsymbol{\mu} \in \mathbb{R}^K$, to represent the mean of arms $\boldsymbol{\mu}_S$ for $|S| = K$. This assumption helps avoiding cases in which specific arms behave inherently different than other arms, such that the problem becomes much easier. For example, for the biased linear function $r(S; \boldsymbol{\mu}) = \sum_{i \in S} (\mu_i + m_i)$ and for any $\boldsymbol{\mu} \in [0, 1]^m$, the optimal action is $S^* = \arg \max_{S \in \mathcal{S}} \sum_{i \in S} i$, regardless of the arm means; therefore, both the upper and lower bounds for this reward function trivially equal zero. In contrast, the lower bound for the linear function are nonzero (see Table 1); thus, without the index-invariance, the lower bounds cannot be characterized solely by the gradient of reward function w.r.t. $\boldsymbol{\mu}$, in contrast to the existing upper bounds. To the best of our knowledge, this assumption holds for all practical applications for CMABs. We also believe that our analysis will hold for reward functions that depend on the *order* of arms inside an action. However, we leave this extension for future work. Besides this assumption, we later move our focus to monotone reward functions, which are defined as follows:

Definition 3 A differential reward function $r(S; \boldsymbol{\mu}) : \mathcal{S} \times [0, 1]^m \rightarrow \mathbb{R}$ is called monotone if for any $S \in \mathcal{S}$, any $\boldsymbol{\mu} \in [0, 1]^m$ and any $i \in [m]$, it holds that $\nabla_i r(S; \boldsymbol{\mu}) \geq 0$.

We remark that in most previous work, the upper bounds only hold for monotone functions, which include most of the practical application, e.g., the linear and PMC problems. We end this part of the preliminaries with an important inequality that was derived for MABs and will enable us to derive our new bounds for CMABs. Let $[m]$ be a set of arms, where each arm $a \in [m]$ is characterized by a distribution ν_a over \mathbb{R}^K , and denote $\underline{\nu} = \{\nu_a\}_{a \in [m]}$.² Assume that at each round, when playing a_t , a sample Y_t is drawn independently at random from ν_{a_t} . Let ψ be a strategy that chooses an arm according to the history and internal i.i.d randomization $U_t \in [0, 1]$. Namely, if $H_t = (U_0, Y_1, U_1, \dots, U_t, Y_t)$, then $a_{t+1} = \psi_t(H_t)$. Also, let $N_{\psi, a}(T)$ be the number of times an arm a was played under strategy ψ up to time T . Under these notations, the following holds:

Lemma 4 (Garivier et al. 2018) For all bandit problems $\underline{\nu}, \underline{\nu}'$, for all $\sigma(H_T)$ -measurable random variables Z with values in $[0, 1]$,

$$\sum_{a=1}^m \mathbb{E}_{\underline{\nu}}[N_{\psi, a}(T)] D_{\text{KL}}(\nu_a, \nu'_a) \geq \text{kl}(\mathbb{E}_{\underline{\nu}}[Z], \mathbb{E}_{\underline{\nu}'}[Z]), \quad (1)$$

1. Previous work on regret upper bounds also allows approximate maximization of r . We focus on the best achievable performance, so we assume we can efficiently maximize r .

2. Garivier et al. (2018) assume that ν_a are distributions over \mathbb{R} , but the exact same proof holds for distributions over \mathbb{R}^K .




Figure 1: Red arrows: confidence intervals on the function parameter (x-axis confidence), due to either Hoeffding or Bernstein inequalities; the latter is tighter near the edges. Blue zone: the resulting confidence intervals on the reward function in the bold curve (y-axis confidence). In Figures 1(a),1(b), these intervals are derived using the global Lipschitz constant of the reward γ_∞ , i.e., $CI(r(\hat{\mu})) \lesssim \gamma_\infty \sum_i CI(\hat{\mu}_i)$. In Figure 1(c), we present the real confidence interval on the reward, that is much tighter than the bound due to γ_∞ . This is since the bound is tight where the gradient is large, which is around the edge of the domain, but loose in other areas, where the gradient is small.

where $\text{kl}(p, q) = p \ln \frac{p}{q} + (1 - p) \ln \frac{1-p}{1-q}$.

In the combinatorial case we use similar notations and denote the action counts by $N_{\psi, S}(T)$.

3.1. Smoothness Measures

We now present the smoothness measures for smooth index-invariant reward functions that govern our lower bounds. The measures are defined for arm parameters $\mu \in \mathbb{R}^K$ and a set $I \subset [m]$ as follows:

1. L2 Gini-weighted smoothness

$$\gamma_{g,2}^2(\mu; I) \triangleq \sum_{i=1}^{K_I} p_i^{\mu, I} (1 - p_i^{\mu, I}) \nabla_i r(p^{\mu, I})^2 = \sum_{i \notin I} \mu_i (1 - \mu_i) \nabla_i r(\mu)^2 \quad (2)$$

2. L1 Gini-weighted smoothness

$$\gamma_{g,1}^2(\mu; I) \triangleq \left(\sum_{i=1}^{K_I} \sqrt{p_i^{\mu, I} (1 - p_i^{\mu, I})} \nabla_i r(p^{\mu, I}) \right)^2 = \left(\sum_{i \notin I} \sqrt{\mu_i (1 - \mu_i)} \nabla_i r(\mu) \right)^2 \quad (3)$$

3. Modified Gini-weighted smoothness

$$\tilde{\gamma}_g^2(\mu; I) = \sum_{i=1}^{K_I} p_i^{\mu, I} (1 - p_i^{\mu, I}) \nabla_i r(p^{\mu, I})^2 + 2 \sum_{i=1}^{K_I} \sum_{j=i+1}^{K_I} p_i^{\mu, I} (1 - p_j^{\mu, I}) \nabla_i r(p^{\mu, I}) \nabla_j r(p^{\mu, I}) \quad (4)$$

Notice that the last equality in Equations (2),(3) is due to the index-invariance assumption. Specifically, the assumption implies that $\nabla_i r(\mu)$ only depends on μ_i and the set of values $\{\mu_i\}_{i=1}^m$. For smooth index-invariant reward functions, $\gamma_{g,2}^2(\mu; I)$ is an extension of the Gini-smoothness, as defined in (Merlis and Mannor, 2019). In particular, their smoothness parameter is defined as $\gamma_g = \max_{I, \mu} \gamma_{g,2}(\mu; I) = \max_{\mu} \gamma_{g,2}(\mu; \emptyset)$. A motivation to this smoothness criterion is presented in Figure 1. Notably, the performance of any algorithm strongly depends on the uncertainty in the reward of actions. However, algorithms only have access to uncertainty in arm parameters, and arm uncertainty must be translated into reward uncertainty. The figure illustrates that doing so using the global Lipschitz constant might lead to loose bounds. Intuitively, if the gradients are small, even

wide confidence intervals do not cause high uncertainty in the reward. Similarly, narrow confidence intervals do not lead to high reward uncertainty even where the gradients are large. Thus, [Merlis and Mannor \(2019\)](#) suggested weighting the gradients according to the confidence intervals of the arms. Specifically, their algorithm (BC-UCB) relies on Empirical-Bernstein concentration-bounds ([Audibert et al., 2009](#)), that depend on the variance of the arms and are proportional to $\sqrt{\mu_i(1 - \mu_i)}$ for Bernoulli arms. Similarly weighting the gradients leads to the Gini-smoothness measures.

In the following sections, we prove lower bounds for the CMAB problem that depend on $\tilde{\gamma}_g^2(\boldsymbol{\mu}; I)$. While complex at first glance, $\tilde{\gamma}_g^2(\boldsymbol{\mu}; I)$ is actually closely related to the other smoothness measures. Notably, observe that the only difference between $\tilde{\gamma}_g^2(\boldsymbol{\mu}; I)$ and $\gamma_{g,1}^2(\boldsymbol{\mu}; I)$ is a small modification to the second (cross) term of (4). In Proposition 3, we indeed prove that $\tilde{\gamma}_g^2(\boldsymbol{\mu}; I) \geq \tilde{\Omega}(\gamma_{g,1}^2(\boldsymbol{\mu}; I))$. When the function is monotone, we later prove that $\gamma_{g,1}(\boldsymbol{\mu}; I)$ can also be related to $\gamma_{g,2}(\boldsymbol{\mu}; \emptyset)$, which leads to tight lower bounds, up to logarithmic factors.

4. Problem-Dependent Lower Bounds

In this section, we prove a problem-dependent lower bound. Specifically, we show that there exist a CMAB instance, such that the asymptotic regret of any consistent strategy on this instance is lower bounded by a logarithmic term that depends on $\tilde{\gamma}_g(\boldsymbol{\mu}; I)$ and the minimal gap $\Delta = \min_{S \in \mathcal{S}, \Delta_S > 0} \Delta_S$. Consistent strategy is defined as follows:

Definition 5 *A bandit strategy ψ is called consistent if for any CMAB problem, any $S \in \mathcal{S}$ such that $\Delta_S > 0$ and any $0 < \alpha \leq 1$, it holds that $\mathbb{E}[N_{\psi,S}(T)] = o(T^\alpha)$.*

To prove the lower bounds, we focus on a subset of CMAB problems which we call *I-disjoint*.

Definition 6 *For a given subset $I \subset [m]$, a CMAB problem is called *I-disjoint* if all arms $i \in I$ are mutually independent of all arms $i \notin I$ and also $S_1 \cap S_2 = I$ for any $S_1 \neq S_2 \in \mathcal{S}$.*

Since $|S| \leq K$, we implicitly assume that $|I| \leq K$ and denote the effective maximal action size by $K_I = K - |I|$. In *I-disjoint* CMAB problems, the base arms $i \in I$ appear in all actions and are mutually independent of the other arms. The rest of the arms can only appear in one action. This notion of CMAB problems actually extends the action sets from previous work – for the linear reward function, [Kveton et al. \(2015c\)](#) divided the arms into $\frac{m}{K}$ disjoint groups, which is equivalent to $I = \emptyset$. In contrast, in the CMAB instance on which the PMC lower bounds were derived, only a single arm per item varied between actions ([Merlis and Mannor, 2019](#)). If there are M items, this is equivalent to $I = [K - M]$. We show that choosing the ‘worst-case’ set I naturally results with tighter lower bounds. We start by deriving general problem-dependent lower bound, using Lemma 4:

Lemma 7 *Let r be a smooth index invariant reward function with $|S| = K$ for all $S \in \mathcal{S}$. Also, let $\underline{\nu}$ be the action distribution of an *I-disjoint* CMAB problem such that there exists an arm $i \notin I$ in S^* with $\Pr\{X_i^{(t)} = \mu_i\} < 1$ and $\nabla_i r(S^*; \boldsymbol{\mu}) \neq 0$. Then, for all consistent strategies ψ and all suboptimal actions S ,*

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_{\underline{\nu}}[N_{\psi,S}(T)]}{\ln T} \geq \frac{1}{D_{\text{KL}}(\nu_S, \nu_{S^*})} .$$

Specifically, it holds that

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_{\underline{\nu}}[R(T)]}{\ln T} \geq \sum_{S: \Delta_S > 0} \frac{\Delta_S}{D_{\text{KL}}(\nu_S, \nu_{S^*})} .$$

The proof is in Appendix A, and partially follows Theorem 1 of (Garivier et al., 2018), with some adjustments due to the nonlinearity of the reward function. Although this lemma gives a general lower bound for I -disjoint CMAB problems, it has no clear dependence on any smoothness measure of the reward. To derive lower bounds that directly depend on such measures, we carefully design the arm distributions ν_S, ν_{S^*} and analyze both Δ_S and $D_{\text{KL}}(\nu_S, \nu_{S^*})$ for these distributions. We do so in the following theorem, which results with the desired lower bound:

Theorem 1 *Let r be a smooth index invariant reward function. For any $\boldsymbol{\mu} \in [0, 1]^K$ and any small enough $\Delta > 0$, there exists an instance of an I -disjoint bandit problem $\underline{\nu}$ with minimal gap Δ and $\mathbb{E}[\nu_{S^*}] = \boldsymbol{\mu}$, such that the expected regret of any consistent algorithm is bounded by*

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_{\underline{\nu}}[R(t)]}{\ln T} \geq \max_I \frac{(m - 2K)\tilde{\gamma}_g^2(\boldsymbol{\mu}; I)}{8K_I\Delta} \triangleq DB^*(\Delta; \boldsymbol{\mu})$$

Before proving the theorem, we start with a short discussion on the tightness of the results and the relation to existing upper bounds. One interesting case is when $p_i^{\boldsymbol{\mu}, I} = p_0$ for all $i \in [K_I]$. Due to the index invariance, the gradient components are also equal, which results with $\tilde{\gamma}_g^2(\boldsymbol{\mu}; I) = K_I^2 p_0(1 - p_0)\nabla_i r(p^{\boldsymbol{\mu}, I})^2$. This choice allows us to easily reproduce the existing lower bounds and leads to the bounds stated in Table 1. Specifically, for the linear reward function, we achieve the bound by choosing $I = \emptyset$ and $p_0 = \frac{1}{2}$; for the PMC problem, we choose the subset I such that it contains all arms except for a single (first) element from each vector $\boldsymbol{\mu}_i = (\frac{1}{2}, 0, \dots, 0)$, i.e., $K_I = M$ and $p_0 = \frac{1}{2}$. We remark that though choosing $K_I = K$ is seemingly optimal, this is not always the case. A notable example for this issue is the reward function $r(S; \boldsymbol{\mu}) = 1 - e^{-\sum_{i \in S} \mu_i^2}$. For this instance, optimizing over p_0 leads to a bound of $\tilde{\gamma}_g^2(\boldsymbol{\mu}; I)/K_I = \Omega(1/\sqrt{K_I})$, and the optimal choice is $K_I = \mathcal{O}(1)$. In this example it also holds that $\gamma_g = \mathcal{O}(1)$, and, therefore, this bound is tight.

We note that this bound only requires the smooth index invariance assumption and holds for non-monotone reward functions. When the reward is also monotone, observe that $\tilde{\gamma}_g^2(\boldsymbol{\mu}; I) \geq \gamma_{g,2}^2(\boldsymbol{\mu}; I)$. Choosing $I = \emptyset$ and maximizing over $\boldsymbol{\mu}$ leads to a lower bound of $\Omega\left(\frac{m\gamma_g^2 \ln T}{\Delta K}\right)$, which differs from the upper bound of (Merlis and Mannor, 2019) by a factor of K . We later prove that when the reward is monotone, a stronger lower bound can be derived, such that it matches the upper bound up to logarithmic factors.

Next, we present the proof of Theorem 1 which is composed of three parts. We first present a carefully designed parametric arm distributions, that allow controlling the suboptimality gap while retaining low KL-divergence. We then bound both the gap and the KL-divergence in terms of the parameters of the distributions and apply Lemma 7 using these bounds. We conclude the proof by optimizing the resulting lower bound over the distribution parameters.

Proof

STEP 1: *Fixing a parametric family of CMAB instances.*

Denote by I^* , the maximizer of I in $DB^*(\Delta; \boldsymbol{\mu})$, and for brevity, let $p = p^{\boldsymbol{\mu}, I^*}$ and $K_I = K_I^* = K - |I^*|$. Note that this also implies that $0 \leq p_1 \leq \dots \leq p_{K_I} \leq 1$. In addition, due to the form of the lower bound, we can deduce that $p_1 > 0$ and $p_{K_I} < 1$, since otherwise, we can move arms with means 0 and 1 into I^* and strictly increase $DB^*(\Delta; \boldsymbol{\mu})$. For now, we also assume that no two arms have the same mean, i.e., $0 < p_1 < \dots < p_{K_I} < 1$, and will return to this assumption later in the proof. Without loss of generality, we also assume that $K_I > 0$ and there exists $i \in [K_I]$ such that $\nabla_i r(p) \neq 0$ since otherwise, $DB^*(\Delta; \boldsymbol{\mu}) = 0$ and trivially holds. Similarly, we assume that $m > 2K$.

We fix the CMAB problem to be I^* -disjoint, and choose the action set \mathcal{S} to be the maximal action set with action sizes K , i.e., $|\mathcal{S}| = \lfloor \frac{m-|I^*|}{K-|I^*|} \rfloor \geq \frac{m-K}{K_I}$. Denote the arm distribution in this problem by $\underline{\nu}$, and fix the distribution of the common arms $i \in I^*$ to any distribution with expectation $\mathbb{E}[\nu_{I^*}] = \boldsymbol{\mu}_{I^*}$, as long as they are mutually independent of the rest of the arms. For a single action, we set the distribution to be ν^* with mean $\boldsymbol{\mu}^{\nu^*} \triangleq \boldsymbol{\mu}$, and for the rest of the actions, we fix it to distribution ν with mean $\boldsymbol{\mu}^\nu$. Both distributions are stated in Table 2. The distribution ν depends on $\epsilon \in \mathbb{R}^{K_I}$, that will be determined later such that ν is strictly suboptimal, and we denote its suboptimality gap by $\Delta_\epsilon = r(\boldsymbol{\mu}^{\nu^*}) - r(\boldsymbol{\mu}^\nu)$. We remark that for $0 < p_1 < \dots < p_{K_I} < 1$, there exists $b_{\epsilon,0}$ such that ν is a valid probability distribution for all $\|\epsilon\|_\infty \leq b_{\epsilon,0}$. We enforce this condition later in the proof.

Note that for all $i \notin I^*$, the arms are Bernoulli random variables with mean $\mu_i \notin \{0, 1\}$, and therefore $\Pr\{X_i^{(t)} = p_i\} = 0$. Also, since the gradient is not zero for some $i \notin I^*$, the conditions of Lemma 7 hold, and we can bound the regret by

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_{\underline{\nu}}[R(T)]}{\ln T} \geq \sum_{i=1}^{|\mathcal{S}|-1} \frac{\Delta_\epsilon}{D_{\text{KL}}(\nu, \nu^*)} \geq \frac{m-2K}{K_I} \frac{\Delta_\epsilon}{D_{\text{KL}}(\nu, \nu^*)}. \quad (5)$$

STEP 2: Deriving lower bounds that depend on the distribution parameters ϵ .

Next, we bound both $D_{\text{KL}}(\nu, \nu^*)$ and Δ_ϵ in terms of ϵ .

Lemma 8 Let $p \triangleq p^{\mu, I} \in \mathbb{R}^K$ such that $0 < p_1 < \dots < p_{K_I} < 1$ and define $p_0 = 0$. Also, let ν, ν^* be the distributions stated in Table 2. Then, there exists a constant $b_{\epsilon,1} > 0$ such that for any $\epsilon \in \mathbb{R}^{K_I}$ with $\|\epsilon\|_\infty \leq b_{\epsilon,1}$, it holds that

$$D_{\text{KL}}(\nu, \nu^*) \leq 2\epsilon^T B_{\text{KL}}(p)\epsilon, \quad (6)$$

where $B_{\text{KL}}(p) = D(p) + \frac{1}{1-p_{K_I}} \mathbf{1}\mathbf{1}^T$, $D(p) \in \mathbb{R}^{K_I \times K_I}$ is a diagonal matrix whose elements are $D_{ii}(p) = \frac{1}{p_i - p_{i-1}}$ and $\mathbf{1} \in \mathbb{R}^{K_I}$ is a vector of ones.

Lemma 9 Let r be a smooth index invariant reward function and let $p \in \mathbb{R}^K$ such that $0 < p_1 \leq \dots \leq p_{K_I} < 1$ and there exists $i \in [K_I]$ with $\nabla_i r(p) \neq 0$. Also, define $c_j = \sum_{i=j}^{K_I} \nabla_i r(p)$ for all $j \in [K_I]$ and let $u \in \mathbb{R}^{K_I}$ be a vector such that $c^T u > 0$. Then, if $\epsilon = \epsilon_0 u$, there exists a constant $b_{\epsilon,2} > 0$ such that for all $0 < \epsilon_0 \leq b_{\epsilon,2}$,

$$\Delta_\epsilon \geq \frac{1}{2} c^T \epsilon > 0. \quad (7)$$

The proofs are presented in Appendix B. Specifically, note that $D_{ii}(p) > 0$, and thus $B_{\text{KL}}(p)$ is positive definite, and the KL bound equals zero only for $\epsilon = 0$. Also, since there exists $i \in [K_I]$ such that $\nabla_i r(p) \neq 0$, it holds that $c \neq 0$. When both lemmas hold, substitution into (5) yields

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_{\underline{\nu}}[R(T)]}{\ln T} \geq \frac{m-2K}{K_I} \frac{\Delta_\epsilon}{D_{\text{KL}}(\underline{\nu}, \underline{\nu}^*)} = \frac{m-2K}{K_I \Delta_\epsilon} \frac{\Delta_\epsilon^2}{D_{\text{KL}}(\underline{\nu}, \underline{\nu}^*)} \geq \frac{m-2K}{8K_I \Delta_\epsilon} \frac{(c^T \epsilon)^2}{\epsilon^T B_{\text{KL}}(p)\epsilon}. \quad (8)$$

Table 2: The probability distributions ν and ν^* of arms outside I^* , with parameters $0 < p_1 < p_2 < \dots < p_{K_I} < 1$ and $\epsilon \in \mathbb{R}^{K_I}$. Arms in I^* have mean $\boldsymbol{\mu}_{I^*}$ and are independent of the rest of the arms.

Observation vector $X^{(t)} = (X_1^{(t)}, \dots, X_{K_I}^{(t)})$	Probability of $X^{(t)}$ in ν^*	Probability of $X^{(t)}$ in ν
$(1, 1, \dots, 1, 1)$	p_1	$p_1 - \epsilon_1$
$(0, 1, \dots, 1, 1)$	$p_2 - p_1$	$p_2 - p_1 - \epsilon_2$
\dots	\dots	\dots
$(0, 0, \dots, 0, 1)$	$p_{K_I} - p_{K_I-1}$	$p_{K_I} - p_{K_I-1} - \epsilon_{K_I}$
$(0, 0, \dots, 0, 0)$	$1 - p_{K_I}$	$1 - p_{K_I} + \sum_{i=j}^{K_I} \epsilon_j$

STEP 3: Finding the worst-case CMAB instance.

We now focus on the function $f_I(\epsilon; p) = \frac{\epsilon^T c c^T \epsilon}{\epsilon^T B_{KL}(p) \epsilon}$, which is defined for any $\epsilon \neq 0$, as $B_{KL}(p)$ is positive definite. $B_{KL}(p)$ is also invertible, and we can therefore apply the invertible transformation $\epsilon = B_{KL}^{-1/2}(p)x$, which results with the following function:

$$\tilde{f}_I(x; p) = \frac{x^T B_{KL}^{-1/2}(p) c c^T B_{KL}^{-1/2}(p) x}{\|x\|_2^2} = \frac{\left(c^T B_{KL}^{-1/2}(p) x\right)^2}{\|x\|_2^2} \stackrel{(*)}{\leq} \frac{\|B_{KL}^{-1/2}(p) c\|_2^2 \|x\|_2^2}{\|x\|_2^2} = c^T B_{KL}^{-1}(p) c ,$$

where $(*)$ is due to Cauchy-Schwarz Inequality, and equality holds for any $\epsilon_0 \neq 0$ and $x = \epsilon_0 B_{KL}^{-1/2}(p) c$. Therefore, the maximal value of $f_I(\epsilon; p)$ is $f_I(\epsilon^*; p) = c^T B_{KL}^{-1}(p) c$ and can be attained with $\epsilon^* = \epsilon_0 B_{KL}^{-1}(p) c$, for any $\epsilon_0 \neq 0$.

Motivated by the maximization property of ϵ^* , we now fix $\epsilon \leftarrow \epsilon^* = \epsilon_0 B_{KL}^{-1}(p) c$, for $0 < \epsilon_0 \leq b_{\epsilon, 2}$ such that $\|\epsilon^*\|_\infty \leq \min\{b_{\epsilon, 0}, b_{\epsilon, 1}\}$. For this choice, $\epsilon^* \neq 0$, since $B_{KL}(p)$ is invertible and $c \neq 0$, and $c^T B_{KL}^{-1}(p) c > 0$ as required for Lemma 9. We explicitly calculate the lower bound in the following lemma (see proof in Appendix B):

Lemma 10 Under the notations of Lemmas 8 and 9, for any $0 < p_1 < \dots < p_{K_I} < 1$ and $c \neq 0$, if $\epsilon^* = \epsilon_0 B_{KL}^{-1}(p) c$, then

$$\epsilon_i^* = \epsilon_0 (p_i - p_{i-1}) \left(c_i - \sum_{j=1}^{K_I} (p_j - p_{i-j}) c_j \right) .$$

Also, if $f_I(\epsilon; p) = \frac{\epsilon^T c c^T \epsilon}{\epsilon^T B_{KL}(p) \epsilon}$, then $f_I(\epsilon^*; p) = \tilde{\gamma}_g^2(\boldsymbol{\mu}; I)$.

An important conclusion is that under the assumptions of the lemma, $\tilde{\gamma}_g^2(\boldsymbol{\mu}; I) > 0$, since $f_I(\epsilon^*; p) = c^T B_{KL}^{-1}(p) c > 0$. A more general result naturally arises from the proof of Lemma 10: for any for any $\boldsymbol{\mu} \in [0, 1]^K$ and any $I \subset [K]$, it holds that $\tilde{\gamma}_g^2(\boldsymbol{\mu}; I) \geq 0$, as expected from a smoothness parameter. We refer the readers to the proof of the lemma for additional details. Substituting back into (8) and recalling that I^* was chosen as the maximizer in $DB_{\boldsymbol{\mu}, I}^*(\Delta_\epsilon)$ we get

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu[R(T)]}{\ln T} \geq \frac{(m - 2K) \tilde{\gamma}_g^2(\boldsymbol{\mu}; I^*)(p)}{8K_{I^*} \Delta_\epsilon} = DB^*(\Delta_\epsilon; \boldsymbol{\mu}).$$

We finally return to our assumption that $p_1 < \dots < p_{K_I}$. If there are equal values in p , i.e., $p_i = p_i = \dots = p_{i+n-1}$, we modify both distributions in Table 2 such that the observations of these arms are identical, namely $X_i^{(t)} = X_{i+1}^{(t)} = \dots = X_{i+n-1}^{(t)}$. Then, we set $\epsilon_{i+1}^* = \dots = \epsilon_{i+n-1}^* = 0$. Notice that this modification does not change the KL divergence, nor the analysis of the gap, and therefore retains the same results. Similarly, Lemma 10 still holds by defining the function f_I over the sub-vector of ϵ with coordinates such that $p_i < p_{i+1}$. Alternatively, note that the existing analysis naturally sets $\epsilon_i = 0$ if $p_i = p_{i-1}$ (Lemma 10), so it is not surprising that this modification does not change the results. We avoided writing the full analysis since it requires indexing the sub-vector of p with strictly increasing values, which would make the notations much more involved.

To conclude the proof, we remark that if $\Delta_0 > 0$ is the gap for some $\epsilon_0 > 0$ such that Lemmas 8 and 9 hold, by tuning ϵ_0 we can achieve any gap $\Delta \leq \Delta_0$, and thus the previous result holds for any small enough gap Δ . ■

5. Problem-Independent Lower Bounds

In this section, we prove a problem-independent regret lower bound. Specifically, we prove that for any fixed strategy and any large enough horizon T , there exists a CMAB instance such that the regret is lower bounded by a gap-independent \sqrt{T} term. We remark that in contrast to problem-dependent bounds, in which the instance is fixed for *all* strategies and time horizons, the instance for problem-independent bounds is designed as the ‘worst-case’ problem for a *specific* strategy and time horizon. Similarly to the previous section, we start by proving a general lower bound for I -disjoint CMAB problems and then apply it with a specific distribution to derive the desired bound.

Lemma 11 *Let r be a smooth index invariant reward function, and let $I \subset [m]$ such that $|I| \leq K$ and $m > K$. Also, let $\boldsymbol{\mu}, \boldsymbol{\mu}^* \in \mathbb{R}^K$ such that $\boldsymbol{\mu}_I = \boldsymbol{\mu}_I^*$ and $\Delta = r(\boldsymbol{\mu}^*) - r(\boldsymbol{\mu}) > 0$. Finally, let ν, ν^* be two distributions with expectations $\boldsymbol{\mu}, \boldsymbol{\mu}^* \in \mathbb{R}^K$, such that arms in I are mutually independent of arms outside I and both distributions are identical for arms in I . Then, for any horizon T and any strategy ψ , there exists an I -disjoint CMAB problem with arm distribution $\underline{\nu}$ such that $\nu_{S^*} = \boldsymbol{\mu}^*$ and its regret under strategy ψ is bounded by*

$$\mathbb{E}_{\underline{\nu}'}[R(T)] \geq T\Delta \left(1 - \frac{K_I}{m-K} - \sqrt{\frac{1}{2} \frac{TK_I}{m-K} D_{\text{KL}}(\nu, \nu^*)} \right).$$

The proof is a variant of Theorem 6 of [Garivier et al. \(2018\)](#) and can be found in [Appendix A](#). With this lemma at hand, and similarly to the problem-dependent bound of [Theorem 1](#), we can also derive a problem-independent lower bound:

Theorem 2 *Let r be a smooth index invariant reward function and assume that $m \geq 3K$. Then, for any $\boldsymbol{\mu} \in [0, 1]^K$, any $T \geq T_0$ and for any strategy ψ , there exists an I -disjoint CMAB problem $\underline{\nu}$ with $\mathbb{E}[\nu_{S^*}'] = \boldsymbol{\mu}$ such that its regret under strategy ψ is bounded by*

$$\mathbb{E}_{\underline{\nu}'}[R(T)] \geq \max_I \frac{\tilde{\gamma}_g(\boldsymbol{\mu}; I)}{32} \sqrt{\frac{T(m-K)}{K_I}} \triangleq IB^*(\boldsymbol{\mu}).$$

We defer the proof to [Appendix C](#). The same discussion from the previous section about the tightness of the bound still holds. We start by noting that $IB^*(\boldsymbol{\mu})$ reproduces the existing lower bounds both for the linear reward function ([Kveton et al., 2015c](#)) and the probabilistic maximum coverage problem ([Merlis and Mannor, 2019](#)). Also, for monotone functions, we can bound $\tilde{\gamma}_g^2(\boldsymbol{\mu}; \emptyset) \geq \gamma_g^2$ and match the problem-independent upper bound of ([Merlis and Mannor, 2019](#)) up to a \sqrt{K} factor. This factor will be improved to a logarithmic factor in the following section.

6. Relations Between Smoothness Measures

To this point, we derived lower bounds that depend on the modified Gini-smoothness $\tilde{\gamma}_g(\boldsymbol{\mu}; I)$. In this section, we show that at a cost of logarithmic factors, we can relate these bounds to the L1 Gini-smoothness. Moreover, for monotone reward functions, we also prove lower bounds that depend on the L2 Gini-smoothness, and thus match the upper bounds of ([Merlis and Mannor, 2019](#)) up to log-factors. We formally state the results in the following propositions:

Proposition 3 *Let r be a smooth index invariant reward function and denote $p = p^{\boldsymbol{\mu}, I}$ for some $\boldsymbol{\mu} \in \mathbb{R}^K$ and $I \subset [K]$. Then,*

$$\tilde{\gamma}_g^2(\boldsymbol{\mu}; I) \geq \frac{\gamma_{g,1}^2(\boldsymbol{\mu}; I)}{3 + \ln \frac{1}{p_1} + \ln \frac{1}{1-p_{K_I}}}, \quad (9)$$

where the r.h.s. is defined as zero if $p_1 = 0$ or $p_{K_I} = 1$.

Proposition 4 *Let r be a monotone smooth index invariant reward function. Then, for any $\boldsymbol{\mu} \in \mathbb{R}^K$, it holds that*

$$\max_I \frac{\gamma_{g,1}^2(\boldsymbol{\mu}; I)}{K_I} \geq \frac{\gamma_{g,2}^2(\boldsymbol{\mu}; \emptyset)}{1 + \ln K} . \quad (10)$$

Furthermore, if $\mu_{\min} = \min_{i:\mu_i > 0} \mu_i$ and $\mu_{\max} = \max_{i:\mu_i < 1} \mu_i$, then

$$\max_I \frac{\tilde{\gamma}_g^2(\boldsymbol{\mu}; I)}{K_I} \geq \frac{\gamma_{g,2}^2(\boldsymbol{\mu}; \emptyset)}{(1 + \ln K) \left(3 + \ln \frac{1}{\mu_{\min}} + \ln \frac{1}{1 - \mu_{\max}} \right)} .$$

We emphasize that Proposition 3 *does not* require the monotonicity assumption. However, when the components of the gradient can be either positive or negative, the bound might equal zero even when the gradient is large. When the function is monotone, Proposition 4 greatly improves the naïve choice of $I = \emptyset$ in the bounds of Theorems 1 and 2, by a factor of K , with only a logarithmic cost. Specifically, by maximizing over $\boldsymbol{\mu}$ and applying Proposition 4, we get a problem-dependent lower bound of $\tilde{\Omega}(m\gamma_g^2 \ln T/\Delta)$ and a problem-independent bound of $\tilde{\Omega}(\gamma_g \sqrt{mT})$, both match the upper bounds of (Merlis and Mannor, 2019) up to log-factors. Thus, for monotone smooth functions, we conclude that the L2 Gini-smoothness parameter characterizes both the upper and lower bounds.

Of the log-factors in the propositions, the more interesting one is that of $\left(\ln \frac{1}{\mu_{\min}} + \ln \frac{1}{1 - \mu_{\max}} \right)$. In cases where the mean values $\boldsymbol{\mu}$ are exponentially close to zero or one, this factor can be of order $1/K$. We suspect that this is due to a proof artefact, but leave the investigation of this factor to future work. Nonetheless, for any practical example, this term is at most of order $\ln K$, which leaves the bound tight up to log-factors in the problem size. Another question that arises is whether this factor is the result of a loose analysis in Proposition 3, and a tighter analysis might yield a better factor (e.g., $\ln K$). Sadly, the answer is negative. In Appendix E.1, we build an example where μ_i are exponentially small and Inequality (9) *does not hold* without a factor of order $1/K$. Therefore, to remove this factor, $DB^*(\Delta; \boldsymbol{\mu})$ and $IB^*(\boldsymbol{\mu})$ also need to be improved.

Due to space limitations, we defer the full proofs to Appendix D and only provide a proof sketch for Proposition 4:

Proof Sketch Notice that $\gamma_{g,1}(\boldsymbol{\mu}; I)$ and $\gamma_{g,2}(\boldsymbol{\mu}; I)$ are closely related to the L1 and L2 norms of a vector whose components are $\sqrt{\mu_i}(i - \mu_i)\nabla_i r(\boldsymbol{\mu})$. Specifically, both are norms of a vector with a subset of these components. We utilize this connection and derive, to the best of our knowledge, a new relation between the norms:

Lemma 12 *Let A be a nonempty subset of indices $A \subset [n]$. For any vector $x \in \mathbb{R}^n$, it holds that*

$$\max_{A \neq \emptyset} \frac{1}{|A|} \|x_A\|_1^2 \geq \frac{1}{1 + \ln n} \|x\|_2^2 .$$

This lemma gives a much stronger relation than the standard L1-L2 inequality and is of independent interest. Specifically, the relation between the norms is logarithmic in the vector dimension, instead of the standard relation $\|x\|_1 \geq \|x\|_2$, that would have resulted in a linear term. This is due to the ability to choose the best subset of the vector on the left-hand side. We remark that this inequality is tight, as demonstrated in Appendix E.2. By applying Lemma 12, we get $\max_I \frac{\gamma_{g,1}^2(\boldsymbol{\mu}; I)}{K_I} \geq \frac{\gamma_{g,2}^2(\boldsymbol{\mu}; \emptyset)}{1 + \ln K}$, and substituting in Proposition 3 concludes the proof. ■

7. Summary and Future Work

In this work, we presented the first lower bounds for the CMAB problem that hold for general reward functions, under very mild assumptions. Specifically, we proved both problem-dependent and problem-independent lower bounds, which depend on the modified Gini-smoothness $\tilde{\gamma}_g(\boldsymbol{\mu}; I)$ and reproduce the existing bounds for specific instances. When the reward function is also monotone, we showed that the upper bounds of (Merlis and Mannor, 2019), which depend on the L2 Gini smoothness of the reward function, are tight up to logarithmic factors. There are a few directions for extending our results that we leave for future work:

Gini-smoothness and non-monotone reward functions: One question that naturally arises is whether the L2 Gini-smoothness γ_g also characterizes the lower bound for non-monotone reward function. If such dependence truly holds, a possible way to derive these bounds is by improving $\tilde{\gamma}_g(\boldsymbol{\mu}; I)$ such that it depends on the absolute values of the gradient components. However, such modification of the analysis is highly nontrivial, and we leave it for future work.

Lower bounds for arbitrary action sets: To derive the lower bounds, we carefully designed the action set of the problem, such that no information is gained on one action by sampling a different one. In practice, different actions might have overlapping arms, which can be sometimes used to achieve better performance. For example, in the linear reward function and when the action set contains all possible subsets of fixed size, superior regret bounds can be attained (Komiyama et al., 2015). To the best of our knowledge, the only similar result is by (Combes et al., 2015) for the specific case of linear rewards and independent arms. They show that the best achievable performance strongly depends on the structure of the action set, and it is interesting to derive such lower bounds for general reward functions and arm distributions.

Distribution-dependent lower bounds: To derive lower bounds that depend on the Gini-smoothness, we designed a family of arm distributions, all with binary support. Similarly, to derive the upper bounds, Merlis and Mannor (2019) bounded the variance of the arms by the variance of Bernoulli arms. We can, therefore, conclude that Bernoulli distribution is the ‘worst-case’ distribution, under which both the upper and the lower bounds depend on the Gini-smoothness. A possible future direction is analyzing both bounds under general distributions and deriving new smoothness criteria that depend on the specific arm distribution, rather than the worst-case distribution.

Other variants and performance measures: In this work, we focused on regret lower bounds for the CMAB problem with semi-bandit feedback. Other interesting problems include the case of full-bandit feedback (Gopalan and Mannor, 2015; Rejwan and Mansour, 2020), where there is no feedback on specific arms, but rather on the reward of the action, or using sample complexity as the performance measure (Chen et al., 2017; Mannor and Tsitsiklis, 2004; Kaufmann et al., 2016). Both variants have mainly been studied for the linear reward functions, and extending the existing upper and lower bounds for general reward functions can be beneficial for many practical settings.

Acknowledgments

This work was partially funded by the Israel Science Foundation under ISF grant number 1380/16.

References

Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):1876–1902,

- 2009.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- Apostolos N Burnetas and Michael N Katehakis. Optimal adaptive policies for sequential allocation problems. *Advances in Applied Mathematics*, 17(2):122–142, 1996.
- Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- Lijie Chen, Jian Li, and Mingda Qiao. Nearly instance optimal sample complexity bounds for top-k arm selection. *arXiv preprint arXiv:1702.03605*, 2017.
- Wei Chen, Yajun Wang, and Yang Yuan. Combinatorial multi-armed bandit: General framework and applications. In *International Conference on Machine Learning*, pages 151–159, 2013.
- Wei Chen, Yajun Wang, Yang Yuan, and Qinshi Wang. Combinatorial multi-armed bandit and its extension to probabilistically triggered arms. *The Journal of Machine Learning Research*, 17(1): 1746–1778, 2016a.
- Wei Chen, Wei Hu, Fu Li, Jian Li, Yu Liu, and Pinyan Lu. Combinatorial multi-armed bandit with general reward functions. In *Advances in Neural Information Processing Systems*, pages 1659–1667, 2016b.
- Richard Combes, Mohammad Sadegh Talebi Mazraeh Shahi, Alexandre Proutiere, et al. Combinatorial bandits revisited. In *Advances in Neural Information Processing Systems*, pages 2116–2124, 2015.
- Rémy Degenne and Vianney Perchet. Combinatorial semi-bandit with known covariance. In *Advances in Neural Information Processing Systems*, pages 2972–2980, 2016.
- Yi Gai, Bhaskar Krishnamachari, and Rahul Jain. Combinatorial network optimization with unknown variables: Multi-armed bandits with linear rewards and individual observations. *IEEE/ACM Transactions on Networking (TON)*, 20(5):1466–1478, 2012.
- Aurélien Garivier, Pierre Ménard, and Gilles Stoltz. Explore first, exploit next: The true shape of regret in bandit problems. *Mathematics of Operations Research*, 2018.
- Aditya Gopalan and Shie Mannor. Thompson sampling for learning parameterized markov decision processes. In *Conference on Learning Theory*, pages 861–898, 2015.
- William W Hager. Updating the inverse of a matrix. *SIAM review*, 31(2):221–239, 1989.
- Alihan Hüyük and Cem Tekin. Thompson sampling for combinatorial network optimization in unknown environments. *arXiv preprint arXiv:1907.04201*, 2019.
- Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *The Journal of Machine Learning Research*, 17(1):1–42, 2016.

- Junpei Komiyama, Junya Honda, and Hiroshi Nakagawa. Optimal regret analysis of thompson sampling in stochastic multi-armed bandit problem with multiple plays. *arXiv preprint arXiv:1506.00779*, 2015.
- Branislav Kveton, Csaba Szepesvari, Zheng Wen, and Azin Ashkan. Cascading bandits: Learning to rank in the cascade model. In *International Conference on Machine Learning*, pages 767–776, 2015a.
- Branislav Kveton, Zheng Wen, Azin Ashkan, and Csaba Szepesvari. Combinatorial cascading bandits. In *Advances in Neural Information Processing Systems*, pages 1450–1458, 2015b.
- Branislav Kveton, Zheng Wen, Azin Ashkan, and Csaba Szepesvari. Tight regret bounds for stochastic combinatorial semi-bandits. In *Artificial Intelligence and Statistics*, pages 535–543, 2015c.
- Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- Shie Mannor and John N Tsitsiklis. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research*, 5(Jun):623–648, 2004.
- Nadav Merlis and Shie Mannor. Batch-size independent regret bounds for the combinatorial multi-armed bandit problem. In *Conference on Learning Theory*, pages 2465–2489, 2019.
- Idan Rejwan and Yishay Mansour. Top- k combinatorial bandits with full-bandit feedback. In *Algorithmic Learning Theory*, pages 752–776, 2020.
- Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952.
- Qinshi Wang and Wei Chen. Improving regret bounds for combinatorial semi-bandits with probabilistically triggered arms and its applications. In *Advances in Neural Information Processing Systems*, pages 1161–1171, 2017.
- Siwei Wang and Wei Chen. Thompson sampling for combinatorial semi-bandits. In *International Conference on Machine Learning*, pages 5101–5109, 2018.

Appendix A. General Lower bound for I-disjoint CMAB Problems

Lemma 7 *Let r be a smooth index invariant reward function with $|S| = K$ for all $S \in \mathcal{S}$. Also, let $\underline{\nu}$ be the action distribution of an I-disjoint CMAB problem such that there exists an arm $i \notin I$ in S^* with $\Pr\{X_i^{(t)} = \mu_i\} < 1$ and $\nabla_i r(S^*; \boldsymbol{\mu}) \neq 0$. Then, for all consistent strategies ψ and all suboptimal actions S ,*

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_{\underline{\nu}}[N_{\psi, S}(T)]}{\ln T} \geq \frac{1}{D_{\text{KL}}(\nu_S, \nu_{S^*})} .$$

Specifically, it holds that

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_{\underline{\nu}}[R(T)]}{\ln T} \geq \sum_{S: \Delta_S > 0} \frac{\Delta_S}{D_{\text{KL}}(\nu_S, \nu_{S^*})} .$$

Proof For any arm distribution $\underline{\nu}$ and any suboptimal action S , consider a modified I-disjoint problem with arm distribution $\underline{\nu}'$, where $\nu'_A = \nu_A$ for all $A \neq S$, and thus $r(A; \boldsymbol{\mu}') = r(A; \boldsymbol{\mu})$, and ν'_S is modified such that $r(S; \boldsymbol{\mu}') > r(S^*; \boldsymbol{\mu})$ and ν'_S has the same support as ν_{S^*} . Specifically, ν'_S can be constructed by initializing $\nu'_S \leftarrow \nu_{S^*}$, and then modifying arm i in the direction of the gradient. Due to the assumption, the arm is not deterministic, and there exist arm distributions with the same support and expectations $\mu'_i > \mu_i$ and $\mu'_i < \mu_i$. Thus, modifying it in the direction of the gradient is valid, and since the gradient is not zero, the reward will increase. Furthermore, as the modification is done for an arm $i \notin I$, it does not affect any other action, and the new instance is still I-disjoint.

We now apply Lemma 4 with $Z = N_{\psi, S}(T)/T$, while noting that $D_{\text{KL}}(\nu_A, \nu'_A) = 0$ for all $A \neq S$:

$$\begin{aligned} \mathbb{E}_{\underline{\nu}}[N_{\psi, S}(T)] D_{\text{KL}}(\nu_S, \nu'_S) &\geq \text{kl}\left(\frac{\mathbb{E}_{\underline{\nu}}[N_{\psi, S}(T)]}{T}, \frac{\mathbb{E}_{\underline{\nu}'}[N_{\psi, S}(T)]}{T}\right) \\ &\geq \left(1 - \frac{\mathbb{E}_{\underline{\nu}}[N_{\psi, S}(T)]}{T}\right) \ln\left(\frac{T}{T - \mathbb{E}_{\underline{\nu}'}[N_{\psi, S}(T)]}\right) - \ln 2 . \end{aligned} \quad (11)$$

where the second inequality uses the following bound which holds for any $p, q \in [0, 1]$

$$\text{kl}(p, q) = \underbrace{p \ln \frac{1}{q}}_{\geq 0} + (1-p) \ln \frac{1}{1-q} + \underbrace{p \ln p + (1-p) \ln(1-p)}_{\geq -\ln 2} \geq (1-p) \ln \frac{1}{1-q} - \ln 2 .$$

Next, as ψ is consistent and all actions $A \neq S$ are strictly suboptimal for bandit problem $\underline{\nu}'$, we get that for any $0 < \alpha \leq 1$,

$$0 \leq T - \mathbb{E}_{\underline{\nu}'}[N_{\psi, S}(T)] = \sum_{A \neq S} \mathbb{E}_{\underline{\nu}'}[N_{\psi, A}] = o(T^\alpha) .$$

In particular, for sufficiently large T , it holds that $T - \frac{\mathbb{E}_{\underline{\nu}'}[N_{\psi, S}(T)]}{T} \leq T^\alpha$, and therefore, for any $0 < \alpha \leq 1$,

$$\liminf_{T \rightarrow \infty} \frac{1}{\ln T} \ln\left(\frac{T}{T - \mathbb{E}_{\underline{\nu}'}[N_{\psi, S}(T)]}\right) \geq \liminf_{T \rightarrow \infty} \frac{1}{\ln T} \ln\left(\frac{T}{T^\alpha}\right) = 1 - \alpha .$$

Furthermore, since ψ is consistent and S is suboptimal in $\underline{\nu}$, $\mathbb{E}_{\underline{\nu}}[N_{\psi,S}(T)]/T \rightarrow 0$, and substituting both inequalities into (11) yields

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_{\underline{\nu}}[N_{\psi,S}(T)]}{\ln T} \geq \frac{1}{D_{\text{KL}}(\nu_S, \nu'_S)} .$$

Next, note that if $D_{\text{KL}}(\nu_S, \nu_{S^*}) = \infty$, the result of the lemma trivially holds. Otherwise, recall that there exist $i \notin I$ in S^* with $\Pr\{X_i^{(t)} = \mu_i\} < 1$ and $\nabla_i r(S^*; \boldsymbol{\mu}) \neq 0$, and without loss of generality assume that $\nabla_i r(S^*; \boldsymbol{\mu}) > 0$. By modifying only this component, we can build a sequence of distributions ν'^n such that the following holds:

1. For all n , $\mu_i'^n > \mu_i^*$ and $\mu_j'^n = \mu_j^*$ for all $j \neq i \in S^*$.
2. $\lim_{n \rightarrow \infty} \mu_i'^n = \mu_i^*$.
3. $\lim_{n \rightarrow \infty} D_{\text{KL}}(\nu_S, \nu'^n) = D_{\text{KL}}(\nu_S, \nu_{S^*})$.

By the index invariance assumption, and since $\nabla_i r(S^*; \boldsymbol{\mu}) > 0$, for large enough n it holds that $r(S; \boldsymbol{\mu}'^n) > r(S^*; \boldsymbol{\mu})$, and by taking the infimum over all of these distributions we get for all suboptimal actions

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_{\underline{\nu}}[N_{\psi,S}(T)]}{\ln T} \geq \frac{1}{D_{\text{KL}}(\nu_S, \nu_{S^*})} .$$

Similarly, if $\nabla_i r(S^*; \boldsymbol{\mu}) < 0$, we require that $\mu_i'^n < \mu_i^*$, and the same arguments hold. To derive the second part of the lemma, notice that $R(T) = \sum_{S \in \mathcal{S}} \Delta_S \mathbb{E}_{\underline{\nu}}[N_{\psi,S}(T)]$, and substituting the previous bound for all of the suboptimal actions yields the desired result. \blacksquare

Lemma 11 *Let r be a smooth index invariant reward function, and let $I \subset [m]$ such that $|I| \leq K$ and $m > K$. Also, let $\boldsymbol{\mu}, \boldsymbol{\mu}^* \in \mathbb{R}^K$ such that $\boldsymbol{\mu}_I = \boldsymbol{\mu}^*_I$ and $\Delta = r(\boldsymbol{\mu}^*) - r(\boldsymbol{\mu}) > 0$. Finally, let ν, ν^* be two distributions with expectations $\boldsymbol{\mu}, \boldsymbol{\mu}^* \in \mathbb{R}^K$, such that arms in I are mutually independent of arms outside I and both distributions are identical for arms in I . Then, for any horizon T and any strategy ψ , there exists an I -disjoint CMAB problem with arm distribution $\underline{\nu}'$ such that $\nu_{S^*} = \boldsymbol{\mu}^*$ and its regret under strategy ψ is bounded by*

$$\mathbb{E}_{\underline{\nu}'}[R(T)] \geq T\Delta \left(1 - \frac{K_I}{m-K} - \sqrt{\frac{1}{2} \frac{TK_I}{m-K} D_{\text{KL}}(\nu, \nu^*)} \right).$$

Proof Define an I -disjoint CMAB problem with the maximal action set such that each action contains K arms, i.e., $|\mathcal{S}| = \lfloor \frac{m-|I|}{K-|I|} \rfloor \geq \frac{m-K}{K_I}$, where all actions are distributed according to $\nu_S = \nu$. We denote this problem by $\underline{\nu}$. Due to the pigeonhole principle, there exists an action S^* such that under strategy ψ , $\mathbb{E}_{\underline{\nu}}[N_{\psi, S^*}(T)] \leq \frac{T}{|\mathcal{S}|}$. Next, define a modified bandit problem $\underline{\nu}'$ such that $\nu'_S = \nu$ for all $S \neq S^*$ and $\nu_{S^*} = \boldsymbol{\mu}^*$. Thus,

$$\mathbb{E}_{\underline{\nu}'}[R(T)] = \sum_{S \neq S^*} \Delta \mathbb{E}_{\underline{\nu}'}[N_{\psi, S}(T)] = T\Delta \left(1 - \frac{\mathbb{E}_{\underline{\nu}'}[N_{\psi, S^*}(T)]}{T} \right). \quad (12)$$

As we only changed S^* , for all $S \neq S^*$, $D_{\text{KL}}(\nu_S, \nu'_S) = 0$. Thus, by applying Lemma 4 on the random variable $Z = N_{\psi, S^*}(T)/T$ and then using Pinsker's Inequality, we get

$$\begin{aligned} \mathbb{E}_{\underline{\nu}}[N_{\psi, S^*}(T)] D_{\text{KL}}(\nu, \nu^*) &\geq \text{kl} \left(\frac{\mathbb{E}_{\underline{\nu}}[N_{\psi, S^*}(T)]}{T}, \frac{\mathbb{E}_{\underline{\nu}'}[N_{\psi, S^*}(T)]}{T} \right) \\ &\geq 2 \left(\frac{\mathbb{E}_{\underline{\nu}}[N_{\psi, S^*}(T)]}{T} - \frac{\mathbb{E}_{\underline{\nu}'}[N_{\psi, S^*}(T)]}{T} \right)^2. \end{aligned}$$

It can be easily verified that when $\mathbb{E}_{\underline{\nu}'}[N_{\psi, S^*}(T)]/T$ is either smaller or larger than $\mathbb{E}_{\underline{\nu}}[N_{\psi, S^*}(T)]/T$, it holds that

$$\frac{\mathbb{E}_{\underline{\nu}'}[N_{\psi, S^*}(T)]}{T} \leq \frac{\mathbb{E}_{\underline{\nu}}[N_{\psi, S^*}(T)]}{T} + \sqrt{\frac{1}{2} \mathbb{E}_{\underline{\nu}}[N_{\psi, S^*}(T)] D_{\text{KL}}(\nu, \nu^*)},$$

Substituting $\mathbb{E}_{\underline{\nu}}[N_{\psi, S^*}(T)] \leq \frac{T}{|\mathcal{S}|}$ and $|\mathcal{S}| \geq \frac{m-K}{K_I}$ into the regret bound (12) leads to the desired result. \blacksquare

Appendix B. Technical Lemmas

Lemma 8 Let $p \triangleq p^{\mu, I} \in \mathbb{R}^K$ such that $0 < p_1 < \dots < p_{K_I} < 1$ and define $p_0 = 0$. Also, let ν, ν^* be the distributions stated in Table 2. Then, there exists a constant $b_{\epsilon, 1} > 0$ such that for any $\epsilon \in \mathbb{R}^{K_I}$ with $\|\epsilon\|_\infty \leq b_{\epsilon, 1}$, it holds that

$$D_{\text{KL}}(\nu, \nu^*) \leq 2\epsilon^T B_{\text{KL}}(p)\epsilon, \quad (6)$$

where $B_{\text{KL}}(p) = D(p) + \frac{1}{1-p_{K_I}} \mathbf{1}\mathbf{1}^T$, $D(p) \in \mathbb{R}^{K_I \times K_I}$ is a diagonal matrix whose elements are $D_{ii}(p) = \frac{1}{p_i - p_{i-1}}$ and $\mathbf{1} \in \mathbb{R}^{K_I}$ is a vector of ones.

Proof Assume that $\|\epsilon\|_\infty \leq b_{\epsilon, 0}$, such that ν is a valid probability distribution. Since all arms $i \in I^*$ are mutually independent of all arms $i \notin I^*$, we can write

$$D_{\text{KL}}(\nu, \nu^*) = \underbrace{D_{\text{KL}}(\nu_{I^*}, \nu_{I^*}^*)}_{=0} + D_{\text{KL}}(\nu_{I^{*c}}, \nu_{I^{*c}}^*) = D_{\text{KL}}(\nu_{I^{*c}}, \nu_{I^{*c}}^*)$$

Next, the KL-divergence between the distributions in Table 2 can be bounded by:

$$\begin{aligned} D_{\text{KL}}(\nu, \nu^*) &= p_1 \ln \frac{p_1}{p_1 - \epsilon_1} + \sum_{j=2}^{K_I} (p_j - p_{j-1}) \ln \frac{p_j - p_{j-1}}{p_j - p_{j-1} - \epsilon_j} + (1 - p_{K_I}) \ln \frac{1 - p_{K_I}}{1 - p_{K_I} + \sum_{j=1}^{K_I} \epsilon_j} \\ &\stackrel{(*)}{\leq} \frac{p_1 \epsilon_1}{p_1 - \epsilon_1} + \sum_{j=2}^{K_I} \frac{(p_j - p_{j-1}) \epsilon_j}{p_j - p_{j-1} - \epsilon_j} - \frac{(1 - p_{K_I}) \sum_{j=1}^{K_I} \epsilon_j}{1 - p_{K_I} + \sum_{j=1}^{K_I} \epsilon_j} \\ &= \epsilon_1 + \frac{\epsilon_1^2}{p_1 - \epsilon_1} + \sum_{j=2}^{K_I} \left(\epsilon_j + \frac{\epsilon_j^2}{p_j - p_{j-1} - \epsilon_j} \right) - \sum_{j=1}^{K_I} \epsilon_j + \frac{\left(\sum_{j=1}^{K_I} \epsilon_j \right)^2}{1 - p_{K_I} + \sum_{j=1}^{K_I} \epsilon_j} \\ &= \frac{\epsilon_1^2}{p_1 - \epsilon_1} + \sum_{j=2}^{K_I} \frac{\epsilon_j^2}{p_j - p_{j-1} - \epsilon_j} + \frac{\left(\sum_{j=1}^{K_I} \epsilon_j \right)^2}{1 - p_{K_I} + \sum_{j=1}^{K_I} \epsilon_j} \end{aligned}$$

where (*) is due to the inequality $\ln x \leq x - 1$. Now let

$$b_{\epsilon, 1} = \frac{1}{2} \min \left\{ p_1, \min_i (p_i - p_{i-1}), \frac{1}{K_I} (1 - p_{K_I}), b_{\epsilon, 0} \right\},$$

and for any $\epsilon \in \mathbb{R}^{K_I}$ such that $\|\epsilon\|_\infty \leq b_{\epsilon, 1}$, it holds that

$$D_{\text{KL}}(\underline{\nu}, \underline{\nu}^*) \leq 2 \left(\frac{\epsilon_1^2}{p_1} + \sum_{j=2}^{K_I} \frac{\epsilon_j^2}{p_j - p_{j-1}} + \frac{\left(\sum_{j=1}^{K_I} \epsilon_j \right)^2}{1 - p_{K_I}} \right).$$

Writing the bound in a matrix formulation yields the desired result and concludes the proof. \blacksquare

Lemma 9 *Let r be a smooth index invariant reward function and let $p \in \mathbb{R}^K$ such that $0 < p_1 \leq \dots \leq p_{K_I} < 1$ and there exists $i \in [K_I]$ with $\nabla_i r(p) \neq 0$. Also, define $c_j = \sum_{i=j}^{K_I} \nabla_i r(p)$ for all $j \in [K_I]$ and let $u \in \mathbb{R}^{K_I}$ be a vector such that $c^T u > 0$. Then, if $\epsilon = \epsilon_0 u$, there exists a constant $b_{\epsilon,2} > 0$ such that for all $0 < \epsilon_0 \leq b_{\epsilon,2}$,*

$$\Delta_\epsilon \geq \frac{1}{2} c^T \epsilon > 0 . \quad (7)$$

Proof Define $a_j = \sum_{i=1}^j u_i$ for all $j \in [K_I]$ and $a_j = 0$ for all $j > K_I$, which also results with

$$a^T \nabla r(p) = \sum_{j=1}^{K_I} \sum_{i=1}^j u_i \nabla_j r(p) = \sum_{j=1}^{K_I} u_j \sum_{i=j}^{K_I} \nabla_i r(p) = \sum_{j=1}^{K_I} u_j c_j = c^T u > 0 .$$

Next, by the definition of the gradient,

$$\lim_{\epsilon_0 \rightarrow 0} \frac{r(p) - r(p - \epsilon_0 a)}{\epsilon_0} = a^T \nabla r(p) .$$

Specifically, for $\delta = c^T u / 2 > 0$, there exists $b_{\epsilon,2} > 0$, such that for all $0 < \epsilon_0 \leq b_{\epsilon,2}$,

$$r(p) - r(p - \epsilon_0 a) \geq \epsilon_0 (a^T \nabla r(p) - \delta) = \frac{1}{2} \epsilon_0 c^T u > 0 .$$

We now show that the l.h.s of the inequality is equal to Δ_ϵ . First, Under the index invariance assumption, it holds that $r(p) = r(\mu^{\nu^*})$. Also, note that when comparing to ν^* , the mean value of base arm i in distribution ν decreases by $\sum_{j=1}^i \epsilon_j = \epsilon_0 a_i$, and therefore $r(p - \epsilon_0 a) = r(\mu^\nu)$. Finally, recall that $\Delta_\epsilon = r(\mu^{\nu^*}) - r(\mu^\nu)$, and substituting into the last inequality yields the desired result. \blacksquare

Lemma 10 *Under the notations of Lemmas 8 and 9, for any $0 < p_1 < \dots < p_{K_I} < 1$ and $c \neq 0$, if $\epsilon^* = \epsilon_0 B_{KL}^{-1}(p)c$, then*

$$\epsilon_i^* = \epsilon_0 (p_i - p_{i-1}) \left(c_i - \sum_{j=1}^{K_I} (p_j - p_{i-j}) c_j \right) .$$

Also, if $f_I(\epsilon; p) = \frac{\epsilon^T c c^T \epsilon}{\epsilon^T B_{KL}(p) \epsilon}$, then $f_I(\epsilon^*; p) = \tilde{\gamma}_g^2(\mu; I)$.

Proof By applying the Sherman-Morrison Formula (Hager, 1989) on $B_{KL}(p)$, we get

$$B_{KL}^{-1}(p) = \left(D(p) + \frac{1}{1 - p_{K_I}} \mathbf{1} \mathbf{1}^T \right)^{-1} = D^{-1}(p) - \frac{1}{1 - p_{K_I}} \frac{D^{-1}(p) \mathbf{1} \mathbf{1}^T D^{-1}(p)}{1 + \frac{1}{1 - p_{K_I}} \mathbf{1}^T D^{-1}(p) \mathbf{1}} .$$

$D^{-1}(p)$ is a diagonal matrix whose elements are $D_{ii}^{-1}(p) = \frac{1}{D_{ii}(p)} = p_i - p_{i-1}$. The elements of the matrix $D^{-1}(p) \mathbf{1} \mathbf{1}^T D^{-1}(p)$ are

$$(D^{-1}(p) \mathbf{1} \mathbf{1}^T D^{-1}(p))_{ij} = D^{-1}(p)_{ii} D^{-1}(p)_{jj} = (p_i - p_{i-1})(p_j - p_{j-1})$$

and denominator can be written as

$$\begin{aligned}
 1 + \frac{1}{1 - p_{K_I}} \mathbf{1}^T D^{-1}(p) \mathbf{1} &= 1 + \frac{1}{1 - p_{K_I}} \sum_{i=1}^{K_I} \frac{1}{D_{ii}(p)} \\
 &= 1 + \frac{1}{1 - p_{K_I}} \sum_{i=1}^{K_I} (p_i - p_{i-1}) \\
 &= 1 + \frac{p_{K_I}}{1 - p_{K_I}} \\
 &= \frac{1}{1 - p_{K_I}} .
 \end{aligned}$$

where we used the fact that $p_0 = 0$. Therefore, the elements of $B_{KL}^{-1}(p)c$ are equal to

$$\begin{aligned}
 (B_{KL}^{-1}(p)c)_i &= (D^{-1}(p)c)_i - \frac{1}{1 - p_{K_I}} \left(\frac{D^{-1}(p) \mathbf{1} \mathbf{1}^T D^{-1}(p) c}{1 + \frac{1}{1 - p_{K_I}} \mathbf{1}^T D^{-1}(p) \mathbf{1}} \right)_i \\
 &= (p_i - p_{i-1})c_i - \sum_{j=1}^{K_I} (p_i - p_{i-1})(p_j - p_{j-1})c_j \\
 &= (p_i - p_{i-1}) \left(c_i - \sum_{j=1}^{K_I} (p_j - p_{j-1})c_j \right) .
 \end{aligned}$$

Recalling that $\epsilon^* = \epsilon_0 B_{KL}^{-1}(p)c$ concludes the first part of the lemma. For the second part of the lemma, we directly substitute ϵ^* :

$$\begin{aligned}
 f_I(\epsilon^*; p) &= c^T B_{KL}^{-1}(p)c \\
 &= \sum_{i=1}^{K_I} (p_i - p_{i-1})c_i \left(c_i - \sum_{j=1}^{K_I} (p_j - p_{j-1})c_j \right) \\
 &= \sum_{i=1}^{K_I} (p_i - p_{i-1})c_i^2 - \left(\sum_{i=1}^{K_I} c_i (p_i - p_{i-1}) \right)^2 \\
 &\stackrel{(*)}{=} \sum_{i=1}^{K_I} p_i (c_i^2 - c_{i+1}^2) - \left(\sum_{i=1}^{K_I} p_i (c_i - c_{i+1}) \right)^2 ,
 \end{aligned} \tag{13}$$

where $(*)$ is due to the following identities and under the notations $p_0 = c_{K_I+1} = 0$:

$$\begin{aligned}
 \sum_{i=1}^{K_I} (p_i - p_{i-1})c_i &= \sum_{i=1}^{K_I} p_i (c_i - c_{i+1}) \\
 \sum_{i=1}^{K_I} (p_i - p_{i-1})c_i^2 &= \sum_{i=1}^{K_I} p_i (c_i^2 - c_{i+1}^2) .
 \end{aligned}$$

This term can be further simplified by substituting $c_i = \sum_{j=i}^{K_I} \nabla_j r(p)$ as follows:

$$\begin{aligned}
 f_I(\epsilon^*; p) &= \sum_{i=1}^{K_I} p_i (c_i^2 - c_{i+1}^2) - \left(\sum_{i=1}^{K_I} p_i (c_i - c_{i+1}) \right)^2 \\
 &= \sum_{i=1}^{K_I} p_i \nabla_i r(p)^2 + 2 \sum_{i=1}^{K_I} p_i \sum_{j=i+1}^{K_I} \nabla_i r(p) \nabla_j r(p) - \sum_{i=1}^{K_I} \sum_{j=1}^{K_I} p_i p_j \nabla_i r(p) \nabla_j r(p) \\
 &= \sum_{i=1}^{K_I} p_i \nabla_i r(p)^2 + 2 \sum_{i=1}^{K_I} p_i \sum_{j=i+1}^{K_I} \nabla_i r(p) \nabla_j r(p) \\
 &\quad - \sum_{i=1}^{K_I} p_i^2 \nabla_i r(p)^2 - 2 \sum_{i=1}^{K_I} \sum_{j=i+1}^{K_I} p_i p_j \nabla_i r(p) \nabla_j r(p) \\
 &= \sum_{i=1}^{K_I} p_i (1 - p_i) \nabla_i r(p)^2 + 2 \sum_{i=1}^{K_I} \sum_{j=i+1}^{K_I} p_i (1 - p_j) \nabla_i r(p) \nabla_j r(p) \\
 &= \tilde{\gamma}_g^2(\boldsymbol{\mu}; I) .
 \end{aligned}$$

We remark that by Equation (13) and under the notations $p_0 = 0$, we observe that

$$\tilde{\gamma}_g^2(\boldsymbol{\mu}; I) = \sum_{i=1}^{K_I} (p_i - p_{i-1}) c_i^2 - \left(\sum_{i=1}^{K_I} c_i (p_i - p_{i-1}) \right)^2 .$$

Denoting $p_{K_I+1} = 1$ and $c_{K_I+1} = 0$, the modified Gini-smoothness can be perceived as the variance of a random variable X such that for all $i \in [K_I + 1]$, $\Pr\{X = c_i\} = p_i - p_{i-1}$. This also holds when $p_1 = 0$ and $p_{K_I} = 1$, or when $p_i = p_{i+1}$ for some $i \in [K_I - 1]$. Thus, we can conclude that for any $\boldsymbol{\mu} \in [0, 1]^K$ and any $I \subset [K]$, it holds that $\tilde{\gamma}_g^2(\boldsymbol{\mu}; I) \geq 0$. \blacksquare

Appendix C. Proof of Problem Independent Lower Bound

Theorem 2 *Let r be a smooth index invariant reward function and assume that $m \geq 3K$. Then, for any $\boldsymbol{\mu} \in [0, 1]^K$, any $T \geq T_0$ and for any strategy ψ , there exists an I -disjoint CMAB problem $\underline{\nu}'$ with $\mathbb{E}[\nu'_{S^*}] = \boldsymbol{\mu}$ such that its regret under strategy ψ is bounded by*

$$\mathbb{E}_{\underline{\nu}'}[R(T)] \geq \max_I \frac{\tilde{\gamma}_g(\boldsymbol{\mu}; I)}{32} \sqrt{\frac{T(m-K)}{K_I}} \triangleq IB^*(\boldsymbol{\mu}) .$$

Proof Without loss of generality, assume that $IB^*(\boldsymbol{\mu}) > 0$, since otherwise the bound trivially holds. As in the proof of Theorem 1, we denote $p = p^{\boldsymbol{\mu}, I^*}$, where I^* is the maximizer in $IB^*(\boldsymbol{\mu})$. We also similarly assume that $0 < p_1 < \dots < p_{K_I} < 1$ and fix the distributions ν and ν^* according to Table 2. Finally, we set $\epsilon^* = \epsilon_0 B_{KL}^{-1} c$, with $B_{KL}(p)$ and c as in Lemmas 8 and 9. Recall that with ϵ_0 small enough, both Lemmas 8 and 9 hold. Combining both inequalities, we get

$$D_{\text{KL}}(\nu, \nu^*) = \Delta^2 \frac{D_{\text{KL}}(\nu, \nu^*)}{\Delta^2} \leq \Delta^2 \frac{2\epsilon^{*T} B_{KL}(p) \epsilon^*}{(\frac{1}{2} c^T \epsilon^*)^2} = \frac{8\Delta^2}{f_I(\epsilon^*; p)} ,$$

where $f_I(\epsilon^*; p)$ is defined in Lemma 10. Specifically for our choice of ϵ^* , Lemma 10 also implies that $f_I(\epsilon^*; p) = \tilde{\gamma}_g^2(\boldsymbol{\mu}; I)$. By Lemma 11, there exists an instance of an I -disjoint CMAB problem with gap Δ and optimal action $\mathbb{E}[\nu'_{S^*}] = \boldsymbol{\mu}$ such that

$$\begin{aligned} \mathbb{E}_{\underline{\nu}'}[R(T)] &\geq T\Delta \left(1 - \frac{K_I}{m-K} - \sqrt{\frac{1}{2} \frac{TK_I}{m-K} D_{\text{KL}}(\nu, \nu^*)} \right) \\ &\geq T\Delta \left(1 - \frac{K_I}{m-K} - 2\Delta \sqrt{\frac{TK_I}{(m-K)\tilde{\gamma}_g^2(\boldsymbol{\mu}; I)}} \right) \end{aligned}$$

Next, recall that there exists Δ_0 such that we can achieve any gap $0 < \Delta \leq \Delta_0$ by tuning ϵ_0 . Therefore, for large enough T , there exists ϵ_0 such that $\Delta = \frac{\tilde{\gamma}_g(\boldsymbol{\mu}; I)}{8} \sqrt{\frac{m-K}{TK_I}}$, for which we get the bound

$$\mathbb{E}_{\underline{\nu}'}[R(T)] \geq \frac{\tilde{\gamma}_g(\boldsymbol{\mu}; I)}{8} \sqrt{\frac{m-K}{TK_I}} \left(1 - \frac{1}{2} - \frac{1}{4} \right) = \frac{\tilde{\gamma}_g(\boldsymbol{\mu}; I)}{32} \sqrt{\frac{T(m-K)}{K_I}} .$$

where in the first inequality we use $m \geq 3K$. For brevity, we omitted the cases where $p_1 = 0$, $p_{K_I} = 1$ or $p_i = p_{i+1}$ for some $i \in [K_I]$, and refer the readers to the proof of Theorem 1 for the required modifications. \blacksquare

Appendix D. Relations Between the Smoothness Measures

Proposition 3 *Let r be a smooth index invariant reward function and denote $p = p^{\mu, I}$ for some $\mu \in \mathbb{R}^K$ and $I \subset [K]$. Then,*

$$\tilde{\gamma}_g^2(\mu; I) \geq \frac{\gamma_{g,1}^2(\mu; I)}{3 + \ln \frac{1}{p_1} + \ln \frac{1}{1-p_{K_I}}} , \quad (9)$$

where the r.h.s. is defined as zero if $p_1 = 0$ or $p_{K_I} = 1$.

Proof Without loss of generality, assume that $p_1 > 0$ and $p_{K_I} < 1$. Otherwise, the r.h.s. is zero, and since $\tilde{\gamma}_g(\mu; I) \geq 0$, the result trivially holds (further details on the nonnegativity of $\tilde{\gamma}_g(\mu; I)$ can be found at the end of the proof of Lemma 10). Similarly to Theorem 1, also assume that $p_1 < \dots < p_{K_I}$ and define B_{KL}, c and $f_I(\epsilon; p)$ as in Lemmas 8, 9 and 10. As in Theorem 1, all results can be modified to the case where different arms have the same mean, by defining ϵ only on indices where $p_i < p_{i+1}$, or equivalently forcing $\epsilon_i = 0$ if $p_i = p_{i+1}$. However, we avoid this case for brevity. Finally, we assume that $K_I > 0$ and $\nabla_i r(p) \neq 0$ for some $i \in [K_I]$, since otherwise, both sides of the inequality equal zero, and the bound trivially holds.

Next, we denote $p_0 = 0$ and set $\epsilon_j = \epsilon_0 \sqrt{p_j(1-p_j)} - \epsilon_0 \sqrt{p_{j-1}(1-p_{j-1})}$ for all $j \in [K_I]$. By construction $\sum_{i=1}^j \epsilon_i = \epsilon_0 \sqrt{p_j(1-p_j)}$, and direct substitution yields

$$c^T \epsilon = \sum_{j=1}^{K_I} \sum_{i=j}^{K_I} \epsilon_j \nabla_i r(p) = \sum_{j=1}^{K_I} \sum_{i=1}^j \epsilon_i \nabla_j r(p) = \epsilon_0 \sum_{j=1}^{K_I} \sqrt{p_j(1-p_j)} \nabla_j r(p) , \quad (14)$$

and

$$\begin{aligned} \epsilon^T B_{KL}(p) \epsilon &= \epsilon_0^2 \left(\sum_{j=1}^{K_I} \frac{(\sqrt{p_j(1-p_j)} - \sqrt{p_{j-1}(1-p_{j-1})})^2}{p_j - p_{j-1}} + p_{K_I} \right) \\ &= \epsilon_0^2 \left(1 - p_1 + p_{K_I} + \sum_{j=2}^{K_I} \frac{(\sqrt{p_j(1-p_j)} - \sqrt{p_{j-1}(1-p_{j-1})})^2}{p_j - p_{j-1}} \right) . \end{aligned} \quad (15)$$

Before we further bound this term, note that for any $0 \leq x \leq y \leq \frac{1}{2}$, it holds that

$$\begin{aligned} \left(\sqrt{y(1-y)} - \sqrt{x(1-x)} \right)^2 &\leq \left(\sqrt{y(1-y)} - \sqrt{x(1-y)} \right)^2 = (1-y)(\sqrt{y} - \sqrt{x})^2 \\ &\leq (\sqrt{y} - \sqrt{x})^2 , \end{aligned}$$

where the first inequality uses the fact that $y(1-y) \geq x(1-x) \geq x(1-y)$. From symmetry, it also holds for any $x, y \in [0, \frac{1}{2}]$. If $x \in [0, \frac{1}{2}]$ and $y \in [\frac{1}{2}, 1]$, we bound

$$\begin{aligned} \left(\sqrt{y(1-y)} - \sqrt{x(1-x)} \right)^2 &= y(1-y) + x(1-x) - 2\sqrt{y(1-y)x(1-x)} \\ &\leq \max\{y(1-y), x(1-x)\} - \min\{y(1-y), x(1-x)\} \\ &= |y(1-y) - x(1-x)| \\ &= (y-x)|1-x-y| \\ &\leq (y-x) \end{aligned}$$

If $x, y \in [\frac{1}{2}, 1]$, from symmetry $(\sqrt{y(1-y)} - \sqrt{x(1-x)})^2 \leq (\sqrt{1-y} - \sqrt{1-x})^2$. Recall that p_j are strictly increasing, and denote the last index in which $p_j \leq \frac{1}{2}$ by n . Applying these inequalities on (15), we get

$$\begin{aligned} \epsilon^T B_{KL}(p)\epsilon &\leq \epsilon_0^2 \left(1 - p_1 + p_{K_I} + \sum_{j=2}^n \frac{(\sqrt{p_j} - \sqrt{p_{j-1}})^2}{p_j - p_{j-1}} + \frac{p_{n+1} - p_n}{p_{n+1} - p_n} + \sum_{j=n+2}^{K_I} \frac{(\sqrt{1-p_j} - \sqrt{1-p_{j-1}})^2}{p_j - p_{j-1}} \right) \\ &\leq \epsilon_0^2 \left(3 + \sum_{j=2}^n \frac{(\sqrt{p_j} - \sqrt{p_{j-1}})^2}{p_j - p_{j-1}} + \sum_{j=n+2}^{K_I} \frac{(\sqrt{1-p_j} - \sqrt{1-p_{j-1}})^2}{p_j - p_{j-1}} \right). \end{aligned} \quad (16)$$

Next, we bound the summands as follows:

$$\begin{aligned} \frac{(\sqrt{p_j} - \sqrt{p_{j-1}})^2}{p_j - p_{j-1}} &= \frac{(\sqrt{p_j} - \sqrt{p_{j-1}})^2}{p_j - p_{j-1}} \cdot \frac{(\sqrt{p_j} + \sqrt{p_{j-1}})^2}{(\sqrt{p_j} + \sqrt{p_{j-1}})^2} = \frac{(p_j - p_{j-1})^2}{(p_j - p_{j-1})(\sqrt{p_j} + \sqrt{p_{j-1}})^2} \\ &\leq \frac{p_j - p_{j-1}}{p_j}. \end{aligned}$$

Similarly, we have

$$\begin{aligned} \frac{(\sqrt{1-p_j} - \sqrt{1-p_{j-1}})^2}{p_j - p_{j-1}} &= \frac{(\sqrt{1-p_j} - \sqrt{1-p_{j-1}})^2}{p_j - p_{j-1}} \cdot \frac{(\sqrt{1-p_j} + \sqrt{1-p_{j-1}})^2}{(\sqrt{1-p_j} + \sqrt{1-p_{j-1}})^2} \\ &= \frac{(p_j - p_{j-1})^2}{(p_j - p_{j-1})(\sqrt{1-p_j} + \sqrt{1-p_{j-1}})^2} \\ &\leq \frac{p_j - p_{j-1}}{1 - p_{j-1}}. \end{aligned}$$

Substitute both into (16) yields:

$$\begin{aligned} \epsilon^T B_{KL}(p)\epsilon &\leq \epsilon_0^2 \left(3 + \sum_{j=2}^n \frac{p_j - p_{j-1}}{p_j} + \sum_{j=n+2}^{K_I} \frac{p_j - p_{j-1}}{1 - p_{j-1}} \right) \\ &\stackrel{(*)}{\leq} \epsilon_0^2 \left(3 + \int_{p_1}^{p_n} \frac{dx}{x} + \int_{p_{n+1}}^{p_{K_I}} \frac{dx}{1-x} \right) \\ &\leq \epsilon_0^2 \left(3 + \ln \frac{1}{p_1} + \ln \frac{1}{1-p_{K_I}} \right), \end{aligned}$$

where $(*)$ utilizes the relation between sums and integrals. Combining with (14) and substituting back into $f_I(\epsilon; p)$, we get

$$f_I(\epsilon; p) \geq \frac{(\epsilon_0 \sum_{j=1}^{K_I} \sqrt{p_j(1-p_j)} \nabla_j r(p))^2}{\epsilon_0^2 \left(3 + \ln \frac{1}{p_1} + \ln \frac{1}{1-p_{K_I}} \right)} = \frac{\gamma_{g,1}^2(\boldsymbol{\mu}; I)}{3 + \ln \frac{1}{p_1} + \ln \frac{1}{1-p_{K_I}}}$$

The proof is concluded by applying Lemma 10 and recalling that ϵ^* is the maximizer of $f_I(\epsilon; p)$; therefore, for any $\epsilon \neq 0$, it holds that $f_I(\epsilon; p) \leq f_I(\epsilon^*; p) \leq \tilde{\gamma}_g^2(\boldsymbol{\mu}; I)$. \blacksquare

Lemma 12 *Let A be a nonempty subset of indices $A \subset [n]$. For any vector $x \in \mathbb{R}^n$, it holds that*

$$\max_{A \neq \emptyset} \frac{1}{|A|} \|x_A\|_1^2 \geq \frac{1}{1 + \ln n} \|x\|_2^2 .$$

Proof Without loss of generality, assume that $x_1 \geq x_2 \geq \dots \geq x_n \geq 0$, as reorganizing and taking absolute values do not affect both sides of the inequality. We bound $\|x\|_2^2$ as follows:

$$\begin{aligned} \|x\|_2^2 &= \sum_{l=1}^n x_l^2 \stackrel{(1)}{\leq} \sum_{l=1}^n \left(\frac{1}{l} \sum_{k=1}^l x_k \right)^2 = \sum_{l=1}^n \frac{1}{l} \left[\frac{1}{l} \left(\sum_{k=1}^l x_k \right)^2 \right] \\ &\stackrel{(2)}{\leq} \sum_{l=1}^n \frac{1}{l} \max_{A \neq \emptyset} \frac{1}{|A|} \|x_A\|_1^2 \stackrel{(3)}{\leq} (1 + \ln n) \max_{A \neq \emptyset} \frac{1}{|A|} \|x_A\|_1^2 \end{aligned}$$

In (1) we use the fact that x_i are decreasing and non-negative, and thus $x_i \leq \frac{1}{i} \sum_{k=1}^i x_k$. For (2), we note that $\frac{1}{l} \left(\sum_{k=1}^l x_k \right)^2 = \frac{1}{|A|} \|x_A\|_1^2$ for $A = \{1, \dots, l\}$, and bound it by the maximum over all possible subsets $A \neq \emptyset$. Finally, (3) uses the well-known property of the harmonic sum $\sum_{l=1}^n \frac{1}{l} \leq 1 + \ln n$. ■

Proposition 4 *Let r be a monotone smooth index invariant reward function. Then, for any $\boldsymbol{\mu} \in \mathbb{R}^K$, it holds that*

$$\max_I \frac{\gamma_{g,1}^2(\boldsymbol{\mu}; I)}{K_I} \geq \frac{\gamma_{g,2}^2(\boldsymbol{\mu}; \emptyset)}{1 + \ln K} . \quad (10)$$

Furthermore, if $\mu_{\min} = \min_{i:\mu_i > 0} \mu_i$ and $\mu_{\max} = \max_{i:\mu_i < 1} \mu_i$, then

$$\max_I \frac{\tilde{\gamma}_g^2(\boldsymbol{\mu}; I)}{K_I} \geq \frac{\gamma_{g,2}^2(\boldsymbol{\mu}; \emptyset)}{(1 + \ln K) \left(3 + \ln \frac{1}{\mu_{\min}} + \ln \frac{1}{1 - \mu_{\max}} \right)} .$$

Proof For the first inequality, note that when the function is monotone, all elements in the summation of $\gamma_{g,1}(\boldsymbol{\mu}; I)$ are nonnegative, and it can be thus conceived as the L1 norm of a vector whose components are $\sqrt{\mu_i(i - \mu_i)} \nabla_i r(\boldsymbol{\mu})$. Furthermore, $\gamma_{g,2}(\boldsymbol{\mu}; I)$ can be conceived as the L2 norm of the same vector. Specifically, both are the respective norm of a vector with a subset of these components. Therefore, we can directly relate these two quantities using standard relations between norms. Nonetheless, we are interested in maximizing the lower bound, which includes an additional ‘penalty’ factor on the number of components in the vector $1/K_I$. As a result, choosing the largest number of elements in the sub-vector is not always optimal. Specifically, we show in Lemma 12 that when optimizing the choice of I , the penalized L1 norm of the sub-vector can be lower bounded by the L2 norm of the *full* vector, up to logarithmic factors. Applying this lemma on $\gamma_{g,1}(\boldsymbol{\mu}; I)$ and $\gamma_{g,2}(\boldsymbol{\mu}; \emptyset)$ leads to the first inequality of the proposition.

Next, we prove the second inequality. By Proposition 3, for any set I , if $p = p^{\mu, I}$, it holds that

$$\tilde{\gamma}_g^2(\boldsymbol{\mu}; I) \geq \frac{\gamma_{g,1}^2(\boldsymbol{\mu}; I)}{3 + \ln \frac{1}{p_1} + \ln \frac{1}{1 - p_{K_I}}} .$$

next, we divide by K_I and maximize over I , which yields

$$\max_I \frac{\tilde{\gamma}_g^2(\boldsymbol{\mu}; I)}{K_I} \geq \max_I \frac{\gamma_{g,1}^2(\boldsymbol{\mu}; I)}{K_I \left(3 + \ln \frac{1}{p_1} + \ln \frac{1}{1 - p_{K_I}} \right)} .$$

If the maximizer on the r.h.s. leads to $p_1 = 0$ or $p_{K_I} = 1$, then the r.h.s. of the inequality equals zero for any $I \subset [K]$. We can then choose $I = i$ for all $i \in [K]$, and thus $\mu_i(1 - \mu_i) \nabla_i r(\boldsymbol{\mu})^2 = 0$ for all $i \in [K]$. Therefore, it also holds that $\gamma_{g,2}^2(\boldsymbol{\mu}; \emptyset) = 0$ and the required inequality trivially holds. Otherwise, $p_1 > 0$ and $p_{K_I} < 1$, and specifically, $p_1 > \mu_{\min}$ and $p_{K_I} < \mu_{\max}$. Then, it also holds that

$$\max_I \frac{\tilde{\gamma}_g^2(\boldsymbol{\mu}; I)}{K_I} \geq \max_I \frac{\gamma_{g,1}^2(\boldsymbol{\mu}; I)}{K_I \left(3 + \ln \frac{1}{\mu_{\min}} + \ln \frac{1}{1 - \mu_{\max}} \right)} ,$$

and applying the first result of the proposition leads to its second result. ■

Appendix E. Tightness of the Relations Between the Smoothness Measures

E.1. Modified Smoothness and L1 Smoothness

In this appendix, we demonstrate the tightness of Proposition 3. Specifically, we show that there exists a CMAB instance and exponentially small arm parameters $\boldsymbol{\mu}$ such that for all $I \subset [K]$, it holds that $\tilde{\gamma}_g^2(\boldsymbol{\mu}; I) = \mathcal{O}\left(\frac{\gamma_{g,1}^2(\boldsymbol{\mu}; I)}{K_I}\right)$. This proves that the logarithmic factor in Inequality (9) cannot be replaced with a better factor of $\Omega\left(\frac{1}{\ln K}\right)$.

We start by fixing $\mu_i = 2^{-2(K-i)-1}$ and choosing the CMAB instance such that $\nabla_{i^*} r(\boldsymbol{\mu}) = 2^{K-i}$. Notice that the elements of $\boldsymbol{\mu}$ are sorted in an increasing order; therefore, for any I , the vector $p^{\boldsymbol{\mu}, I}$ contains all the elements of $\boldsymbol{\mu}$ outside I in their original order. Thus, we can directly bound $\tilde{\gamma}_g^2(\boldsymbol{\mu}; I)$ by:

$$\begin{aligned}
\tilde{\gamma}_g^2(\boldsymbol{\mu}; I) &= \sum_{i \notin I} \mu_i (1 - \mu_i) \nabla_{i^*} r(\boldsymbol{\mu})^2 + 2 \sum_{i \notin I} \sum_{j > i, j \notin I} \mu_i (1 - \mu_j) \nabla_{i^*} r(\boldsymbol{\mu}) \nabla_{j^*} r(\boldsymbol{\mu}) \\
&\stackrel{(1)}{\leq} \sum_{i \notin I} \mu_i \nabla_{i^*} r(\boldsymbol{\mu})^2 + 2 \sum_{i \notin I} \sum_{j > i, j \notin I} \mu_i \nabla_{i^*} r(\boldsymbol{\mu}) \nabla_{j^*} r(\boldsymbol{\mu}) \\
&\stackrel{(2)}{\leq} 2 \sum_{i \notin I} \mu_i \nabla_{i^*} r(\boldsymbol{\mu}) \sum_{j \geq i, j \notin I} \nabla_{j^*} r(\boldsymbol{\mu}) \\
&\stackrel{(3)}{\leq} 2 \sum_{i \notin I} 2^{-2(K-i)-1} 2^{K-i} \sum_{j \geq i} 2^{K-j} \\
&= \sum_{i \notin I} 2^{-(K-i)} \sum_{j \geq i} 2^{K-j} \\
&\leq \sum_{i \notin I} 2^{-(K-i)} 2^{K-i+1} \\
&= 2(K - |I|) = 2K_I
\end{aligned}$$

In (1) we removed the terms $1 - \mu_i, 1 - \mu_j$, which increases the expression since they are smaller than 1 and the gradients are positive. Similarly, in (2) we multiplied the first term by 2 and combined the sums into a single term. In (3) we substituted the values of the parameters and increased the internal sum by summing over all $j \geq i$ (including elements in I). Next, we lower bound $\gamma_{g,1}^2(\boldsymbol{\mu}; I)$:

$$\begin{aligned}
\gamma_{g,1}^2(\boldsymbol{\mu}; I) &= \left(\sum_{i \notin I} \sqrt{\mu_i (1 - \mu_i)} \nabla_{i^*} r(\boldsymbol{\mu}) \right)^2 \stackrel{(1)}{\geq} \frac{1}{2} \left(\sum_{i \notin I} \sqrt{\mu_i} \nabla_{i^*} r(\boldsymbol{\mu}) \right)^2 \\
&= \frac{1}{2} \left(\sum_{i \notin I} 2^{-(K-i)-\frac{1}{2}} 2^{K-i} \right)^2 = \frac{1}{4} \left(\sum_{i \notin I} 1 \right)^2 = \frac{K_I^2}{4}
\end{aligned}$$

In (1) we used the fact that for all i , $1 - \mu_i \geq \frac{1}{2}$. Using both bounds, we conclude that for all $I \subset [K]$, it holds that

$$\frac{\tilde{\gamma}_g^2(\boldsymbol{\mu}; I)}{\gamma_{g,1}^2(\boldsymbol{\mu}; I)} \leq \frac{2K_I}{\frac{1}{4}K_I^2} = \mathcal{O}\left(\frac{1}{K_I}\right).$$

Another conclusion from this example is that

$$\frac{\max_I \tilde{\gamma}_g^2(\boldsymbol{\mu}; I)}{\max_I \gamma_{g,1}^2(\boldsymbol{\mu}; I)} \leq \frac{2K}{\frac{1}{4}K^2} = \mathcal{O}\left(\frac{1}{K}\right).$$

Thus, in contrast to the relations between the L1 and L2 Gini-smoothness measure, we cannot improve the inequality by maximizing over the set I . We end this section by remarking that in this example, one can easily observe that $p_1 \leq 2^{-2K_I+1}$, and thus $\ln \frac{1}{p_1} \approx K_I$. Therefore, for this instance, Proposition 3 is tight.

E.2. L1 Smoothness and L2 Smoothness

In this appendix we prove the tightness of the relation between the L1 and L2 smoothness measures. We do so by proving that Lemma 12 is tight up to a constant factor. Let $x \in \mathbb{R}^n$ such that $x_i = \sqrt{i} - \sqrt{i-1}$. Specifically, x is positive and sorted in a decreasing order, and therefore for any $d \in [n]$ and any A such that $|A| = d$, the maximal value of $\frac{\|x_A\|_1^2}{|A|}$ is obtained for $A = [d]$. Moreover, for this specific set, it also holds that

$$\frac{1}{|A|} \|x_A\|_1^2 = \frac{1}{d} \left(\sum_{i=1}^d x_i \right)^2 = \frac{1}{d} \left(\sum_{i=1}^d \sqrt{i} - \sqrt{i-1} \right)^2 = \frac{1}{d} (\sqrt{d})^2 = 1.$$

Finally, we can write the l.h.s. of the norm inequality by

$$\max_A \frac{1}{|A|} \|x_A\|_1^2 = \max_d \max_{|A|=d} \frac{1}{|A|} \|x_A\|_1^2 = \max_d 1 = 1.$$

Next, we bound the r.h.s. of the inequality by

$$\begin{aligned} \|x_A\|_2^2 &= \sum_{i=1}^n (\sqrt{i} - \sqrt{i-1})^2 \\ &= \sum_{i=1}^n \frac{(\sqrt{i} - \sqrt{i-1})^2 (\sqrt{i} + \sqrt{i-1})^2}{(\sqrt{i} + \sqrt{i-1})^2} \\ &= \sum_{i=1}^n \frac{1}{(\sqrt{i} + \sqrt{i-1})^2} \\ &\geq \sum_{i=1}^n \frac{1}{4i} \\ &\geq \frac{\ln(n+1)}{4} \end{aligned}$$

Thus, for this example, $\max_A \frac{1}{|A|} \|x_A\|_1^2 \leq \frac{4}{\ln(n+1)} \|x_A\|_2^2$, and Lemma 12 is tight up to a constant factor.