

Optimal Strategies for Graph-Structured Bandits

Hassan Saber

HASSAN.SABER@INRIA.FR

SequeL Research Group

Inria Lille-Nord Europe & CRIStAL

Villeneuve-d'Ascq, Parc scientifique de la Haute-Borne, France

Pierre Ménard

PIERRE.MENARD@INRIA.FR

SequeL Research Group

Inria Lille-Nord Europe & CRIStAL

Villeneuve-d'Ascq, Parc scientifique de la Haute-Borne, France

Odalric-Ambrym Maillard

ODALRIC.MAILLARD@INRIA.FR

SequeL Research Group

Inria Lille-Nord Europe & CRIStAL

Villeneuve-d'Ascq, Parc scientifique de la Haute-Borne, France

Editor:

Abstract

We study a structured variant of the multi-armed bandit problem specified by a set of Bernoulli distributions $\nu = (\nu_{a,b})_{a \in \mathcal{A}, b \in \mathcal{B}}$ with means $(\mu_{a,b})_{a \in \mathcal{A}, b \in \mathcal{B}} \in [0, 1]^{\mathcal{A} \times \mathcal{B}}$ and by a given weight matrix $\omega = (\omega_{b,b'})_{b, b' \in \mathcal{B}}$, where \mathcal{A} is a finite set of arms and \mathcal{B} is a finite set of users. The weight matrix ω is such that for any two users $b, b' \in \mathcal{B}$, $\max_{a \in \mathcal{A}} |\mu_{a,b} - \mu_{a,b'}| \leq \omega_{b,b'}$. This formulation is flexible enough to capture various situations, from highly-structured scenarios ($\omega \in \{0, 1\}^{\mathcal{B} \times \mathcal{B}}$) to fully unstructured setups ($\omega \equiv 1$). We consider two scenarios depending on whether the learner chooses only the actions to sample rewards from or both users and actions. We first derive problem-dependent lower bounds on the regret for this generic graph-structure that involves a structure dependent linear programming problem. Second, we adapt to this setting the Indexed Minimum Empirical Divergence (IMED) algorithm introduced by Honda and Takemura (2015), and introduce the IMED-GS* algorithm. Interestingly, IMED-GS* does not require computing the solution of the linear programming problem more than about $\log(T)$ times after T steps, while being provably asymptotically optimal. Also, unlike existing bandit strategies designed for other popular structures, IMED-GS* does not resort to an explicit forced exploration scheme and only makes use of local counts of empirical events. We finally provide numerical illustration of our results that confirm the performance of IMED-GS*.

Keywords: Graph-structured stochastic bandits, regret analysis, asymptotic optimality, Indexed Minimum Empirical Divergence (IMED) algorithm.

1. Introduction

The multi-armed bandit problem is a popular framework to formalize sequential decision making problems. It was first introduced in the context of medical trials (Thompson, 1933, 1935) and later formalized by Robbins (1952). In this paper, we consider a contextual and structured variant of the problem, specified by a set of distributions $\nu = (\nu_{a,b})_{a \in \mathcal{A}, b \in \mathcal{B}}$ with means $(\mu_{a,b})_{a \in \mathcal{A}, b \in \mathcal{B}}$, where \mathcal{A} is a finite set of arms and \mathcal{B} is a finite set of users. Such ν is called a (bandit) configuration where each $\nu_b = (\nu_{a,b})_{a \in \mathcal{A}}$ can be seen as a classical multi-armed bandit problem. The streaming

protocol is the following: at each time $t \geq 1$, the learner deals with a user $b_t \in \mathcal{B}$ and chooses an arm $a_t \in \mathcal{A}$, based only on the past. We consider two scenarios: either the sequence of users is deterministic (uncontrolled scenario) or the learner has the possibility to choose the user (controlled scenario), see Section 1.1. The learner then receives and observes a reward X_t sampled according to ν_{a_t, b_t} conditionally independent from the past. We assume binary rewards: each $\nu_{a,b}$ is a Bernoulli distribution $\text{Bern}(\mu_{a,b})$ with mean $\mu_{a,b} \in (0, 1)$ and we denote by \mathcal{D} the set of such configurations. The goal of the learner is then to maximize its expected cumulative reward over T rounds, or equivalently minimize regret given by

$$R(\nu, T) = \mathbb{E}_\nu \left[\sum_{t=1}^T \max_{a \in \mathcal{A}} \mu_{a, b_t} - X_t \right].$$

For this problem one can run, for example, a separate instance of a bandit algorithm for each user b , but we would like to exploit a *known structure* among the users (which we detail below).

Unstructured bandits The classical bandit problem (when $|\mathcal{B}| = 1$) received increased attention in the middle of the 20th century. The seminal paper [Lai and Robbins \(1985\)](#) established the first lower bounds on the cumulative regret, showing that designing a strategy that is optimal uniformly over a given set of configurations \mathcal{D} comes with a price. The study of the lower performance bounds in multi-armed bandits successfully lead to the development of asymptotically optimal strategies for specific configuration sets, such as the KL-UCB strategy ([Lai, 1987](#); [Cappé et al., 2013](#); [Maillard, 2018](#)) for exponential families, or alternatively the DMED and IMED strategies from [Honda and Takemura \(2011, 2015\)](#). The lower bounds from [Lai and Robbins \(1985\)](#), later extended by [Burnetas and Katehakis \(1997\)](#) did not cover all possible configurations, and in particular *structured* configuration sets were not handled until [Agrawal et al. \(1989\)](#) and then [Graves and Lai \(1997\)](#) established generic lower bounds. Here, structure refers to the fact that pulling an arm may reveal information that enables to refine estimation of other arms. Unfortunately, designing efficient strategies that are provably optimal remains a challenge for many structures at the cost of a high computational complexity.

Structured configurations Motivated by the growing popularity of bandits in a number of industrial and societal application domains, the study of structured configuration sets has received increasing attention over the last few years: The linear bandit problem is one typical illustration ([Abbasi-Yadkori et al., 2011](#); [Srinivas et al., 2010](#); [Durand et al., 2017](#)), for which the linear structure considerably modifies the achievable lower bound, see [Lattimore and Szepesvari \(2017\)](#). The study of a unimodal structure naturally appears in the context of wireless communications, and has been considered in [Combes and Proutiere \(2014\)](#) from a bandit perspective, providing an explicit lower bound together with a strategy exploiting this structure. Other structures include Lipschitz bandits ([Magureanu et al., 2014](#)), and we refer to the manuscript [Magureanu \(2018\)](#) for other examples, such as cascading bandits that are useful in the context of recommender systems. [Combes et al. \(2017\)](#) introduced a generic strategy called OSSB (Optimal Structured Stochastic Bandit), stepping the path towards generic multi-armed bandit strategies that are adaptive to a given structure.

Graph-structure In this paper, we consider the following structure: For a *given* weight matrix $\omega = (\omega_{b,b'})_{b,b' \in \mathcal{B}} \in [0, 1]^{\mathcal{B} \times \mathcal{B}}$ inducing a metric on \mathcal{B} , we assume that for any two users $b, b' \in \mathcal{B}$, $\|\mu_b - \mu_{b'}\|_\infty := \max_{a \in \mathcal{A}} |\mu_{a,b} - \mu_{a,b'}| \leq \omega_{b,b'}$. We see the matrix ω as an adjacency matrix of a fully connected weighted graph where each vertex represents a user and each weight $\omega_{b,b'}$ measures proximity between two users, hence we call this a “*graph structure*”. The motivation to study such a

structure is two-fold. On the one hand, in view of paving the way to solving generic structured bandits, the graph structure yields nicely interpretable lower bounds that show how ω effectively modifies the achievable optimal regret and suggests a natural strategy, while being flexible enough to *interpolate* between a fully unstructured and a highly structured setup. On the other hand, multi-armed bandits have been extensively applied to recommender systems: In such systems it is natural to assume that users may not react arbitrarily differently from each other, but that two users that are "close" in some sense will also react similarly when presented with the same item (action). Now, the similarity between any two users may be loosely or accurately known (by studying for instance activities of users on various social networks and refining this knowledge once in a while): The weight matrix ω enables to summarize such imprecise knowledge. Indeed $\omega_{b,b'} = 0$ means that two users behave identically, while $\omega_{b,b'} = 1$ is not informative on the true similarity $\|\mu_b - \mu_{b'}\|_\infty$ that can be anything from arbitrarily small to 1. Hence, studying this structure is both motivated by a theoretical challenge and more applied considerations. To our knowledge this is the first work on *graph structure*. Other structured problems such as Clustered bandits (Gentile et al., 2014), Latent bandits (Maillard and Mannor, 2014), or Spectral bandits (Valko et al., 2014) do not deal with this particular setting.

Goal The primary goal of this paper is to build a provably optimal strategy for this *flexible* notion of structure. To do so, we derive lower bounds and use them to build intuition on how to handle structure, which enables us to establish a novel bandit strategy, that we prove to be optimal. Although specialized to this structure, the mechanisms leading to the strategy and introduced in the proof technique are novel and are of independent interest.

Outline and contributions We formally introduce the graph-structure model in Section 1.2. Graph structure is simple enough while interpolating between a fully unstructured case and highly-structured settings such as clustered bandits (see Figure 1): This makes it a convenient setting to study *structured* multi-armed bandits. In Section 2, we first establish in Proposition 5 a lower bound on the asymptotic number of times a sub-optimal couple must be pulled by any consistent strategy (see Definition 3), together with its corresponding lower bound on the regret (see Corollary 8) involving an optimization problem. In Section 3, we revisit the Indexed Minimum Empirical Divergence (IMED) strategy from Honda and Takemura (2011) introduced for unstructured multi-armed bandits, and adapt it to the graph-structured setting, making use of the lower bounds of Section 2. The resulting strategy is called IMED-GS in the controlled scenario and IMED-GS₂ in the uncontrolled scenario. Our analysis reveals that in view of asymptotic optimality, these strategies may still not optimally exploit the graph-structure in order to trade-off information gathering and low regret. In order to address this difficulty, we introduce the modified IMED-GS* strategy for the controlled scenario (and IMED-GS*₂ in the uncontrolled one). We show in Theorem 11, which is the main result of this paper, that both IMED-GS* and IMED-GS*₂ are *asymptotically optimal* consistent strategies. Interestingly, IMED-GS* does not compute a solution to the optimization problem appearing in the lower bound *at each time step*, unlike for instance OSSB introduced for generic structures, but only about $\log(T)$ times after T steps. Also, if *forced exploration* does not seem to be avoidable for this problem, IMED-GS* does not make use of an explicit *forced exploration* scheme but a more implicit one, based on local counters of empirical events. Up to our knowledge, IMED-GS* is the first strategy with such properties, in the context of a structure requiring to solve an optimization problem, that is provably asymptotically optimal. On a broader perspective, we believe the mechanism used in IMED-GS* as well as the proof techniques could be extended beyond the considered graph-structure, thus opening promising perspective in order to build structure-adaptive optimal strategies for generic structures. Last, we

provide in Section 4 numerical illustrations on synthetic data. They show that IMED-GS* is also *numerically efficient* in practice, both in terms of regret minimization and computation time; this contrasts with some bandit strategies introduced for other structures (as in Combes et al. (2017), Lattimore and Szepesvari (2017)), that in practice suffer from a prohibitive burn-in phase.

1.1 Setting

Let us recall that the goal of the learner is to maximize its expected cumulative reward over T rounds, or equivalently minimize regret given by

$$R(\nu, T) = \mathbb{E}_\nu \left[\sum_{t=1}^T \max_{a \in \mathcal{A}} \mu_{a, b_t} - X_t \right].$$

As mentioned, for this problem one can run, for example, a separate instance of bandit algorithms for each user b , but we would like to exploit the graph structure. We consider two typical scenarios.

Uncontrolled scenario The sequence of users $(b_t)_{t \geq 1}$ is assumed deterministic and does not depend on the strategy of the learner. At each time step $t \geq 1$, the user b_t is revealed to the learner.

Controlled scenario The sequence of users $(b_t)_{t \geq 1}$ is strategy-dependent and at each time step $t \geq 1$, the learner has to choose a user b_t to deal with, based only on the past.

Both scenarios are motivated by practical considerations: *uncontrolled scenario* is the most common setup for recommender systems, while *controlled scenario* is more natural in case the learner interacts actively with available users as in advertisement campaigns. In an *uncontrolled scenario*, the frequencies of user-arrivals are imposed and may be arbitrary, while in a *controlled scenario* all users are available and the learner has to deal with them with similar frequency (even if this means considering a subset of users). We formalize the notion of frequency in the following definition.

Definition 1 (Log-frequency of a user) A sequence of user $(b_t)_{t \geq 1}$ has log-frequencies $\beta \in [0, 1]^{\mathcal{B}}$ if, almost surely, the number of times the learner has dealt with user $b \in \mathcal{B}$ is $N_b(T) = \Theta(T^{\beta_b})$ ¹. In this case, almost surely we have

$$\forall b \in \mathcal{B}, \lim_{T \rightarrow \infty} \frac{\log(N_b(T))}{\log(T)} = \beta_b.$$

In an *uncontrolled scenario*, we assume that the sequence of users $(b_t)_{t \geq 1}$ has positive log-frequencies $\beta \in (0, 1]^{\mathcal{B}}$, with β unknown to the learner. In a *controlled scenario*, we focus only on strategies that induce sequences of users with same log-frequencies, hence all equal to 1, independently on the considered configuration, that is strategies such that, almost surely, $N_b(T) = \Theta(T)$ for all user $b \in \mathcal{B}$.

1.2 Graph Structure

In this section, we introduce the graph structure. We assume that all bandit configurations ν belong to a set of the form:

$$\overline{\mathcal{D}}_\omega := \left\{ \nu \in \mathcal{D} : \forall b, b' \in \mathcal{B}, \max_{a \in \mathcal{A}} |\mu_{a, b} - \mu_{a, b'}| \leq \omega_{b, b'} \right\},$$

1. We say that $u_T = \Theta(v_T)$, if the two sequences u_T and v_T are equivalent.

where $\omega = (\omega_{b,b'})_{b,b' \in \mathcal{B}} \in [0, 1]^{\mathcal{B} \times \mathcal{B}}$ is a weight matrix known to the learner. Intuitively, when the weights are close to 1, we expect no change to the agnostic situation. But, when the weights are close to $\|\mu_b - \mu_{b'}\|_\infty := \max_{a \in \mathcal{A}} |\mu_{a,b} - \mu_{a,b'}|$, we expect significantly lower achievable regret.

Remark 2 For the specific case where $\omega_{b,b'} = 0$, $\overline{\mathcal{D}}_\omega$ corresponds to user b and b' known to be perfectly clustered. The weight matrix given in Figure 1 models three smooth clusters of users. Each cluster is included in a ball of diameter α for the infinite norm $\|\cdot\|_\infty$.

In the sequel we assume the following properties on the weights.

Assumption 1 (Metric weight property) The weight matrix ω satisfies:

- $\omega_{b,b} = 0$ and $\omega_{b,b'} > 0$ for all $b \neq b' \in \mathcal{B}$,
- $\omega_{b,b'} = \omega_{b',b}$ and $\omega_{b,b'} \leq \omega_{b,b''} + \omega_{b'',b'}$ for all $b, b', b'' \in \mathcal{B}$.

This comes without loss of generality, since for the first property, if two users share exactly the same distribution we can see them as one unique user. For the second property, considering $\tilde{\omega}_{b,b'} = \sup_{a \in \mathcal{A}, \nu \in \overline{\mathcal{D}}_\omega} |\mu_{a,b} - \mu_{a,b'}|$ leads to the same set of configuration $\overline{\mathcal{D}}_\omega = \overline{\mathcal{D}}_{\tilde{\omega}}$ and it holds $\tilde{\omega}_{b,b'} = \tilde{\omega}_{b',b}$, $\tilde{\omega}_{b,b'} \leq \tilde{\omega}_{b,b''} + \tilde{\omega}_{b'',b'}$. Such a weight matrix ω naturally induces a metric on \mathcal{B}^2 .




Figure 1: A cluster structure. *Left*: weight matrix of three clusters. *Right*: range of two armed bandit problems included in clusters with center (ν_1, ν_2, ν_3) for various α (the larger α the lighter and larger the box). The value $\alpha = 0$ corresponds to perfect clusters.

1.3 Notations

Let $\mu_b^* = \max_{a \in \mathcal{A}} \mu_{a,b}$ denote the optimal mean for user b and $\mathcal{A}_b^* = \operatorname{argmax}_{a \in \mathcal{A}} \mu_{a,b}$ the set of optimal arms for this user. We define for a couple $(a, b) \in \mathcal{A} \times \mathcal{B}$ its gap $\Delta_{a,b} = \mu_b^* - \mu_{a,b}$. Thus a couple is optimal if its gap is equal to zero and sub-optimal if it is positive. We denote by $\mathcal{O}^* = \{(a, b) \in \mathcal{A} \times \mathcal{B} : \mu_{a,b} = \mu_b^*\}$ the set of optimal couples. Thanks to the chain rule we can rewrite the regret as follows:

$$R(\nu, T) = \sum_{a,b \in \mathcal{A} \times \mathcal{B}} \Delta_{a,b} \mathbb{E}_\nu [N_{a,b}(T)], \quad \text{where } N_{a,b}(t) = \sum_{s=1}^t \mathbb{I}_{\{(a_s, b_s) = (a,b)\}}$$

is the number of pulls of arm a and user b up to time t .

2. Regret Lower bound

In this subsection, we establish lower bounds on the regret for the structure $\overline{\mathcal{D}}_\omega$. In order to obtain non trivial lower bounds we consider, as in the classical bandit problem, strategies that are *consistent* (uniformly good) on $\overline{\mathcal{D}}_\omega$.

Definition 3 (Consistent strategy) *A strategy is consistent on $\overline{\mathcal{D}}_\omega$ if for all configuration $\nu \in \overline{\mathcal{D}}_\omega$, for all sub-optimal couple (a, b) , for all $\alpha > 0$,*

$$\lim_{T \rightarrow \infty} \mathbb{E}_\nu \left[\frac{N_{a,b}(T)}{N_b(T)^\alpha} \right] = 0.$$

Remark 4 *When $\mathcal{B} = \{b\}$, $N_b(T) = T$ and we recover the usual notion of consistency (Lai and Robbins, 1985).*

Before we provide below the lower bound on the cumulative regret, let us give some intuition: To that end, we fix a configuration $\nu \in \overline{\mathcal{D}}_\omega$ and a sub-optimal couple (a, b) . One key observation is that if for all $b' \in \mathcal{B}$ it holds $\mu_b^* - \mu_{a,b'} < \omega_{b,b'}$, this means we can form an environment $\tilde{\nu} \in \overline{\mathcal{D}}_\omega$ such that $\tilde{\mu}_{a',b'} = \mu_{a',b'}$ for all couples (a', b') except (a, b) , and such that $\tilde{\mu}_{a,b}$ satisfies $\mu_b^* < \tilde{\mu}_{a,b} < \mu_{a,b'} + \omega_{b,b'}$. Indeed, in this novel environment, $\tilde{\mu}_{a,b} - \tilde{\mu}_{a,b'} < \omega_{b,b'}$ still holds but (a, b) is now optimal. Hence, we can transform the sub-optimal couple (a, b) in an optimal one without moving the means of the other users. Thanks to this remarkable property, and introducing $\text{kl}(\mu|\mu')$ to denote the Kullback-Leibler divergence between two Bernoulli distributions $\text{Bern}(\mu)$ and $\text{Bern}(\mu')$ with the usual conventions, one can prove then that for all consistent strategy

$$\liminf_{T \rightarrow \infty} \mathbb{E}_\nu \left[\frac{N_{a,b}(T)}{\log(N_b(T))} \right] \geq \frac{1}{\text{kl}(\mu_{a,b} | \mu_b^*)},$$

which is the lower bound that we get *without* graph structure. This suggests that only the users b' such that $\mu_b^* - \mu_{a,b'} > \omega_{b,b'}$ provide information about the behavior of user b . This justifies to introduce for each couple (a, b) the fundamental set

$$\mathcal{B}_{a,b} := \{b' \in \mathcal{B} : \mu_{a,b'} < \mu_b^* - \omega_{b,b'}\}.$$

It is also convenient to introduce its frontier, denoted $\partial \mathcal{B}_{a,b} := \{b' \in \mathcal{B} : \mu_{a,b'} = \mu_b^* - \omega_{b,b'}\}$. Now, in order to report the lower bounds while avoiding tedious technicalities, we slightly restrict the set $\overline{\mathcal{D}}_\omega$. To this end, we introduce the set

$$\mathcal{D}_\omega := \{\nu \in \overline{\mathcal{D}}_\omega : \forall (a, b) \in \mathcal{A} \times \mathcal{B}, \partial \mathcal{B}_{a,b} = \emptyset\}.$$

This definition is justified since the closure of \mathcal{D}_ω is indeed $\overline{\mathcal{D}}_\omega$ (we only remove from $\overline{\mathcal{D}}_\omega$ sets of empty interior). We can now state the following proposition.

Proposition 5 (Graph-structured lower bounds on pulls) *Let us consider a consistent strategy. Then, for all configuration $\nu \in \mathcal{D}_\omega$, almost surely it holds for all sub-optimal couple $(a, b) \notin \mathcal{O}^*$,*

$$\lim_{T \rightarrow \infty} N_b(T) < +\infty \quad \text{or} \quad \liminf_{T \rightarrow \infty} \frac{1}{\log(N_b(T))} \sum_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) N_{a,b'}(T) \geq 1. \quad (1)$$

We then introduce the notion of *Pareto-optimality* based on the lower bounds given in Proposition 5.

Definition 6 (Pareto-optimality) *A strategy is asymptotically Pareto-optimal if for all $\nu \in \mathcal{D}_\omega$,*

$$\forall a \in \mathcal{A}, \quad \limsup_{T \rightarrow \infty} \min_{b: (a,b) \notin \mathcal{O}^*} \frac{1}{\log(N_b(T))} \sum_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) N_{a,b'}(T) \leq 1,$$

with the convention $\min_\emptyset = -\infty$.

Remark 7 *This proposition reveals that the set $\mathcal{B}_{a,b} = \{b' \in \mathcal{B}: \mu_{a,b'} < \mu_b^* - \omega_{b,b'}\}$ plays a crucial role in the graph structure. The definition of \mathcal{D}_ω excludes specific situations when there exists $b, b' \in \mathcal{B}$, $a \in \mathcal{A}$, $\omega_{b,b'} = \mu_b^* - \mu_{a,b'} = \Delta_{a,b} + \mu_{a,b} - \mu_{a,b'}$, that belong to the close set $\overline{\mathcal{D}_\omega}$. Extending the result to $\overline{\mathcal{D}_\omega}$ seems possible but at the price of clarity due to the need to handle degenerate cases.*

In order to derive an asymptotic lower bound on the regret from these asymptotic lower bounds, we have to characterize the growth of the counts $(N_b(\cdot))_{b \in \mathcal{B}}$.

Corollary 8 (Lower bounds on the regret) *Let us consider a consistent strategy and sequences of users with log-frequencies $\beta \in (0, 1]^\mathcal{B}$ independently of the considered configuration in \mathcal{D}_ω . Then, for all configuration $\nu \in \mathcal{D}_\omega$*

$$\liminf_{T \rightarrow \infty} \frac{R(\nu, T)}{\log(T)} \geq C_\omega^*(\beta, \nu) := \min \left\{ \sum_{a,b \notin \mathcal{O}^*} \Delta_{a,b} n_{a,b} : n \in \mathbb{R}_+^{\mathcal{A} \times \mathcal{B}} \right. \quad (2)$$

$$\left. \text{s.t. } \forall (a,b) \notin \mathcal{O}^*, \quad \sum_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) n_{a,b'} \geq \beta_b \right\}.$$

Hence such a strategy is asymptotically optimal if for all $\nu \in \mathcal{D}_\omega$

$$\limsup_{T \rightarrow \infty} \frac{R(\nu, T)}{\log(T)} \leq C_\omega^*(\beta, \nu).$$

Remark 9 *In the previous corollary, log-frequencies β may be either strategy dependent or independent. In an uncontrolled scenario, β is imposed by the setting and does not depend on the followed strategy, while in a controlled scenario we consider strategies that impose $\beta = 1_\mathcal{B} := (1)_{b \in \mathcal{B}}$.*

Like other structured bandit problems (as in Combes et al. (2017), Lattimore and Szepesvari (2017)) this lower bound is characterized by a problem-dependent constant $C_\omega^*(\beta, \nu)$ solution to an *optimization problem*. In the agnostic case we recover the lower bound of the classical multi-armed bandit problem. Indeed, let us introduce for $\alpha \in [0, 1]$ the weight matrix ω_α where all the weights are equal to α (except for the zero diagonal). ω_α is the same weight matrix as in Figure 1 but only for one cluster. Then when there is no structure ($\omega \equiv \omega_1$), we obtain the explicit constant

$$C_{\omega_1}^*(\beta, \nu) = \sum_{b \in \mathcal{B}} \beta_b \sum_{a \in \mathcal{A}: (a,b) \notin \mathcal{O}^*} \frac{\Delta_{a,b}}{\text{kl}(\mu_{a,b} | \mu_b^*)}, \quad (3)$$

that corresponds to solving $|\mathcal{B}|$ bandit problems in parallel (independently the ones from the others). Thus the graph structure allows to interpolate smoothly between $|\mathcal{B}|$ independent bandit problems and a unique one when all the users share the same distributions. In order to illustrate the gain of information due to the graph structure we plot in Figure 2 the expectation $\mathbb{E}_{\nu \sim \mathcal{U}(\mathcal{D}_{\omega_\alpha})} [C_{\omega_\alpha}^*(1_{\mathcal{B}}, \nu) / C_{\omega_1}^*(1_{\mathcal{B}}, \nu)]$ of the ratio between the constant in the structured case (2) and that in the agnostic case (3), where $\mathcal{U}(\mathcal{D}_{\omega_\alpha})$ denotes the uniform distribution over $\mathcal{D}_{\omega_\alpha}$, $\alpha \in [0, 1]$.




Figure 2: Plot of $\alpha \mapsto \mathbb{E}_{\nu \sim \mathcal{U}(\mathcal{D}_{\omega_\alpha})} [C_{\omega_\alpha}^*(1_{\mathcal{B}}, \nu) / C_{\omega_1}^*(1_{\mathcal{B}}, \nu)]$ where ω_α is a matrix where all the weights are equal to α (except for the zero diagonal) and ν is sampled uniformly at random in $\mathcal{D}_{\omega_\alpha}$.

3. IMED type strategies for Graph-structured Bandits

In this section, we present for both the controlled and uncontrolled scenarios, two strategies: IMED-GS* that matches the asymptotic lower bound of Corollary 8 and IMED-GS with a lower computational complexity but weaker guaranty. Both are inspired by the Indexed Minimum Empirical Divergence (IMED) proposed by [Honda and Takemura \(2011\)](#). The general idea behind this algorithm is to enforce, via a well chosen index, the constraints (1) that appears in the optimization problem (2) of the asymptotic lower bound. These constraints intuitively serve as tests to assert whether or not a couple is optimal.

3.1 IMED type strategies for the *controlled scenario*

We consider the *controlled scenario* where the sequence of users $(b_t)_{t \geq 1}$ is strategy-dependent and at each time step $t \geq 1$, the learner has to choose a user b_t and an arm a_t , based only on the past.

3.1.1 THE IMED-GS STRATEGY.

We denote by $\hat{\mu}_{a,b}(t) = \frac{1}{N_{a,b}(t)} \sum_{s=1}^t \mathbb{I}_{\{(a_s, b_s) = (a,b)\}} X_s$ if $N_{a,b}(t) > 0$, 0 otherwise, the empirical mean of the rewards from couple (a, b) . Guided by the lower bound (1) we generalize the IMED index to

take into account the graph structure as follows. For a couple (a, b) and at time t we define

$$I_{a,b}(t) = \begin{cases} \log(N_{a,b}(t)) & \text{if } (a, b) \in \hat{\mathcal{O}}^*(t) \\ \sum_{b' \in \hat{\mathcal{B}}_{a,b}(t)} \text{kl}(\hat{\mu}_{a,b'}(t) | \hat{\mu}_b^*(t) - \omega_{b,b'}) N_{a,b'}(t) + \log(N_{a,b}(t)) & \text{otherwise} \end{cases}, \quad (4)$$

where $\hat{\mu}_b^*(t) = \max_{a \in \mathcal{A}} \hat{\mu}_{a,b}(t)$ is the current best mean for user b , the current set of optimal couple is

$$\hat{\mathcal{O}}^*(t) := \left\{ (a, b) \in \mathcal{A} \times \mathcal{B} : \hat{\mu}_{a,b}(t) = \hat{\mu}_b^*(t) \right\}$$

and the current set of informative users for an empirical sub-optimal couple (a, b) is

$$\hat{\mathcal{B}}_{a,b}(t) := \left\{ b' \in \mathcal{B} : N_{a,b'}(t) > 0 \text{ and } \hat{\mu}_{a,b'}(t) < \hat{\mu}_b^*(t) - \omega_{b,b'} \right\}.$$

This quantity can be seen as a transportation cost for “moving²” a sub-optimal couple to an optimal one, plus exploration terms (the logarithms of the numbers of pulls). When an optimal couple is considered, the transportation cost is null and only the exploration part remains. Note that, as stated in [Honda and Takemura \(2011\)](#), $I_{a,b}(t)$ is an index in the weaker sense since it is not determined only by samples from the couple (a, b) but also uses empirical means of current optimal arms. We define IMED-GS (Indexed Minimum Empirical Divergence for Graph Structure) to be the strategy consisting of pulling a couple with minimum index in Algorithm 1. It works well in practice, see Section 4, and has a low computational complexity (proportional to the number of couples). However, it is known for other structures, see [Lattimore and Szepesvari \(2017\)](#), that such greedy strategy *does not exploit optimally* the structure of the problem. Indeed, at a high level, pulling an apparently sub-optimal couple (a, b) allows to gather information not only about this particular couple but also about other couples due to the structure. In order to attain optimality one needs to find couples that provide the best trade-off between information and low regret. This is exactly what is done in the optimization problem (2).

3.1.2 THE IMED-GS* STRATEGY

In order to address this difficulty we first, thanks to the (weak) indexes, decide whether we need to exploit or explore. In the second case, in order to explore optimally according to the graph structure we solve the optimization problem (2) parametrized by the current estimates of the means and then track the optimal numbers of pulls given by the solution of this problem. More precisely at each round we *choose* but not immediately pull a couple with minimum index

$$(\underline{a}_t, \underline{b}_t) \in \underset{(a,b) \in \mathcal{A} \times \mathcal{B}}{\operatorname{argmin}} I_{a,b}(t).$$

Exploitation: If this couple is currently optimal, $(\underline{a}_t, \underline{b}_t) \in \hat{\mathcal{O}}^*(t)$, we exploit, that is pull this couple.

Exploration: Else we explore arm $a_{t+1} = \underline{a}_t$. To this end, let $n^{\text{opt}}(t)$ be a solution of the empirical

2. This notion refers to the generic proof technique used to derive regret lower bounds. It involves a change-of-measure argument, from the initial configuration in which the couple is sub-optimal to another one chosen to make it optimal.

version of (2) with $\beta = 1_{\mathcal{B}}$, that is

$$n^{\text{opt}}(t) \in \underset{n \in \mathbb{R}_+^{\mathcal{A} \times \mathcal{B}}}{\operatorname{argmin}} \left\{ \sum_{(a,b) \in \mathcal{A} \times \mathcal{B}} (\hat{\mu}_b^*(t) - \hat{\mu}_{a,b}(t)) n_{a,b} \right. \quad (5)$$

$$\left. s.t. \quad \forall (a,b) \notin \hat{\mathcal{O}}^*(t), \hat{\mathcal{B}}_{a,b}(t) \neq \emptyset : \sum_{b' \in \hat{\mathcal{B}}_{a,b}(t)} \operatorname{kl}(\hat{\mu}_{a,b'}(t) | \hat{\mu}_b^*(t) - \omega_{b,b'}) n_{a,b'} \geq 1 \right\}.$$

The current optimal numbers of pulls given by

$$N_{a,b}^{\text{opt}}(t) = n_{a,b}^{\text{opt}}(t) \min_{b' \in \mathcal{B}} I_{a,b'}(t). \quad (6)$$

We then *track*

$$b_{t+1} \in \underset{b \in \hat{\mathcal{B}}_{\underline{a}_t, b_t}(t) \cup \{b_t\}}{\operatorname{argmax}} N_{\underline{a}_t, b}^{\text{opt}}(t) - N_{\underline{a}_t, b}(t). \quad (7)$$

Asymptotically, we expect that all the sub-optimal couples are pulled roughly $\log(T)$ times. Therefore, for all sub-optimal couple (a, b) , the index $I_{a,b}(T)$ should be of order $\log(T)$. Thus we asymptotically recover in the definition of $N_{a,b}^{\text{opt}}(\cdot)$ the optimal number of pulls of couple (a, b) , that is $n_{a,b}' \log(T)$ as suggested in Corollary 8. Finally we pull the selected couple (a_{t+1}, b_{t+1}) . In order to ensure optimality, however, such a direct tracking of the current optimal number of pulls is still a bit too aggressive and we need to force exploration in *some* exploration rounds. We proceed as follows: when we explore arm \underline{a}_t we automatically pull a couple (\underline{a}_t, b) if its number of pulls $N_{\underline{a}_t, b}$ is lower than the logarithm of the number of time we decided to explore this arm. See Algorithm 2 for details. This does not hurt the asymptotic optimality because we expect to explore a sub-optimal arm not more than $\log(T)$ times. On the bright side, this is still different than the traditional forced exploration. Indeed, only few rounds are dedicated to exploration thanks to the first selection with the indexes and among them only a logarithmic number will consist of pure exploration: Thus, we expect an overall $\log \log(T)$ rounds of forced exploration. Note also that all the quantities involved in this forced exploration use empirical counters. Putting all together we end up with strategy IMED-GS* described in Algorithm 2.

Comparison with other strategies IMED-GS* combines ideas from IMED introduced by [Honda and Takemura \(2011\)](#) and from OSSB by [Combes et al. \(2017\)](#). More precisely, it generalizes the index from IMED to the graph structure. From OSSB it borrows the tracking of the optimal counts given by the asymptotic lower bound (see also [Lattimore and Szepesvari \(2017\)](#)) and the way to force exploration sparingly. The main difference with OSSB is that IMED-GS* leverages the indexes to deal with the exploitation-exploration trade-off. In particular IMED-GS* *does not need* to solve at each round the optimization problem (2). This greatly improves the computational complexity. Also, note that OSSB requires choosing a tuning parameter that must be positive to ensure theoretical guarantees but that must be set equal to 0 to work well in practice. This is not the case for IMED-GS* that requires no parameter tuning and that works well both in theory and in practice (see Section 4).

3.2 IMED type strategies for the *uncontrolled scenario*

In this section, an *uncontrolled scenario* is considered where the sequence of users $(b_t)_{t \geq 1}$ is assumed deterministic and does not depend on the strategy of the learner. We adapt the two previous strategies IMED-GS and IMED-GS* to this scenario.

Algorithm 1 IMED-GS (controlled scenario)

Require: Weight matrix $(\omega_{b,b'})_{b,b' \in \mathcal{B}}$.
for $t = 1 \dots T$ **do**
 Pull $(a_{t+1}, b_{t+1}) \in \underset{(a,b) \in \mathcal{A} \times \mathcal{B}}{\operatorname{argmin}} I_{a,b}(t)$
end for

Algorithm 2 IMED-GS^{*} (controlled scenario)

Require: Weight matrix $(\omega_{b,b'})_{b,b' \in \mathcal{B}}$.
 $\forall a \in \mathcal{A}, \quad c_a, c_a^+ \leftarrow 1$
for For $t = 1 \dots T$ **do**
 Choose $(\underline{a}_t, \underline{b}_t) \in \underset{(a,b) \in \mathcal{A} \times \mathcal{B}}{\operatorname{argmin}} I_{a,b}(t)$
 if $(\underline{a}_t, \underline{b}_t) \in \hat{\mathcal{O}}^*(t)$ **then**
 Choose $(a_{t+1}, b_{t+1}) = (\underline{a}_t, \underline{b}_t)$
 else
 Set $a_{t+1} = \underline{a}_t$
 if $c_{a_{t+1}} = c_{a_{t+1}}^+$ **then**
 $c_{a_{t+1}}^+ \leftarrow 2c_{a_{t+1}}^+$
 Choose $b_{t+1} \in \underset{b \in \mathcal{B}}{\operatorname{argmin}} N_{a,b}(t)$
 else
 Choose $b_{t+1} \in \underset{b \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t) \cup \{\underline{b}_t\}}{\operatorname{argmax}} N_{a_{t+1}, b}^{\operatorname{opt}}(t) - N_{a_{t+1}, b}(t)$
 end if
 $c_{a_{t+1}} \leftarrow c_{a_{t+1}} + 1$
 end if
 Pull (a_{t+1}, b_{t+1})
end for

IMED-GS₂ strategy At time step $t \geq 1$ the choice of user b_t is no longer strategy-dependent but is imposed by the sequence of users $(b_t)_{t \geq 1}$ which is assumed to be deterministic in the *uncontrolled scenario*. The learner only chooses an arm to pull a_t knowing user b_t . We define IMED-GS₂ to be the strategy consisting of pulling an arm with minimum index in Algorithm 3 of Appendix C. IMED-GS₂ suffers the same advantages and shortcomings as IMED-GS. It does not exploit optimally the structure of the problem but it works well in practice, see Section 4, and has a low computational complexity.

IMED-GS₂^{*} strategy In order to explore optimally according to the graph structure in the *uncontrolled scenario*, we also track the optimal numbers of pulls. β may be at first glance different from $1_{\mathcal{B}}$. This requires some normalizations. First, for all time step $t \geq 1$, $n^{\operatorname{opt}}(t)$ now denotes a solution of the empirical version of (2) with $\beta = (\hat{\beta}_b(t))_{b \in \mathcal{B}}$ where $\hat{\beta}_b(t) = \log(N_b(t)) / \log(t)$ estimates log-frequency β_b of user $b \in \mathcal{B}$. Second, we have to consider normalized indexes $\tilde{I}_{a,b}(t) = I_{a,b}(t) / \hat{\beta}_b(t)$ for couples $(a, b) \in \mathcal{A} \times \mathcal{B}$ in order to have $\tilde{I}_{a,b}(T) \sim \log(T)$ as in the *controlled scenario*. An additional difficulty is that at a given time step $t \geq 1$, while the indexes indicate to explore, the current tracked user (see Equation 7) given is likely to be different from user b_t with whom the learner deals. This difficulty is easy to circumvent by postponing and prioritizing the exploration until the learner deals with the tracked user. Priority in exploration phases is given to first delayed forced-exploration and delayed exploration based on solving optimization problem (2), then exploration based on current indexes (see Algorithm 4 in Appendix C). IMED-GS₂^{*} corresponds essentially to IMED-GS^{*} with some delays due to the fact that the tracked and the current users may be different. This has no impact on the optimality of IMED-GS₂^{*} since log-frequencies of users are enforced to be positive.

3.3 Asymptotic optimality of IMED type strategies

In order to prove the asymptotic optimality of IMED-GS^{*} we introduce the following mild assumptions on the configuration considered.

Definition 10 (Non-peculiar configuration) A configuration $\nu \in \mathcal{D}_\omega$ is non-peculiar if the optimization problem (2) admits a unique solution and each user b admits a unique optimal arm a_b^* .

In Theorem 11 we state the main result of this paper, namely, the asymptotic optimality of IMED-GS^{*} and IMED-GS₂^{*}. We prove this result for IMED-GS^{*} in Appendix E and adapt this proof in Appendix G for IMED-GS₂^{*}. Please refer to Proposition 20 (Appendix D) for more refined finite-time upper bounds. As a byproduct of this analysis we deduce the Pareto-optimality of IMED-GS and IMED-GS₂ stated in Proposition 12 and proved in Appendix G.

Theorem 11 (Asymptotic optimality) Both IMED-GS^{*} and IMED-GS₂^{*} are consistent strategies. Further, they are asymptotically optimal on the set of non-peculiar configurations, that is, for all $\nu \in \mathcal{D}_\omega$ non-peculiar, under IMED-GS^{*} the sequence of users has log-frequencies $\mathbf{1}_B$ and we have

$$\limsup_{T \rightarrow \infty} \frac{R(\nu, T)}{\log(T)} \leq C_\omega^*(\mathbf{1}_B, \nu),$$

and, under IMED-GS₂^{*}, assuming a sequence of users with log-frequencies $\beta \in (0, 1]^B$, we have

$$\limsup_{T \rightarrow \infty} \frac{R(\nu, T)}{\log(T)} \leq C_\omega^*(\beta, \nu).$$

Proposition 12 (Asymptotic Pareto-optimality) Both IMED-GS and IMED-GS₂ are consistent strategies. Further, they are asymptotically Pareto-optimal on the set of non-peculiar configurations, that is, under IMED-GS or IMED-GS₂, for all $\nu \in \mathcal{D}_\omega$ non-peculiar,

$$\forall a \in \mathcal{A}, \quad \limsup_{T \rightarrow \infty} \min_{b: (a,b) \notin \mathcal{C}^*} \frac{1}{\log(N_b(T))} \sum_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) N_{a,b'}(T) \leq 1.$$

Discussion Removing forced exploration remains the most challenging task for structured bandit problems. In the context of this structure, forced exploration would have been to use criteria like "if $N_{a,b}(T) < \hat{c}_{a,b}(T) \log(T)$, then pull couple (a, b) " for some constants $\hat{c}_{a,b}(T)$ that depends on the minimization problem coming from the lower bound and where $\hat{c}_{a,b}(T) \log(T)$ can be interpreted as an estimator of the theoretical asymptotic lower bound on the numbers of pulls of couple (a, b) . In stark contrast, in IMED-GS^{*} there is no forced exploration in choosing the arm to explore and, in choosing the user to explore, the used criteria is more intrinsic as it reads "if $N_{a,b}(T) < \hat{c}_{a,b}(T) I_a(T)$, then pull couple (a, b) ", where $I_a(T) \sim \log(T)$ but really depends on $(N_{a,b}(T))_{(a,b) \notin \mathcal{C}^*}$. Thus, the used criteria are not asymptotic, and do not depend on the time t but on the current numbers of pull of sub-optimal arms. Since theoretical asymptotic lower bounds on the numbers of pulls are significantly larger than the current numbers of pulls in finite horizon (see Figure 3), IMED-GS^{*} strategy is also expected to behave better than strategies based on usual (conservative) forced exploration. Although entirely removing forced exploration would be nicer, in IMED-GS^{*}, forced exploration is only done in a sparing way.

4. Numerical experiments

In this section, we compare empirically the following strategies introduced beforehand: IMED-GS and IMED-GS^{*} described respectively in Algorithms 1, 2, IMED-GS₂ and IMED-GS₂^{*} described respectively

in Algorithms 3, 4 and the baseline IMED by Honda and Takemura (2011) that does not exploit the structure. We compare these strategies on two setups, each with $|\mathcal{B}| = 10$ users and $|\mathcal{A}| = 5$ arms. For the *uncontrolled scenario* we consider the round-robin sequence of users. As expected the strategies leveraging the graph structure perform better than the baseline IMED that does not exploit it. Furthermore, the plots suggest that IMED-GS and IMED-GS* (respectively IMED-GS₂ and IMED-GS₂*) perform similarly in practice.

Fixed configuration Figure 3: *Left* – For these experiments we investigate these strategies on a *fixed* configuration. The weight matrix ω and the configuration $\nu \in \mathcal{D}_\omega$ are given in Appendix I. This enables us to plot also the asymptotic lower bound on the regret for reference: We plot the unstructured lower bound (LB_agnostic) in dashed red line, and the structured lower bound (LB_struct) in dashed blue line.

Random configurations Figure 3: *Right* – In these experiments we average regrets over random configurations. We proceed as follows: At each run we sample uniformly at random a weight matrix ω and then sample uniformly at random a configuration $\nu \in \mathcal{D}_\omega$.

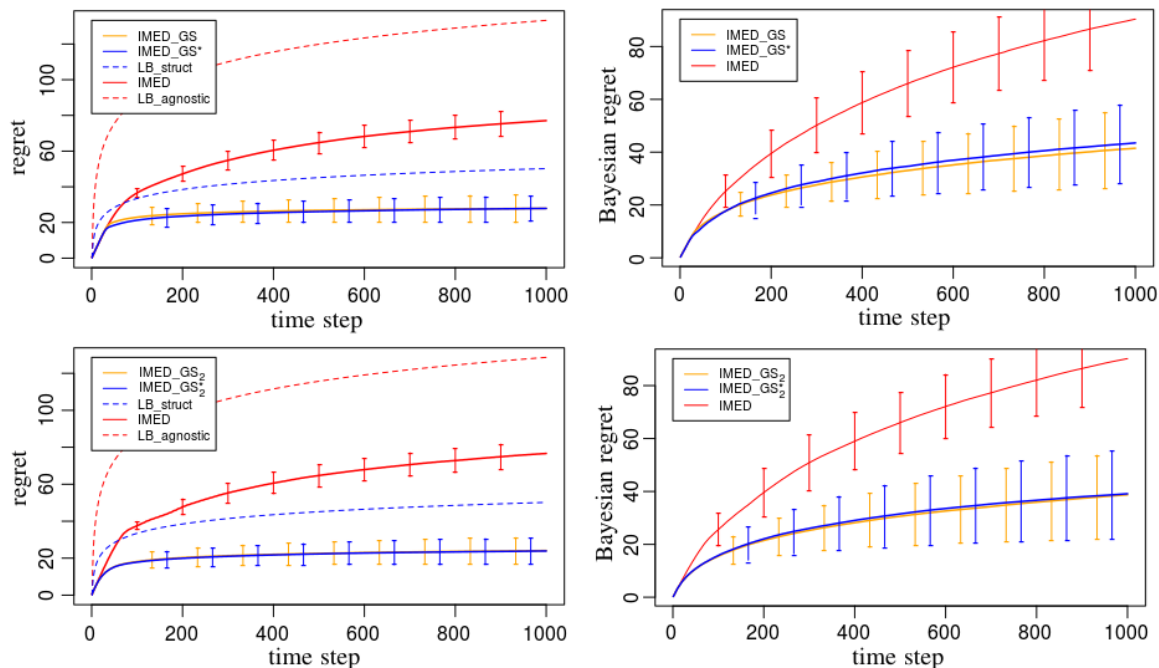


Figure 3: Regret approximated over 1000 runs. *Top: controlled scenario. Bottom: uncontrolled scenario, $(b_t)_{t \geq 1}$ is the round-robin sequence of users. Left: Fixed configuration. Right: Random configurations.*

Additional experiments in Appendix I confirm that both IMED-GS and IMED-GS* induce sequences of users with log-frequencies all equal to 1.

Acknowledgments

This work was supported by CPER Nord-Pas de Calais/FEDER DATA Advanced data science and technologies 2015-2020, the French Ministry of Higher Education and Research, the French Agence Nationale de la Recherche (ANR) under grant ANR-16- CE40-0002 (project BADASS), and Inria Lille – Nord Europe.

References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320, 2011.
- Rajeev Agrawal, Demosthenis Teneketzis, and Venkatachalam Anantharam. Asymptotically efficient adaptive allocation schemes for controlled iid processes: Finite parameter space. *IEEE Transactions on Automatic Control*, 34(3), 1989.
- Apostolos N. Burnetas and Michael N. Katehakis. Optimal adaptive policies for Markov decision processes. *Mathematics of Operations Research*, 22(1):222–255, 1997.
- Olivier Cappé, Aurélien Garivier, Odalric-Ambrym Maillard, Rémi Munos, and Gilles Stoltz. Kullback–Leibler upper confidence bounds for optimal sequential allocation. *Annals of Statistics*, 41(3):1516–1541, 2013.
- Richard Combes and Alexandre Proutiere. Unimodal bandits: Regret lower bounds and optimal algorithms. In *International Conference on Machine Learning*, 2014.
- Richard Combes, Stefan Magureanu, and Alexandre Proutiere. Minimal exploration in structured stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 1763–1771, 2017.
- Audrey Durand, Odalric-Ambrym Maillard, and Joelle Pineau. Streaming kernel regression with provably adaptive mean, variance, and regularization. *arXiv preprint arXiv:1708.00768*, 2017.
- Claudio Gentile, Shuai Li, and Giovanni Zappella. Online clustering of bandits. In *International Conference on Machine Learning*, pages 757–765, 2014.
- Todd L Graves and Tze Leung Lai. Asymptotically efficient adaptive choice of control laws in controlled markov chains. *SIAM journal on control and optimization*, 35(3):715–743, 1997.
- Junya Honda and Akimichi Takemura. An asymptotically optimal policy for finite support models in the multiarmed bandit problem. *Machine Learning*, 85(3):361–391, 2011.
- Junya Honda and Akimichi Takemura. Non-asymptotic analysis of a new bandit algorithm for semi-bounded rewards. *Machine Learning*, 16:3721–3756, 2015.
- Tze Leung Lai. Adaptive treatment allocation and the multi-armed bandit problem. *The Annals of Statistics*, pages 1091–1114, 1987.
- Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- Tor Lattimore and Csaba Szepesvári. The end of optimism? an asymptotic analysis of finite-armed linear bandits. In *Artificial Intelligence and Statistics*, pages 728–737, 2017.
- Stefan Magureanu. *Efficient Online Learning under Bandit Feedback*. PhD thesis, KTH Royal Institute of Technology, 2018.

- Stefan Magureanu, Richard Combes, and Alexandre Proutiere. Lipschitz bandits: Regret lower bounds and optimal algorithms. *Machine Learning*, 35:1–25, 2014.
- O-A Maillard. Boundary crossing probabilities for general exponential families. *Mathematical Methods of Statistics*, 27(1):1–31, 2018.
- Odalric-Ambrym Maillard and Shie Mannor. Latent bandits. In *International Conference on Machine Learning*, pages 136–144, 2014.
- H. Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematics Society*, 58:527–535, 1952.
- Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: no regret and experimental design. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, pages 1015–1022. Omnipress, 2010.
- William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- William R Thompson. On a criterion for the rejection of observations and the distribution of the ratio of deviation to sample standard deviation. *The Annals of Mathematical Statistics*, 6(4):214–219, 1935.
- Michal Valko, Rémi Munos, Branislav Kveton, and Tomáš Kocák. Spectral bandits for smooth graph functions. In *International Conference on Machine Learning*, 2014.

Contents

1	Introduction	1
1.1	Setting	4
1.2	Graph Structure	4
1.3	Notations	5
2	Regret Lower bound	6
3	IMED type strategies for Graph-structured Bandits	8
3.1	IMED type strategies for the <i>controlled scenario</i>	8
3.1.1	The IMED-GS strategy.	8
3.1.2	The IMED-GS* strategy	9
3.2	IMED type strategies for the <i>uncontrolled scenario</i>	10
3.3	Asymptotic optimality of IMED type strategies	11
4	Numerical experiments	12
A	Notations and reminders	18
B	Proof related to the regret lower bound (Section 2)	18
B.1	Almost sure asymptotic lower bounds under consistent strategies	18
B.1.1	$\mathbb{E}_{\tilde{\nu}}[N_b(T)^{1-\alpha}\mathbb{1}_{\{\Omega_T\}}]$ tends to 0 when T tends to infinity	20
B.1.2	$\mathbb{P}_{\nu}(\Omega_T \cap E_T^c)$ tends to 0 when T tends to infinity	21
B.2	Asymptotic lower bounds on the regret	23
C	Algorithms for the <i>uncontrolled scenario</i>	24
D	IMED-GS*: Finite-time analysis	25
D.1	Strategy-based empirical bounds	25
D.2	Reliable current best arm and means	28
D.3	Pareto-optimality and upper bounds on the numbers of pulls of sub-optimal arms	29
D.4	IMED-GS* is consistent and induces sequences of users with log-frequencies $1_{\mathcal{B}}$	31
D.5	The counters c_a and c_a^+ coincide at most $O(\log(\log(T)))$ times	33
D.6	All couples $(a, b) \in \mathcal{A} \times \mathcal{B}$ are asymptotically pulled an infinite number of times	35
D.7	Concentration lemmas	36
D.8	Proof of Lemma 18	37
E	IMED-GS*: Proof of Theorem 11 (main result)	38
E.1	Almost surely $n^{\text{opt}}(T)$ tends to n^{ν}	39
E.2	Almost surely and on expectation, for all sub-optimal couple $\frac{N_{a,b}(T)}{\log(T)}$ tends to $n_{a,b}^{\nu}$	40
F	Concentration lemmas: Proofs	41

G	IMED-GS, IMED-GS₂, IMED-GS₂[*]: Finite-time analysis	44
G.1	IMED-GS finite-time analysis	44
G.2	<i>Uncontrolled scenario</i> : Finite-time analysis	45
G.2.1	Empirical bounds on the numbers of pulls	45
G.2.2	Concentration inequality with bounded time delays	46
H	Continuity of solutions to parametric linear programs	47
I	Details on numerical experiments	47

Appendix A. Notations and reminders

For $\nu \in \mathcal{D}_\nu$, we define

$$\varepsilon_\nu := \min_{\substack{(a,b) \notin \mathcal{O}^*, b' \in \mathcal{B} \\ \mu_{a,b'} - \mu_b^* - \omega_{b,b'} \neq 0}} \left\{ \frac{|\mu_{a,b'} - \mu_b^* - \omega_{b,b'}|}{4}, \frac{\mu_{a,b}}{4}, \frac{1 - \mu_b^*}{4} \right\}.$$

Then, there exists $\alpha_\nu : \mathbb{R}_+^* \rightarrow \mathbb{R}_+^*$ such that $\lim_{\varepsilon \rightarrow 0} \alpha_\nu(\varepsilon) = 0$ and such that for all $0 < \varepsilon < \varepsilon_\nu$, for all $(a, b) \notin \mathcal{O}^*$, for all $b' \in \mathcal{B}$,

$$\mu_{a,b'} \neq \mu_b^* - \omega_{b,b'} \quad \Rightarrow \quad \frac{\text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'})}{1 + \alpha_\nu(\varepsilon)} \leq \text{kl}(\mu_{a,b'} \pm \varepsilon | \mu_b^* - \omega_{b,b'} \pm \varepsilon) \leq (1 + \alpha_\nu(\varepsilon)) \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}).$$

We also introduce the following constant of interest

$$E_\nu := 6e \max_{a \in \mathcal{A}, b \in \mathcal{B}} \left(1 - \frac{\log(1 - \mu_{a,b} - \varepsilon_\nu)}{\log(1 - \mu_{a,b})} \right)^{-1} \left(1 - e^{-\left(1 - \frac{\log(1 - \mu_{a,b} - \varepsilon_\nu)}{\log(1 - \mu_{a,b})} \right) \text{kl}(\mu_{a,b} | \mu_{a,b} - \varepsilon_\nu)} \right)^{-3}.$$

Lastly, for all couple $(a, b) \in \mathcal{A} \times \mathcal{B}$, for all $n \geq 1$, we consider the stopping times

$$\tau_{a,b}^n := \inf \{ t \geq 1 : N_{a,b}(t) = n \}$$

and define

$$\hat{\mu}_{a,b}^n := \hat{\mu}_{a,b}(\tau_{a,b}^n).$$

Appendix B. Proof related to the regret lower bound (Section 2)

In this section we regroup the proofs related to the lower bounds.

B.1 Almost sure asymptotic lower bounds under consistent strategies

In this section we prove Proposition 5.

Let us consider a consistent strategy on \mathcal{D}_ω . Let $\nu \in \mathcal{D}_\omega$ and let us consider $(a, b) \notin \mathcal{O}^*$. We show that almost surely $\lim_{T \rightarrow \infty} N_b(T) = +\infty$ implies

$$\liminf_{T \rightarrow \infty} \frac{1}{\log(N_b(T))} \sum_{b' \in \mathcal{B}_{a,b}} N_{a,b'}(T) \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) \geq 1$$

where $\mathcal{B}_{a,b} = \{b' \in \mathcal{B} : (a, b') \notin \mathcal{O}^* \text{ and } \mu_{a,b'} < \mu_b^* - \omega_{b,b'}\}$.

Proof Let us consider $\tilde{\nu} \in \mathcal{D}_\omega$, a maximal confusing distribution for the sub-optimal couple (a, b) such that a is the unique optimal arm of user b for $\tilde{\nu}$, defined as follows:

- $\forall a' \neq a, b' \in \mathcal{B}, \quad \tilde{\mu}_{a',b'} = \mu_{a',b'}$
- $\forall b'' \notin \mathcal{B}_{a,b}, \quad \tilde{\mu}_{a,b''} = \mu_{a,b''}$
- $\forall b' \in \mathcal{B}_{a,b}, \quad \tilde{\mu}_{a,b'} = \mu_b^* - \omega_{b,b'} + \varepsilon$

where $0 < \varepsilon < \varepsilon_0 = \min_{b' \in \mathcal{B}} |\mu_b^* - \omega_{b,b'} - \mu_{a,b'}|$. Our assumption on $\mathcal{D}_\omega \subset \overline{\mathcal{D}}_\omega$ ensures that $\varepsilon_0 > 0$.

Note that ε is chosen in such a way that for all $b', b'' \in \mathcal{B}, \max_{a \in \mathcal{A}} |\tilde{\mu}_{a,b'} - \tilde{\mu}_{a,b''}| \leq \omega_{b',b''}$. Indeed we have:

- for $b', b'' \in \mathcal{B}_{a,b} : \quad |\tilde{\mu}_{a,b'} - \tilde{\mu}_{a,b''}| = |\omega_{b,b'} - \omega_{b,b''}| \leq \omega_{b',b''}$
- for $b', b'' \notin \mathcal{B}_{a,b} : \quad |\tilde{\mu}_{a,b'} - \tilde{\mu}_{a,b''}| = |\mu_{a,b'} - \mu_{a,b''}| \leq \omega_{b',b''}$
- for $b' \in \mathcal{B}_{a,b}$ and $b'' \notin \mathcal{B}_{a,b}$, we have $\tilde{\mu}_{a,b'} - \tilde{\mu}_{a,b''} = \mu_b^* - \omega_{b,b'} + \varepsilon - \mu_{a,b''}$. Since in this case $b' \in \mathcal{B}_{a,b}$ it implies $\mu_{a,b'} \leq \mu_b^* - \omega_{b,b'}$ and since $b'' \notin \mathcal{B}_{a,b} : \mu_{a,b''} \geq \mu_b^* - \omega_{b,b''} + \varepsilon_0$. Therefore on one hand we get

$$\mu_b^* - \omega_{b,b'} + \varepsilon - \mu_{a,b''} \geq \mu_{a,b'} - \mu_{a,b''} \geq -\omega_{b',b''},$$

and on the other hand

$$\mu_b^* - \omega_{b,b'} + \varepsilon - \mu_{a,b''} \leq \mu_b^* - \omega_{b,b'} + \varepsilon - (\mu_b^* - \omega_{b,b''} + \varepsilon_0) = \varepsilon - \varepsilon_0 + \omega_{b,b''} - \omega_{b,b'} \leq \omega_{b',b''}.$$

Actually, we can choose $0 < \varepsilon < \varepsilon_\nu$ so that :

$$\forall b' \in \mathcal{B}_{a,b}, \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'}) \leq (1 + \alpha_\nu(\varepsilon)) \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}).$$

We refer to Appendix A for the definitions of ε_ν and $\alpha_\nu(\cdot)$. Note that $\alpha_\nu(\cdot)$ is such that $\lim_{\varepsilon \rightarrow 0} \alpha_\nu(\varepsilon) = 0$.

Let $0 < c < 1$. We will show that almost surely $\lim_{T \rightarrow \infty} N_b(T) = +\infty$ implies

$$\liminf_{T \rightarrow \infty} \frac{1}{\log(N_b(T))} \sum_{b' \in \mathcal{B}_{a,b}} N_{a,b'}(T) \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'}) \geq c.$$

We start with the following inequality

$$\begin{aligned} & \mathbb{P}_\nu \left(\liminf_{T \rightarrow \infty} \frac{1}{\log(N_b(T))} \sum_{b' \in \mathcal{B}_{a,b}} N_{a,b'}(T) \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'}) < c, \lim_{T \rightarrow \infty} N_b(T) = \infty \right) \\ & \leq \liminf_{T \rightarrow \infty} \mathbb{P}_\nu \left(\frac{1}{\log(N_b(T))} \sum_{b' \in \mathcal{B}_{a,b}} N_{a,b'}(T) \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'}) < c, \lim_{T \rightarrow \infty} N_b(T) = \infty \right). \end{aligned}$$

Let us consider an horizon $T \geq 1$ and let us introduce the event

$$\Omega_T = \left\{ \sum_{b' \in \mathcal{B}_{a,b}} N_{a,b'}(T) \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'}) < c \log(N_b(T)), \lim_{T \rightarrow \infty} N_b(T) = \infty \right\}.$$

We want to provide an upper bound on $\mathbb{P}_\nu(\Omega_T)$ to ensure $\lim_{T \rightarrow \infty} \mathbb{P}_\nu(\Omega_T) = 0$. We start by taking advantage of the following lemma.

Lemma 13 (Change of measure) *For every event Ω and random variable Z both measurable with respect to ν and $\tilde{\nu}$,*

$$\mathbb{P}_\nu(\Omega \cap E) = \mathbb{E}_{\tilde{\nu}} \left[\frac{d\nu}{d\tilde{\nu}}(\psi) \mathbb{I}_{\{\Omega \cap E\}} \right] \leq \mathbb{E}_{\tilde{\nu}} [e^Z \mathbb{I}_{\{\Omega\}}]$$

where $E = \left\{ \log \left(\frac{d\nu}{d\tilde{\nu}}(\psi) \right) \leq Z \right\}$ and $\psi = ((a_t, b_t), X_t)_{t=1..T}$ is the sequence of pulled couples and rewards.

Let $\alpha \in (0, 1)$ and let us introduce the event

$$E_T = \left\{ \log \left(\frac{d\nu}{d\tilde{\nu}}(\psi) \right) \leq (1 - \alpha) \log(N_b(T)) \right\}.$$

Then we can decompose the probability as follows

$$\mathbb{P}_\nu(\Omega_T) = \mathbb{P}_\nu(\Omega_T \cap E_T) + \mathbb{P}_\nu(\Omega_T \cap E_T^c) \leq \mathbb{E}_{\tilde{\nu}} [N_b(T)^{1-\alpha} \mathbb{I}_{\{\Omega_T\}}] + \mathbb{P}_\nu(\Omega_T \cap E_T^c)$$

Now, we control successively $\mathbb{E}_{\tilde{\nu}} [N_b(T)^{1-\alpha} \mathbb{I}_{\{\Omega_T\}}]$ and $\mathbb{P}_\nu(\Omega_T \cap E_T^c)$ and show that they both tend to 0 as T tends to ∞ .

B.1.1 $\mathbb{E}_{\tilde{\nu}} [N_b(T)^{1-\alpha} \mathbb{I}_{\{\Omega_T\}}]$ TENDS TO 0 WHEN T TENDS TO INFINITY

We first provide an upper bound on $\mathbb{I}_{\{\Omega_T\}}$ as follows, denoting $c' = c/\text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'})$,

$$\begin{aligned} \Omega_T &\subset \left\{ N_{a,b}(T) < c' \log(N_b(T)), \lim_{T \rightarrow \infty} N_b(T) = \infty \right\} \\ &= \left\{ N_b(T) < c' \log(N_b(T)) + \sum_{a' \neq a} N_{a',b}(T), \lim_{T \rightarrow \infty} N_b(T) = \infty \right\}. \end{aligned}$$

Thus, we have

$$\mathbb{I}_{\{\Omega_T\}} \leq c' \mathbb{I}_{\{N_b(T) \geq 1, N_b(T) \rightarrow \infty\}} \frac{\log(N_b(T))}{N_b(T)} + \sum_{a' \neq a} \frac{N_{a',b}(T)}{N_b(T)}$$

Considering $f_\alpha : x \geq 1 \mapsto \log(x)/x^\alpha$, we have $f_\alpha \leq e^{-1}/\alpha$. Then, the dominated convergence theorem ensures

$$\mathbb{E}_{\tilde{\nu}} \left[\mathbb{I}_{\{N_b(T) \geq 1, N_b(T) \rightarrow \infty\}} \frac{\log(N_b(T))}{N_b(T)^\alpha} \right] = o(1).$$

Furthermore, since the considered strategy is assumed consistent we know that for $a' \neq a$, since a' is a sub-optimal arm for user b and configuration $\tilde{\nu}$,

$$\mathbb{E}_{\tilde{\nu}} \left[\frac{N_{a,b'}(T)}{N_b(T)^\alpha} \right] = o(1),$$

therefore we get

$$\mathbb{E}_{\tilde{\nu}} [N_b(T)^{1-\alpha} \mathbb{I}_{\{\Omega_T\}}] = o(1).$$

B.1.2 $\mathbb{P}_\nu(\Omega_T \cap E_T^c)$ TENDS TO 0 WHEN T TENDS TO INFINITY

For each time $t = 1, \dots, T$, the reward X_t is sampled independently from the past and according to ν_{a_t, b_t} . Hence the likelihood ratio rewrites

$$\frac{d\nu}{d\tilde{\nu}}(\psi) = \prod_{t=1}^T \frac{d\nu_{a_t, b_t}}{d\tilde{\nu}_{a_t, b_t}}(X_t)$$

where, for all $(a, b) \in \mathcal{A} \times \mathcal{B}$ and for all $x \in \{0, 1\}$, we have : $\frac{d\nu_{a,b}}{d\tilde{\nu}_{a,b}}(x) = \frac{\mu_{a,b}^x (1 - \mu_{a,b})^{1-x}}{\tilde{\mu}_{a,b}^x (1 - \tilde{\mu}_{a,b})^{1-x}}$.

Thus, since for all $b' \notin \mathcal{B}_{a,b}$, $\mu_{a,b} = \tilde{\mu}_{a,b}$, the log-likelihood ratio is

$$\log \left(\frac{d\nu}{d\tilde{\nu}}(\psi) \right) = \sum_{b' \in \mathcal{B}_{a,b}} \sum_{t=1}^T \mathbb{I}_{\{(a_t, b_t) = (a, b')\}} \log \left(\frac{d\nu_{a,b'}}{d\tilde{\nu}_{a,b'}}(X_t) \right).$$

Let us introduce, for $(a, b) \in \mathcal{A} \times \mathcal{B}$, $X_{a,b}^n = X_{\tau_{a,b}^n}$ where $\tau_{a,b}^n = \min \{t \geq 1 \text{ s.t. } N_{a,b}(t) = n\}$. Note that the random variables $\tau_{a,b}^n$ are predictable stopping times, since $\{\tau_{a,b}^n = t\}$ is measurable with respect to the filtration generated by $((a_1, b_1), X_1, \dots, (a_{t-1}, b_{t-1}), X_{t-1})$. Hence we can rewrite the event E_T

$$E_T = \left\{ \sum_{b' \in \mathcal{B}_{a,b}} \sum_{t=1}^T \mathbb{I}_{\{(a_t, b_t) = (a, b')\}} \frac{d\nu_{a,b'}}{d\tilde{\nu}_{a,b'}}(X_t) \leq (1 - \alpha) \log(N_b(T)) \right\}$$

and, since $\Omega_T = \left\{ \sum_{b' \in \mathcal{B}_{a,b}} N_{a,b'}(T) \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'}) < c \log(N_b(T)), \lim_{T \rightarrow \infty} N_b(T) = \infty \right\}$, we have

$$\Omega_T \cap E_T^c \subset \left\{ \begin{aligned} &\exists (n_{b'})_{b' \in \mathcal{B}_{a,b}} : \sum_{b' \in \mathcal{B}_{a,b}} n_{b'} \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'}) < c \log(N_b(T)), \lim_{T \rightarrow \infty} N_b(T) = \infty \\ &\text{and } \sum_{b' \in \mathcal{B}_{a,b}} \sum_{n=1..n_{b'}} \frac{d\nu_{a,b'}}{d\tilde{\nu}_{a,b'}}(X_{a,b'}^n) > (1 - \alpha) \log(N_b(T)) \end{aligned} \right\}.$$

For $b' \in \mathcal{B}_{a,b}$ and $n \geq 1$, let us consider $Z_{b'}^n = \frac{d\nu_{a,b'}}{d\tilde{\nu}_{a,b'}}(X_{a,b'}^n)$. Then $Z_{b'}^n$ is positive and bounded by $B_{b'} = \frac{1}{\tilde{\mu}_{a,b'}(1 - \tilde{\mu}_{a,b'})}$, with mean $\mathbb{E}_\nu[Z_{b'}^n] = \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'})$. Furthermore, the random variables $Z_{b'}^n$, for $b' \in \mathcal{B}_{a,b}$ and $n \geq 1$, are independent. Thus, it holds

$$\Omega_T \cap E_T^c \subset \left\{ \max_{(n'_b) \in \mathcal{N}_{a,b}} \sum_{b' \in \mathcal{B}_{a,b}} \sum_{n=1..n_{b'}} Z_{b'}^n - \mathbb{E}_\nu[Z_{b'}^n] > \left(\frac{1 - \alpha}{c} - 1 \right) c \log(N_b(T)), N_b(T) \rightarrow \infty \right\},$$

where

$$\mathcal{N}_{a,b} := \left\{ (n_{b'})_{b' \in \mathcal{B}_{a,b}} : \sum_{b' \in \mathcal{B}_{a,b}} n_{b'} \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'}) < c \log(T) \right\}.$$

In the following, we apply Doob's maximal inequality. For $b' \in \mathcal{B}_{a,b}$ and $\lambda > 0$, let us introduce the super-martingale

$$M_n^{b'} = \exp\left(\lambda \sum_{k=1}^n (Z_{b'}^k - E[Z_{b'}^k]) - n\lambda^2 \frac{B_{b'}^2}{8}\right).$$

Then noting that

$$\frac{\sum_{b' \in \mathcal{B}_{a,b}} \lambda^2 n_{b'} \frac{B_{b'}^2}{8}}{c \log(T)} < \frac{\sum_{b' \in \mathcal{B}_{a,b}} \lambda^2 n_{b'} \frac{B_{b'}^2}{8}}{\sum_{b' \in \mathcal{B}_{a,b}} n_{b'} \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'})} \leq \lambda^2 \frac{\max_{b' \in \mathcal{B}_{a,b}} B_{b'}^2}{8 \min_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'})},$$

we obtain

$$\begin{aligned} \Omega_T \cap E_T^c &\subset \left\{ \max_{(n'_b) \in \mathcal{N}_{a,b}} \prod_{b' \in \mathcal{B}_{a,b}} M_{n'_b}^{b'} > T^{\left[\lambda \left(\frac{1-\alpha}{c} - 1 \right) - \lambda^2 \frac{\max_{b' \in \mathcal{B}_{a,b}} B_{b'}^2}{8 \min_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'})} \right] c}, N_b(T) \rightarrow \infty \right\} \\ &\subset \left\{ \exists b' \in \mathcal{B}_{a,b} : \max_{n \leq n_{\max}} M_n^{b'} > N_b(T)^\gamma, N_b(T) \rightarrow \infty \right\}, \end{aligned}$$

where $n_{\max} = \frac{c \log(T)}{\min_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'})}$ and $\gamma = \left[\lambda \left(\frac{1-\alpha}{c} - 1 \right) - \lambda^2 \frac{\max_{b' \in \mathcal{B}_{a,b}} B_{b'}^2}{8 \min_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'})} \right] \frac{c}{|\mathcal{B}_{a,b}|}$. In

order to have $\gamma > 0$, we impose:

- $0 < \alpha < 1 - c$ (this implies $\frac{1-\alpha}{c} - 1 > 0$)
- $\lambda \in \operatorname{argmax}_{\lambda' \geq 0} \left\{ \lambda' \left(\frac{1-\alpha}{c} - 1 \right) - \lambda'^2 \frac{\max_{b' \in \mathcal{B}_{a,b}} B_{b'}^2}{8 \min_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \tilde{\mu}_{a,b'})} \right\} > 0$.

Thus for $A > 0$, we have

$$\begin{aligned} \mathbb{P}_\nu(\Omega_T \cap E_T^c) &\leq \sum_{b' \in \mathcal{B}_{a,b}} \mathbb{P}_\nu \left(\max_{n \leq n_{\max}} M_n^{b'} > N_b(T)^\gamma, N_b(T) \rightarrow \infty \right) \quad (\text{Union bound}) \\ &\leq \sum_{b' \in \mathcal{B}_{a,b}} \mathbb{P}_\nu(N_b(T)^\gamma \leq A, N_b(T) \rightarrow \infty) + \mathbb{P}_\nu \left(\max_{n \leq n_{\max}} M_n^{b'} > A \right) \\ &\leq |\mathcal{B}_{a,b}| \mathbb{P}_\nu(N_b(T)^\gamma \leq A, N_b(T) \rightarrow \infty) + \sum_{b' \in \mathcal{B}_{a,b}} \frac{\mathbb{E}_\nu[M_0^{b'}]}{A} \quad (\text{Doob's maximal inequality}) \\ &= |\mathcal{B}_{a,b}| \mathbb{P}_\nu(N_b(T)^\gamma \leq A, N_b(T) \rightarrow \infty) + \frac{|\mathcal{B}_{a,b}|}{A}. \end{aligned}$$

Furthermore, we have

$$\begin{aligned} \lim_{T \rightarrow \infty} \mathbb{P}_{\tilde{\nu}}(N_b(T)^\gamma \leq A, N_b(T) \rightarrow \infty) &\leq \mathbb{P}_{\tilde{\nu}} \left(\limsup_{T \rightarrow \infty} (N_b(T)^\gamma \leq A), N_b(T) \rightarrow \infty \right) \\ &\leq \mathbb{P}_{\tilde{\nu}} \left(\limsup_{T \rightarrow \infty} N_b(T) < \infty, N_b(T) \rightarrow \infty \right) = 0. \end{aligned}$$

Thus we have shown

$$\forall A > 0, \quad \limsup_{T \rightarrow \infty} \mathbb{P}_\nu(\Omega_T \cap E_T^c) \leq \frac{|\mathcal{B}_{a,b}|}{A},$$

that is $\mathbb{P}_\nu(\Omega_T \cap E_T^c) = o(1)$. ■

B.2 Asymptotic lower bounds on the regret

Here, we explain how we obtain the lower bounds on the regret given in Corollary 8.

Proof [Proof of Corollary 8.] Let us consider a consistent strategy on \mathcal{D}_ω and let $\nu \in \mathcal{D}_\omega$. Let $(T_k)_{k \in \mathbb{N}}$ be a sub-sequence such that

$$\liminf_{T \rightarrow \infty} \frac{R(T, \nu)}{\log(T)} = \lim_{k \rightarrow \infty} \frac{R(T_k, \nu)}{\log(T_k)}.$$

We assume that this limit is finite otherwise the result is straightforward. This implies in particular that for all $(a, b) \notin \mathcal{O}^*$

$$\limsup_{k \rightarrow \infty} \frac{\mathbb{E}_\nu [N_{a,b}(T_k)]}{\log(T_k)} < +\infty.$$

By Cantor's diagonal argument there exists an extraction of $(T_k)_{k \in \mathbb{N}}$ denoted by $(T'_k)_{k \in \mathbb{N}}$ such that for all $(a, b) \notin \mathcal{O}^*$, there exist $N_{a,b} \notin \mathcal{O}^*$ such that

$$\lim_{k' \rightarrow \infty} \frac{\mathbb{E}_\nu [N_{a,b}(T'_{k'})]}{\log(T'_{k'})} = N_{a,b}.$$

Hence we get

$$\liminf_{T \rightarrow \infty} \frac{R(T, \nu)}{\log(T)} = \sum_{(a,b) \notin \mathcal{O}^*} N_{a,b} \Delta_{a,b}.$$

But thanks to Proposition 5 we have for all $(a, b) \notin \mathcal{O}^*$, since user b has a log-frequency β_b ,

$$\begin{aligned} & \sum_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) N_{a,b'} \\ &= \lim_{k \rightarrow \infty} \sum_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) \frac{\mathbb{E}_\nu [N_{a,b'}(T'_k)]}{\log(T'_k)} \\ &\geq \liminf_{k \rightarrow \infty} \sum_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) \frac{\mathbb{E}_\nu [N_{a,b'}(T'_k)]}{\log(N_b(T'_k))} \times \liminf_{k \rightarrow \infty} \frac{\log(N_b(T'_k))}{\log(T'_k)} \geq \beta_b. \end{aligned}$$

Therefore we obtain the lower bound

$$\begin{aligned} \liminf_{T \rightarrow \infty} \frac{R(\nu, T)}{\log(T)} &\geq C_\omega^*(\beta, \nu) := \min_{n \in \mathbb{R}_+^{A \times B}} \sum_{a,b \notin \mathcal{O}^*} n_{a,b} \Delta_{a,b} \\ &\quad \text{s.t.} \quad \forall (a, b) \notin \mathcal{O}^* : \sum_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) n_{a,b'} \geq \beta_b. \end{aligned}$$
■

Appendix C. Algorithms for the *uncontrolled scenario*

We regroup in this section the algorithms IMED-GS₂ and IMED-GS₂^{*} for the *uncontrolled scenario*.

Algorithm 3 IMED-GS ₂	Algorithm 4 IMED-GS ₂ [*]
<p>Require: Weight matrix $(\omega_{b,b'})_{b,b' \in \mathcal{B}}$.</p> <p>for $t = 1 \dots T$ do</p> <p style="padding-left: 2em;">Pull $a_{t+1} \in \operatorname{argmin}_{a \in \mathcal{A}} \tilde{I}_{a,b_{t+1}}(t)$</p> <p>end for</p>	<p>Require: Weight matrix $(\omega_{b,b'})_{b,b' \in \mathcal{B}}$.</p> <p>$\forall a \in \mathcal{A}, c_a, c_a^+ \leftarrow 1$</p> <p>$\forall b \in \mathcal{B}, E(b) \leftarrow \emptyset, FE(b) \leftarrow \emptyset$</p> <p>for For $t = 1 \dots T$ do</p> <p style="padding-left: 2em;">Choose $\bar{a}_t \in \operatorname{argmin}_{a \in \mathcal{A}} \tilde{I}_{a,b_{t+1}}(t)$</p> <p>if $(\bar{a}_t, b_{t+1}) \in \hat{\mathcal{O}}^*(t)$ then</p> <p style="padding-left: 2em;">Choose $a_{t+1} = \bar{a}_t$</p> <p>else</p> <p style="padding-left: 2em;">Choose $(\underline{a}_t, \underline{b}_t) \in \operatorname{argmin}_{(a,b) \in \mathcal{A} \times \mathcal{B}} \tilde{I}_{a,b}(t)$</p> <p>if $(\underline{a}_t, \underline{b}_t) \notin \hat{\mathcal{O}}^*(t)$ then</p> <p style="padding-left: 2em;">if $c_{\underline{a}_t} = c_{\underline{a}_t}^+$ then</p> <p style="padding-left: 4em;">$c_{\underline{a}_t}^+ \leftarrow 2c_{\underline{a}_t}^+$</p> <p style="padding-left: 2em;">Choose $\bar{b}_t \in \operatorname{argmin}_{b \in \mathcal{B}} N_{a,b}(t)$</p> <p style="padding-left: 2em;">$FE(\bar{b}_t) \leftarrow \underline{a}_t$</p> <p>else</p> <p style="padding-left: 2em;">Choose $\bar{b}_t \in \operatorname{argmax}_{b \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t) \cup \{\underline{b}_t\}} N_{a_{t+1}, b}^{\text{opt}}(t) - N_{a_{t+1}, b}(t)$</p> <p style="padding-left: 2em;">$E(\bar{b}_t) \leftarrow \underline{a}_t$</p> <p>end if</p> <p style="padding-left: 2em;">$c_{a_{t+1}} \leftarrow c_{a_{t+1}} + 1$</p> <p>end if</p> <p>Priority rule in exploration phases:</p> <p>if $FE(b_{t+1}) \neq \emptyset$ then</p> <p style="padding-left: 2em;">Choose $a_{t+1} = FE(b_{t+1})$ (delayed forced-exploration)</p> <p style="padding-left: 2em;">$FE(b_{t+1}) \leftarrow \emptyset$</p> <p>else if $E(b_{t+1}) \neq \emptyset$ then</p> <p style="padding-left: 2em;">Choose $a_{t+1} = E(b_{t+1})$ (delayed exploration)</p> <p style="padding-left: 2em;">$E(b_{t+1}) \leftarrow \emptyset$</p> <p>else</p> <p style="padding-left: 2em;">Choose $a_{t+1} = \bar{a}_t$ (current exploration)</p> <p>end if</p> <p>end if</p> <p style="padding-left: 2em;">Pull a_{t+1}</p> <p>end for</p>

Appendix D. IMED-GS*: Finite-time analysis

IMED-GS* strategy implies empirical lower and empirical upper bounds on the numbers of pulls (Lemma 14, Lemma 15). Based on concentration lemmas (see Appendix D.7), the strategy-based empirical lower bounds ensure the reliability of the estimators of interest (Lemma 19). Then, combining the reliability of these estimators with the obtained strategy-base empirical upper bounds, we get upper bounds on the average numbers of pulls (Proposition 20). We first show that IMED-GS* strategy is Pareto-optimal (for minimization problem 2) and that it is a consistent strategy which induces sequences of users with log-frequencies all equal to 1 (independently from the considered bandit configuration). From an asymptotic analysis, we then prove that IMED-GS* strategy is asymptotically optimal.

D.1 Strategy-based empirical bounds

IMED-GS* strategy implies inequalities between the indexes that can be rewritten as inequalities on the numbers of pulls. While asymptotic analysis suggests lower bounds involving $\log(N_{b_{t+1}}(t))$ might be expected, we show in this non-asymptotic context lower bounds on the numbers of pulls involving instead the logarithm of the number of pulls of the current chosen arm, $\log(N_{a_{t+1}, b_{t+1}}(t))$. In contrast, we provide upper bounds involving $\log(N_{b_{t+1}}(t))$ on $N_{a_{t+1}, b_{t+1}}(t)$.

We believe that establishing these empirical lower and upper bounds is a key element of our proof technique, that is of independent interest and not *a priori* restricted to the graph structure.

Lemma 14 (Empirical lower bounds) *Under IMED-GS*, at each step time $t \geq 1$, for all couple $(a, b) \notin \hat{\mathcal{O}}^*(t)$,*

$$\log(N_{a_{t+1}, b_{t+1}}(t)) \leq \sum_{b' \in \hat{\mathcal{B}}_{a, b}(t)} N_{a, b'}(t) \text{kl}(\hat{\mu}_{a, b'}(t) | \hat{\mu}_b^*(t) - \omega_{b, b'}) + \log(N_{a, b'}(t)) .$$

Furthermore, for all couple $(a, b) \in \hat{\mathcal{O}}^*(t)$,

$$N_{a_{t+1}, b_{t+1}}(t) \leq N_{a, b}(t) .$$

Proof According to IMED-GS* strategy (see Algorithm 2), $a_{t+1} = \underline{a}_t$ and for all couple $(a, b) \in \mathcal{A} \times \mathcal{B}$

$$I_{a, b}(t) \geq I_{\underline{a}_t, \underline{b}_t}(t) .$$

There is three possible cases.

Case 1: $(a_{t+1}, b_{t+1}) = (\underline{a}_t, \underline{b}_t) \in \hat{\mathcal{O}}^*(t)$ and $I_{\underline{a}_t, \underline{b}_t}(t) = \log(N_{a_{t+1}, b_{t+1}}(t))$.

Case 2: $b_{t+1} \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t) \cup \{\underline{b}_t\}$ and

$$I_{\underline{a}_t, \underline{b}_t}(t) = \sum_{b' \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t)} N_{\underline{a}_t, b'}(t) \text{kl}(\hat{\mu}_{\underline{a}_t, b'}(t) | \hat{\mu}_{\underline{b}_t}^*(t) - \omega_{\underline{b}_t, b'}) + \log(N_{\underline{a}_t, b'}(t)) .$$

Note that $\underline{b}_t \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t)$

except if $N_{\underline{a}_t, \underline{b}_t}(t) = 0$. Thus $I_{\underline{a}_t, \underline{b}_t}(t) \geq \log(N_{a_{t+1}, b_{t+1}}(t))$.

Case 3: $b_{t+1} \in \underset{b \in \mathcal{B}}{\text{argmin}} N_{\underline{a}_t, b}(t)$ and $I_{\underline{a}_t, \underline{b}_t}(t) \geq \min_{b \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t)} \log(N_{a_{t+1}, b}(t)) \geq \log(N_{a_{t+1}, b_{t+1}}(t))$.

This implies for all couple $(a, b) \in \mathcal{A} \times \mathcal{B}$,

$$I_{a, b}(t) \geq \log(N_{a_{t+1}, b_{t+1}}(t)) .$$

Thus, according to the definition of the indexes (Eq. 4) for all couple $(a, b) \notin \hat{\mathcal{O}}^*(t)$ we obtain

$$\log(N_{a_{t+1}, b_{t+1}}(t)) \leq \sum_{b' \in \hat{\mathcal{B}}_{a,b}(t)} N_{a,b'}(t) \text{kl}(\hat{\mu}_{a,b'}(t) | \hat{\mu}_b^*(t) - \omega_{b,b'}) + \log(N_{a,b'}(t)) ,$$

and for all couple $(a, b) \in \hat{\mathcal{O}}^*(t)$,

$$\log(N_{a_{t+1}, b_{t+1}}(t)) \leq \log(N_{a,b}(t)) .$$

Taking the exponential in the last inequality allows us to conclude. ■

Lemma 15 (Empirical upper bounds) *Under IMED-GS*, at each step time $t \geq 1$ such that $(a_{t+1}, b_{t+1}) \notin \hat{\mathcal{O}}^*(t)$ we have*

$$\sum_{b' \in \hat{\mathcal{B}}_{a_{t+1}, b_t}(t)} N_{a_{t+1}, b'}(t) \text{kl}(\hat{\mu}_{a_{t+1}, b'}(t) | \hat{\mu}_{b_t}^*(t) - \omega_{b_t, b'}) \leq \log(N_{b_t}(t))$$

and

$$\frac{N_{a_{t+1}, b_{t+1}}(t)}{\log(t)} \leq \begin{cases} \frac{1}{\text{kl}(\hat{\mu}_{a_{t+1}, b_t}(t) | \hat{\mu}_{b_t}^*(t))} , & \text{if } c_{a_{t+1}} = c_{a_{t+1}}^+ \\ \min \left(\frac{1}{\text{kl}(\hat{\mu}_{a_{t+1}, b_{t+1}}(t) | \hat{\mu}_{b_t}^*(t) - \omega_{b_t, b_{t+1}})} , n_{a_{t+1}, b_{t+1}}^{\text{opt}}(t) \right) & \text{otherwise.} \end{cases}$$

Proof For all current optimal couple $(a, b) \in \hat{\mathcal{O}}^*(t)$, we have $I_{a,b}(t) = \log(N_{a,b}(t)) \leq \log(N_b(t))$. This implies

$$I_{a_t, b_t}(t) \leq \log(N_{b_t}(t)) .$$

Furthermore, since $(a_{t+1}, b_{t+1}) \notin \hat{\mathcal{O}}^*(t)$, we have $a_{t+1} = a_t$ and the previous inequality implies

$$\sum_{b' \in \hat{\mathcal{B}}_{a_{t+1}, b_t}(t)} N_{a_{t+1}, b'}(t) \text{kl}(\hat{\mu}_{a_{t+1}, b'}(t) | \hat{\mu}_{b_t}^*(t) - \omega_{b_t, b'}) \leq \log(N_{b_t}(t)) . \quad (8)$$

In the following we study separately the two cases either $c_{a_{t+1}} = c_{a_{t+1}}^+$ or $c_{a_{t+1}} < c_{a_{t+1}}^+$.

Case 1: $c_{a_{t+1}} = c_{a_{t+1}}^+$

Then $b_{t+1} \in \underset{b \in \mathcal{B}}{\text{argmin}} N_{a_{t+1}, b}(t)$ and from Eq. 8 we get

$$N_{a_{t+1}, b_{t+1}}(t) \leq N_{a_{t+1}, b_t}(t) \leq \frac{\log(t)}{\text{kl}(\hat{\mu}_{a_{t+1}, b_t}(t) | \hat{\mu}_{b_t}^*(t))} .$$

Case 2: $c_{a_{t+1}} < c_{a_{t+1}}^+$

Then $b_{t+1} \in \underset{b \in \hat{\mathcal{B}}_{a_t, b_t}(t) \cup \{b_t\}}{\text{argmax}} N_{a_t, b}^{\text{opt}}(t) - N_{a_t, b}(t)$ and we have

$$N_{a_{t+1}, b_{t+1}}(t) \leq \frac{\log(t)}{\text{kl}(\hat{\mu}_{a_{t+1}, b_{t+1}}(t) | \hat{\mu}_{b_t}^*(t) - \omega_{b_t, b_{t+1}})} .$$

Last, Lemma 16 stated below implies

$$N_{a_{t+1}, b_{t+1}}(t) \leq N_{a_{t+1}, b_{t+1}}^{\text{opt}}(t) = n_{a_{t+1}, b_{t+1}}^{\text{opt}}(t) \min_{b \in \mathcal{B}} I_{a_{t+1}, b}(t) \leq n_{a_{t+1}, b_{t+1}}^{\text{opt}}(t) \log(t).$$

■

Lemma 16 (N^{opt} dominates N) *Under IMED-GS^{*}, at each time step $t \geq 1$ such that $(\underline{a}_t, \underline{b}_t) \notin \hat{\mathcal{O}}^*(t)$ we have*

$$\max_{b \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t) \cup \{\underline{b}_t\}} N_{\underline{a}_t, b}^{\text{opt}}(t) - N_{\underline{a}_t, b}(t) \geq 0.$$

Proof If $\underline{b}_t \notin \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t)$, then $N_{\underline{a}_t, \underline{b}_t}(t) = 0$ and $N_{\underline{a}_t, \underline{b}_t}^{\text{opt}}(t) - N_{\underline{a}_t, \underline{b}_t}(t) \geq 0$. In the following we assume that $\underline{b}_t \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t) \neq \emptyset$.

From Eq. 6, since $\min_{b' \in \mathcal{B}} I_{\underline{a}_t, b'}(t) = I_{\underline{a}_t, \underline{b}_t}(t)$, for all $b \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t)$ we have

$$N_{\underline{a}_t, b}^{\text{opt}}(t) = n_{\underline{a}_t, b}^{\text{opt}}(t) I_{\underline{a}_t, \underline{b}_t}(t), \quad (9)$$

and, since $(\underline{a}_t, \underline{b}_t) \notin \hat{\mathcal{O}}^*(t)$, from Eq. 5 we get

$$\sum_{b' \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t)} \text{kl}\left(\hat{\mu}_{\underline{a}_t, b'}(t) \middle| \hat{\mu}_{\underline{b}_t}^*(t) - \omega_{\underline{b}_t, b'}\right) n_{\underline{a}_t, b'}^{\text{opt}}(t) \geq 1. \quad (10)$$

Then Eq. 9 and 10 imply

$$\sum_{b' \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t)} N_{\underline{a}_t, b'}^{\text{opt}}(t) \text{kl}\left(\hat{\mu}_{\underline{a}_t, b'}(t) \middle| \hat{\mu}_{\underline{b}_t}^*(t) - \omega_{\underline{b}_t, b'}\right) \geq I_{\underline{a}_t, \underline{b}_t}(t).$$

Hence from the definitions of the indexes (Eq. 4) this implies

$$\begin{aligned} \sum_{b' \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t)} N_{\underline{a}_t, b'}^{\text{opt}}(t) \text{kl}\left(\hat{\mu}_{\underline{a}_t, b'}(t) \middle| \hat{\mu}_{\underline{b}_t}^*(t) - \omega_{\underline{b}_t, b'}\right) \\ \geq \sum_{b' \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t)} N_{\underline{a}_t, b'}(t) \text{kl}\left(\hat{\mu}_{\underline{a}_t, b'}(t) \middle| \hat{\mu}_{\underline{b}_t}^*(t) - \omega_{\underline{b}_t, b'}\right). \end{aligned} \quad (11)$$

Since we assume $\underline{b}_t \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t) \neq \emptyset$, previous Eq. 11 implies

$$\max_{b \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t) \cup \{\underline{b}_t\}} N_{\underline{a}_t, b}^{\text{opt}}(t) - N_{\underline{a}_t, b}(t) = \max_{b \in \hat{\mathcal{B}}_{\underline{a}_t, \underline{b}_t}(t)} N_{\underline{a}_t, b}^{\text{opt}}(t) - N_{\underline{a}_t, b}(t) \geq 0.$$

■

D.2 Reliable current best arm and means

In this subsection, we consider the subset $\mathcal{T}_{\varepsilon,\gamma}$ of times where everything behaves well, that is: The current best couples correspond to the true ones, and, the empirical means of the best couples and the couples at least pulled proportionally (with a coefficient $\gamma \in (0, 1)$) to the number of pulls of the current chosen couple are ε -accurate for $0 < \varepsilon < \varepsilon_\nu$, i.e.

$$\mathcal{T}_{\varepsilon,\gamma} := \left\{ t \geq 1 : \begin{array}{l} \hat{\mathcal{O}}^*(t) = \mathcal{O}^* \\ \forall (a,b) \text{ s.t. } N_{a,b}(t) \geq \gamma N_{a_{t+1},b_{t+1}}(t) \text{ or } (a,b) \in \mathcal{O}^*, |\hat{\mu}_{a,b}(t) - \mu_{a,b}| < \varepsilon \end{array} \right\}.$$

We will show that its complementary is finite on average. In order to prove this we decompose the set $\mathcal{T}_{\varepsilon,\gamma}$ in the following way. Let $\mathcal{E}_{\varepsilon,\gamma}$ be the set of times where the means are well estimated,

$$\mathcal{E}_{\varepsilon,\gamma} := \left\{ t \geq 1 : \forall (a,b) \text{ s.t. } N_{a,b}(t) \geq \gamma N_{a_{t+1},b_{t+1}}(t) \text{ or } (a,b) \in \hat{\mathcal{O}}^*(t), |\hat{\mu}_{a,b}(t) - \mu_{a,b}| < \varepsilon \right\},$$

and Λ_ε the set of times where the mean of a couple that is not the current optimal neither pulled is underestimated

$$\Lambda_\varepsilon := \left\{ t \geq 1 \mid \exists a \in \mathcal{A}, \mathcal{B}' \subset \mathcal{B} : \begin{array}{l} \forall b \in \mathcal{B}', \hat{\mu}_{a,b}(t) < \mu_{a,b} - \varepsilon \\ \log(N_{a_{t+1},b_{t+1}}(t)) \leq \sum_{b \in \mathcal{B}'} N_{a,b}(t) \text{kl}(\hat{\mu}_{a,b}(t) | \mu_{a,b} - \varepsilon_\nu) + \log(N_{a,b}(t)) \end{array} \right\}.$$

Then we prove below the following inclusion.

Lemma 17 (Relations between the subsets of times) For $0 < \varepsilon < \varepsilon_\nu$ and $\gamma \in (0, 1)$,

$$\mathcal{T}_{\varepsilon,\gamma}^c \setminus \mathcal{E}_{\varepsilon,\gamma}^c \subset \Lambda_{\varepsilon_\nu}. \quad (12)$$

Refer to Appendix A for the definition of ε_ν .

Proof For all user $b \in \mathcal{B}$ it is assumed that there exists a unique optimal arm $a_b^* \in \mathcal{A}$ such that $(a_b^*, b) \in \mathcal{O}^*$. We have $\mathcal{O}^* = \bigcup_{b \in \mathcal{B}} \{(a_b^*, b)\}$. In particular, for all time step $t \geq 1$, if $\hat{\mathcal{O}}^*(t) \neq \mathcal{O}^*$ then there exists $b \in \mathcal{B}$ and $\hat{a}_b^* \in \mathcal{A}$ such that $(\hat{a}_b^*, b) \in \hat{\mathcal{O}}^*(t) \setminus \mathcal{O}^*$ (and $\hat{a}_b^* \neq a_b^*$).

Let $t \in \mathcal{T}_{\varepsilon,\gamma}^c \setminus \mathcal{E}_{\varepsilon,\gamma}^c$. Then $\hat{\mathcal{O}}^*(t) \neq \mathcal{O}^*$ and there exists $b \in \mathcal{B}$, $\hat{a}_b^* \in \mathcal{A}$ such that $(\hat{a}_b^*, b) \in \hat{\mathcal{O}}^*(t) \setminus \mathcal{O}^*$. Thus we know that $\hat{a}_b^* \neq a_b^*$. In particular, we have $\mu_{a_b^*,b} = \mu_b^* > \mu_{\hat{a}_b^*,b} + 2\varepsilon_\nu$. Since $t \in \mathcal{E}_{\varepsilon,\gamma}$ and $\varepsilon < \varepsilon_\nu$, this implies

$$\mu_{a_b^*,b} > \hat{\mu}_{\hat{a}_b^*,b}(t) + \varepsilon_\nu = \hat{\mu}_b^*(t) + \varepsilon_\nu \geq \hat{\mu}_{a_b^*}(t) + \varepsilon \quad (13)$$

and $(a_b^*, b) \notin \hat{\mathcal{O}}^*(t)$. From Lemma 14 we have the following empirical lower bound

$$\log(N_{a_{t+1},b_{t+1}}(t)) \leq \sum_{b' \in \hat{\mathcal{B}}_{a_b^*,b}(t)} N_{a_b^*,b'}(t) \text{kl}(\hat{\mu}_{a_b^*,b'}(t) | \hat{\mu}_b^*(t) - \omega_{b,b'}) + \log(N_{a_b^*,b'}(t)). \quad (14)$$

In particular, for all $b' \in \hat{\mathcal{B}}_{a_b^*,b}(t)$ we have $\hat{\mu}_{a_b^*,b'}(t) < \hat{\mu}_b^*(t) - \omega_{b,b'}$ and Eq. 13 implies

$$\hat{\mu}_{a_b^*,b'}(t) < \hat{\mu}_b^*(t) - \omega_{b,b'} < \mu_{a_b^*,b} - \varepsilon_\nu - \omega_{b,b'} < \mu_{a_b^*,b'} - \varepsilon_\nu, \quad (15)$$

and the monotonic properties of $\text{kl}(\cdot|\cdot)$ implies

$$\text{kl}\left(\hat{\mu}_{a_b^*,b'}(t) \middle| \hat{\mu}_b^*(t) - \omega_{b,b'}\right) \leq \text{kl}\left(\hat{\mu}_{a_b^*,b'}(t) \middle| \mu_{a_b^*,b'} - \varepsilon_\nu\right). \quad (16)$$

Therefore, by combining Eq. 14, 15 and 16, we have for such t

$$\forall b' \in \widehat{\mathcal{B}}_{a_b^*,b}(t), \hat{\mu}_{a_b^*,b'}(t) < \mu_{a_b^*,b'} - \varepsilon_\nu$$

and

$$\log(N_{a_{t+1},b_{t+1}}(t)) \leq \sum_{b' \in \widehat{\mathcal{B}}_{a_b^*,b}(t)} N_{a_b^*,b'}(t) \text{kl}\left(\hat{\mu}_{a_b^*,b'}(t) \middle| \mu_{a_b^*,b'} - \varepsilon_\nu\right) + \log(N_{a_b^*,b'}(t)),$$

which concludes the proof. ■

Using classical concentration arguments we prove in Appendix D.8 the following upper bounds.

Lemma 18 (Bounded subsets of times) For $0 < \varepsilon < \varepsilon_\nu$ and $\gamma \in (0, 1/2)$,

$$\mathbb{E}_\nu[|\mathcal{E}_{\varepsilon,\gamma}^c|] \leq \frac{17}{\gamma\varepsilon^4} |\mathcal{A}|^2 |\mathcal{B}|^2 \quad \mathbb{E}_\nu[|\Lambda_{\varepsilon_\nu}|] \leq 2 |\mathcal{A}|^2 |\mathcal{B}| (1 + E_\nu)^{|\mathcal{B}|}.$$

Refer to Appendix A for the definitions of ε_ν and E_ν .

Thus combining them with (12) we obtain

$$\mathbb{E}_\nu[|\mathcal{T}_{\varepsilon,\gamma}^c|] \leq \mathbb{E}_\nu[|\mathcal{E}_{\varepsilon,\gamma}^c|] + \mathbb{E}_\nu[|\Lambda_{\varepsilon_\nu}|] \leq \frac{17}{\gamma\varepsilon^4} |\mathcal{A}|^2 |\mathcal{B}|^2 + 2 |\mathcal{A}|^2 |\mathcal{B}| (1 + E_\nu)^{|\mathcal{B}|}.$$

Hence, we just proved the following lemma.

Lemma 19 (Reliable estimators) For $0 < \varepsilon < \varepsilon_\nu$ and $\gamma \in (0, 1/2)$,

$$\mathbb{E}_\nu[|\mathcal{T}_{\varepsilon,\gamma}^c|] \leq \frac{17}{\gamma\varepsilon^4} |\mathcal{A}|^2 |\mathcal{B}|^2 + 2 |\mathcal{A}|^2 |\mathcal{B}| (1 + E_\nu)^{|\mathcal{B}|}.$$

Refer to Appendix A for the definitions of ε_ν and E_ν .

D.3 Pareto-optimality and upper bounds on the numbers of pulls of sub-optimal arms

In this section, we combine the different results of the previous sections to prove the following proposition.

Proposition 20 (Upper bounds) Let $\nu \in \mathcal{D}_\omega$. Let $0 < \varepsilon < \varepsilon_\nu$ and $\gamma \in (0, 1/2)$. Let us consider

$$\mathcal{T}_{\varepsilon,\gamma} := \left\{ t \geq 1 : \begin{array}{l} \hat{\mathcal{O}}^*(t) = \mathcal{O}^* \\ \forall (a,b) \text{ s.t. } N_{a,b}(t) \geq \gamma N_{a_{t+1},b_{t+1}}(t) \text{ or } (a,b) \in \mathcal{O}^*, |\hat{\mu}_{a,b}(t) - \mu_{a,b}| < \varepsilon \end{array} \right\}.$$

Then under IMED-GS* strategy,

$$\mathbb{E}_\nu[|\mathcal{T}_{\varepsilon,\gamma}^c|] \leq \frac{17}{\gamma\varepsilon^4} |\mathcal{A}|^2 |\mathcal{B}|^2 + 2 |\mathcal{A}|^2 |\mathcal{B}| (1 + E_\nu)^{|\mathcal{B}|}$$

and for all horizon time $T \geq 1$,

$$\forall a \in \mathcal{A}, \min_{b: (a,b) \notin \mathcal{O}^*} \frac{1}{\log(N_b(T))} \sum_{b' \in \mathcal{B}_{a,b}} N_{a,b'}(T) \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) \leq (1 + \alpha_\nu(\varepsilon)) \left[1 + \gamma \frac{M_\nu}{m_\nu} \right] + \frac{M_\nu |\mathcal{T}_{\varepsilon,\gamma}|}{\min_{b \in \mathcal{B}} \log(N_b(T))}$$

where m_ν and M_ν are defined as follows:

$$m_\nu = \min_{\substack{(a,b) \notin \mathcal{O}^* \\ b' \in \mathcal{B}_{a,b}}} \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}), \quad M_\nu = \max_{(a,b) \notin \mathcal{O}^*} \sum_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}).$$

Furthermore, we have

$$\forall (a,b) \notin \mathcal{O}^*, \quad N_{a,b}(T) \leq \frac{1 + \alpha_\nu(\varepsilon)}{m_\nu} \log(N_b(T)) + |\mathcal{T}_{\varepsilon,\gamma}^c|.$$

Refer to Appendix A for the definitions of ε_ν , $\alpha_\nu(\cdot)$ and E_ν .

Proof From Lemma 19, we have:

$$\mathbb{E}_\nu[|\mathcal{T}_{\varepsilon,\gamma}^c|] \leq \frac{17}{\gamma \varepsilon^4} |\mathcal{A}|^2 |\mathcal{B}|^2 + 2 |\mathcal{A}|^2 |\mathcal{B}| (1 + E_\nu)^{|\mathcal{B}|}.$$

Let $a \in \mathcal{A}$. Let us consider $1 \leq t \leq T$ such that $a_{t+1} = a$, $(a_{t+1}, b_{t+1}) \notin \mathcal{O}^*$ and $t \in \mathcal{T}_{\varepsilon,\gamma}$. Then, according to IMED-GS* strategy (see Algorithm 2), we have

$$(a, \underline{b}_t) = (\underline{a}_t, \underline{b}_t) \notin \hat{\mathcal{O}}^*(t).$$

From Lemma 15 this implies

$$\sum_{b' \in \hat{\mathcal{B}}_{a, \underline{b}_t}(t)} N_{a,b'}(t) \text{kl}(\hat{\mu}_{a,b'}(t) | \hat{\mu}_{\underline{b}_t}^*(t) - \omega_{\underline{b}_t, b'}) \leq \log(N_b(T)). \quad (17)$$

Since $t \in \mathcal{T}_{\varepsilon,\gamma}$ and $\varepsilon < \varepsilon_\nu$, we have

$$\{b' \in \mathcal{B}_{a, \underline{b}_t} : N_{a,b'}(t) \geq \gamma N_{a_{t+1}, b_{t+1}}(t)\} \subset \hat{\mathcal{B}}_{a, \underline{b}_t}(t). \quad (18)$$

Combining inequality (17) with inclusion (18), it comes

$$\sum_{b' \in \mathcal{B}_{a, \underline{b}_t} : N_{a,b'}(t) \geq \gamma N_{a_{t+1}, b_{t+1}}(t)} N_{a,b'}(t) \text{kl}(\hat{\mu}_{a,b'}(t) | \hat{\mu}_{\underline{b}_t}^*(t) - \omega_{\underline{b}_t, b'}) \leq \log(N_b(T)). \quad (19)$$

Since $t \in \mathcal{T}_{\varepsilon,\gamma}$, we have

$$\left| \hat{\mu}_{\underline{b}_t}^*(t) - \mu_{\underline{b}_t}^* \right| < \varepsilon \quad \text{and} \quad \forall b' \in \mathcal{B}_{a, \underline{b}_t} \text{ s.t. } N_{a,b'}(t) \geq \gamma N_{a_{t+1}, b_{t+1}}(t), \quad \left| \hat{\mu}_{a,b'}(t) - \mu_{a,b'} \right| < \varepsilon. \quad (20)$$

By construction of $\alpha_\nu(\cdot)$ (see Section A), since $\varepsilon < \varepsilon_\nu$, inequalities (19) and (20) give us

$$\sum_{b' \in \mathcal{B}_{a, \underline{b}_t} : N_{a,b'}(t) \geq \gamma N_{a_{t+1}, b_{t+1}}(t)} N_{a,b'}(t) \text{kl}(\mu_{a,b'} | \mu_{\underline{b}_t}^* - \omega_{\underline{b}_t, b'}) \leq (1 + \alpha_\nu(\varepsilon)) \log(N_b(T)). \quad (21)$$

This implies

$$\sum_{b' \in \mathcal{B}_{a,b_t}} N_{a,b'}(t) \text{kl}(\mu_{a,b'} | \mu_{b_t}^* - \omega_{b_t,b'}) \leq (1 + \alpha_\nu(\varepsilon)) \log(N_b(T)) + \gamma M_\nu N_{a_{t+1},b_{t+1}}(t).$$

Furthermore, using inequality (21), we get

$$N_{a_{t+1},b_{t+1}}(t) \leq \begin{cases} N_{a,b_t}(t) \leq \frac{(1 + \alpha_\nu(\varepsilon)) \log(N_b(T))}{\text{kl}(\mu_{a,b_t} | \mu_{b_t}^*)} & \text{if } c_a = c_a^+, \\ \frac{(1 + \alpha_\nu(\varepsilon)) \log(N_b(T))}{\text{kl}(\mu_{a_{t+1},b_{t+1}} | \mu_{b_t}^* - \omega_{b_t,b_{t+1}})} & \text{if } c_a < c_a^+. \end{cases} \quad (22)$$

Thus, we have shown that for all arm $a \in \mathcal{A}$, for all time step $1 \leq t \leq T$ such that $a_{t+1} = a$, $(a_{t+1}, b_{t+1}) \notin \mathcal{O}^*$ and $t \in \mathcal{T}_{\varepsilon,\gamma}$:

$$\min_{b: (a,b) \notin \mathcal{O}^*} \frac{1}{\log(N_b(T))} \sum_{b' \in \mathcal{B}_{a,b}} N_{a,b'}(t) \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) \leq (1 + \alpha_\nu(\varepsilon)) \left(1 + \gamma \frac{M_\nu}{m_\nu}\right)$$

and

$$N_{a_{t+1},b_{t+1}}(t) \leq \frac{(1 + \alpha_\nu(\varepsilon)) \log(N_b(T))}{m_\nu}.$$

This implies for all arm $a \in \mathcal{A}$ and for all time step $1 \leq t \leq T$,

$$\min_{b: (a,b) \notin \mathcal{O}^*} \frac{1}{\log(N_b(T))} \sum_{b' \in \mathcal{B}_{a,b}} N_{a,b'}(T) \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) \leq (1 + \alpha_\nu(\varepsilon)) \left[1 + \gamma \frac{M_\nu}{m_\nu}\right] + \frac{M_\nu |\mathcal{T}_{\varepsilon,\gamma}|}{\min_{b \in \mathcal{B}} \log(N_b(T))}$$

and

$$\forall b : (a, b) \notin \mathcal{O}^*, \quad N_{a,b}(T) \leq \frac{(1 + \alpha_\nu(\varepsilon)) \log(N_b(T))}{m_\nu} + |\mathcal{T}_{\varepsilon,\gamma}^c|.$$

■

It can be easily proved that under IMED-GS* $N_b(T) \rightarrow \infty$ for all $b \in \mathcal{B}$ (see Lemma 27). From previous Proposition 20, we deduce the following corollary by doing $T \rightarrow \infty$, then $\varepsilon, \gamma \rightarrow 0$.

Corollary 21 (Pareto optimality) *Let $\nu \in \mathcal{D}_\omega$. Let $a \in \mathcal{A}$ such that $\{b \in \mathcal{B} : (a, b) \notin \mathcal{O}^*\} \neq \emptyset$. Then, we have*

$$\limsup_{T \rightarrow \infty} \min_{b: (a,b) \notin \mathcal{O}^*} \frac{1}{\log(N_b(T))} \sum_{b' \in \mathcal{B}_{a,b}} N_{a,b'}(T) \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) \leq 1.$$

D.4 IMED-GS* is consistent and induces sequences of users with log-frequencies $1_{\mathcal{B}}$

In this section we show that IMED-GS* is a consistent strategy that induces sequences of users with log-frequencies all equal to 1, independently from the considered bandit configuration in \mathcal{D} .

Lemma 22 (Consistency, log-frequencies $1_{\mathcal{B}}$) *IMED-GS^{*} is a consistent strategy and induces sequences of users with log-frequencies all equal to 1.*

Proof We first show that IMED-GS^{*} induces sequences of users with log-frequencies all equal to 1.

Let $\nu \in \mathcal{D}_\omega$ and let us consider an horizon $T \geq 1$. Let $0 < \varepsilon < \varepsilon_\nu$ and $\gamma \in (0, 1/2)$. Let us consider again the set of times

$$\mathcal{T}_{\varepsilon, \gamma} = \left\{ T \geq t \geq 1: \begin{array}{l} \hat{\mathcal{O}}^*(t) = \mathcal{O}^* \\ \forall (a, b) \text{ s.t. } N_{a,b}(t) \geq \gamma N_{a_{t+1}, b_{t+1}}(t) \text{ or } (a, b) \in \mathcal{O}^*, |\hat{\mu}_{a,b}(t) - \mu_{a,b}| < \varepsilon \end{array} \right\}.$$

Then, according to Proposition 20, under IMED-GS^{*} strategy,

$$\mathbb{E}_\nu[|\mathcal{T}_{\varepsilon, \gamma}^c|] \leq \frac{17}{\gamma \varepsilon^4} |\mathcal{A}|^2 |\mathcal{B}|^2 + 2 |\mathcal{A}|^2 |\mathcal{B}| (1 + E_\nu)^{|\mathcal{B}|} < \infty. \quad (23)$$

and for all horizon time $T \geq 1$, for all $(a, b) \notin \mathcal{O}^*$,

$$N_{a,b}(T) \leq \frac{1 + \alpha_\nu(\varepsilon)}{m_\nu} \log(N_b(T)) + |\mathcal{T}_{\varepsilon, \gamma}^c| \leq \frac{1 + \alpha_\nu(\varepsilon)}{m_\nu} \log(T) + |\mathcal{T}_{\varepsilon, \gamma}^c|, \quad (24)$$

where $m_\nu = \min_{\substack{(a,b) \notin \mathcal{O}^* \\ b' \in \mathcal{B}_{a,b}}} \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'})$, and $\varepsilon_\nu, \alpha_\nu(\cdot), E_\nu$ defined in Appendix A.

Note that, under IMED-GS^{*}, for all $t \geq 1$ such that $(a_{t+1}, b_{t+1}) \in \hat{\mathcal{O}}^*(t)$ we have $(a_{t+1}, b_{t+1}) \in \underset{(a,b) \in \hat{\mathcal{O}}^*(t)}{\text{argmin}} N_{a,b}(t)$. This implies by definition of $\mathcal{T}_{\varepsilon, \gamma}$ that

$$\forall (a, b), (a', b') \in \mathcal{O}^*, \quad |N_{a,b}(T) - N_{a',b'}(T)| \leq |\mathcal{T}_{\varepsilon, \gamma}^c| + 1. \quad (25)$$

Indeed the difference of pulls between two optimal couples is non-decreasing only at times $t \geq 1$ such that the difference is greater than 1 and $\hat{\mathcal{O}}^*(t) \neq \mathcal{O}^*$. Combining Eq. 24 and 25 we get

$$\begin{aligned} \min_{b \in \mathcal{B}} N_b(T) &\geq \min_{(a,b) \in \mathcal{O}^*} N_{a,b}(T) - 1 \\ &\geq \max_{(a,b) \in \mathcal{O}^*} N_{a,b}(T) - |\mathcal{T}_{\varepsilon, \gamma}^c| - 1 \quad (\text{Eq. 24}) \\ &\geq \frac{1}{|\mathcal{B}|} \sum_{(a,b) \in \mathcal{O}^*} N_{a,b}(T) - |\mathcal{T}_{\varepsilon, \gamma}^c| - 1 \\ &= \frac{1}{|\mathcal{B}|} \left(T - \sum_{(a,b) \notin \mathcal{O}^*} N_{a,b}(T) \right) - |\mathcal{T}_{\varepsilon, \gamma}^c| - 1 \\ &\geq \frac{1}{|\mathcal{B}|} \left(T - (|\mathcal{A}| - 1) |\mathcal{B}| \left[\frac{1 + \alpha_\nu(\varepsilon)}{m_\nu} \log(T) + |\mathcal{T}_{\varepsilon, \gamma}^c| \right] \right) - |\mathcal{T}_{\varepsilon, \gamma}^c| - 1 \quad (\text{Eq. 25}) \\ &\geq \frac{T}{|\mathcal{B}|} - |\mathcal{A}| \frac{1 + \alpha_\nu(\varepsilon)}{m_\nu} \log(T) - |\mathcal{A}| |\mathcal{T}_{\varepsilon, \gamma}^c| - |\mathcal{A}|. \end{aligned}$$

Since $\mathbb{E}_\nu[|\mathcal{T}_{\varepsilon, \gamma}^c|] < \infty$ (see Eq. 23), this implies that IMED-GS^{*} induces sequences of users with log-frequencies all equal to 1.

We show the consistency of IMED-GS^{*} in the following. Let $(a, b) \notin \mathcal{O}^*$ and $\alpha \in (0, 1)$. According to Proposition 20,

$$N_{a,b}(T) \leq \frac{1 + \alpha_\nu(\varepsilon)}{m_\nu} \log(N_b(T)) + \limsup_{T \rightarrow \infty} |\mathcal{T}_{\varepsilon,\gamma}^c|,$$

and monotone convergence theorem ensures

$$\mathbb{E}_\nu[\limsup_{T \rightarrow \infty} |\mathcal{T}_{\varepsilon,\gamma}^c|] = \limsup_{T \rightarrow \infty} \mathbb{E}_\nu[|\mathcal{T}_{\varepsilon,\gamma}^c|] \leq \frac{17}{\gamma \varepsilon^4} |\mathcal{A}|^2 |\mathcal{B}|^2 + 2 |\mathcal{A}|^2 |\mathcal{B}| (1 + E_\nu)^{|\mathcal{B}|} < \infty.$$

This implies

$$\frac{N_{a,b}(T)}{N_b(T)^\alpha} \leq \frac{1 + \alpha_\nu(\varepsilon)}{m_\nu} \frac{\log(N_b(T))}{N_b(T)^\alpha} + \frac{\limsup_{T \rightarrow \infty} |\mathcal{T}_{\varepsilon,\gamma}^c|}{N_b(T)^\alpha},$$

and, taking the expectation, dominated convergence theorem implies

$$\mathbb{E}_\nu \left[\frac{N_{a,b}(T)}{N_b(T)^\alpha} \right] \leq \mathbb{E}_\nu \left[\frac{1 + \alpha_\nu(\varepsilon)}{m_\nu} \frac{\log(N_b(T))}{N_b(T)^\alpha} + \frac{\limsup_{T \rightarrow \infty} |\mathcal{T}_{\varepsilon,\gamma}^c|}{N_b(T)^\alpha} \right] \rightarrow 0.$$

Indeed, it can be easily shown that under IMED-GS^{*} $N_b(T) \rightarrow \infty$ (see Lemma 27). This implies

$$\limsup_{T \rightarrow \infty} \mathbb{E}_\nu \left[\frac{N_{a,b}(T)}{N_b(T)^\alpha} \right] = 0.$$

■

D.5 The counters c_a and c_a^+ coincide at most $O(\log(\log(T)))$ times

Let us consider $0 < \varepsilon < \varepsilon_\nu$ and $\gamma \in (0, 1/2)$. Let us introduce

$$\mathcal{T}_c(T) := \left\{ t \in \mathcal{T}_{\varepsilon,\gamma} : (a_{t+1}, b_{t+1}) \notin \mathcal{O}^* \text{ and } c_{a_{t+1}}(t) = c_{a_{t+1}}^+(t) \right\},$$

where $\mathcal{T}_{\varepsilon,\gamma}$ is define as in Appendix D.2.

In this section, we want to bound $|\mathcal{T}_c(T)|$.

Lemma 23 *Let $0 < \varepsilon < \varepsilon_\nu$ and $\gamma \in (0, 1/2)$. Let us consider an horizon $T \geq 1$. Then, it holds*

$$|\mathcal{T}_c(T)| \leq 2 |\mathcal{A}| + |\mathcal{A}| \log_2 \left(\frac{(1 + \alpha_\nu(\varepsilon)) |\mathcal{B}|}{m_\nu} \log(T) + |\mathcal{B}| |\mathcal{T}_{\varepsilon,\gamma}^c| \right).$$

where $m_\nu = \min_{\substack{(a,b) \notin \mathcal{O}^* \\ b' \in \mathcal{B}_{a,b}}} \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'})$, and $\varepsilon_\nu, \alpha_\nu(\cdot)$ defined in Appendix A.

Proof From Lemma 24, we get:

$$|\mathcal{T}_c(T)| \leq 2|\mathcal{A}| + \sum_{a \in \mathcal{A}} \log_2(c_a(T)).$$

Then applying Lemma 25, it comes:

$$|\mathcal{T}_c(T)| \leq 2|\mathcal{A}| + \sum_{a \in \mathcal{A}} \log_2 \left(\sum_{b: (a,b) \notin \mathcal{O}^*} N_{a,b}(T) + |\mathcal{T}_{\varepsilon,\gamma}^c| \right).$$

We end the proof by combining the previous inequality with Proposition 20 that ensures

$$\forall (a,b) \notin \mathcal{O}^*, \quad N_{a,b}(T) \leq \frac{(1 + \alpha_\nu(\varepsilon))}{m_\nu} \log(T) + |\mathcal{T}_{\varepsilon,\gamma}^c|.$$

■

Lemma 24 Let $0 < \varepsilon < \varepsilon_\nu$ and $\gamma \in (0, 1/2)$. Let us consider an horizon $T \geq 1$. Then, it holds

$$|\mathcal{T}_c(T)| \leq 2|\mathcal{A}| + \sum_{a \in \mathcal{A}} \log_2(c_a(T)).$$

Refer to Appendix A for the definition of ε_ν .

Proof Let $a \in \mathcal{A}$. By construction of $(c_a(t))_{1 \leq t \leq T}$ and $(c_a^+(t))_{1 \leq t \leq T}$, we have

$$c_a^+(T) = 2^{\sum_{t=1}^T \mathbb{I}\{c_a(t) = c_a^+(t)\}}^{-1}. \quad (26)$$

Furthermore, the following inequalities are satisfied

$$|\mathcal{T}_c(T)| \leq \sum_{a \in \mathcal{A}} \sum_{t=1}^T \mathbb{I}\{c_a(t) = c_a^+(t)\} \quad \text{and} \quad \forall a \in \mathcal{A}, \quad c_a^+(T) \leq 2c_a(T). \quad (27)$$

Then Eq. 26 and 27 imply

$$2^{|\mathcal{T}_c(T)| - |\mathcal{A}|} \leq 2^{|\mathcal{A}|} \prod_{a \in \mathcal{A}} c_a(T).$$

■

Lemma 25 Let $0 < \varepsilon < \varepsilon_\nu$ and $\gamma \in (0, 1/2)$. Let us consider an horizon $T \geq 1$. Then, it holds

$$\forall a \in \mathcal{A}, \quad \left| c_a(T) - \sum_{b: (a,b) \notin \mathcal{O}^*} N_{a,b}(T) \right| \leq |\mathcal{T}_{\varepsilon,\gamma}^c|.$$

Refer to Appendix A for the definition of ε_ν .

Proof Let $a \in \mathcal{A}$. At each time step $t \geq 1$ we increment $c_a(t)$ only if $(a_{t+1}, b_{t+1}) \notin \hat{\mathcal{O}}^*(t)$ and $a_{t+1} = a$. Then, if $t \in \mathcal{T}_{\varepsilon,\gamma}$, we have $\hat{\mathcal{O}}^*(t) = \mathcal{O}^*$ and we increment $c_a(t)$ only if we increment one of the $N_{a,b}(t)$ for $b \in \mathcal{B}$ such that $(a,b) \notin \mathcal{O}^*$. ■

D.6 All couples $(a, b) \in \mathcal{A} \times \mathcal{B}$ are asymptotically pulled an infinite number of times

Let $0 < \varepsilon < \varepsilon_\nu$ (defined in Appendix A) and $\gamma \in (0, 1/2)$. Let us consider

$$\mathcal{T}_{\varepsilon, \gamma} = \left\{ t \geq 1 : \begin{array}{l} \hat{\mathcal{O}}^*(t) = \mathcal{O}^* \\ \forall (a, b) \text{ s.t. } N_{a,b}(t) \geq \gamma N_{a_{t+1}, b_{t+1}}(t) \text{ or } (a, b) \in \mathcal{O}^*, |\hat{\mu}_{a,b}(t) - \mu_{a,b}| < \varepsilon \end{array} \right\}.$$

Then, according to Proposition 20, under IMED-GS* strategy, $\mathbb{E}_\nu[|\mathcal{T}_{\varepsilon, \gamma}^c|] < \infty$. In particular, almost surely $|\mathcal{T}_{\varepsilon, \gamma}^c| < \infty$.

Lemma 26 (The indexes tend to infinity) *For all strategy we have $\lim_{t \rightarrow \infty} N_{a_{t+1}, b_{t+1}}(t) = \infty$ and, under IMED-GS*,*

$$\forall (a, b) \in \mathcal{A} \times \mathcal{B}, \lim_{t \rightarrow \infty} I_{a,b}(t) = \infty.$$

Proof

For all couple $(a, b) \in \mathcal{A} \times \mathcal{B}$ such that $N_{a,b}(\infty) < \infty$, we have $\mathbb{I}_{\{(a_{t+1}, b_{t+1})=(a,b)\}} \rightarrow 0$. Then

$$\sum_{(a,b) \in \mathcal{A} \times \mathcal{B}: N_{a,b}(\infty) = \infty} \mathbb{I}_{\{(a_{t+1}, b_{t+1})=(a,b)\}} \rightarrow 1.$$

This implies

$$\begin{aligned} & \sum_{(a,b) \in \mathcal{A} \times \mathcal{B}: N_{a,b}(\infty) = \infty} \mathbb{I}_{\{(a_{t+1}, b_{t+1})=(a,b)\}} N_{a,b}(t) \\ & \geq \min_{(a,b) \in \mathcal{A} \times \mathcal{B}: N_{a,b}(\infty) = \infty} N_{a,b}(t) \sum_{(a,b) \in \mathcal{A} \times \mathcal{B}: N_{a,b}(\infty) = \infty} \mathbb{I}_{\{(a_{t+1}, b_{t+1})=(a,b)\}} \rightarrow \infty. \end{aligned}$$

Thus, since

$$\begin{aligned} N_{a_{t+1}, b_{t+1}}(t) &= \sum_{(a,b) \in \mathcal{A} \times \mathcal{B}: N_{a,b}(\infty) < \infty} \mathbb{I}_{\{(a_{t+1}, b_{t+1})=(a,b)\}} N_{a,b}(t) \\ &+ \sum_{(a,b) \in \mathcal{A} \times \mathcal{B}: N_{a,b}(\infty) = \infty} \mathbb{I}_{\{(a_{t+1}, b_{t+1})=(a,b)\}} N_{a,b}(t), \end{aligned}$$

we have

$$N_{a_{t+1}, b_{t+1}}(t) \rightarrow \infty.$$

Furthermore, under IMED-GS* strategy we have

$$\forall (a, b) \in \mathcal{A} \times \mathcal{B}, \quad I_{a,b}(t) \geq \log(N_{a_{t+1}, b_{t+1}}(t)),$$

which ends the proof. ■

Lemma 27 (The numbers of pulls tend to infinity) *Under IMED-GS* the numbers of pulls almost surely satisfy*

$$\forall (a, b) \in \mathcal{A} \times \mathcal{B}, N_{a,b}(T) \rightarrow \infty.$$

In particular, almost surely for all $(a, b) \in \mathcal{A} \times \mathcal{B}$, $\lim_{T \rightarrow \infty} \hat{\mu}_{a,b}(T) = \mu_{a,b}$.

Proof Lemma 26 ensures $\lim_{T \rightarrow \infty} I_{a,b}(T) = \infty$, for all $(a, b) \in \mathcal{A} \times \mathcal{B}$.

Let $(a, b) \in \mathcal{O}^*$. Since $|\mathcal{T}_{\varepsilon, \gamma}^c| < \infty$ and $\forall T \in \mathcal{T}_{\varepsilon, \gamma}$, $\hat{\mathcal{O}}^*(T) = \mathcal{O}^*$, we have for all $(a, b) \in \mathcal{O}^*$

$$\lim_{T \rightarrow \infty} \log(N_{a,b}(T)) = \lim_{T \rightarrow \infty} I_{a,b}(T) = \infty.$$

Then, for all $(a, b) \in \mathcal{O}^*$, $\lim_{T \rightarrow \infty} N_{a,b}(T) = \infty$ and $\lim_{T \rightarrow \infty} \hat{\mu}_{a,b}(T) = \mu_{a,b}$.

Let $(a, b) \notin \mathcal{O}^*$ and let $T \in \mathcal{T}_{\varepsilon, \gamma}$. Then, the following inequalities occur

$$I_{a,b}(T) \leq \sum_{b': (a,b') \notin \mathcal{O}^*} N_{a,b'}(t) \text{kl}(\hat{\mu}_{a,b'}(t) | \hat{\mu}_b^*(t) - \omega_{b,b'}) + \log(\max(1, N_{a,b'}(t))) \quad (28)$$

and

$$\forall b \in \mathcal{B}, \hat{\mu}_b^*(t) < 1 - \varepsilon_\nu. \quad (29)$$

Since $|\mathcal{T}_{\varepsilon, \gamma}^c| < \infty$, Eq. 28 and 29 imply

$$\lim_{T \rightarrow \infty} \sum_{b': (a,b') \notin \mathcal{O}^*} N_{a,b'}(T) \rightarrow \infty.$$

Then, since $|\mathcal{T}_{\varepsilon, \gamma}^c| < \infty$, from Lemma 25 we get $\lim_{T \rightarrow \infty} c_a(T) = \infty$. This implies

$$\text{argmax} \{t \in \llbracket 1, T \rrbracket : c_a(t) = c_a^+(t)\} \rightarrow \infty \text{ and } \lim_{T \rightarrow \infty} \min_{b \in \mathcal{B}} N_{a,b}(T) = \infty.$$

■

D.7 Concentration lemmas

We state two concentration lemmas that do not depend on the followed strategy. Lemma 28 comes from Lemma B.1 in [Combes and Proutiere \(2014\)](#) and Lemma 29 comes from Lemma 14 in [Honda and Takemura \(2015\)](#). Proofs are provided in Appendix F.

Lemma 28 (Concentration inequalities) *Let $\nu \in \mathcal{D}_\omega$. For all $0 < \varepsilon, \gamma \leq 1/2$ and for all couples $(a, b), (a', b') \in \mathcal{A} \times \mathcal{B}$,*

$$\mathbb{E}_\nu \left[\sum_{t \geq 1} \mathbb{I}_{\{(a_{t+1}, b_{t+1}) = (a, b), N_{a', b'}(t) \geq \gamma N_{a, b}(t), |\hat{\mu}_{a', b'}(t) - \mu_{a', b'}| \geq \varepsilon\}} \right] \leq \frac{17}{\gamma \varepsilon^4}.$$

Lemma 29 (Large deviation probabilities) *Let $\nu \in \mathcal{D}_\omega$. For all couple $(a, b) \in \mathcal{A} \times \mathcal{B}$, for all $0 < \mu < \mu_{a,b}$,*

$$\mathbb{E}_\nu \left[\sum_{n \geq 1} \mathbb{I}_{\{\hat{\mu}_{a,b}^n < \mu\}} n \exp(n \text{kl}(\hat{\mu}_{a,b}^n | \mu)) \right] \leq 6e \left(1 - \frac{\log(1 - \mu)}{\log(1 - \mu_{a,b})} \right)^{-1} \left(1 - e^{-\left(1 - \frac{\log(1 - \mu)}{\log(1 - \mu_{a,b})}\right) \text{kl}(\mu_{a,b} | \mu)} \right)^{-3},$$

where $\hat{\mu}_{a,b}^n$ estimates $\mu_{a,b}$ after n pulls of couple (a, b) (see Appendix A).

D.8 Proof of Lemma 18

Using Lemma 14, for all time step $t \geq 1$, we have

$$\forall (a', b') \in \hat{\mathcal{O}}^*(t), \quad N_{a', b'}(t) \geq N_{a_{t+1}, b_{t+1}}(t) \geq \gamma N_{a_{t+1}, b_{t+1}}(t).$$

Then, based on the concentration inequalities from Lemma 28, we obtain

$$\begin{aligned} \mathbb{E}_\nu [|\mathcal{E}_{\varepsilon, \gamma}^c|] &\leq \sum_{(a,b), (a', b') \in \mathcal{A} \times \mathcal{B}} \mathbb{E}_\nu \left[\sum_{t \geq 1} \mathbb{I}_{\{(a_{t+1}, b_{t+1}) = (a,b), N_{a', b'}(t) \geq \gamma N_{a,b}(t), |\hat{\mu}_{a', b'}(t) - \mu_{a', b'}| \geq \varepsilon\}} \right] \\ &\leq \sum_{(a,b), (a', b') \in \mathcal{A} \times \mathcal{B}} \frac{17}{\gamma \varepsilon^4} \\ &\leq \frac{17}{\gamma \varepsilon^4} |\mathcal{A}|^2 |\mathcal{B}|^2. \end{aligned}$$

Furthermore, for $t \geq 1$, $a \in \mathcal{A}$ and $\mathcal{B}' \subset \mathcal{B}$, we have

$$\begin{aligned} \log(N_{a_{t+1}, b_{t+1}}(t)) &\leq \sum_{b \in \mathcal{B}'} N_{a,b}(t) \text{kl}(\hat{\mu}_{a,b}(t) | \lambda_{a,b}) + \log(N_{a,b}(t)) \\ \Leftrightarrow N_{a_{t+1}, b_{t+1}}(t) &\leq \prod_{b \in \mathcal{B}'} N_{a,b}(t) e^{N_{a,b}(t) \text{kl}(\hat{\mu}_{a,b}(t) | \lambda_{a,b})}, \end{aligned}$$

where $\lambda_{a,b} = \mu_{a,b} - \varepsilon_\nu$ for all couple $(a, b) \in \mathcal{A} \times \mathcal{B}$. Thus, considering estimators of means based

on the numbers of pulls $(\hat{\mu}_{a,b}^n)_{(a,b) \in \mathcal{A} \times \mathcal{B}, n \geq 1}$ (see Appendix A), we have

$$\begin{aligned}
|\Lambda_{\varepsilon_\nu}| &\leq \sum_{t \geq 1} \sum_{\substack{a \in \mathcal{A} \\ \mathcal{B}' \subset \mathcal{B}}} \mathbb{I} \left\{ \forall b \in \mathcal{B}', \hat{\mu}_{a,b}(t) < \lambda_{a,b} \text{ and } N_{a,t+1,b,t+1}(t) \leq \prod_{b \in \mathcal{B}'} N_{a,b}(t) e^{N_{a,b}(t) \text{kl}(\hat{\mu}_{a,b}(t) | \lambda_{a,b})} \right\} \\
&= \sum_{\substack{t \geq 1 \\ (a',b') \in \mathcal{A} \times \mathcal{B}}} \sum_{\substack{a \in \mathcal{A} \\ \mathcal{B}' \subset \mathcal{B}}} \sum_{\substack{n_b \geq 0 \\ b \in \mathcal{B}'}} \mathbb{I} \{(a_{t+1}, b_{t+1}) = (a', b'), N_{a,b}(t) = n_b\} \mathbb{I} \left\{ \forall b \in \mathcal{B}', \hat{\mu}_{a,b}^{n_b} < \lambda_{a,b}, N_{a',b'}(t) \leq \prod_{b \in \mathcal{B}'} n_b e^{n_b \text{kl}(\hat{\mu}_{a,b}^{n_b} | \lambda_{a,b})} \right\} \\
&\leq \sum_{\substack{t \geq 1 \\ (a',b') \in \mathcal{A} \times \mathcal{B}}} \sum_{\substack{a \in \mathcal{A} \\ \mathcal{B}' \subset \mathcal{B}}} \sum_{\substack{n_b \geq 1 \\ b \in \mathcal{B}'}} \mathbb{I} \{(a_{t+1}, b_{t+1}) = (a', b')\} \mathbb{I} \left\{ \forall b \in \mathcal{B}', \hat{\mu}_{a,b}^{n_b} < \lambda_{a,b} \right\} \mathbb{I} \left\{ 1 \leq N_{a',b'}(t) \leq \prod_{b \in \mathcal{B}'} n_b e^{n_b \text{kl}(\hat{\mu}_{a,b}^{n_b} | \lambda_{a,b})} \right\} \\
&+ \sum_{\substack{t \geq 1 \\ (a',b') \in \mathcal{A} \times \mathcal{B}}} \mathbb{I} \{(a_{t+1}, b_{t+1}) = (a', b')\} \mathbb{I} \{N_{a',b'}(t) = 0\} \\
&\leq \sum_{\substack{(a',b') \in \mathcal{A} \times \mathcal{B} \\ \mathcal{B}' \subset \mathcal{B}}} \sum_{\substack{a \in \mathcal{A} \\ \mathcal{B}' \subset \mathcal{B}}} \sum_{\substack{n_b \geq 1 \\ b \in \mathcal{B}'}} \mathbb{I} \left\{ \forall b \in \mathcal{B}', \hat{\mu}_{a,b}^{n_b} < \lambda_{a,b} \right\} \sum_{t \geq 1} \mathbb{I} \{(a_{t+1}, b_{t+1}) = (a', b')\} \mathbb{I} \left\{ 1 \leq N_{a',b'}(t) \leq \prod_{b \in \mathcal{B}'} n_b e^{n_b \text{kl}(\hat{\mu}_{a,b}^{n_b} | \lambda_{a,b})} \right\} \\
&+ |\mathcal{A}| |\mathcal{B}| \\
&\leq \sum_{\substack{(a',b') \in \mathcal{A} \times \mathcal{B} \\ \mathcal{B}' \subset \mathcal{B}}} \sum_{\substack{a \in \mathcal{A} \\ \mathcal{B}' \subset \mathcal{B}}} \sum_{\substack{n_b \geq 1 \\ b \in \mathcal{B}'}} \mathbb{I} \left\{ \forall b \in \mathcal{B}', \hat{\mu}_{a,b}^{n_b} < \lambda_{a,b} \right\} \prod_{b \in \mathcal{B}'} n_b e^{n_b \text{kl}(\hat{\mu}_{a,b}^{n_b} | \lambda_{a,b})} + |\mathcal{A}| |\mathcal{B}| \\
&= |\mathcal{A}| |\mathcal{B}| \sum_{\substack{a \in \mathcal{A} \\ \mathcal{B}' \subset \mathcal{B}}} \sum_{\substack{n_b \geq 1 \\ b \in \mathcal{B}'}} \prod_{b \in \mathcal{B}'} \mathbb{I} \left\{ \hat{\mu}_{a,b}^{n_b} < \lambda_{a,b} \right\} n_b e^{n_b \text{kl}(\hat{\mu}_{a,b}^{n_b} | \lambda_{a,b})} + |\mathcal{A}| |\mathcal{B}| \\
&= |\mathcal{A}| |\mathcal{B}| \left[1 + \sum_{\substack{a \in \mathcal{A} \\ \mathcal{B}' \subset \mathcal{B}}} \prod_{b \in \mathcal{B}'} \sum_{n \geq 1} \mathbb{I} \left\{ \hat{\mu}_{a,b}^n < \lambda_{a,b} \right\} n e^{n \text{kl}(\hat{\mu}_{a,b}^n | \lambda_{a,b})} \right]
\end{aligned}$$

and

$$\mathbb{E}_\nu [|\Lambda_{\varepsilon_\nu}|] \leq |\mathcal{A}| |\mathcal{B}| \left(1 + \sum_{\substack{a \in \mathcal{A} \\ \mathcal{B}' \subset \mathcal{B}}} \prod_{b \in \mathcal{B}'} \mathbb{E}_\nu \left[\sum_{n \geq 1} \mathbb{I} \left\{ \hat{\mu}_{a,b}^n < \lambda_{a,b} \right\} n e^{n \text{kl}(\hat{\mu}_{a,b}^n | \lambda_{a,b})} \right] \right). \quad (30)$$

Then, by applying Lemma 29 based on large deviation inequalities, we have

$$\forall (a, b) \in \mathcal{A} \times \mathcal{B}, \quad \mathbb{E}_\nu \left[\sum_{n \geq 1} \mathbb{I} \left\{ \hat{\mu}_{a,b}^n < \lambda_{a,b} \right\} n e^{n \text{kl}(\hat{\mu}_{a,b}^n | \lambda_{a,b})} \right] \leq E_\nu, \quad (31)$$

where $E_\nu = 6e \max_{a \in \mathcal{A}, b \in \mathcal{B}} \left(1 - \frac{\log(1-\lambda_{a,b})}{\log(1-\mu_{a,b})} \right)^{-1} \left(1 - e^{-\left(1 - \frac{\log(1-\lambda_{a,b})}{\log(1-\mu_{a,b})}\right) \text{kl}(\mu_{a,b} | \lambda_{a,b})} \right)^{-3}$.

By combining Eq. 30 and 31, we conclude that

$$\mathbb{E}_\nu [|\Lambda_{\varepsilon_\nu}|] \leq |\mathcal{A}| |\mathcal{B}| \left(1 + |\mathcal{A}| (1 + E_\nu)^{|\mathcal{B}|} \right) \leq 2 |\mathcal{A}|^2 |\mathcal{B}| (1 + E_\nu)^{|\mathcal{B}|}.$$

Appendix E. IMED-GS*: Proof of Theorem 11 (main result)

In this section we prove the asymptotic optimality of IMED-GS* strategy. The proof is based on the finite time analysis detailed in Appendix D.

E.1 Almost surely $n^{\text{opt}}(T)$ tends to n^ν

For $a \in \mathcal{A}$ such that $\mathcal{B}_a = \{b \in \mathcal{B} : (a, b) \notin \mathcal{O}^*\} \neq \emptyset$, let us define the linear programming

$$\begin{aligned} C_{\omega, a}^*(\nu) &:= \min_{n \in \mathbb{R}_+^{\mathcal{B}_a}} \sum_{b \in \mathcal{B}_a} n_b \Delta_{a, b} \\ \text{s.t.} \quad \forall b \in \mathcal{B}_a &: \sum_{b' \in \mathcal{B}_a} \text{kl}^+(\mu_{a, b'} | \mu_b^* - \omega_{b, b'}) n_{a, b'} \geq 1. \end{aligned}$$

Then $(n_{a, b}^\nu)_{b \in \mathcal{B}_a}$ is the unique optimal solution of the previous minimization problem. Furthermore, we can state the following lemma.

Lemma 30 $\lim_{T \rightarrow \infty} (n_{a, b}^{\text{opt}}(T))_{b \in \mathcal{B}_a} = (n_{a, b}^\nu)_{b \in \mathcal{B}_a}$.

Proof This is a direct application of Lemma 31 and Lemma 43 stated below. ■

Lemma 31 Let $0 < \varepsilon < \varepsilon_\nu$ (see Appendix A) and $\gamma \in (0, 1)$. Let $a \in \mathcal{A}$ such that $\mathcal{B}_a = \{b \in \mathcal{B} : (a, b) \notin \mathcal{O}^*\}$. Let us consider for $T \geq 1$, $\hat{K}_a(T) = (\text{kl}^+(\hat{\mu}_{a, b'}(T) | \hat{\mu}_b^*(T) - \omega_{b, b'}))_{b, b' \in \mathcal{B}_a}$, the vector $\hat{\Delta}_a(T) = (\hat{\mu}_b^*(T) - \hat{\mu}_{a, b}(T))_{b \in \mathcal{B}_a}$ and the parameter $\hat{h}_a(T) = (\hat{K}_a(T), \hat{\Delta}_a(T))$. We also consider

$$\mathcal{H}_a := \left\{ \hat{h}_a(T) \quad , T \in \mathcal{T}_{\varepsilon, \gamma} \right\},$$

where $\mathcal{T}_{\varepsilon, \gamma}$, defined in Appendix D.2, satisfies $|\mathcal{T}_{\varepsilon, \gamma}^c| < \infty$. Then, we have

$$\forall h = (K, \Delta) \in \mathcal{H}_a, K \neq 0 \text{ and } \min_{h=(K, \Delta) \in \mathcal{H}_a} \min_{b \in \mathcal{B}_a} \Delta_b > 0.$$

Proof Let $h = (K, \Delta) \in \mathcal{H}_a$. There exists $T \in \mathcal{T}_{\varepsilon, \gamma}$ such that $h = (K, \Delta) = \hat{h}_a(T) = (\hat{K}_a(T), \hat{\Delta}_a(T))$. Since $T \in \mathcal{T}_{\varepsilon, \gamma}$, we have $\hat{\mathcal{O}}^*(T) = \mathcal{O}^*$. In particular for all $b \in \mathcal{B}_a$, $K_{b, b} = \hat{K}_{a, b, b}(T) = \text{kl}^+(\hat{\mu}_{a, b}(T) | \hat{\mu}_b^*(T)) > 0$. Furthermore, we have

$$\min_{b \in \mathcal{B}_a} \Delta_b = \min_{b \in \mathcal{B}_a} \hat{\Delta}_{a, b}(T) = \min_{b \in \hat{\mathcal{B}}_a(T)} \hat{\mu}_b^*(T) - \hat{\mu}_{a, b}(T) > 0.$$

Lastly since $\forall b \in \mathcal{B}_a, \hat{\mu}_b^*(\infty) = \mu_b^*, \hat{\mu}_{a, b}(\infty) = \mu_{a, b}$, we have

$$\min_{b \in \mathcal{B}_a} \hat{\Delta}_{a, b}(T) \rightarrow \min_{b \in \mathcal{B}_a} \mu_b^* - \mu_{a, b} > 0$$

and

$$\min_{h=(K, \Delta) \in \mathcal{H}_a} \min_{b \in \mathcal{B}_a} \Delta_b = \min_{T \notin \mathcal{E} \cup \mathcal{T}_0} \min_{b \in \mathcal{B}_a} \hat{\Delta}_{a, b}(T) > 0. \quad \blacksquare$$

E.2 Almost surely and on expectation, for all sub-optimal couple $\frac{N_{a,b}(T)}{\log(T)}$ tends to $n_{a,b}^\nu$

Combining the upper bounds from the finite analysis and the asymptotic behaviour of $n^{\text{opt}}(\cdot)$, we prove the asymptotic optimality of IMED-GS*.

Lemma 32 (Asymptotic upper bounds) *For all sub-optimal couple $(a, b) \notin \mathcal{O}^*$,*

$$\limsup_{T \rightarrow \infty} \frac{N_{a,b}(T)}{\log(T)} \leq n_{a,b}^\nu.$$

Proof Let $0 < \varepsilon < \varepsilon_\nu$ (see Appendix A) and $\gamma \in (0, 1/2)$. Let $(a, b) \notin \mathcal{O}^*$ and let us consider an horizon $T \geq 1$. Let us introduce the random variable

$$\tau := \min \{t \in \llbracket 1, T \rrbracket \text{ s.t. } t \in \mathcal{T}_{\varepsilon, \gamma} \setminus \mathcal{T}_c(T) \text{ and } (a_{t+1}, b_{t+1}) = (a, b)\},$$

where $\mathcal{T}_{\varepsilon, \gamma}$ and $\mathcal{T}(T)$ are respectively introduced in Appendix D.2 and D.5. Then, by definition of τ and since $|\mathcal{T}_{\varepsilon, \gamma}^c| < \infty$, from Lemma 23 we have

$$N_{a,b}(T) \leq N_{a,b}(\tau) + |\mathcal{T}_{\varepsilon, \gamma}^c| + |\mathcal{T}_c(T)| = N_{a,b}(\tau) + O(\log(\log(T))). \quad (32)$$

Furthermore, since $\tau \notin \mathcal{T}_c(T)$ we have $c_a(\tau) \neq c_a^+(\tau)$. In addition, since $\tau \in \mathcal{T}_{\varepsilon, \gamma}$ and $(a, b) = (a_{\tau+1}, b_{\tau+1}) \notin \mathcal{O}^*$, Lemma 15 implies the following empirical upper bound

$$N_{a,b}(\tau) \leq \log(\tau) n_{a,b}^{\text{opt}}(\tau). \quad (33)$$

In particular, since $\log(\tau) \leq \log(T)$, Eq. 32 and 33 imply

$$a.s. \quad \frac{N_{a,b}(T)}{\log(T)} \leq \frac{N_{a,b}(\tau)}{\log(\tau)} + \frac{O(\log(\log(T)))}{\log(T)} \leq n^{\text{opt}}(\tau) + \frac{O(\log(\log(T)))}{\log(T)}$$

and, since $a.s. \lim_{T \rightarrow \infty} \tau = \infty$, from Lemma 30 we get

$$a.s. \quad \limsup_{T \rightarrow \infty} \frac{N_{a,b}(T)}{\log(T)} \leq \limsup_{T \rightarrow \infty} n_{a,b}^{\text{opt}}(\tau) + \limsup_{T \rightarrow \infty} \frac{O(\log(\log(T)))}{\log(T)} = n_{a,b}^\nu. \quad \blacksquare$$

Lemma 33 (Asymptotic optimality) *For all sub-optimal couple $(a, b) \notin \mathcal{O}^*$, we have*

$$a.s. \quad \lim_{T \rightarrow \infty} \frac{N_{a,b}(T)}{\log(T)} = n_{a,b}^\nu \quad \text{and} \quad \lim_{T \rightarrow \infty} \mathbb{E}_\nu \left[\frac{N_{a,b}(T)}{\log(T)} \right] = n_{a,b}^\nu.$$

Proof Since IMED-GS* is a consistent strategy that induces sequences of users with log-frequencies equal to 1, we have

$$\forall (a, b) \neq \mathcal{O}^*, \quad \liminf_{T \rightarrow \infty} \frac{1}{\log(T)} \sum_{b' \in \mathcal{B}_a} N_{a,b'}(T) \text{kl}^+(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) \geq 1.$$

Then, Pareto-optimality of IMED-GS* combined with asymptotic upper bounds given in Lemma 32 ensures that for all $(a, b) \notin \mathcal{O}^*$, $N_{a,b}(T)/\log(T) \rightarrow n_{a,b}^\nu$. Since, the $N_{a,b}(T)/\log(T)$ are dominated by an integrable variable (see Proposition 20 in Appendix D), we also have these convergences on average. \blacksquare

Appendix F. Concentration lemmas: Proofs

Lemma Let $\nu \in \mathcal{D}_\omega$. For all $0 < \varepsilon, \gamma \leq 1/2$ and for all couples $(a, b), (a', b') \in \mathcal{A} \times \mathcal{B}$,

$$\mathbb{E}_\nu \left[\sum_{t \geq 1} \mathbb{I}_{\{(a_{t+1}, b_{t+1}) = (a, b), N_{a', b'}(t) \geq \gamma N_{a, b}(t), |\hat{\mu}_{a', b'}(t) - \mu_{a', b'}| \geq \varepsilon\}} \right] \leq \frac{17}{\gamma \varepsilon^4}.$$

Proof Considering the stopping times $\tau_{a, b}^n = \inf \{t \geq 1, N_{a, b}(t) = n\}$ we will rewrite the sum $\sum_{t \geq 1} \mathbb{I}_{\{(a_{t+1}, b_{t+1}) = (a, b), N_{a', b'}(t) \geq \gamma N_{a, b}(t), |\hat{\mu}_{a', b'}(t) - \mu_{a', b'}| \geq \varepsilon\}}$ and use an Hoeffding's type argument.

$$\begin{aligned} & \sum_{t \geq 1} \mathbb{I}_{\{(a_{t+1}, b_{t+1}) = (a, b), N_{a', b'}(t) \geq \gamma N_{a, b}(t), |\hat{\mu}_{a', b'}(t) - \mu_{a', b'}| \geq \varepsilon\}} \\ & \leq \sum_{t \geq 1} \sum_{n \geq 1, m \geq 0} \mathbb{I}_{\{\tau_{a, b}^n = t+1, N_{a', b'}(t) = m\}} \mathbb{I}_{\{m \geq \gamma(n-1), |\hat{\mu}_{a', b'}^m - \mu_{a', b'}| \geq \varepsilon\}} \\ & = \sum_{m \geq 0} \sum_{n \geq 1} \mathbb{I}_{\{m \geq \gamma(n-1), |\hat{\mu}_{a', b'}^m - \mu_{a', b'}| \geq \varepsilon\}} \sum_{t \geq 1} \mathbb{I}_{\{\tau_{a, b}^n = t+1, N_{a', b'}(t) = m\}} \\ & \leq \sum_{m \geq 0} \sum_{n \geq 1} \mathbb{I}_{\{m \geq \gamma(n-1), |\hat{\mu}_{a', b'}^m - \mu_{a', b'}| \geq \varepsilon\}} \sum_{t \geq 1} \mathbb{I}_{\{\tau_{a, b}^n = t+1\}} \\ & \leq \sum_{m \geq 0} \sum_{n \geq 1} \mathbb{I}_{\{m \geq \gamma(n-1), |\hat{\mu}_{a', b'}^m - \mu_{a', b'}| \geq \varepsilon\}} \end{aligned}$$

Taking the expectation, it comes:

$$\begin{aligned} & \mathbb{E}_\nu \left[\sum_{t \geq 1} \mathbb{I}_{\{(a_{t+1}, b_{t+1}) = (a, b), N_{a', b'}(t) \geq \gamma N_{a, b}(t), |\hat{\mu}_{a', b'}(t) - \mu_{a', b'}| \geq \varepsilon\}} \right] \\ & \leq \sum_{m \geq 0} \sum_{n \geq 1} \mathbb{I}_{\{m \geq \gamma(n-1)\}} \mathbb{P}_\nu (|\hat{\mu}_{a'}^m - \mu_{a'}| \geq \varepsilon) \\ & \leq \sum_{m \geq 0} \sum_{n \geq 1} \mathbb{I}_{\{m \geq \gamma(n-1)\}} 2e^{-2m\varepsilon^2} \quad (\text{Hoeffding inequality}) \\ & = 2 \sum_{m \geq 0} \left(\frac{m}{\gamma} + 1 \right) e^{-2m\varepsilon^2} \\ & = 2 \sum_{m \geq 1} \frac{m}{\gamma} e^{-2m\varepsilon^2} + 2 \sum_{m \geq 0} e^{-2m\varepsilon^2} \\ & = \frac{1}{\gamma} \frac{2e^{-2\varepsilon^2}}{(1 - e^{-2\varepsilon^2})^2} + \frac{1}{1 - e^{-2\varepsilon^2}} \\ & = \frac{1}{\gamma} \frac{2e^{2\varepsilon^2}}{(e^{2\varepsilon^2} - 1)^2} + \frac{e^{2\varepsilon^2}}{e^{2\varepsilon^2} - 1} \\ & \leq \frac{1}{\gamma} \frac{8e^{1/2}}{\varepsilon^4} + \frac{e^{2\varepsilon^2}}{2\varepsilon^2} \leq \frac{17}{\gamma \varepsilon^4}. \end{aligned}$$

■

Lemma Let $\nu \in \mathcal{D}_\omega$. For all couple $(a, b) \in \mathcal{A} \times \mathcal{B}$, for all $0 < \mu < \mu_{a,b}$,

$$\mathbb{E}_\nu \left[\sum_{n \geq 1} \mathbb{I}_{\{\hat{\mu}_{a,b}^n < \mu\}} n \exp(n \text{kl}(\hat{\mu}_{a,b}^n | \mu)) \right] \leq \frac{6e}{\left(1 - \frac{\log(1-\mu)}{\log(1-\mu_{a,b})}\right) \left(1 - e^{-\left(1 - \frac{\log(1-\mu)}{\log(1-\mu_{a,b})}\right) \text{kl}(\mu_{a,b} | \mu)}\right)^3},$$

where $\hat{\mu}_{a,b}^n$ estimates $\mu_{a,b}$ after n pulls of couple (a, b) (see Appendix A).

Proof The proof is based on a Chernoff type inequality and a calculation by measurement change. The proof comes from [Honda and Takemura \(2015\)](#). We explicit here the particular case of Bernoulli distributions for completeness.

Let us rephrase Proposition 11 from [Honda and Takemura \(2015\)](#). Since we consider Bernoulli distributions, we get a more explicit formulation.

Proposition 34 Let $\nu \in \mathcal{D}_\omega$. Let $(a, b) \in \mathcal{A} \times \mathcal{B}$ and $0 < \mu < \mu_{a,b}$. Then, for all $n \geq 0$ and $u \in \mathbb{R}$, we have

$$\mathbb{P}_\nu(\text{kl}(\hat{\mu}_{a,b}^n | \mu) \geq u, \hat{\mu}_{a,b}^n \leq \mu) \leq \begin{cases} e^{-n \text{kl}(\mu_{a,b} | \mu)} & \text{if } u \leq \frac{\log(1-\mu)}{\log(1-\mu_{a,b})} \text{kl}(\mu_{a,b} | \mu) \\ 2e\left(1 + \frac{\log(1-\mu_{a,b})}{\log(1-\mu)}\right)n e^{-n \frac{\log(1-\mu_{a,b})}{\log(1-\mu)} u} & \text{otherwise.} \end{cases}$$

We know rewrite equality (27) from [Honda and Takemura \(2015\)](#) with our notations. Let $n \geq 1$. We have from [Proposition 34](#) that :

$$\begin{aligned}
& \mathbb{E}_\nu \left[\mathbb{I}_{\{\hat{\mu}_{a,b}^n \leq \mu\}} n e^{n \text{kl}(\hat{\mu}_{a,b}^n | \mu)} \right] \\
&= \int_0^\infty \mathbb{P}_\nu \left(\mathbb{I}_{\{\hat{\mu}_{a,b}^n \leq \mu\}} n e^{n \text{kl}(\hat{\mu}_{a,b}^n | \mu)} > x \right) dx \\
&= \int_0^\infty \mathbb{P}_\nu \left(n e^{n \text{kl}(\hat{\mu}_{a,b}^n | \mu)} > x, \hat{\mu}_{a,b}^n \leq \mu \right) dx \\
&= \int_{-\infty}^\infty n^2 e^{nu} \mathbb{P}_\nu \left(\text{kl}(\hat{\mu}_{a,b}^n | \mu) > u, \hat{\mu}_{a,b}^n \leq \mu \right) du \quad (x = n e^{nu}, dx = n^2 e^{nu} du) \\
&= \int_{-\infty}^{\frac{\text{kl}(\mu_{a,b} | \mu) \log(1-\mu)}{\log(1-\mu_{a,b})}} n^2 e^{nu} \mathbb{P}_\nu \left(\text{kl}(\hat{\mu}_{a,b}^n | \mu) > u, \hat{\mu}_{a,b}^n \leq \mu \right) du \\
&+ \int_{\frac{\text{kl}(\mu_{a,b} | \mu) \log(1-\mu)}{\log(1-\mu_{a,b})}}^\infty n^2 e^{nu} \mathbb{P}_\nu \left(\text{kl}(\hat{\mu}_{a,b}^n | \mu) > u, \hat{\mu}_{a,b}^n \leq \mu \right) du \\
&\leq \int_{-\infty}^{\frac{\text{kl}(\mu_{a,b} | \mu) \log(1-\mu)}{\log(1-\mu_{a,b})}} n^2 e^{nu} e^{-n \text{kl}(\mu_{a,b} | \mu)} du \\
&+ \int_{\frac{\text{kl}(\mu_{a,b} | \mu) \log(1-\mu)}{\log(1-\mu_{a,b})}}^\infty n^2 e^{nu} 2e \left(1 + \frac{\log(1-\mu_{a,b})}{\log(1-\mu)} n \right) e^{-n \frac{\log(1-\mu_{a,b})}{\log(1-\mu)} u} du \quad (\text{Proposition 34}) \\
&= n e^{-n \text{kl}(\mu_{a,b} | \mu)} \int_{-\infty}^{\frac{\text{kl}(\mu_{a,b} | \mu) \log(1-\mu)}{\log(1-\mu_{a,b})}} n e^{nu} du \\
&+ 2n e \left(1 + \frac{\log(1-\mu_{a,b})}{\log(1-\mu)} n \right) \int_{\frac{\text{kl}(\mu_{a,b} | \mu) \log(1-\mu)}{\log(1-\mu_{a,b})}}^\infty n e^{-\left(\frac{\log(1-\mu_{a,b})}{\log(1-\mu)} - 1 \right) nu} du \\
&= n e^{-n \left(1 - \frac{\log(1-\mu)}{\log(1-\mu_{a,b})} \right) \text{kl}(\mu_{a,b} | \mu)} + 2n e \left(1 + \frac{\log(1-\mu_{a,b})}{\log(1-\mu)} n \right) \frac{e^{-n \left(1 - \frac{\log(1-\mu)}{\log(1-\mu_{a,b})} \right) \text{kl}(\mu_{a,b} | \mu)}}{\frac{\log(1-\mu_{a,b})}{\log(1-\mu)} - 1} \\
&= \left(1 + \frac{2e}{\frac{\log(1-\mu_{a,b})}{\log(1-\mu)} - 1} \right) n e^{-n \left(1 - \frac{\log(1-\mu)}{\log(1-\mu_{a,b})} \right) \text{kl}(\mu_{a,b} | \mu)} \\
&+ \frac{2e}{1 - \frac{\log(1-\mu)}{\log(1-\mu_{a,b})}} n^2 e^{-n \left(1 - \frac{\log(1-\mu)}{\log(1-\mu_{a,b})} \right) \text{kl}(\mu_{a,b} | \mu)}
\end{aligned}$$

To ends the proof, we use the following inequalities for $r > 0$:

$$\begin{aligned}
\sum_{n \geq 1} n e^{-nr} &\leq \frac{1}{(1-e^{-r})^2} \leq \frac{1}{(1-e^{-r})^3} \\
\sum_{n \geq 1} n^2 e^{-nr} &\leq \frac{2}{(1-e^{-r})^3}.
\end{aligned}$$

■

Appendix G. IMED-GS, IMED-GS₂, IMED-GS₂^{*}: Finite-time analysis

In this subsection we rewrite and adapt the results established in Sections D, E for IMED-GS^{*} strategy to the other considered strategies. Mainly, we rewrite the empirical lower bounds and upper bounds detailed in Lemmas 14, and 15. These inequalities form the basis of the analysis of IMED-GS^{*} strategy. For the sake of brevity and clarity, proofs are not given.

G.1 IMED-GS finite-time analysis

Under IMED-GS strategy we do not solve empirical versions of optimisation problem 2 and pull the couples with minimum (pseudo) indexes. This leads to the following empirical bounds.

Lemma 35 (IMED-GS: Empirical lower bounds) *Under IMED-GS, at each step time $t \geq 1$, for all couple $(a, b) \notin \hat{\mathcal{O}}^*(t)$,*

$$\log(N_{a_{t+1}, b_{t+1}}(t)) \leq \sum_{b' \in \hat{\mathcal{B}}_{a,b}(t)} N_{a,b'}(t) \text{kl}(\hat{\mu}_{a,b'}(t) | \hat{\mu}_b^*(t) - \omega_{b,b'}) + \log(N_{a,b'}(t)) .$$

Furthermore, for all couple $(a, b) \in \hat{\mathcal{O}}^*(t)$,

$$N_{a_{t+1}, b_{t+1}}(t) \leq N_{a,b}(t) .$$

Lemma 36 (IMED-GS: Empirical upper bounds) *Under IMED-GS, at each step time $t \geq 1$ such that $(a_{t+1}, b_{t+1}) \notin \hat{\mathcal{O}}^*(t)$ we have*

$$\sum_{b' \in \hat{\mathcal{B}}_{a_{t+1}, b_{t+1}}(t)} N_{a_{t+1}, b'}(t) \text{kl}(\hat{\mu}_{a_{t+1}, b'}(t) | \hat{\mu}_{b_{t+1}}^*(t) - \omega_{b_{t+1}, b'}) \leq \log(N_{b_{t+1}}(t)) .$$

In particular

$$N_{a_{t+1}, b_{t+1}}(t) \text{kl}(\hat{\mu}_{a_{t+1}, b_{t+1}}(t) | \hat{\mu}_{b_{t+1}}^*(t)) \leq \log(N_{b_{t+1}}(t)) .$$

Based on this lemmas, one can prove IMED-GS Pareto-optimality in a similar way as for IMED-GS^{*} strategy.

Proposition 37 (IMED-GS: Upper bounds) *Let $\nu \in \mathcal{D}_\omega$. Let $0 < \varepsilon < \varepsilon_\nu$ and $\gamma \in (0, 1/2)$. Let us introduce*

$$\mathcal{T}_{\varepsilon, \gamma} := \left\{ t \geq 1 : \begin{array}{l} \hat{\mathcal{O}}^*(t) = \mathcal{O}^* \\ \forall (a, b) \text{ s.t. } N_{a,b}(t) \geq \gamma N_{a_{t+1}, b_{t+1}}(t) \text{ or } (a, b) \in \mathcal{O}^*, |\hat{\mu}_{a,b}(t) - \mu_{a,b}| < \varepsilon \end{array} \right\} .$$

Then under IMED-GS strategy,

$$\mathbb{E}_\nu[|\mathcal{T}_{\varepsilon, \gamma}^c|] \leq \frac{17}{\gamma \varepsilon^4} |\mathcal{A}|^2 |\mathcal{B}|^2 + 2 |\mathcal{A}|^2 |\mathcal{B}| (1 + E_\nu)^{|\mathcal{B}|}$$

and for all horizon time $T \geq 1$, for all arm $a \in \mathcal{A}$,

$$\min_{b: (a,b) \notin \mathcal{O}^*} \frac{1}{\log(N_b(T))} \sum_{b' \in \mathcal{B}_{a,b}} N_{a,b'}(T) \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) \leq (1 + \alpha_\nu(\varepsilon)) \left[1 + \gamma \frac{M_\nu}{m_\nu} \right] + \frac{M_\nu |\mathcal{T}_{\varepsilon,\gamma}|}{\min_{b \in \mathcal{B}} \log(N_b(T))}$$

where m_ν and M_ν are defined as follows:

$$m_\nu = \min_{\substack{(a,b) \notin \mathcal{O}^* \\ b' \in \mathcal{B}_{a,b}}} \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}), \quad M_\nu = \max_{(a,b) \notin \mathcal{O}^*} \sum_{b' \in \mathcal{B}_{a,b}} \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}).$$

Furthermore, we have

$$\forall (a,b) \notin \mathcal{O}^*, \quad \frac{N_{a,b}(T)}{\log(N_b(T))} \leq \frac{1 + \alpha_\nu(\varepsilon)}{\text{kl}(\mu_{a,b} | \mu_b^*)} + \frac{|\mathcal{T}_{\varepsilon,\gamma}^c|}{\log(N_b(T))}.$$

Refer to Appendix A for the definitions of ε_ν , $\alpha_\nu(\cdot)$ and E_ν .

From the previous proposition, we deduce the following corollary by doing $T \rightarrow \infty$, then $\varepsilon, \gamma \rightarrow 0$.

Corollary 38 (IMED-GS: Pareto optimality) Let $\nu \in \mathcal{D}_\omega$. Under IMED-GS strategy we have

$$\forall a \in \mathcal{A}, \quad \limsup_{T \rightarrow \infty} \min_{b: (a,b) \notin \mathcal{O}^*} \frac{1}{\log(N_b(T))} \sum_{b' \in \mathcal{B}_{a,b}} N_{a,b'}(T) \text{kl}(\mu_{a,b'} | \mu_b^* - \omega_{b,b'}) \leq 1.$$

G.2 Uncontrolled scenario: Finite-time analysis

When *uncontrolled scenario* is considered, the learner does not choose the users to deal with and the exploration phases may be performed with some delay. This can be formalized within the empirical bounds induced by IMED-GS₂ and IMED-GS₂^{*} strategies.

G.2.1 EMPIRICAL BOUNDS ON THE NUMBERS OF PULLS

For time step $t \geq 1$, let us introduce the last return time of couple $(a_{t+1}, b_{t+1}) \in \mathcal{B}$ as

$$\tau_t := \min \{ t - t' : b_{t'} = b_{t+1}, t \geq t' \geq 1 \}.$$

By definition of τ_t we have

$$N_{a_{t+1}, b_{t+1}}(t) = N_{a_{t+1}, b_{t+1}}(t - \tau_t).$$

Now, empirical bounds on the numbers of pull can be formulated for the *uncontrolled scenario*. These inequalities are the same as those for the *controlled scenario* up to (random) time-delays.

Lemma 39 (Uncontrolled scenario: Empirical lower bounds) Under IMED-GS₂ and IMED-GS₂^{*}, at each step time $t \geq 1$ there exists a random time delay σ_t such that $0 \leq \sigma_t \leq \tau_t$ and for all couple $(a, b) \notin \hat{\mathcal{O}}^*(t - \sigma_t)$

$$\log(N_{a_{t+1}, b_{t+1}}(t - \sigma_t)) \leq \sum_{b' \in \hat{\mathcal{B}}_{a,b}(t - \sigma_t)} N_{a,b'}(t - \sigma_t) \text{kl}(\hat{\mu}_{a,b'}(t - \sigma_t) | \hat{\mu}_b^*(t - \sigma_t) - \omega_{b,b'}) + \log(N_{a,b'}(t - \sigma_t)).$$

Furthermore, for all couple $(a, b) \in \hat{\mathcal{O}}^*(t - \sigma_t)$,

$$N_{a_{t+1}, b_{t+1}}(t - \sigma_t) \leq N_{a, b}(t - \sigma_t).$$

Lemma 40 (Empirical upper bounds) Under IMED-GS₂ and IMED-GS₂^{*}, at each step time $t \geq 1$ such that $(a_{t+1}, b_{t+1}) \notin \hat{\mathcal{O}}^*(t)$, we have

$$\sum_{b' \in \hat{\mathcal{B}}_{a_{t+1}, b_t}(t - \sigma_t)} N_{a_{t+1}, b'}(t - \sigma_t) \text{kl}\left(\hat{\mu}_{a_{t+1}, b'}(t - \sigma_t) \middle| \hat{\mu}_{\underline{b}_t}^*(t - \sigma_t) - \omega_{\underline{b}_t, b'}\right) \leq \log(N_b(t - \sigma_t)),$$

where σ_t is a random time delay such that $0 \leq \sigma_t \leq \tau_t$. Furthermore, we have under IMED-GS₂

$$b_{t+1} = \underline{b}_t, \quad N_{a_{t+1}, b_{t+1}}(t - \sigma_t) \text{kl}\left(\hat{\mu}_{a_{t+1}, b_{t+1}}(t - \sigma_t) \middle| \hat{\mu}_{\underline{b}_{t+1}}^*(t - \sigma_t)\right) \leq \log(N_{b_{t+1}}(t - \sigma_t))$$

and under IMED-GS₂^{*}

$$\frac{N_{a_{t+1}, b_{t+1}}(t - \sigma_t)}{\log(t - \sigma_t)} \leq \begin{cases} \frac{1}{\text{kl}\left(\hat{\mu}_{a_{t+1}, b_t}(t - \sigma_t) \middle| \hat{\mu}_{\underline{b}_t}^*(t - \sigma_t)\right)}, & \text{if } c_{a_{t+1}}(t - \sigma_t) = c_{a_{t+1}}^+(t - \sigma_t) \\ \min\left(\frac{1}{\text{kl}\left(\hat{\mu}_{a_{t+1}, b_{t+1}}(t - \sigma_t) \middle| \hat{\mu}_{\underline{b}_t}^*(t - \sigma_t) - \omega_{\underline{b}_t, b_{t+1}}\right)}, n_{a_{t+1}, b_{t+1}}^{\text{opt}}(t - \sigma_t)\right), & \text{else.} \end{cases}$$

Thus, we prove respectively the Pareto-optimality and optimality of IMED-GS₂ and IMED-GS₂^{*} since we show that the empirical means $\hat{\mu}_{a, b}(t - \sigma_t)$ of couples (a, b) involved in the previous inequalities concentrate as in the case of the *controlled scenario*. This is the case as it is stated in Lemmas 41 of the next subsection.

G.2.2 CONCENTRATION INEQUALITY WITH BOUNDED TIME DELAYS

We prove a concentration lemma that does not depend on the followed strategy. It is a rewriting for the case of *controlled scenario* of Lemma 28.

Lemma 41 (Concentration inequalities) Let $\nu \in \mathcal{D}_\omega$, $0 < \varepsilon$, $\gamma \leq 1/2$ and $(a, b), (a', b') \in \mathcal{A} \times \mathcal{B}$. Then for all sequence of stopping times $(\sigma_t)_{t \geq 1}$ such that $0 \leq \sigma_t \leq \tau_t$ for all $t \geq 1$, we have

$$\mathbb{E}_\nu \left[\sum_{t \geq 1} \mathbb{I}_{\{(a_{t+1}, b_{t+1}) = (a, b), N_{a', b'}(t - \sigma_t) \geq \gamma N_{a, b}(t - \sigma_t), |\hat{\mu}_{a', b'}(t - \sigma_t) - \mu_{a', b'}| \geq \varepsilon\}} \right] \leq \frac{17}{\gamma \varepsilon^4}.$$

Remark 42 There is no need to adapt Lemma 29 for the case of *controlled scenario* since this concentration lemma does not involve the current time steps explicitly.

Proof It is pointed out that for all time step $t \geq 1$, $N_{a_{t+1}, b_{t+1}}(t - \sigma_t) \geq N_{a_{t+1}, b_{t+1}}(t - \tau_t) = N_{a_{t+1}, b_{t+1}}(t)$, then we proceed as in Appendix F.

Considering the stopping times $\tau_{a, b}^n = \inf\{t \geq 1, N_{a, b}(t) = n\}$ we will rewrite the sum

$$\sum_{t \geq 1} \mathbb{I}_{\{(a_{t+1}, b_{t+1}) = (a, b), N_{a', b'}(t - \sigma_t) \geq \gamma N_{a, b}(t - \sigma_t), |\hat{\mu}_{a', b'}(t - \sigma_t) - \mu_{a', b'}| \geq \varepsilon\}}$$

and use an Hoeffding's type argument.

$$\begin{aligned}
& \sum_{t \geq 1} \mathbb{I}\{(a_{t+1}, b_{t+1}) = (a, b), N_{a', b'}(t - \sigma_t) \geq \gamma N_{a, b}(t - \sigma_t), |\hat{\mu}_{a', b'}^m(t - \sigma_t) - \mu_{a', b'}| \geq \varepsilon\} \\
& \leq \sum_{t \geq 1} \mathbb{I}\{(a_{t+1}, b_{t+1}) = (a, b), N_{a', b'}(t - \sigma_t) \geq \gamma N_{a, b}(t), |\hat{\mu}_{a', b'}^m(t - \sigma_t) - \mu_{a', b'}| \geq \varepsilon\} \\
& \leq \sum_{t \geq 1} \sum_{n \geq 1, m \geq 0} \mathbb{I}\{\tau_{a, b}^n = t + 1, N_{a', b'}(t - \sigma_t) = m\} \mathbb{I}\{m \geq \gamma(n - 1), |\hat{\mu}_{a', b'}^m - \mu_{a', b'}| \geq \varepsilon\} \\
& = \sum_{m \geq 0} \sum_{n \geq 1} \mathbb{I}\{m \geq \gamma(n - 1), |\hat{\mu}_{a', b'}^m - \mu_{a', b'}| \geq \varepsilon\} \sum_{t \geq 1} \mathbb{I}\{\tau_{a, b}^n = t + 1, N_{a', b'}(t - \sigma_t) = m\} \\
& \leq \sum_{m \geq 0} \sum_{n \geq 1} \mathbb{I}\{m \geq \gamma(n - 1), |\hat{\mu}_{a', b'}^m - \mu_{a', b'}| \geq \varepsilon\} \sum_{t \geq 1} \mathbb{I}\{\tau_{a, b}^n = t + 1\} \\
& \leq \sum_{m \geq 0} \sum_{n \geq 1} \mathbb{I}\{m \geq \gamma(n - 1), |\hat{\mu}_{a', b'}^m - \mu_{a', b'}| \geq \varepsilon\},
\end{aligned}$$

where the $\hat{\mu}_{a', b'}^m$ are defined in Appendix A. The proof ends the same way as in Appendix F. \blacksquare

Appendix H. Continuity of solutions to parametric linear programs

In this section we recall Lemma 13 established in Magureanu et al. (2014) on the continuity of solutions to parametric linear programs.

Lemma 43 Consider $K \in \mathbb{R}_+^{B \times B}$, $\Delta \in \mathbb{R}_+^B$, and $\mathcal{H} \subset \mathbb{R}_+^{B \times B} \times \mathbb{R}_+^B$. Define $h = (K, \Delta)$. Consider the function Q and the set-valued map Q^*

$$Q(h) = \inf_{x \in \mathbb{R}_+^B} \{\Delta \cdot x \mid K \cdot x \geq 1\}$$

$$Q^*(h) = \{x \geq 0 : \Delta \cdot x \leq Q(h) \mid K \cdot x \geq 1\}.$$

Assume that:

- (i) For all $h \in \mathcal{H}$, all rows and columns of K are non-identically 0
- (ii) $\min_{h \in \mathcal{H}} \min_{b \in B} \Delta_b > 0$.

Then:

- (a) Q is continuous on \mathcal{H}
- (b) Q^* is upper hemicontinuous on \mathcal{H} .

Appendix I. Details on numerical experiments

For the fixed configuration experiments we used the weight matrix ω of Table 1 and the configuration ν described in Table 2. ω and ν have been chosen at random in such a way that the regret under IMED exceeds the structured lower bound on the regret. This means the structure ω is informative for the bandit configuration ν and not taking it into account hinders optimality.

user\user	b_1	b_2	b_3	b_4	b_5	b_6	b_7	b_8	b_9	b_{10}
b_1	0	0.07	0.07	0.12	0.20	0.05	0.16	0.14	0.28	0.03
b_2	0.07	0	0.14	0.13	0.21	0.12	0.09	0.07	0.21	0.04
b_3	0.07	0.14	0	0.19	0.27	0.12	0.11	0.13	0.25	0.10
b_4	0.12	0.13	0.19	0	0.26	0.17	0.22	0.20	0.34	0.09
b_5	0.20	0.21	0.27	0.26	0	0.25	0.18	0.20	0.32	0.17
b_6	0.05	0.12	0.12	0.17	0.25	0	0.21	0.19	0.33	0.08
b_7	0.16	0.09	0.11	0.22	0.18	0.21	0	0.02	0.14	0.13
b_8	0.14	0.07	0.13	0.20	0.20	0.19	0.02	0	0.16	0.11
b_9	0.28	0.21	0.25	0.34	0.32	0.33	0.14	0.16	0	0.25
b_{10}	0.03	0.04	0.10	0.09	0.17	0.08	0.13	0.11	0.25	0

Table 1: weight matrix ω used in the fixed configuration experiment.

arm \user	b_1	b_2	b_3	b_4	b_5	b_6	b_7	b_8	b_9	b_{10}
a_1	0.15	0.11	0.19	0.19	0.08	0.15	0.09	0.08	0.13	0.13
a_2	0.70	0.73	0.71	0.71	0.70	0.71	0.78	0.79	0.64	0.70
a_3	0.13	0.14	0.13	0.15	0.02	0.08	0.17	0.17	0.04	0.14
a_4	0.02	0.04	0.09	0.05	0.16	0.02	0.11	0.11	0.06	0.01
a_5	0.95	0.98	1.00	0.97	0.98	0.98	0.90	0.91	0.84	0.97

Table 2: configuration ν used in the fixed configuration experiment.




Figure 4: $\min_{b \in \mathcal{B}} N_b(\cdot)$ approximated over 1000 runs. At each run we sample uniformly at random a weight matrix ω and then sample uniformly at random a configuration $\nu \in \mathcal{D}_\omega$.