
Covariance adaptive best arm identification

El Mehdi Saad

Université Paris-Saclay, CentraleSupélec
Laboratoire des Signaux et Systèmes
Paris, France
el-mehdi.saad@centralesupelec.fr

Gilles Blanchard

Institut de Mathématique d'Orsay
Université Paris-Saclay
Paris, France
gilles.blanchard@universite-paris-saclay.fr

Nicolas Verzelen

Mistea, INRAE
Institut Agro, Université de Montpellier
Montpellier, France
nicolas.verzelen@inrae.fr

Abstract

We consider the problem of best arm identification in the multi-armed bandit model, under fixed confidence. Given a confidence input δ , the goal is to identify the arm with the highest mean reward with a probability of at least $1 - \delta$, while minimizing the number of arm pulls. While the literature provides solutions to this problem under the assumption of independent arms distributions, we propose a more flexible scenario where arms can be dependent and rewards can be sampled simultaneously. This framework allows the learner to estimate the covariance among the arms distributions, enabling a more efficient identification of the best arm. The relaxed setting we propose is relevant in various applications, such as clinical trials, where similarities between patients or drugs suggest underlying correlations in the outcomes. We introduce new algorithms that adapt to the unknown covariance of the arms and demonstrate through theoretical guarantees that substantial improvement can be achieved over the standard setting. Additionally, we provide new lower bounds for the relaxed setting and present numerical simulations that support their theoretical findings.

1 Introduction and setting

Best arm identification (BAI) is a sequential learning and decision problem that refers to finding the arm with the largest mean (average reward) among a finite number of arms in a stochastic multi-armed bandit (MAB) setting. An MAB model ν is a set of K distributions in \mathbb{R} : ν_1, \dots, ν_K , with means μ_1, \dots, μ_K . An "arm" is identified with the corresponding distribution index. The observation consists in sequential draws ("queries" or "arm pulls") from these distributions, and each such outcome is a "reward". The learner's goal is to identify the optimal arm $i^* := \text{Arg Max}_{i \in \llbracket K \rrbracket} \mu_i$, efficiently. There are two main variants of BAI problems: The *fixed budget setting* [1, 8], where given a fixed number of queries T , the learner allocates queries to candidates arms and provides a

guess for the optimal arm. The theoretical guarantee in this case takes the form of an upper bound on the probability p_T of selecting a sub-optimal arm. The second variant is the *fixed confidence* setting [14, 20], where a confidence parameter $\delta \in (0, 1)$ is given as an input to the learner, and the objective is to output an arm $\psi \in \llbracket K \rrbracket$, such that $\mathbb{P}(\psi = i^*) \geq 1 - \delta$, using the least number of arm pulls. The complexity in this case corresponds to the total number of queries made before the algorithm stops and gives a guess for the best arm that is valid with probability at least $1 - \delta$ according to a specified stopping rule. In this paper we specifically focus on the fixed confidence setting.

The problem of best-arm identification with fixed confidence was extensively studied and is well understood in the literature [16, 19, 20, 14]. However, in these previous works, the problem was considered under the assumption that all observed rewards are independent. More precisely, in each round t , a fresh sample (reward vector), independent of the past, $(X_{1,t}, \dots, X_{K,t})$ is secretly drawn from ν by the environment, and the learner is only allowed to choose one arm A_t out of K (and observe its reward $X_{A_t,t}$). We relax this setting by allowing simultaneous queries. Specifically, we consider the MAB model ν as a *joint probability distribution* of the K arms, and in each round t the learner chooses a subset $C_t \subset \llbracket K \rrbracket$ and observes the rewards $(X_{i,t})_{i \in C_t}$ (see the Game Protocol 1). The high-level idea of our work is that allowing multiple queries per round opens up opportunities to estimate and leverage the underlying structure of the arms distribution, which would otherwise remain inaccessible with one-point feedback. This includes estimating the covariance between rewards at the same time point. It is important to note that our proposed algorithms do not require any prior knowledge of the covariance between arms.

Throughout this paper, we consider two cases: bounded rewards (in Section 4) and Gaussian rewards (in Section 5), with the following assumptions:

Assumption 1 . *Suppose that:*

- *IID assumption with respect to t : $(X_t)_{t \geq 1} = (X_{1,t}, \dots, X_{K,t})_{t \geq 1}$ are independent and identically distributed variables following ν .*
- *There is only one optimal arm: $|\text{Arg Max}_{i \in \llbracket K \rrbracket} \mu_i| = 1$.*

Assumption 2 – \mathcal{B} . *Bounded rewards: The support of ν is in $[0, 1]^K$.*

Assumption 2 – \mathcal{G} . *Gaussian rewards: ν is a multivariate normal distribution.*

Protocol 1 The Game Protocol

Parameters: δ .
while [condition] do
 Choose a subset $C \subseteq \llbracket K \rrbracket$.
 The environment reveals the rewards $(X_i)_{i \in C}$.
end while
Output the selected arm: ψ .

A round corresponds to an iteration in Protocol 1. Denote by $i^* \in \llbracket K \rrbracket$ the optimal arm. The learner algorithm consists of: a sequence of queried subsets $(C_t)_t$ of $\llbracket K \rrbracket$, such that subset C_t at round t is chosen based on past observations, a halting condition to stop sampling (i.e. a stopping time written τ) and an arm ψ to output after stopping the sampling procedure. The theoretical guarantees take the form of a high probability upper-bound on the total number N of queries made through the game:

$$N := \sum_{t=1}^{\tau} |C_t|. \quad (1)$$

2 Motivation and main contributions

The query complexity of best arm identification with independent rewards, when arms distributions are σ -sub-Gaussian is characterized by the following quantity:

$$H(\nu) := \sum_{i \in \llbracket K \rrbracket \setminus \{i^*\}} \frac{\sigma^2}{(\mu_{i^*} - \mu_i)^2}. \quad (2)$$

Some δ -PAC algorithms [16] guarantee a total number of queries satisfying $\tau = \tilde{O}(H(\nu) \log(1/\delta))$ (where $\tilde{O}(\cdot)$ hides a logarithmic factor in the problem parameters).

Recall that the number of queries required for comparing the means of two variables, namely $X_1 \sim \mathcal{N}(\mu_1, 1)$ and $X_2 \sim \mathcal{N}(\mu_2, 1)$, within a tolerance error δ , is approximately $\mathcal{O}(\log(1/\delta)/(\mu_1 - \mu_2)^2)$. An alternative way to understand the sample complexity stated in equation (2) is to view the task of identifying the optimal arm i^* as synonymous with determining that the remaining arms $\llbracket K \rrbracket \setminus \{i^*\}$ are suboptimal. Consequently, the cost of identifying the best arm corresponds to the sum of the comparison costs between each suboptimal arm and the optimal one.

In practical settings, the arms distributions are not independent. This lack of independence can arise in various scenarios, such as clinical trials, where the effects of drugs on patients with similar traits or comparing drugs with similar components may exhibit underlying correlations. These correlations provide additional information that can potentially expedite the decision-making process [18].

In such cases, Protocol 1 allows the player to estimate the means and the covariances of arms. This additional information naturally raises the following question:

*Can we accelerate best arm identification by leveraging
the (unknown) covariance between the arms?*

We give two arguments for a positive answer to the last question: When allowed simultaneous queries (more than one arm per round as in Protocol 1), the learner can adapt to the covariance between variables. To illustrate, consider the following toy example for 2 arms comparison: $X_1 \sim \mathcal{N}(\mu_1, 1)$ and $X_2 = X_1 + Y$ where $Y \sim \epsilon \mathcal{N}(1, 1)$ for some $\epsilon > 0$. BAI algorithms in one query per round framework require $\mathcal{O}((1 + \epsilon)^2 \log(1/\delta)/\epsilon^2)$ queries, which gives $\mathcal{O}(\log(1/\delta)/\epsilon^2)$ for small ϵ . In contrast, when two queries per round are possible, the learner can perform the following test $\mathcal{H}_0 : \mathbb{E}[X_1 - X_2] > 0$ against $\mathcal{H}_1 : \mathbb{E}[X_1 - X_2] \leq 0$. Therefore using standard test algorithms adaptive to unknown variances, such as Student's t -test, leads to a number of queries of the order $\mathcal{O}(\text{Var}(X_1 - X_2) \log(1/\delta)/\epsilon^2) = \mathcal{O}(\log(1/\delta))$. Hence a substantial improvement in the sample complexity can be achieved when leveraging the covariance. We can also go one step further to even reduce the sample complexity in some settings. Indeed, we can establish the sub-optimality of an arm i by comparing it with another sub-optimal arm j faster than when comparing it to the optimal arm i^* . To illustrate consider the following toy example with $K = 3$: let $X_1 \sim \mathcal{N}(\mu_1, 1)$, $X_2 = X_1 + Y$ where $Y \sim \epsilon \mathcal{N}(1, 1)$ and $X_3 \sim \mathcal{N}(\mu_1 + 2\epsilon, 1)$, with X_1, Y and X_3 independent. X_3 is clearly the optimal arm. Eliminating X_1 with a comparison with X_3 requires $\mathcal{O}(\log(1/\delta)/\epsilon^2)$ queries, while comparing X_1 with the sub-optimal arm X_2 requires $\mathcal{O}(\text{Var}(X_1 - X_2) \log(1/\delta)/\epsilon^2) = \mathcal{O}(\log(1/\delta))$.

The aforementioned arguments can be adapted to accommodate bounded variables, such as when comparing correlated Bernoulli variables. These arguments suggest that by utilizing multiple queries per round (as shown in Protocol 1), we can expedite the identification of the best arm compared to the standard one-query-per-setting approach. Specifically, for variables bounded by 1, our algorithm ensures that the cost of eliminating a sub-optimal arm $i \in \llbracket K \rrbracket$ is given by:

$$\min_{j \in \llbracket K \rrbracket: \mu_j > \mu_i} \left\{ \frac{\text{Var}(X_j - X_i)}{(\mu_j - \mu_i)^2} + \frac{1}{\mu_j - \mu_i} \right\}. \quad (3)$$

The additional term $1/(\mu_j - \mu_i)$ arises due to the sub-exponential tail behavior of the sum of bounded variables. To compare the quantity in (3) with its counterpart in the independent case, it is important to note that the minimum is taken over a set that includes the best arm, and that $\text{Var}(X_{i^*} - X_i) \leq 2(\text{Var}(X_{i^*}) + \text{Var}(X_i))$. Consequently, when the variables are bounded, the quantity (3) is no larger than a numerical constant times $H(\nu)$ (its independent case counterpart in 2), with potentially a significant improvements if there is positive correlation between an arm $j \in \llbracket K \rrbracket$ with a higher mean and arm i (so that $\text{Var}(X_j - X_i)$ is small). As a result, our algorithm for bounded variables has a sample complexity, up to a logarithmic factor, that corresponds to the sum of quantities (3) over all sub-optimal arms.

We expand our analysis to encompass the scenario where arms follow a Gaussian distribution. In this context, our procedure ensures that the cost of eliminating a sub-optimal arm $i \in \llbracket K \rrbracket$ is given by:

$$\min_{\substack{j \in \llbracket K \rrbracket: \\ \mu_j > \mu_i}} \left\{ \frac{\text{Var}(X_j - X_i)}{(\mu_j - \mu_i)^2} \vee 1 \right\}.$$

Similar to the setting with bounded variables, the aforementioned quantity is always smaller than its counterpart for independent arms when all variables X_j have a unit variance.

In Section 6, we present two lower bound results for the sample complexity of best arm identification in our multiple query setting, specifically when the arms are correlated. Notably, these lower bounds pinpoint how the optimal sample complexity may decrease with the covariance between arms, making them the first of their kind in the context of best arm identification, to the best of our knowledge. The presented lower bounds are sharp, up to a logarithmic factor, in the case where all arms are positively correlated with the optimal arm. However, it remains an open question to determine a sharper lower bound applicable to a more general class of distributions with arbitrary covariance matrices.

In Section B of the appendix, we introduce a new algorithm that differs from the previous ones by performing comparisons between the candidate arm and a convex combination of the remaining arms, rather than pairwise comparisons. We provide theoretical guarantees for the resulting algorithm.

Lastly, we conduct numerical experiments using synthetic data to assess the practical relevance of our approach.

3 Related work

Best arm identification: BAI in the fixed confidence setting was studied by [11], [22], and [12], where the objective is to find ϵ -optimal arms under the PAC (“probably approximately correct”) model. A summary of various optimal bounds for this problem is presented in [8, 20]. Prior works on fixed confidence BAI [25] and [17] developed strategies adaptive to the unknown variances of the arms. In contrast, our proposed algorithms demonstrate adaptability to all entries of the covariance matrix of the arms. In particular, we establish that in the worst-case scenario, where the arms are independent, our guarantees align with the guarantees provided by these previous approaches.

Covariance in the Multi-Armed Bandits model: Recently, the concept of leveraging arm dependencies in the multi-armed bandit (MAB) model for best-arm identification was explored in [15]. However, their framework heavily relies on prior knowledge, specifically upper bounds on the conditional expectation of rewards from unobserved arms given the chosen arm’s reward. In a similar vein, a game protocol that allows simultaneous queries was examined in [21]. However, their objective differs from ours as their focus is on identifying the most correlated arms, whereas our primary goal is to identify the arm with the highest mean reward.

The extension of the standard multi-armed bandit setting to multiple-point bandit feedback was also considered in the stochastic combinatorial semi-bandit problem (2, 9, 10 and 13). At each round $t \geq 1$, the learner pulls m out of K arms and receives the sum of the pulled arms rewards. The objective is to minimize the cumulative regret with respect to the best choice of arms subset. [27] proposed an algorithm that adapts to the covariance of the covariates within the same arm. While this line of research shares the intuition of exploiting the covariance structure with our paper, there are essential differences between the two settings. In the combinatorial semi-bandit problem, the learner receives the sum of rewards from all selected arms in each round and aims to minimize cumulative regret, necessitating careful exploration during the game. In contrast, our approach does not impose any constraint on the number of queried arms per round, and our focus is purely on exploration.

Simultaneous queries of multiple arms was also considered in the context of graph-based bandit problems [7]. However, in these studies, it is assumed that the distributions of arms are independent.

Model selection racing: Racing algorithms for model selection refers to the problem of selecting the best model out of a finite set efficiently. The main idea consists of early elimination of poorly performing models and concentrating the selection effort on good models. This idea was seemingly first exploited in [23] through Hoeffding Racing. It consists of sequentially constructing a confidence interval for the generalization error of each (non-eliminated) model. Once two intervals become disjoint, the corresponding sub-optimal model is discarded. Later [25] presented an adaptive stopping algorithm using confidence regions derived with empirical Bernstein concentration inequality (3, 24). The resulting algorithm is adaptive to the unknown marginal variances of the models. Similarly, [4] presented a procedure centered around sequential hypothesis testing to make decisions between two possibilities. In their setup, they assume independent samples and use Bernstein concentration inequality tailored for bounded variables to adapt to the variances of the two variables under consideration.

While the idea of exploiting the possible dependence between models was shown [6, 26] to empirically outperform methods based on individual performance monitoring, there is an apparent lack of theoretical guarantees. This work aims to develop a control on the number of sufficient queries for

reliable best arm identification, while being adaptive to the unknown correlation of the candidate arms.

4 Algorithm and main theorem for bounded variables

In this section, we focus on the scenario where the arms are bounded by 1. Let us establish the notation we will be using throughout. For each arm $i \in \llbracket K \rrbracket$ associated with the reward variable X_i , we denote by μ_i the mean of X_i . Additionally, for any pair of arms $i, j \in \llbracket K \rrbracket$, we denote V_{ij} as the variance of the difference between X_i and X_j , i.e., $V_{ij} := \text{Var}(X_i - X_j)$. We remind the reader that the quantities V_{ij} are unknown to the learner.

Algorithm 2 is developed based on the ideas introduced in Section 2, which involves conducting sequential tests between every pair $(i, j) \in \llbracket K \rrbracket \times \llbracket K \rrbracket$ of arms. The key element to adapt to the covariance between arms is the utilization of the empirical Bernstein's inequality [24] for the sequence of differences $(X_{i,t} - X_{j,t})_t$ for $i, j \in \llbracket K \rrbracket$. To that end, we introduce the following quantity:

$$\hat{\Delta}_{ij}(t, \delta) := \hat{\mu}_{j,t} - \hat{\mu}_{i,t} - \frac{3}{2} \alpha(t, \delta) \sqrt{2\widehat{V}_{ij,t}} - 9 \alpha^2(t, \delta),$$

where, $\hat{\mu}_{i,t}$ represents the empirical mean of the samples obtained from arm i up to round t , and $\widehat{V}_{ij,t}$ denotes the empirical variance associated with the difference variable $(X_i - X_j)$ up to round t . The term $\alpha(t, \delta)$ is defined as $\alpha(t, \delta) := \sqrt{\log(1/\delta_t)/(t-1)}$, and δ_t is given by $\delta/(2K^2t(t+1))$.

By leveraging the empirical Bernstein's inequality (restated in Theorem K.1), we can establish that if $\hat{\Delta}_{ij}(t, \delta) > 0$ at any time t , then with a probability of at least $1 - \delta$, we have $\mu_i < \mu_j$. This observation indicates that arm i is sub-optimal. Furthermore, when $\mu_j > \mu_i$ for $i, j \in \llbracket K \rrbracket$, a sufficient sample sizes ensuring that the quantity $\hat{\Delta}_{ij}(t, \delta)$ is positive is proportional to:

$$\Lambda_{ij} := \left(\frac{V_{ij}}{(\mu_i - \mu_j)^2} + \frac{1}{\mu_j - \mu_i} \right) \log \left(\frac{K}{(\mu_j - \mu_i)\delta} \right). \quad (4)$$

Moreover, we show in Lemma E.5 in the appendix that the last quantity is necessary, up to a smaller numerical constant factor, in order to have $\hat{\Delta}_{ij}(t, \delta) > 0$.

Algorithm 2 follows a successive elimination approach based on the tests $\hat{\Delta}_{ij}(t, \delta) > 0$. Our objective is to ensure that any sub-optimal arm i is queried at most $\min_j \Lambda_{ij}$ times, where the minimum is taken over arms with means larger than μ_i . However, this approach poses a challenge: the arm j^* that achieves the minimum of Λ_{ij} may be eliminated early in the process, and the algorithm is then constrained to compare arm i with other arms costing larger complexity Λ_{ij} . To address this limitation, it is useful to continue querying arms even after deciding their sub-optimality through the pairwise tests.

In Algorithm 2, we introduce two sets at round t : S_t , which contains arms that are candidates to be optimal, initialized as $S_1 = \llbracket K \rrbracket$, and C_t , which represents the set of arms queried at round t . Naturally, C_t contains S_t and also includes arms that were freshly eliminated from S_t , as we hope that these arms will help in further eliminating candidate arms from S_t more quickly.

An important consideration is how long the algorithm should continue sampling an arm that has been eliminated from S_t . In Theorem E.2, we prove that when arm j is eliminated at round t , it is sufficient to keep sampling it up to round $82t$ (the constant 82 is discussed in Remark 1). The rationale behind this number of additional queries is explained in the sketch of the proof of Theorem 4.1. It suggests that if arm j fails to eliminate another arm after round $82t$, then the arm that eliminated j from S_t can ensure faster eliminations for the remaining arms in S_t .

Theorem 4.1. *Suppose Assumption 1 and 2 – \mathcal{B} holds. Consider Algorithm 2, with input $\delta \in (0, 1)$. We have with probability at least $1 - \delta$: the algorithm identifies the best arm and the total number of queries N satisfies:*

$$N \leq c \log(K\Lambda\delta^{-1}) \sum_{i \in \llbracket K \rrbracket \setminus \{i^*\}} \min_{\substack{j \in \llbracket K \rrbracket: \\ \mu_j > \mu_i}} \left\{ \frac{V_{ij}}{(\mu_i - \mu_j)^2} + \frac{1}{\mu_j - \mu_i} \right\}, \quad (5)$$

where $\Lambda = \max_{i \in \llbracket K \rrbracket} \min_{\substack{j \in \llbracket K \rrbracket: \\ \mu_j > \mu_i}} \left\{ \frac{V_{ij}}{(\mu_i - \mu_j)^2} + \frac{1}{\mu_j - \mu_i} \right\}$ and c is a numerical constant.

Algorithm 2 Pairwise-BAI

```
1: Input  $\delta$ .
2: Initialization:
3:   Query all arms for 2 rounds and compute the empirical means vector  $\hat{\mu}_t, t \leftarrow 3$ .
4:    $S \leftarrow \llbracket K \rrbracket$ , /*Set of candidate arms*/
5:    $C \leftarrow \llbracket K \rrbracket$ , /*Set of queried arms*/
6: while  $|S| > 1$  do
7:   Jointly query all arms in  $C$ .
8:   Update  $\hat{\mu}_t$  and compute  $\max_{j \in C} \hat{\Delta}_{ij}(t, \delta)$  for each  $i \in S$ .
9:   for  $i \in S$  do
10:    if  $\max_{j \in C} \hat{\Delta}_{ij}(t, \delta) > 0$  then
11:      Eliminate  $i$  from  $S$ :  $S \leftarrow S \setminus \{i\}$ . /*  $i$  is sub-optimal */
12:      Mark  $i$  for elimination from  $C$  at round:  $82t$ .
13:    end if
14:  end for
15:  Increment  $t$ .
16: end while
17: Return  $S$ .
```

Moreover, if we omit line 12 from Algorithm 2, that is we do not query non-candidate arms, we have with probability at least $1 - \delta$: the algorithm identifies the best arm and the total number of queries N satisfies:

$$N \leq c \log(K\Lambda\delta^{-1}) \sum_{i \in \llbracket K \rrbracket \setminus \{i^*\}} \left\{ \frac{V_{ii^*}}{(\mu_{i^*} - \mu_i)^2} + \frac{1}{\mu_{i^*} - \mu_i} \right\}, \quad (6)$$

where Λ is defined above and c is a numerical constant.

Summary of the proof for bound (5): Let i denote a sub-optimal arm and $\Upsilon_i := \text{Arg Min}_j \Lambda_{ij}$. First, we show that at round t , if $i \in S_t$ then necessarily we have $\Upsilon_i \cap C_t \neq \emptyset$. We proceed by a contradiction argument: assume $\Upsilon_i \cap C_t = \emptyset$, let j denote the element of Υ_i with the largest mean. Let k denote the arm that eliminated j from S at a round $s < t$. Lemma E.3 shows that, since $\hat{\Delta}_{jk}(s, \delta) > 0$, we necessarily have $\log(1/\delta)\Lambda_{jk}/4 \lesssim s$. Moreover, j was kept up to round $82s$ and in this last round we had $\hat{\Delta}_{ij}(82s, \delta) \leq 0$, which gives by Lemma E.3: $82s \lesssim (25/2) \log(1/\delta_{82s})\Lambda_{ij}$. Combining the two bounds on s , gives: $\Lambda_{jk} < \Lambda_{ij}$. Finally, we use the ultra-metric property satisfied by the quantities Λ_{uv} (Lemma E.6), stating that: $\Lambda_{ik} \leq \max\{\Lambda_{ij}, \Lambda_{jk}\}$. Combining the last inequality with the latter, we get $\Lambda_{ik} \leq \Lambda_{ij}$, which means that $k \in \Upsilon_i$. The contradiction arises from the fact that j is the element of the largest mean in Υ_i and arm k eliminated j (hence $\mu_k > \mu_j$). We conclude that necessarily $\Upsilon_i \cap C_t \neq \emptyset$, therefore by Lemma E.3, at a round of the order of $\min_j \Lambda_{ij}$, we will necessarily have $\hat{\Delta}_{ij}(t, \delta) > 0$ for some $j \in C_t$.

Remark 1. From a practical standpoint, to keep sampling an arm eliminated at t for additional 81 t rounds is a conservative approach. The stated value of 81 is determined by specifics of the proof, but we believe that it can be optimized to a smaller constant based on the insights gained from numerical simulations. Furthermore, even if we omit the oversampling step, as presented in Theorem 4.1, the algorithm is still guaranteed to identify the best arm with probability $1 - \delta$. Only the query complexity guarantee is weaker, but the algorithm may still lead to effective arm elimination, although with potentially slightly slower convergence.

Remark 2. The idea of successive elimination based on evaluating the differences between variables was previously introduced in [28], in the context of model selection aggregation. Their analysis allows having a bound slightly looser than bound (6) (with the distances between variables instead of variances). On the other hand, Theorem 4.1 in our paper provides a sharper bound in (5).

First, it is important to note that bound (5) is always sharper (up to a numerical constant factor) than bound (6) because the optimal arm is included in the set over which the minimum in bound (5) is taken. On the other hand, bound (5) can be smaller than bound (6) by a factor of $1/K$. This situation can arise in scenarios where the $K - 1$ sub-optimal arms, which have close means, are highly correlated, while the optimal arm is independent of the rest. In such a situation, the terms in

the sum in (5) corresponding to the correlated sub-optimal arms are relatively small, except for the arm (denoted as i) with the largest mean among the $K - 1$ correlated suboptimal rewards. That last remaining arm will be eliminated by the optimal arm, incurring a potentially large cost, but for that arm only. On the other hand, each term in the second bound (6) would be of the order of the last cost.

To provide perspective on our guarantees compared to those developed in the independent one query per round setting, it is worth noting that the standard guarantees in that setting provide a sample complexity corresponding to $\sum_{i \neq i^*} \log(1/\delta) / (\mu_{i^*} - \mu_i)^2$. Since variables are bounded by 1, their variances are bounded by 1. Therefore, in the worst case, we recover the previous guarantees with a numerical constant factor of 2.

A refined adaptive algorithm presented in [25] also utilizes a successive elimination approach using confidence intervals for the arm means based on the empirical Bernstein's inequality. However, unlike our algorithm, they use the concentration inequality to evaluate each arm's mean independently of the other arms. Their approach allows for adaptability to individual arm variances, and is particularly beneficial are small, i.e., $\text{Var}(X_i) \ll 1$. The sample complexity of their algorithm is of the order of:

$$\log(K\Gamma\delta^{-1}) \sum_{i \neq i^*} \frac{\text{Var}(X_i) + \text{Var}(X_{i^*})}{(\mu_{i^*} - \mu_i)^2},$$

where $\Gamma := \max_{i \neq i^*} (\text{Var}(X_i) + \text{Var}(X_{i^*}) / (\mu_{i^*} - \mu_i)^2)$. Neglecting numerical constant factors, the last bound is larger than both our bounds (5) and (6). This is because the variance of the differences can be bounded as follows: $V_{ii^*} \leq 2(\text{Var}(X_i) + \text{Var}(X_{i^*}))$.

5 Algorithms and main theorem for Gaussian distributions

In this section, we address the scenario where arms are assumed to follow a Gaussian distribution. We consider a setting where the learner has no prior knowledge about the arms' distribution parameters, and we continue using the notation introduced in the previous section. The main difference between this case and the bounded variables setting, other than the form of the sample complexity obtained, is the behavior of the algorithm when variances of differences between variables tend to 0, displayed by the second bound in Theorem 5.1.

Our algorithm relies on the empirical Bernstein inequality, which was originally designed in the literature for bounded variables [3, 24]. We have extended this inequality to accommodate Gaussian variables by leveraging existing Gaussian concentration results. Note that extending such inequalities more generally for sub-Gaussian variables is a non-trivial task. One possible direction is to suppose that arms follow a sub-Gaussian distribution and satisfy a Bernstein moment assumption (such extensions were pointed by works on bounded variables e.g., ?). Given the last class of distributions, we can build on the standard Bernstein inequality with known variance, then plug in an estimate of the empirical variance leveraging the concentration of quadratic forms (see 5). However, it remains uncertain whether an extension for sub-Gaussian variables (without additional assumptions) is practically feasible.

We extend the previous algorithm to the Gaussian case by performing sequential tests between pairs of arms $(i, j) \in \llbracket K \rrbracket$. We establish a confidence bound for the difference variables $(X_i - X_j)$ (Lemma D.1 in the appendix) and introduce the following quantity:

$$\hat{\Delta}'_{ij}(t, \delta) := \hat{\mu}_{j,t} - \hat{\mu}_{i,t} - \frac{3}{2}\alpha(t, \delta)\sqrt{2f(\alpha(t, \delta))\hat{V}_{ij,t}}, \quad (7)$$

where the f is defined by: $f(x) = \exp(2x + 1)$ if $x \geq 1/3$, and $f(x) = 1/(1 - 2x)$ otherwise.

Using the empirical confidence bounds on the differences between arms, we apply the same procedure as in the case of bounded variables. We make one modification to Algorithm 2: we use the quantities $\hat{\Delta}'_{ij}(t, \delta)$ instead of $\hat{\Delta}_{ij}(t, \delta)$.

By following the same analysis as in the bounded setting, we establish that the sample complexity required for the comparison tests between arms i and j is characterized by the quantity $\frac{V_{ij}}{(\mu_i - \mu_j)^2} \vee 1$. The following theorem provides guarantees on the algorithm presented above:

Theorem 5.1. *Suppose Assumptions 1 and 2 – \mathcal{G} holds. Consider the algorithm described above, with input $\delta \in (0, 1)$. We have with probability at least $1 - \delta$: the algorithm identifies the best arm*

and the total number of queries N satisfies:

$$N \leq c \log(K\Lambda\delta^{-1}) \sum_{i \in \llbracket K \rrbracket \setminus \{i^*\}} \min_{\substack{j \in \llbracket K \rrbracket: \\ \mu_j > \mu_i}} \left\{ \frac{V_{ij}}{(\mu_i - \mu_j)^2} \vee 1 \right\}, \quad (8)$$

where $\Lambda = \max_{i \in \llbracket K \rrbracket} \min_{\substack{j \in \llbracket K \rrbracket: \\ \mu_j > \mu_i}} \left\{ \frac{V_{ij}}{(\mu_i - \mu_j)^2} \vee 1 \right\}$ and c is a numerical constant.

Moreover, if we omit line 12 from Algorithm 2, that is we do not query non-candidate arms, we have with probability at least $1 - \delta$: the algorithm identifies the best arm and the total number of queries N satisfies:

$$N \leq 3K + c \log\left(\frac{K\delta^{-1}}{\log(1 + 1/\Lambda)}\right) \sum_{i \in \llbracket K \rrbracket \setminus \{i^*\}} \frac{1}{\log\left(1 + \frac{(\mu_{i^*} - \mu_i)^2}{V_{ii^*}}\right)}, \quad (9)$$

where $\Lambda = \max_{i \neq i^*} V_{ii^*}/(\mu_{i^*} - \mu_i)^2$ is defined above and c is a numerical constant.

Remarks. On the sample complexity cost of being adaptive to the variance: If the sample size is larger than $\log(1/\delta)$, the cost of plugging in the empirical variance estimate into the Chernoff's concentration inequality is only a multiplicative constant slightly larger than one (nearly $1 + 2\sqrt{\log(1/\delta)/n}$). However, in the case of a small sample regime ($n < \log(1/\delta)$), the cost is a multiplicative factor of $\exp(\sqrt{\log(1/\delta)/n} + 1/2)$ due to the nature of the left tail of the chi-squared distribution (see Sections D and K of the appendix for detailed calculations). For most natural regimes, the number of queries made for each arm is larger than $\log(1/\delta)$, hence the last described effect does not arise. However, in some specific regimes (such as the case of very small variances of the arms) an optimal algorithm should query less than $\log(1/\delta)$ samples, which necessitates introducing the exponential multiplicative term above into the concentration upper bound. This translates into a different form of guarantee presented in Theorem 5.1, inequality (9). It is important to note the cost in this regime cannot be avoided as highlighted by our lower bound presented in Theorem 6.2.

Theorem 5.1 above shows that Algorithm 2 is applied to Gaussian distribution guarantees that each sub-optimal arm i is eliminated after roughly $\min_j V_{ij}^2/(\mu_i - \mu_j)^2$ queries. Bounds (8) and (9) derived from our algorithm are smaller than the standard complexity bound in the independent case, which is given by $\sum_{i \neq i^*} \sigma^2/(\mu_i - \mu_{i^*})^2$ when all arms have variances smaller than σ^2 . It is important to note that unlike most existing procedures in the literature that achieve the standard complexity bound, our algorithm does not require knowledge of the parameter σ^2 .

About the upper bound (9): This bound can be further bounded by $\log(\delta^{-1}) \sum_{i \neq i^*} \frac{V_{ii^*}}{(\mu_i - \mu_{i^*})^2} \vee 1$. The form presented in Theorem 5.1 is particularly sharp when the variances V_{ii^*} tend to zero, resulting in a constant upper bound. Recently, in the independent setting where variances are unknown, [17] analyzed the comparison of two arms i and j . They derived a complexity of the order $1/\log\left(1 + \frac{(\mu_i - \mu_j)^2}{(\sigma_i^2 + \sigma_j^2)}\right)$, where σ_i^2 and σ_j^2 are the variances of arms i and j , respectively. This result is reflected in our bound (9), where instead of the sum of variances, our bound considers the variance of the difference, leading to a sharper bound. In Section 6, we provide a lower bound that nearly matches (9).

6 Lower bounds

In this section, we present lower bounds for the problem of best arm identification with multiple queries per round, following Protocol 1. We provide lower bounds for both the bounded distributions setting (Theorem 6.1) and the Gaussian distributions setting (Theorem 6.2). It is important to note that our lower bounds are derived considering a class of correlated arm distributions. Therefore, the results obtained in the standard one-query-per-round setting [20] do not hold in our setting.

Our first lower bound is derived for the case where arms follow a Bernoulli distribution. Let $\boldsymbol{\mu} = (\mu_i)_{i \in \llbracket K \rrbracket}$ be a sequence of means in $[1/4, 3/4]$, denote by i^* the index of the largest mean. Consider a sequence of positive numbers $\mathbf{V} = (V_{ii^*})_{i \in \llbracket K \rrbracket}$. We define $\mathbb{B}_K(\boldsymbol{\mu}, \mathbf{V})$ as the set of Bernoulli arm distributions such that: (i) $\mathbb{E}[X_i] = \mu_i$; (ii) $\text{Var}(X_i - X_{i^*}) \leq V_{ii^*}$ for all $i \in \llbracket K \rrbracket$.

For Bernoulli distributions, the variances and means are linked. Specifically, Lemma K.8 demonstrates that for any pair of Bernoulli random variables (B_1, B_2) with means (b_1, b_2) , the following inequality holds: $(b_1 - b_2) - (b_1 - b_2)^2 \leq \text{Var}(B_1 - B_2) \leq \min\{2 - (b_1 + b_2); b_1 + b_2\} - (b_1 - b_2)^2$.

We assume that the sequences \mathbf{V} and $\boldsymbol{\mu}$ satisfy the aforementioned condition for $(X_i, X_{i^*})_{i \in \llbracket K \rrbracket}$, as otherwise the class $\mathbb{B}_K(\boldsymbol{\mu}, \mathbf{V})$ would be an empty set. Theorem 6.1 provides a lower bound on the sample complexity required for best arm identification in the worst-case scenario over the class $\mathbb{B}_K(\boldsymbol{\mu}, \mathbf{V})$.

Theorem 6.1. *Let $K \geq 2$ and $\delta \in (0, 1)$. For any δ -sound algorithm, we have*

$$\max_{\mathcal{B} \in \mathbb{B}_K(\boldsymbol{\mu}, \mathbf{V})} \mathbb{E}_{\mathcal{B}}[N] \geq \frac{1}{8} \log(1/4\delta) \sum_{i \in \llbracket K \rrbracket \setminus \{i^*\}} \max \left\{ \frac{V_{ii^*}}{(\mu_{i^*} - \mu_i)^2}, \frac{1}{\mu_{i^*} - \mu_i} \right\},$$

where N is the total number of queries.

The presented lower bound takes into account the correlation between arms by incorporating the quantities V_{ii^*} as upper bounds for the variances of the differences between arm i and the optimal arm i^* . This lower bound indicates that for class $\mathbb{B}_K(\boldsymbol{\mu}, \mathbf{V})$, Algorithm 2 is nearly optimal.

Next, we provide a lower bound in the Gaussian case. Let $\boldsymbol{\mu} = (\mu_i)_{i \in \llbracket K \rrbracket}$ be a sequence of means, where i^* denotes the index of the largest mean. We also consider a sequence of positive numbers $\mathbf{V} = (V_{ii^*})_{i \in \llbracket K \rrbracket}$. We define $\mathbb{G}_K(\boldsymbol{\mu}, \mathbf{V})$ as the set of Gaussian arm distributions satisfying the following conditions: (i) $\mathbb{E}[X_i] = \mu_i$, (ii) $\text{Var}(X_i - X_{i^*}) = V_{ii^*}$, and (iii) $\text{Var}(X_i) \geq 1$ for all $i \in \llbracket K \rrbracket$. Theorem 6.2 provides a lower bound on the sample complexity required for best arm identification in the worst-case scenario over the class $\mathbb{G}_K(\boldsymbol{\mu}, \mathbf{V})$.

Theorem 6.2. *Let $K \geq 2$ and $\delta \in (0, 1)$. For any δ -sound algorithm, we have*

$$\max_{\mathcal{G} \in \mathbb{G}_K(\boldsymbol{\mu}, \mathbf{V})} \mathbb{E}_{\mathcal{G}}[N] \geq 2 \log(1/4\delta) \sum_{i \in \llbracket K \rrbracket \setminus \{i^*\}} \frac{1}{\log \left(1 + \frac{(\mu_i - \mu_{i^*})^2}{V_{ii^*}^2} \right)},$$

where N is the total number of queries.

Theorem 6.2 demonstrates that our algorithm achieves near-optimal performance over $\mathbb{G}_K(\boldsymbol{\mu}, \mathbf{V})$, up to a logarithmic factor.

7 Numerical simulations

We consider the Gaussian rewards scenario. We compare our algorithm Pairwise-BAI (Algorithm 2) to 3 benchmark algorithms: Hoeffding race [23], adapted to the Gaussian setting (consisting of successive elimination based on Chernoff's bounds) and LUCB [19], which is an instantiation of the upper confidence bound (UCB) method. We assume that the last two algorithms have a prior knowledge on the variances of the arms. The third benchmark algorithm consists of using a successive elimination approach using the empirical estimates of the variances. We evaluated two variations of our algorithm. The first one, Pairwise-BAI+, implemented Algorithm 2 for Gaussian variables. In this instance, we modified line 12 by continuing to sample sub-optimal arms that were eliminated at round t until round $2t$ instead of $82t$. We stress that both variants guarantee a δ -sound decision on the optimal arm (see Theorem 5.1). The second instance involved removing the last instruction, meaning we directly stopped querying sub-optimal arms. Figure 1 displays the average sample complexities for each considered algorithm. As expected, the larger the correlation between arms, the better Pairwise-BAI performs.

The second experiment aims to demonstrate that Algorithm 2 adapts to the covariance between sub-optimal arms, as indicated by bound (8) in Theorem 4.1, rather than solely adapting to the correlation with the optimal arm as shown by (9). In this experiment, we consider that the arms are organized into clusters, where each pair of arms within the same cluster exhibits a high correlation (close to 1), while arms from different clusters are independent. According to bound (8) (see also discussion after Theorem 4.1) we expect to observe a gain of up to a factor corresponding to the number of arms per cluster compared to algorithms that do not consider covariance. (Since the total number of arms is kept fixed, the number of arms per cluster scales as the inverse of the number of clusters.)

Figure 1 illustrates the results, showing that the performance of Pairwise-BAI improves as the number of clusters decreases (indicating a larger number of correlated arms). This suggests that increasing

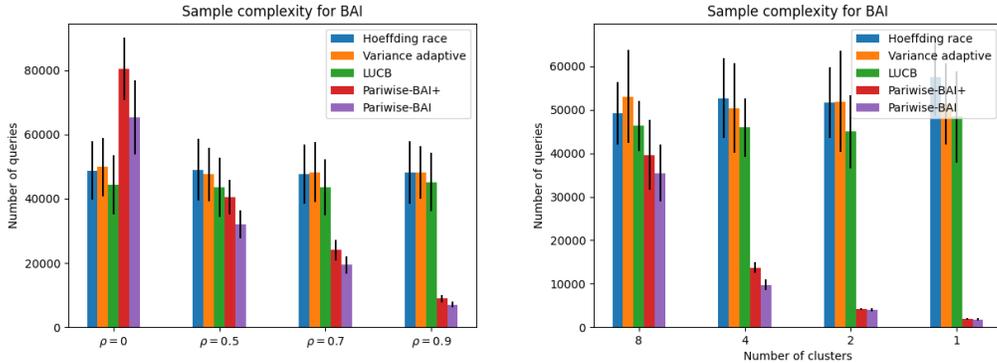


Figure 1: Left figure: the average sample complexities, in 4 scenarios with $K = 10$ arms with means $\mu_i = i/10$. The covariance matrix C is defined as follows: for $i \in \llbracket K \rrbracket$, $C_{ii} = 1$ and for $i \neq j$: $C_{ij} = \rho$. We consider 4 scenarios with the correlation: $\rho \in \{0, 0.5, 0.7, 0.9\}$. Right figure: the average sample complexities with $K = 16$ arms with means $\mu_i = i/10$. Arms in the same cluster have a correlation of 0.99, and arms from different clusters are independent. We consider 4 scenarios with different number of clusters: $n_{cl} \in \{8, 4, 2, 1\}$. All clusters are of the same size.

the number of correlated sub-optimal arms, which are independent from the optimal arm, still leads to significant performance improvement. These findings support the idea that Pairwise-BAI and Pairwise-BAI+ exhibit behavior that aligns more closely with bound (8) in Theorem 4.1, rather than bound (9).

In both experiments, we observe that Pairwise-BAI+ performs worse compared to Pairwise-BAI, indicating that, empirically, in the given scenarios, continuing to sample sub-optimal arms does not contribute to improved performance. While this modification provides better theoretical guarantees, it may not lead to empirical performance improvements in general scenarios.

8 Conclusion and future directions

This work gives rise to several open questions. Firstly, the presented lower bounds take into account partially the covariance between arms. It would be interesting to explore the development of a more precise lower bound that can adapt to any covariance matrix of the arms. Additionally, in terms of the upper bound guarantees, our focus has been on pairwise comparisons, along with an algorithm that compares candidate arms with convex combinations of the remaining arms (Section B). An interesting direction for further research would involve extending this analysis to an intermediate setting, involving comparisons with sparse combinations.

Acknowledgements: This work is supported by ANR-21-CE23-0035 (ASCAI) and ANR-19-CHIA-0021-01 (BISCOTTE). This work was conducted while EM Saad was in INRAE Montpellier.

References

- [1] Jean-Yves Audibert and Sébastien Bubeck. Best arm identification in multi-armed bandits. In *Conference on Learning Theory*, 2010.
- [2] Jean-Yves Audibert, Sébastien Bubeck, and Gábor Lugosi. Regret in online combinatorial optimization. *Mathematics of Operations Research*, 39(1):31–45, 2014.
- [3] Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Tuning bandit algorithms in stochastic environments. In *International conference on algorithmic learning theory*, pages 150–165. Springer, 2007.
- [4] Akshay Balsubramani and Aaditya Ramdas. Sequential nonparametric testing with the law of the iterated logarithm. In *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence*, pages 42–51, 2016.
- [5] Pierre C Bellec. Concentration of quadratic forms under a Bernstein moment assumption. *arXiv preprint arXiv:1901.08736*, 2019.
- [6] Mauro Birattari, Zhi Yuan, Prasanna Balaprakash, and Thomas Stützel. F-race and iterated F-race: An overview. *Experimental methods for the analysis of optimization algorithms*, pages 311–336, 2010.
- [7] Stephane Caron, Branislav Kveton, Marc Lelarge, and S Bhagat. Leveraging side observations in stochastic bandits. In *UAI*, 2012.
- [8] Alexandra Carpentier and Andrea Locatelli. Tight (lower) bounds for the fixed budget best arm identification bandit problem. In *Conference on Learning Theory*, pages 590–604. PMLR, 2016.
- [9] Nicolo Cesa-Bianchi and Gábor Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422, 2012.
- [10] Wei Chen, Yajun Wang, and Yang Yuan. Combinatorial multi-armed bandit: General framework and applications. In *International conference on machine learning*, pages 151–159. PMLR, 2013.
- [11] Eyal Even-Dar, Shie Mannor, and Yishay Mansour. PAC bounds for multi-armed bandit and Markov decision processes. In *International Conference on Computational Learning Theory*, pages 255–270. Springer, 2002.
- [12] Eyal Even-Dar, Shie Mannor, Yishay Mansour, and Sridhar Mahadevan. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of machine learning research*, 7(6), 2006.
- [13] Yi Gai, Bhaskar Krishnamachari, and Rahul Jain. Combinatorial network optimization with unknown variables: Multi-armed bandits with linear rewards and individual observations. *IEEE/ACM Transactions on Networking*, 20(5):1466–1478, 2012.
- [14] Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *Conference on Learning Theory*, pages 998–1027. PMLR, 2016.
- [15] Samarth Gupta, Gauri Joshi, and Osman Yağan. Best-arm identification in correlated multi-armed bandits. *IEEE Journal on Selected Areas in Information Theory*, 2(2):549–563, 2021.
- [16] Kevin Jamieson, Matthew Malloy, Robert Nowak, and Sébastien Bubeck. Lil’ucb: An optimal exploration algorithm for multi-armed bandits. In *Conference on Learning Theory*, pages 423–439. PMLR, 2014.
- [17] Marc Jourdan, Rémy Degenne, and Emilie Kaufmann. Dealing with unknown variances in best-arm identification. In *International Conference on Algorithmic Learning Theory*, pages 776–849. PMLR, 2023.
- [18] Brennan C Kahan and Tim P Morris. Improper analysis of trials randomised using stratified blocks or minimisation. *Statistics in medicine*, 31(4):328–340, 2012.

- [19] Shivaram Kalyanakrishnan, Ambuj Tewari, Peter Auer, and Peter Stone. PAC subset selection in stochastic multi-armed bandits. In *ICML*, volume 12, pages 655–662, 2012.
- [20] Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *The Journal of Machine Learning Research*, 17(1):1–42, 2016.
- [21] Che-Yu Liu and Sébastien Bubeck. Most correlated arms identification. In *Conference on Learning Theory*, pages 623–637. PMLR, 2014.
- [22] Shie Mannor and John N Tsitsiklis. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research*, 5(Jun):623–648, 2004.
- [23] Oded Maron and Andrew Moore. Hoeffding races: Accelerating model selection search for classification and function approximation. *Advances in neural information processing systems*, 6, 1993.
- [24] Andreas Maurer and Massimiliano Pontil. Empirical Bernstein bounds and sample-variance penalization. In *COLT 2009 - The 22nd Conference on Learning Theory, Montreal, Quebec, Canada, June 18-21, 2009*, 2009.
- [25] Volodymyr Mnih, Csaba Szepesvári, and Jean-Yves Audibert. Empirical Bernstein stopping. In *Proceedings of the 25th international conference on Machine learning*, pages 672–679, 2008.
- [26] Andrew W Moore and Mary S Lee. Efficient algorithms for minimizing cross validation error. In *Machine Learning Proceedings 1994*, pages 190–198. Elsevier, 1994.
- [27] Pierre Perrault, Michal Valko, and Vianney Perchet. Covariance-adapting algorithm for semi-bandits with application to sparse outcomes. In *Conference on Learning Theory*, pages 3152–3184. PMLR, 2020.
- [28] El Mehdi Saad and Gilles Blanchard. Fast rates for prediction with limited expert advice. *Advances in Neural Information Processing Systems*, 34, 2021.
- [29] Alexandre B Tsybakov. Optimal rates of aggregation. In *Learning theory and kernel machines*, pages 303–313. Springer, 2003.
- [30] Nicolas Verzelen. Minimax risks for sparse regressions: Ultra-high dimensional phenomena. *Electronic Journal of Statistics*, 6:38 – 90, 2012.
- [31] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.

Supplementary material for: Covariance adaptive best arm identification.

A Notation

- Let $\mathbf{X} = (X_1, \dots, X_K)$ denote the vector of variables associated to the arms.
- Let $\boldsymbol{\mu} = (\mu_i)_{i \in \llbracket K \rrbracket}$ denote the vector of arms' means.
- For each $i, j \in \llbracket K \rrbracket$, let $V_{ij} = \text{Var}(X_i - X_j)$.
- For each round $t \geq 1$ let $\mathbf{X}_t = (X_{1,t}, \dots, X_{K,t})$ denote the rewards sampled by the environment at round t .
- Let $\hat{\mu}_{i,t}$ denote empirical mean of samples pulled from arm i up to round t :

$$\hat{\mu}_{i,t} := \frac{1}{t} \sum_{s=1}^t X_{i,s}.$$

Denote $\hat{\boldsymbol{\mu}}_t = (\hat{\mu}_{1,t}, \dots, \hat{\mu}_{K,t})$.

- Let $(Y_t)_t$ denote a sequence of random variables distributed following Y :

$$\widehat{V}_t(Y) = \frac{1}{t(t-1)} \sum_{1 \leq u < v \leq t} (Y_u - Y_v)^2.$$

- For $i, j \in \llbracket K \rrbracket$, define the empirical variance for $(X_i - X_j)$ as follows:

$$\widehat{V}_{ij,t} := \frac{1}{t(t-1)} \sum_{1 \leq u < v \leq t} ((X_{i,u} - X_{j,u}) - (X_{i,v} - X_{j,v}))^2. \quad (10)$$

- Define $\delta_t := \delta / (2K^2 t(t+1))$ and $\alpha(t, \delta) := \sqrt{\frac{\log(\delta_t^{-1})}{t-1}}$.
- For bounded variables, define:

$$\widehat{\Delta}_{ij}(t, \delta) := \hat{\mu}_{j,t} - \hat{\mu}_{i,t} - \frac{3}{2} \alpha(t, \delta) \sqrt{2\widehat{V}_{ij,t} - 9\alpha^2(t, \delta)}.$$

- For Gaussian variables, define:

$$\widehat{\Delta}'_{ij}(t, \delta) := \hat{\mu}_{j,t} - \hat{\mu}_{i,t} - \frac{3}{2} \alpha(t, \delta) \sqrt{2f(\alpha(t, \delta)) \widehat{V}_{ij,t}}.$$

- For bounded variables, define:

$$\widehat{\Gamma}_i(w, t, \delta) := \langle w, \hat{\boldsymbol{\mu}}_t \rangle - \hat{\mu}_{i,t} - 2\sqrt{2K\widehat{V}_t(X_i, \langle w, \mathbf{X} \rangle)} \alpha(t, \delta) - 14K \|w - e_i\|_1 \alpha^2(t, \delta),$$

where $w \in \mathbf{B}_1^+$ and $\mathbf{B}_1^+ := \{w, w \in [0, 1]^K \text{ and } \|w\|_1 = 1\}$.

- Define for $S \subseteq \llbracket K \rrbracket$ and $t \geq 1$

$$\mathbf{B}_1^+(S) := \{w \in \mathbf{B}_1^+ \text{ such that: } \text{supp}(w) \subseteq S\}. \quad (11)$$

- For bounded arms problem, define for $i, j \in \llbracket K \rrbracket$:

$$\Lambda_{ij} := \begin{cases} +\infty & \text{if } \mu_j \leq \mu_i \\ \frac{V_{ij}}{(\mu_j - \mu_i)^2} + \frac{3}{\mu_j - \mu_i} & \text{otherwise} \end{cases}$$

- For Gaussian arms problem, define for $i, j \in \llbracket K \rrbracket$:

$$\Lambda'_{ij} := \begin{cases} +\infty & \text{if } \mu_j \leq \mu_i \\ \frac{V_{ij}}{(\mu_j - \mu_i)^2} & \text{otherwise} \end{cases}$$

- Observe the quantities Λ_{ij} and Λ'_{ij} defined in the previous sections are slightly different from the above quantities.
- For bounded arms problem, define for $i \in \llbracket K \rrbracket$ and $w \in \mathbf{B}_1^+$:

$$\Xi_i(w) := \begin{cases} +\infty & \text{if } \langle w, \boldsymbol{\mu} \rangle \leq \mu_i \\ \max \left\{ \frac{\text{Var}(\langle \mathbf{X}, w \rangle - X_i)}{(\langle w, \boldsymbol{\mu} \rangle - \mu_i)^2}, \frac{3\|w - e_i\|_1}{\langle w, \boldsymbol{\mu} \rangle - \mu_i} \right\} & \text{otherwise} \end{cases}$$

- Notation for Algorithms 3 and 2: In round t , let S_t denote the set of candidate arms and C_t the set of queried arms at round t .

B Additional Results for bounded variables: comparison to convex combination

In this section we consider that arms distributions are bounded by 1. We adopt the same notation introduced in the main body and add the following: Let $\mathbf{B}_1^+ \subset \mathbb{R}^K$ denote the set of vectors w with non-negative entries such that $\|w\|_1 = 1$. Let $\mathbf{X} := (X_1, \dots, X_K)$ and $\langle \cdot, \cdot \rangle$ denote the euclidean scalar product in \mathbb{R}^K . For a subset $A \subset \llbracket K \rrbracket$, we denote by $\mathbf{B}_1^+(A)$ the set of elements in \mathbf{B}_1^+ with support in A .

While in the previous sections the main idea of the presented procedures is to perform pairwise comparisons between arms, we consider here that for some classes of covariance matrices between the arms, it may be beneficial to perform sequential tests comparing the candidate arms with convex combinations of the non-eliminated arms. For example, for a sub-optimal arm i , it is possible to have for some weights vector w , with support in $\llbracket K \rrbracket \setminus \{i\}$: $\text{Var}(X_i - \langle w, \mathbf{X} \rangle) \ll \min_{j \in \llbracket K \rrbracket} \text{Var}(X_i - X_j)$. In this case, it is advantageous to eliminate arm i through a comparison with the combination $\langle w, \mathbf{X} \rangle$ instead of pairwise comparisons, as concluding that $\mathbb{E}[X_i - \langle w, \mathbf{X} \rangle] < 0$ for some $w \in \mathbf{B}_1^+$ signifies that arm i is sub-optimal.

The approach used in this section shares similarities with the preceding methodology. More precisely, we develop an empirical second-order concentration inequality over the differences $(X_i - \langle w, \mathbf{X} \rangle)$ for $i \in \llbracket K \rrbracket$ and $w \in \mathbf{B}_1^+$, based on empirical Bernstein inequality and a covering argument over \mathbf{B}_1^+ . We define the following quantity: for $i \in \llbracket K \rrbracket$ and $w \in \mathbf{B}_1^+$.

$$\hat{\Gamma}_i(w, t, \delta) := \langle w, \hat{\boldsymbol{\mu}}_t \rangle - \hat{\mu}_{i,t} - 2\sqrt{2K\hat{V}_t(X_i - \langle w, \mathbf{X} \rangle)\alpha(t, \delta)} - 14K\|w - e_i\|_1 \alpha^2(t, \delta).$$

Lemma C.2 shows that if $\hat{\Gamma}_i(w, t, \delta) > 0$, then $\mathbb{E}[\langle w, \mathbf{X} \rangle] > \mu_i$.

Algorithm 3 Convex-BAI

- 1: **Input** δ .
 - 2: **Initialization:**
 - 3: Query all arms for 2 rounds and compute the empirical means vector $\hat{\boldsymbol{\mu}}_t$, $t \leftarrow 3$.
 - 4: $S \leftarrow \llbracket K \rrbracket$, /*Set of candidate arms*/
 - 5: $C \leftarrow \llbracket K \rrbracket$, /*Set of queried arms*/
 - 6: **while** $|S| > 1$ **do**
 - 7: Jointly query all the experts in C .
 - 8: Update $\hat{\boldsymbol{\mu}}_t$ and compute $\sup_{w \in \mathbf{B}_1^+(C)} \hat{\Gamma}_i(w, t, \delta)$ for each $i \in S$.
 - 9: **for** $i \in S$ **do**
 - 10: **if** $\sup_{w \in \mathbf{B}_1^+(C)} \hat{\Gamma}_i(w, t, \delta) > 0$ **then**
 - 11: Eliminate i from S : $S \leftarrow S \setminus \{i\}$. /* i is sub-optimal */
 - 12: Activate a counter to eliminate i from C at round: $98t$.
 - 13: **end if**
 - 14: **end for**
 - 15: **Increment** t .
 - 16: **end while**
 - 17: **Return** S .
-

Remarks. In Algorithm 3, we did not specify a method to perform the test: $\sup_{w \in \mathbf{B}_1^+(C_t)} \hat{\Gamma}_i(w, t, \delta) > 0$. Several developments can be envisioned, such that using methods for convex optimization over a simplex.

Finally, Theorem B.1 below provides guarantees on the strategy of Algorithm 3. Where tests are performed for each expert against convex combination of all arms.

Theorem B.1. *Suppose Assumption 2 – \mathcal{B} hold. Consider Algorithm 3, with input $\delta \in (0, 1)$. With probability at least $1 - \delta$: the procedure selects the best arm i^* , and the total number of queries N satisfies:*

$$N \leq c \log(K\Lambda\delta^{-1}) K \sum_{i \in \llbracket K \rrbracket \setminus \{i^*\}} \min_{\substack{w \in \mathbf{B}_1^+ \\ \langle w, \mu \rangle > \mu_i}} \left\{ \frac{\text{Var}(X_i - \langle w, \mathbf{X} \rangle)}{(\langle w, \mu \rangle - \mu_i)^2} + \frac{\|w - e_i\|_1}{\langle w, \mu \rangle - \mu_i} \right\}.$$

C Concentration lemmas for bounded variables

Define the event (\mathcal{B}_1) for pairwise comparisons:

$$\forall t \geq 2, \forall i, j \in \llbracket K \rrbracket :$$

$$\left\{ \begin{array}{l} |(\hat{\mu}_{i,t} - \hat{\mu}_{j,t}) - (\mu_i - \mu_j)| \leq \alpha(t, \delta) \sqrt{2\hat{V}_{ij,t}} + 6\alpha^2(t, \delta) \\ \left| \sqrt{\hat{V}_{ij,t}} - \sqrt{V_{ij}} \right| \leq 2\sqrt{2}\alpha(t, \delta), \end{array} \right. \quad (12a)$$

$$\left| \sqrt{\hat{V}_{ij,t}} - \sqrt{V_{ij}} \right| \leq 2\sqrt{2}\alpha(t, \delta), \quad (12b)$$

where $\hat{V}_{ij,t}$ is the empirical variance of the sequence $(X_{i,t} - X_{j,t})_{t \geq 1}$, $V_{ij} = \text{Var}(X_i - X_j)$, $\alpha(t, \delta) = \sqrt{\log(\delta_t^{-1})/(t-1)}$ and $\delta_t = 2K^2\delta/(t(t+1))$.

Define the event (\mathcal{B}_2) for comparisons with convex combinations:

$$\forall t \geq 2, \forall i \in \llbracket K \rrbracket, \forall w \in \mathbf{B}_1^+(C_t):$$

$$\left\{ \begin{array}{l} |(\langle w, \hat{\mu}_t \rangle - \hat{\mu}_{i,t}) - (\langle w, \mu \rangle - \mu_i)| \leq \sqrt{2K\hat{V}_t(X_i - \langle w, \mathbf{X} \rangle)} \alpha(t, \delta) + 7K\|w - e_i\|_1 \alpha^2(t, \delta) \\ \left| \sqrt{\hat{V}_t(X_i - \langle w, \mathbf{X} \rangle)} - \sqrt{\text{Var}(X_i - \langle w, \mathbf{X} \rangle)} \right| \leq 3\sqrt{K}\|w - e_i\|_1 \alpha(t, \delta), \end{array} \right. \quad (13a)$$

$$\left| \sqrt{\hat{V}_t(X_i - \langle w, \mathbf{X} \rangle)} - \sqrt{\text{Var}(X_i - \langle w, \mathbf{X} \rangle)} \right| \leq 3\sqrt{K}\|w - e_i\|_1 \alpha(t, \delta), \quad (13b)$$

where $\mathbf{B}_1^+(C_t)$ is defined in Section A (11).

We show that events (\mathcal{B}_1) and (\mathcal{B}_2) , defined in (13a), (13b) and (12a), (12b) respectively, hold with high probability.

Lemma C.1. *We have $\mathbb{P}(\mathcal{B}_1) \geq 1 - 3\delta$.*

Proof. The first inequality is a direct consequence of empirical Bernstein's inequality (Theorem K.1) applied to the sequence of i.i.d variables $(X_{i,s} - X_{j,s})_{s \leq t}$, and using a union bound over $i, j \in \llbracket K \rrbracket$ and $t \geq 2$. The second inequality of event (\mathcal{B}_1) is a direct consequence of Theorem K.2. \square

Lemma C.2. *We have $\mathbb{P}(\mathcal{B}_2) \geq 1 - 4\delta$.*

Proof. Let \mathbf{B}_1 denote the unit ball with respect to $\|\cdot\|_1$ in \mathbb{R}^K . We will show that: $\forall t \geq 2, \forall i \in \llbracket K \rrbracket, \forall u \in \mathbf{B}_1$:

$$\left| \langle u, \hat{\mu}_t - \mu \rangle \right| \leq \sqrt{2K\hat{V}_t(\langle u, \mathbf{X} \rangle)} \alpha(t, \delta) + 7K\alpha^2(t, \delta) \quad (14a)$$

$$\left| \sqrt{\hat{V}_t(\langle u, \mathbf{X} \rangle)} - \sqrt{\text{Var}(\langle u, \mathbf{X} \rangle)} \right| \leq 3\sqrt{K} \alpha(t, \delta), \quad (14b)$$

The result follows by taking $u = (w - e_i)/\|w - e_i\|_1$ or $u = e_i$.

We use a standard covering argument. Recall that the ϵ -covering number for \mathbf{B}_1 , with respect to $\|\cdot\|_1$, is upper bounded by $(3/\epsilon)^K$ (Lemma 5.7 in 31).

Fix $\delta \in (0, 1)$. For each $i \in \llbracket K \rrbracket$, $t \geq 1$, let $\epsilon_t > 0$ be a parameter to be specified later. Let \mathcal{N}_t be an ϵ_t -cover of the set of \mathbf{B}_1 , with respect to $\|\cdot\|_1$. We will first prove that the event defined in the beginning of the proof is true for all $u \in \mathcal{N}_t$, then using the triangle inequality, we will prove the inequality for any $u \in \mathbf{B}_1$.

Let $i \in \llbracket K \rrbracket$ and $u \in \mathcal{N}_t$. Applying Theorem K.1 to the sequence of i.i.d variables $(\langle u, \mathbf{X}_s \rangle)_{s \leq t}$ bounded by 1, we have with probability at least $1 - \delta$,

$$|\langle u, \hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu} \rangle| \leq \sqrt{\frac{2 \log(3\delta^{-1}) \hat{V}_t(\langle u, \mathbf{X} \rangle)}{t}} + 6 \frac{\log(3\delta^{-1})}{t},$$

where $\hat{V}_t(\langle u, \mathbf{X} \rangle)$ denotes the empirical variance of $(\langle u, \mathbf{X} \rangle)$ at round t . Using a union bound over $t \geq 1$, $i \in \llbracket K \rrbracket$ and $u \in \mathcal{N}_t$, we have with probability at least $1 - \delta$: $\forall t \geq 1, i \in \llbracket K \rrbracket, u \in \mathcal{N}_t$:

$$\begin{aligned} |\langle u, \hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu} \rangle| &\leq \sqrt{2} \sqrt{\frac{\log(|\mathcal{N}_t| \delta_t^{-1})}{t-1} \hat{V}_t(\langle u, \mathbf{X} \rangle)} + 6 \frac{\log(|\mathcal{N}_t| \delta_t^{-1})}{t-1} \\ &\leq \alpha(t, \delta/|\mathcal{N}_t|) \sqrt{2 \hat{V}_t(\langle u, \mathbf{X} \rangle)} + 6\alpha^2(t, \delta/|\mathcal{N}_t|), \end{aligned} \quad (15)$$

where $\delta_t = \delta/(2K^2t(t+1))$.

Applying Theorem K.2 to the sequence $(\langle u, \mathbf{X} \rangle)$ at round t , we have with probability at least $1 - 2\delta$:

$$\left| \sqrt{\hat{V}_t(\langle u, \mathbf{X} \rangle)} - \sqrt{\text{Var}(\langle u, \mathbf{X} \rangle)} \right| \leq 2 \sqrt{\frac{2 \log(\delta^{-1})}{t-1}}.$$

Now, we use a union bound over $t \geq 1$, $i \in \llbracket K \rrbracket$ and $u \in \mathcal{N}_t$ to obtain with probability at least $1 - \delta$: $\forall t \geq 1, i \in \llbracket K \rrbracket, u \in \mathcal{N}_t$

$$\left| \sqrt{\hat{V}_t(\langle u, \mathbf{X} \rangle)} - \sqrt{\text{Var}(\langle u, \mathbf{X} \rangle)} \right| \leq 2\sqrt{2}\alpha(t, \delta/|\mathcal{N}_t|). \quad (16)$$

To wrap up, fix $t \geq 1$ and let $u' \in \mathbf{B}_1$. Since \mathcal{N}_t is a covering for \mathbf{B}_1 , we have: $\exists u \in \mathcal{N}_t$ such that $\|u' - u\|_1 \leq \epsilon_t$.

Hence

$$\begin{aligned} |\langle u', \hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu} \rangle| &\leq |\langle u, \hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu} \rangle| + |\langle u - u', \hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu} \rangle| \\ &\leq \sqrt{2 \hat{V}_t(\langle u, \mathbf{X} \rangle)} \alpha(t, \delta/|\mathcal{N}_t|) + 6\alpha^2(t, \delta/|\mathcal{N}_t|) + \epsilon_t, \end{aligned}$$

where we used (15) and $\|u - u'\|_1 \leq \epsilon_t$. Moreover, we have

$$\begin{aligned} \sqrt{\hat{V}_t(\langle u, \mathbf{X} \rangle)} &\leq \sqrt{\hat{V}_t(\langle u', \mathbf{X} \rangle)} + \sqrt{\hat{V}_t(\langle u - u', \mathbf{X} \rangle)} \\ &\leq \sqrt{\hat{V}_t(\langle u', \mathbf{X} \rangle)} + \|u - u'\|_1 \\ &\leq \sqrt{\hat{V}_t(\langle u', \mathbf{X} \rangle)} + \epsilon_t. \end{aligned}$$

Therefore

$$\begin{aligned} |\langle u', \hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu} \rangle| &\leq \sqrt{2 \hat{V}_t(\langle u', \mathbf{X} \rangle)} \alpha(t, \delta/|\mathcal{N}_t|) \\ &\quad + 6\alpha^2(t, \delta/|\mathcal{N}_t|) + \epsilon_t \left(1 + \sqrt{2}\alpha(t, \delta/|\mathcal{N}_t|)\right). \end{aligned} \quad (17)$$

We choose $\epsilon_t = 3\delta_t^{\frac{K-1}{K}} = 3\left(\frac{\delta}{2K^2t(t+1)}\right)^{\frac{K-1}{K}}$, therefore

$$\begin{aligned}\alpha(t, \delta/|\mathcal{N}_t|) &= \sqrt{\frac{\log((3/\epsilon_t)^K \delta_t^{-1})}{t-1}} \\ &= \sqrt{\frac{\log(\delta_t^{-(K-1)} \delta_t^{-1})}{t-1}} \\ &= \sqrt{K} \alpha(t, \delta).\end{aligned}\tag{18}$$

Furthermore, we have

$$\begin{aligned}\epsilon_t \left(1 + \sqrt{2}\alpha(t, \delta/|\mathcal{N}_t|)\right) &\leq 3(\delta_t)^{\frac{K-1}{K}} \left(1 + \sqrt{K} \alpha(t, \delta)\right) \\ &\leq 3(\delta_t)^{1/2} \left(1 + \sqrt{K} \alpha(t, \delta)\right) \\ &\leq 3\delta^{1/2} \left(\frac{1}{K \log(\delta_t^{-1})} + \frac{2\sqrt{K}}{K \sqrt{(t-1) \log(\delta_t^{-1})}}\right) \frac{\log(\delta_t^{-1})}{t-1} \\ &\leq \alpha^2(t, \delta).\end{aligned}\tag{19}$$

We plug (18) and (19) into (17), and obtain that with probability at least $1 - \delta$

$$|\langle u', \hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu} \rangle| \leq \sqrt{2K \widehat{V}_t(\langle u', \mathbf{X} \rangle)} \alpha(t, \delta) + 7K \alpha^2(t, \delta).\tag{20}$$

We proceed similarly for the second concentration inequality. We have with probability at least $1 - \delta$

$$\begin{aligned}\left| \sqrt{\widehat{V}_t(\langle u', \mathbf{X} \rangle)} - \sqrt{\text{Var}(\langle u', \mathbf{X} \rangle)} \right| &\leq \left| \sqrt{\widehat{V}_t(\langle u, \mathbf{X} \rangle)} - \sqrt{\text{Var}(\langle u, \mathbf{X} \rangle)} \right| + \epsilon_t \\ &\leq 2\sqrt{2}\alpha(t, \delta/|\mathcal{N}_t|) + \epsilon_t \\ &\leq 3\sqrt{K}\alpha(t, \delta).\end{aligned}\tag{21}$$

We conclude by combining (20) and (21). \square

D Concentration lemmas for Gaussian variables

Recall the definition $\alpha(t, \delta) = \sqrt{\log(2K^2t(t+1)\delta^{-1})/(t-1)}$. Define the function f for positive numbers as follows:

$$f(x) := \begin{cases} \exp(2x+1) & \text{if } x \geq 1/3 \\ \frac{1}{1-2x} & \text{otherwise,} \end{cases}$$

Define the event (\mathcal{A}_1) : $\forall t \geq 2, \forall i, j \in \llbracket K \rrbracket$:

$$\left\{ \begin{aligned} |(\hat{\mu}_{i,t} - \hat{\mu}_{j,t}) - (\mu_i - \mu_j)| &\leq \sqrt{2f(\alpha(t, \delta)) \widehat{V}_{i,j,t}} \alpha(t, \delta) \end{aligned} \right. \tag{22a}$$

$$\left\{ \begin{aligned} \widehat{V}_{i,j,t} &\leq V_{ij} (1 + 2\alpha(t, \delta) + 2\alpha^2(t, \delta)) \end{aligned} \right. \tag{22b}$$

$$\left\{ \begin{aligned} V_{ij} &\leq \widehat{V}_{i,j,t} f(\alpha(t, \delta)). \end{aligned} \right. \tag{22c}$$

Lemma below shows that event (\mathcal{A}_1) defined above holds with high probability.

Lemma D.1. *We have $\mathbb{P}(\mathcal{A}_1) \geq 1 - 4\delta$.*

Proof. We start by proving inequalities (22b) and (22c). We use Lemma K.7 with $Y_t = (X_{i,t} - X_{j,t} - (\mu_i - \mu_j))/V_{ij}$. A union bound over $t \geq 2$ and $i, j \in \llbracket K \rrbracket$ gives with probability at least $1 - \delta$

$$\widehat{V}_{i,j,t} \leq V_{ij} (1 + 2\alpha(t, \delta) + 2\alpha^2(t, \delta)).$$

For inequality (22c), we apply the first result of Lemma K.7 to the (Y_t) defined above. Using a union bound, we have with probability at least $1 - \delta$

$$\widehat{V}_{ij,t} \geq V_{ij} \max\{1 - 2\alpha(t, \delta); e^{-1} \exp(-2\alpha^2(t, \delta))\}.$$

Inverting the inequality above leads to (22c).

To prove (22a), we use Chernoff's concentration bound for Gaussian variables (and a union bound), with probability at least $1 - \delta$: for all $t \geq 2$ and $i, j \in \llbracket K \rrbracket$:

$$|(\hat{\mu}_{i,t} - \hat{\mu}_{j,t}) - (\mu_i - \mu_j)| \leq \sqrt{2\widehat{V}_{ij,t}}\alpha(t, \delta). \quad (23)$$

We plug-in the bound (22c) to obtain the result. □

E Key lemmas

Lemma E.1. *If (\mathcal{B}_1) defined in (12) holds, we have the following:*

For any $i \in \llbracket K \rrbracket$, if there exists $t \geq 1$ and $j \in \llbracket K \rrbracket$ such that $\widehat{\Delta}_{ij}(t, \delta) > 0$, then $i \neq i^$.*

Moreover, if (\mathcal{A}_1) defined in (12) holds, we have the following:

For any $i \in \llbracket K \rrbracket$, if there exists $t \geq 2$ and $j \in \llbracket K \rrbracket$ such that $\widehat{\Delta}'_{ij}(t, \delta) > 0$, then $i \neq i^$.*

Proof. Suppose that (\mathcal{B}_1) is true. Let $t \geq 2$, $i, j \in \llbracket K \rrbracket$. We have

$$\begin{aligned} \mu_j - \mu_i &= \widehat{\Delta}_{ij}(t, \delta) + \mu_j - \mu_i - (\hat{\mu}_{j,t} - \hat{\mu}_{i,t}) + \frac{3}{2}\sqrt{2\widehat{V}_{ij,t}}\alpha(t, \delta) + 9\alpha^2(t, \delta) \\ &\geq \widehat{\Delta}_{ij}(t, \delta), \end{aligned}$$

where we used (22a). Finally, if $\widehat{\Delta}_{ij}(t, \delta) > 0$, we have $\mu_j > \mu_i$.

Following the exact same steps we have the result for Gaussian variables (the second claim). □

Lemma E.2. *If (\mathcal{B}_2) defined in (13) holds, we have the following:*

For any $i \in \llbracket K \rrbracket$, if there exists $t \geq 1$ and $w \in \mathbf{B}_1^+(C_t)$ such that: $\widehat{\Gamma}_i(w, t, \delta) > 0$, then $i \neq i^$.*

Proof. Suppose that (\mathcal{B}_2) is true. Let $t \geq 1$, $i \in \llbracket K \rrbracket$ and $w \in \mathbf{B}_1^+(C_t)$. We have

$$\begin{aligned} \langle w, \boldsymbol{\mu} \rangle - \mu_i &= \widehat{\Gamma}_i(w, t, \delta) + \langle w, \boldsymbol{\mu} \rangle - \mu_i - (\langle w, \hat{\boldsymbol{\mu}}_t \rangle - \hat{\mu}_{i,t}) + 2\sqrt{2K\widehat{V}_t(X_i - \langle w, \mathbf{X} \rangle)}\alpha(t, \delta) \\ &\quad + 14K\|w - e_i\|_1\alpha^2(t, \delta) \\ &\geq \widehat{\Gamma}_i(w, t, \delta), \end{aligned}$$

where we used (13a). If $\widehat{\Gamma}_i(w, t, \delta) > 0$, we have $\langle w, \boldsymbol{\mu} \rangle > \mu_i$. Since w is a vector of convex weights, we conclude that $\max_{j \in \text{supp}(w)} \mu_j \geq \langle w, \boldsymbol{\mu} \rangle > \mu_i$. □

Lemma E.3. *If (\mathcal{B}_1) defined in (13) holds, then for any $t \geq 1$, $i, j \in C_t$: such that $\mu_j > \mu_i$*

If $\widehat{\Delta}_{ij}(t, \delta) > 0$, then

$$t - 1 \geq \frac{1}{4} \log(\delta_t^{-1}) \Lambda_{ij}.$$

Furthermore, if $\widehat{\Delta}_{ij}(t, \delta) \leq 0$, then

$$t - 1 \leq \frac{25}{2} \log(\delta_t^{-1}) \Lambda_{ij}.$$

Proof. Suppose that (\mathcal{B}_1) is true. Let $t \geq 1, i, j \in \llbracket K \rrbracket$. Suppose that $\hat{\Delta}_{ij}(t, \delta) > 0$. We have

$$\begin{aligned}
\mu_j - \mu_i &= \hat{\Delta}_{ij}(t, \delta) - (\hat{\mu}_{j,t} - \hat{\mu}_{i,t}) + \mu_j - \mu_i + \frac{3}{2} \sqrt{2\widehat{V}_{ij,t}} \alpha(t, \delta) + 9\alpha^2(t, \delta) \\
&\geq \hat{\Delta}_{ij}(t, \delta) - \sqrt{2V_{ij}} \alpha(t, \delta) - \frac{4}{3} \alpha^2(t, \delta) + \frac{3}{2} \sqrt{2\widehat{V}_{ij,t}} \alpha(t, \delta) + 9\alpha^2(t, \delta) \\
&\geq \hat{\Delta}_{ij}(t, \delta) + \sqrt{\frac{V_{ij}}{2}} \alpha(t, \delta) + \frac{3}{2} \sqrt{2\widehat{V}_{ij,t}} \alpha(t, \delta) - \frac{3}{2} \sqrt{2V_{ij}} \alpha(t, \delta) + \frac{23}{3} \alpha^2(t, \delta) \\
&> \sqrt{\frac{V_{ij}}{2}} \alpha(t, \delta) + \frac{5}{3} \alpha^2(t, \delta), \tag{24}
\end{aligned}$$

where we used Bennett's inequality (Theorem 3 in [24]) in the second line, (12b) with $\hat{\Delta}_{ij}(t, \delta) > 0$ in the third line.

Solving inequality(24), gives

$$\begin{aligned}
\alpha(t, \delta) &\leq \frac{3}{14} \left(\sqrt{\frac{V_{ij}}{2} + \frac{20}{3}(\mu_j - \mu_i)} - \sqrt{\frac{V_{ij}}{2}} \right) \\
&= \frac{2(\mu_j - \mu_i)}{\sqrt{\frac{V_{ij}}{2} + \frac{20}{3}(\mu_j - \mu_i)} + \sqrt{\frac{V_{ij}}{2}}}.
\end{aligned}$$

Therefore, we have

$$\begin{aligned}
t - 1 &\geq \log(\delta_t^{-1}) \left(\frac{V_{ij}}{4(\mu_j - \mu_i)^2} + \frac{5/3}{\mu_j - \mu_i} \right) \\
&\geq \frac{1}{4} \log(\delta_t^{-1}) \Lambda_{ij}.
\end{aligned}$$

Which gives the first result.

For the second bound, we proceed similarly. Suppose that $\hat{\Delta}_{ij}(t, \delta) \leq 0$, we have:

$$\begin{aligned}
\mu_j - \mu_i &= \hat{\Delta}_{ij}(t, \delta_t) - (\hat{\mu}_{j,t} - \hat{\mu}_{i,t}) + \mu_j - \mu_i + \frac{3}{2} \sqrt{2\widehat{V}_{ij,t}} \alpha(t, \delta) + 9\alpha^2(t, \delta) \\
&\leq \hat{\Delta}_{ij}(t, \delta_t) + \sqrt{2V_{ij}} \alpha(t, \delta) + \frac{4}{3} \alpha^2(t, \delta) + \frac{3}{2} \sqrt{2\widehat{V}_{ij,t}} \alpha(t, \delta) + 9\alpha^2(t, \delta) \tag{25}
\end{aligned}$$

$$\begin{aligned}
&\leq \hat{\Delta}_{ij}(t, \delta_t) + \frac{5}{2} \sqrt{2V_{ij}} \alpha(t, \delta) + \frac{49}{3} \alpha^2(t, \delta) \\
&\leq \frac{5}{2} \sqrt{2V_{ij}} \alpha(t, \delta) + \frac{49}{3} \alpha^2(t, \delta), \tag{26}
\end{aligned}$$

Similarly to the previous case, we have:

$$\begin{aligned}
\alpha(t, \delta) &\geq \frac{\sqrt{\frac{25}{2} V_{ij} + \frac{196}{3}(\mu_j - \mu_i)} - \sqrt{\frac{25}{2} V_{ij}}}{98/3} \\
&= \frac{2(\mu_j - \mu_i)}{\sqrt{\frac{25}{2} V_{ij} + \frac{196}{3}(\mu_j - \mu_i)} + \sqrt{\frac{25}{2} V_{ij}}},
\end{aligned}$$

which gives

$$t - 1 \leq \log(\delta_t^{-1}) \left(\frac{25}{2} \frac{V_{ij}}{(\mu_i - \mu_j)^2} + \frac{98/3}{\mu_j - \mu_i} \right)$$

We conclude by inverting (26) leading to:

$$t - 1 \leq \frac{25}{2} \log(\delta_t^{-1}) \Lambda_{ij}.$$

□

Lemma E.4. If (\mathcal{A}_1) defined in (13) holds, then for any $t \geq 2$, $i, j \in C_t$, such that $\mu_j > \mu_i$

If $\hat{\Delta}'_{ij}(t, \delta) > 0$, then

$$t - 1 \geq \frac{1}{2} \log(1/\delta_t) \Lambda'_{ij}.$$

Furthermore, if $\hat{\Delta}'_{ij}(t, \delta) \leq 0$, then

$$t - 1 \leq \frac{25 \log(1/\delta_t)}{\log(1 + 1/\Lambda'_{ij})}.$$

Moreover, if $\hat{\Delta}'_{ij}(t, \delta) \leq 0$, then: if $\Lambda'_{ij} < \frac{1}{5}$

$$t - 1 \leq \frac{3 \log(1/\delta_t)}{2 \log\left(\frac{\sqrt{2}}{5\Lambda'_{ij}}\right)},$$

while if $\Lambda'_{ij} \geq \frac{1}{5}$

$$t - 1 \leq 25 \log(1/\delta_t) \Lambda'_{ij}.$$

Proof. Suppose that (\mathcal{A}_1) is true. Let $t \geq 2$, $i, j \in \llbracket K \rrbracket$. Suppose that $\hat{\Delta}'_{ij}(t, \delta_t) > 0$. We have

$$\begin{aligned} \mu_j - \mu_i &= \hat{\Delta}'_{ij}(t, \delta) - (\hat{\mu}_{j,t} - \hat{\mu}_{i,t}) + \mu_j - \mu_i + \frac{3}{2} \sqrt{2f(\alpha(t, \delta)) \widehat{V}_{ij,t}} \alpha(t, \delta) \\ &\geq \hat{\Delta}'_{ij}(t, \delta) + \frac{1}{2} \sqrt{2V_{ij}} \alpha(t, \delta) \\ &> \frac{1}{2} \sqrt{2V_{ij}} \alpha(t, \delta), \end{aligned} \tag{27}$$

where we used (12a) in the second line and (12b) with $\hat{\Delta}'_{ij}(t, \delta_t) > 0$ in the third line.

Therefore inequality(27), gives

$$\alpha(t, \delta) < \frac{\mu_j - \mu_i}{\sqrt{V_{ij}/2}}.$$

Therefore, we have

$$\begin{aligned} t - 1 &\geq \frac{1}{2} \log(1/\delta_t) \frac{V_{ij}}{(\mu_i - \mu_j)^2} \\ &= \frac{1}{2} \log(1/\delta_t) \Lambda'_{ij} \end{aligned}$$

Which gives the first result.

For the second bound, suppose that $\hat{\Delta}'_{ij}(t, \delta) \leq 0$, we have:

$$\begin{aligned} \mu_j - \mu_i &= \hat{\Delta}'_{ij}(t, \delta_t) - (\hat{\mu}_{j,t} - \hat{\mu}_{i,t}) + \mu_j - \mu_i + \frac{3}{2} \sqrt{2f(\alpha(t, \delta)) \widehat{V}_{ij,t}} \alpha(t, \delta) \\ &\leq \hat{\Delta}'_{ij}(t, \delta_t) + \frac{5}{2} \sqrt{2f(\alpha(t, \delta)) \widehat{V}_{ij,t}} \alpha(t, \delta) \\ &\leq \frac{5}{2} \sqrt{V_{ij}} \sqrt{2f(\alpha(t, \delta)) (1 + 2\alpha(t, \delta) + 2\alpha^2(t, \delta))} \alpha(t, \delta), \end{aligned} \tag{28}$$

Hence:

$$f(\alpha(t, \delta)) (\alpha^2(t, \delta) + 2\alpha^3(t, \delta) + 2\alpha^4(t, \delta)) \geq \frac{2(\mu_j - \mu_i)^2}{25V_{ij}} \tag{29}$$

We consider two distinct cases:

Case 1: $|\mu_i - \mu_j| > 5\sqrt{V_{ij}}$. Observe that in this case, inequality (29) implies in particular that:

$$f(\alpha(t, \delta))(\alpha^2(t, \delta) + 2\alpha^3(t, \delta) + 2\alpha^4(t, \delta)) \geq 2.$$

Recall that for a positive number $x < 1/3$, we have $(x^2 + 2x^3 + 2x^4)/(1 - 2x^2) < 2$. Hence, for the latter inequality to hold, we necessarily have that: $\alpha(t, \delta) \geq 1/3$. Therefore, using the definition of f we have $f(\alpha(t, \delta)) = \exp(2\alpha^2(t, \delta) + 1)$, taking the logarithms in inequality (29):

$$2\alpha^2(t, \delta) + \log(\alpha^2(t, \delta) + 2\alpha^3(t, \delta) + 2\alpha^4(t, \delta)) \geq 2 \log\left(\frac{\sqrt{2}|\mu_i - \mu_j|}{5\sqrt{V_{ij}}}\right),$$

Observe that for any $x > 0$, we have $3x^2 > 2x^2 + \log(x^2 + 2x^3 + 2x^4)$. Therefore

$$\alpha^2(t, \delta) > \frac{2}{3} \log\left(\frac{\sqrt{2}|\mu_i - \mu_j|}{5\sqrt{V_{ij}}}\right).$$

We conclude that

$$t - 1 < \frac{3/2 \log(\delta_t^{-1})}{\log\left(\frac{\sqrt{2}|\mu_i - \mu_j|}{5\sqrt{V_{ij}}}\right)} \quad (30)$$

Case 2: $|\mu_i - \mu_j| \leq 5\sqrt{V_{ij}}$. If $\alpha(t, \delta) > 1/3$, then

$$t - 1 < 9 \log(\delta_t^{-1}).$$

Otherwise, if $\alpha(t, \delta) \leq 1/3$, we have using the definition of $f(\alpha(t, \delta))$ and inequality (29)

$$\begin{aligned} \frac{2(\mu_j - \mu_i)^2}{25V_{ij}} &\leq \alpha^2(t, \delta) \frac{1 + 2\alpha(t, \delta) + 2\alpha^2(t, \delta)}{1 - 2\alpha(t, \delta)} \\ &< 2 \frac{\alpha^2(t, \delta)}{1 - 2\alpha(t, \delta)}. \end{aligned}$$

Inverting the inequality above in t , we obtain

$$\alpha(t, \delta) \geq \sqrt{\frac{(\mu_j - \mu_i)^2}{25V_{ij}} \left(\frac{(\mu_j - \mu_i)^2}{25V_{ij}} + 1 \right)} - \frac{(\mu_j - \mu_i)^2}{25V_{ij}}.$$

Therefore, we have

$$\alpha^2(t, \delta) \geq \frac{(\mu_j - \mu_i)^2}{25V_{ij}}.$$

We conclude that if $|\mu_i - \mu_j| \leq 5\sqrt{V_{ij}}$:

$$t - 1 < 25 \log(1/\delta_t) \frac{V_{ij}}{(\mu_i - \mu_j)^2}. \quad (31)$$

Now in order to unify the bounds obtained in Cases 1 and 2, observe that the function f defined for positive numbers by

$$f(x) := \begin{cases} \frac{3}{\log(2/(25x^2))} & \text{if } x \leq 1/5 \\ 25x^2 & \text{otherwise,} \end{cases}$$

satisfies for any $0 < x < 1/5$

$$f(x) \leq \frac{10}{\log(1 + \frac{1}{x^2})},$$

and for $x > 1/5$

$$f(x) \leq \frac{25}{\log(1 + \frac{1}{x^2})},$$

Therefore, we conclude that if $\hat{\Delta}'_{ij}(t, \delta) \leq 0$, then

$$t - 1 \leq \frac{25 \log(1/\delta_t)}{\log(1 + 1/\Lambda'_{ij})}.$$

□

Lemma E.5. If (\mathcal{B}_2) defined in (13) holds, then for any $i \in S_t$, $t \geq 1$ and $w \in \mathbf{B}_1^+(C_t)$:

If $\hat{\Gamma}_i(w, t, \delta) > 0$, then

$$t \geq K \log(\delta_t^{-1}) \Xi_i(w).$$

Furthermore, if $\hat{\Gamma}_i(w, t, \delta) \leq 0$, then

$$t \leq 36K \log(\delta_t^{-1}) \Xi_i(w).$$

Proof. Suppose that (\mathcal{B}_2) is true. Let $t \geq 1$, $i \in S_t$ and $w \in \mathbf{B}_1^+(C_t)$. Suppose that $\hat{\Gamma}_i(w, t, \delta) > 0$. We have

$$\begin{aligned} \langle w, \boldsymbol{\mu} \rangle - \mu_i &= \hat{\Gamma}_i(w, t, \delta) - (\langle w, \hat{\boldsymbol{\mu}}_t \rangle - \hat{\mu}_{i,t}) + \langle w, \boldsymbol{\mu} \rangle - \mu_i + 2\sqrt{2K\hat{V}_t(X_i - \langle w, \mathbf{X} \rangle)} \alpha(t, \delta) \\ &\quad + 14K\|w - e_i\|_1 \alpha^2(t, \delta) \\ &\geq \hat{\Gamma}_i(w, t, \delta) - \sqrt{2K\text{Var}(X_i - \langle w, \mathbf{X} \rangle)} - \frac{4}{3}K\alpha^2(t, \delta) + 2\sqrt{2K\hat{V}_t(X_i - \langle w, \mathbf{X} \rangle)} \alpha(t, \delta) \\ &\quad + 14K\|w - e_i\|_1 \alpha^2(t, \delta) \\ &\geq \hat{\Gamma}_i(w, t, \delta) + \sqrt{2K\text{Var}(X_i - \langle w, \mathbf{X} \rangle)} \alpha(t, \delta) + 7K\|w - e_i\|_1 \alpha^2(t, \delta) \\ &> \sqrt{2K\text{Var}(X_i - \langle w, \mathbf{X} \rangle)} \alpha(t, \delta) + 7K\|w - e_i\|_1 \alpha^2(t, \delta), \end{aligned}$$

where we used Bennett's inequality in the second line and (13b) with $\hat{\Gamma}_i(w, t, \delta) > 0$ in the third line and fourth lines.

Solving the inequality above in $\alpha(t, \delta)$, gives

$$\begin{aligned} \alpha(t, \delta) &\leq \frac{\sqrt{2K\text{Var}(X_i - \langle w, \mathbf{X} \rangle)} + 28K\|w - e_i\|_1(\langle w, \boldsymbol{\mu} \rangle - \mu_i) - \sqrt{2K\text{Var}(X_i - \langle w, \mathbf{X} \rangle)}}{14K\|w - e_i\|_1} \\ &= \frac{2(\langle w, \boldsymbol{\mu} \rangle - \mu_i)}{\sqrt{2K\text{Var}(X_i - \langle w, \mathbf{X} \rangle)} + 28K\|w - e_i\|_1(\langle w, \boldsymbol{\mu} \rangle - \mu_i) + \sqrt{2K\text{Var}(X_i - \langle w, \mathbf{X} \rangle)}}. \end{aligned}$$

Therefore, we have

$$\begin{aligned} t &\geq K \log(\delta_t^{-1}) \left(\frac{\text{Var}(X_i - \langle w, \mathbf{X} \rangle)}{(\langle w, \boldsymbol{\mu} \rangle - \mu_i)^2} + \frac{7\|w - e_i\|_1}{\langle w, \boldsymbol{\mu} \rangle - \mu_i} \right) \\ &\geq K \log(\delta_t^{-1}) \Xi_i(w). \end{aligned}$$

Which gives the first result.

Now let us prove the second claim. Suppose that $\hat{\Gamma}_i(w, t, \delta) \leq 0$. We have

$$\begin{aligned} \langle w, \boldsymbol{\mu} \rangle - \mu_i &= \hat{\Gamma}_i(w, t, \delta) - (\langle w, \hat{\boldsymbol{\mu}}_t \rangle - \hat{\mu}_{i,t}) + \langle w, \boldsymbol{\mu} \rangle - \mu_i + 2\sqrt{2K\hat{V}_t(X_i, \langle w, \mathbf{X} \rangle)} \alpha(t, \delta) \\ &\quad + 14K\|w - e_i\|_1 \alpha^2(t, \delta) \\ &\leq \hat{\Gamma}_i(w, t, \delta) + \sqrt{2K\text{Var}(X_i - \langle w, \mathbf{X} \rangle)} \alpha(t, \delta) + 2\sqrt{2K\hat{V}_t(X_i - \langle w, \mathbf{X} \rangle)} \alpha(t, \delta) \\ &\quad + \frac{42K}{3}\|w - e_i\|_1 \alpha^2(t, \delta) \\ &\leq \hat{\Gamma}_i(w, t, \delta) + 3\sqrt{2K\text{Var}(X_i - \langle w, \mathbf{X} \rangle)} \alpha(t, \delta) + 9K\alpha^2(t, \delta) + \frac{42K}{3}\|w - e_i\|_1 \alpha^2(t, \delta) \\ &\leq 3\sqrt{2K\text{Var}(X_i - \langle w, \mathbf{X} \rangle)} \alpha(t, \delta) + 25K\|w - e_i\|_1 \alpha^2(t, \delta), \end{aligned}$$

where we used (13a) in the second line and (13b) with $\hat{\Gamma}_i(w, t, \delta) \leq 0$ in the third line. Suppose that $\langle w, \boldsymbol{\mu} \rangle > \mu_i$. Solving the inequality above in $\alpha(t, \delta)$, gives

$$\begin{aligned} \alpha(t, \delta) &\geq \frac{\sqrt{18K\text{Var}(X_i - \langle w, \mathbf{X} \rangle)} + 100K\|w - e_i\|_1(\langle w, \boldsymbol{\mu} \rangle - \mu_i) - 3\sqrt{2K\text{Var}(X_i - \langle w, \mathbf{X} \rangle)}}{50K\|w - e_i\|_1} \\ &= \frac{2(\langle w, \boldsymbol{\mu} \rangle - \mu_i)}{\sqrt{18K\text{Var}(X_i - \langle w, \mathbf{X} \rangle)} + 100K\|w - e_i\|_1(\langle w, \boldsymbol{\mu} \rangle - \mu_i) + 3\sqrt{2K\text{Var}(X_i - \langle w, \mathbf{X} \rangle)}}. \end{aligned}$$

Therefore, we have

$$\begin{aligned} t &\leq K \log(\delta_t^{-1}) \left(\frac{18 \text{Var}(X_i - \langle w, \mathbf{X} \rangle)}{(\langle w, \boldsymbol{\mu} \rangle - \mu_i)^2} + \frac{50 \|w - e_i\|_1}{\langle w, \boldsymbol{\mu} \rangle - \mu_i} \right) \\ &\leq 36K \log(\delta_t^{-1}) \Xi_i(w). \end{aligned}$$

If $\langle w, \boldsymbol{\mu} \rangle \leq \mu_i$, then $\Xi_i(w) = +\infty$ and the inequality above is straightforward. □

Lemma E.6. *Let i, j and $k \in \llbracket K \rrbracket$, we have:*

$$\begin{aligned} \Lambda'_{ij} &\leq \max\{\Lambda'_{ik}, \Lambda'_{kj}\} \\ \Lambda_{ij} &\leq \max\{\Lambda_{ik}, \Lambda_{kj}\} \end{aligned}$$

Proof. Let i, j and $k \in \llbracket K \rrbracket$. Suppose that $\mu_j \leq \mu_i$. Hence, for any $k \in \llbracket K \rrbracket$, $\mu_k \leq \mu_i$ or $\mu_k \geq \mu_j$. Therefore

$$\max\{\Lambda'_{ik}, \Lambda'_{kj}\} = +\infty,$$

which proves the result. The same argument applies to the quantities Λ_{ik} and Λ_{kj} .

Now suppose that $\mu_j > \mu_i$ (hence $\Lambda'_{ij} < +\infty$ and $\Lambda_{ij} < +\infty$). Let us start by proving the first claim. Recall the definition in Section A:

$$\Lambda'_{ij} := \frac{V_{ij}}{(\mu_j - \mu_i)^2}.$$

We have

$$\begin{aligned} \frac{\sqrt{V_{ij}}}{\mu_j - \mu_i} &\leq \frac{\sqrt{V_{ik}} + \sqrt{V_{kj}}}{(\mu_j - \mu_k) + (\mu_k - \mu_i)} \\ &\leq \max\left\{ \frac{\sqrt{V_{kj}}}{\mu_j - \mu_k}, \frac{\sqrt{V_{ik}}}{\mu_k - \mu_i} \right\}, \end{aligned}$$

where the first line follows by the triangle inequality and the second is a consequence of the inequality $\frac{a+b}{c+d} \leq \max\{\frac{a}{c}, \frac{b}{d}\}$ (Lemma K.3), which proves the first claim.

Moreover, we have using the result above: for any k such that $\mu_i < \mu_k < \mu_j$

$$\begin{aligned} \Lambda_{ij} &= \frac{V_{ij}}{(\mu_j - \mu_i)^2} + \frac{3}{\mu_j - \mu_i} \\ &\leq \max\left\{ \frac{V_{ik}}{(\mu_k - \mu_i)^2} + \frac{3}{\mu_j - \mu_i}, \frac{V_{kj}}{(\mu_j - \mu_k)^2} + \frac{3}{\mu_j - \mu_i} \right\} \\ &\leq \max\left\{ \frac{V_{ik}}{(\mu_k - \mu_i)^2} + \frac{3}{\mu_k - \mu_i}, \frac{V_{kj}}{(\mu_j - \mu_k)^2} + \frac{3}{\mu_j - \mu_k} \right\} \\ &= \max\{\Lambda_{ik}, \Lambda_{kj}\}. \end{aligned}$$

If k is such that $\mu_k < \mu_i$ or $\mu_k > \mu_j$, we have $\max\{\Lambda_{ik}, \Lambda_{kj}\} = +\infty$, which proves the statement. □

F Proof of Theorem 4.1

The proof of Theorem 4.1 follows the same steps as the proof of Theorem 5.1.

For any $i \in \llbracket K \rrbracket \setminus \{i^*\}$, let us define Υ_i by

$$\Upsilon_i := \underset{j \in \llbracket K \rrbracket}{\text{Arg Min}} \Lambda_{ij}. \tag{32}$$

Lemma F.1. *Consider Algorithm 2 with inputs $\delta \in (0, 1)$. If (\mathcal{B}_1) defined in (13) holds, then for any $i \in \llbracket K \rrbracket \setminus \{i^*\}$ and $t \geq 1$:*

If $i \in S_t$, then $\Upsilon_i \cap C_t \neq \emptyset$, where Υ_i is defined in (32).

Proof. Suppose that (\mathcal{B}_1) holds. Let $t \geq 1$, $i \in \llbracket K \rrbracket \setminus \{i^*\}$. Proceeding by proof via contradiction, suppose that $\Upsilon_i \cap C_t = \emptyset$. This implies in particular that all elements in Υ_i were eliminated prior to t . Let j denote the element of Υ_i with the largest mean:

$$j \in \underset{l \in \Upsilon_i}{\text{Arg Max}} \{\mu_l\}.$$

Let s denote the round where j has failed the test (i.e. $\exists k \in C_s, \hat{\Delta}_{jk}(s, \delta) > 0$).

Hence, using Lemma E.3, we have

$$\frac{1}{4} \log(\delta_s^{-1}) \Lambda_{jk} \leq s. \quad (33)$$

Moreover, j was kept for testing up to round $82s$ (i.e. $j \in C_{82s}$) and $82s < t$ (since $j \notin C_t$). At round $82s$ we necessarily had $\hat{\Delta}_{ij}(82s, \delta) \leq 0$.

Therefore, using Lemma E.3

$$82s \leq 25 \log(\delta_{82s}^{-1}) \Lambda_{ij}. \quad (34)$$

Combining (33) and (34) gives

$$\frac{82}{4} \log(\delta_s^{-1}) \Lambda_{jk} \leq 25 \log(\delta_{82s}^{-1}) \Lambda_{ij}.$$

Therefore

$$\Lambda_{jk} \leq \frac{50}{82} \left(1 + \frac{\log(82)}{\log(1/\delta_s)} \right) \Lambda_{ij} \leq \Lambda_{ij}, \quad (35)$$

recall that $\delta < 1/4$, $s \geq 3$, $K \geq 3$, therefore: $\frac{50}{82} \left(1 + \frac{\log(82)}{\log(1/\delta_s)} \right) \leq 1$. Using Lemma E.6, we have

$$\Lambda_{ik} \leq \max\{\Lambda_{ij}, \Lambda_{jk}\}. \quad (36)$$

We plug the bound $\Lambda_{jk} \leq \Lambda_{ij}$ from (35) into (36) and obtain $\Lambda_{ik} \leq \Lambda_{ij}$. Therefore $k \in \Upsilon_i$.

To conclude, recall that k eliminates j , hence $\mu_k > \mu_j$. The contradiction arises from $k \in \Upsilon_i$ and the definition of j as the element with largest mean in Υ_i . \square

We introduce the following notation. For $i \in \llbracket K \rrbracket$ and $t \geq 1$ let $N_{i,t}$ denote the number of queries made for arm i up to round t

$$N_{i,t} := \sum_{s=1}^t \mathbf{1}(i \in C_s). \quad (37)$$

Lemma F.2. Consider Algorithm 2 with inputs $\delta \in (0, 1)$. If (\mathcal{B}_1) defined in (13) holds, then we have for each $i \in \llbracket K \rrbracket \setminus \{i^*\}$:

$$\forall t \geq 1: \quad N_{i,t} \leq 1368 \min_{j \in \llbracket K \rrbracket} \{\Lambda_{ij}\} \log(216K^2 \delta^{-1} \min_{j \in \llbracket K \rrbracket} \{\Lambda_{ij}\}),$$

where $N_{i,t}$ is defined in (37).

Proof. Suppose (\mathcal{B}_1) holds. Let $i \in \llbracket K \rrbracket \setminus \{i^*\}$ and $t \geq 1$. Let u denote the last round such that $i \in S_u$. Lemma F.1 states that $\Upsilon_i \cap C_u \neq \emptyset$, where Υ_i is defined in (45). Let $j \in \Upsilon_i \cap C_u$, since $i \in S_u$, we necessarily have

$$\hat{\Delta}_{ij}(u-1, \delta) \leq 0.$$

Using Lemma E.3, we have

$$u-1 \leq \frac{25}{2} \log(\delta_{u-1}^{-1}) \Lambda_{ij}.$$

Recall that u is the last round such that $i \in S_u$, hence $i \notin C_{19u+1}$. Therefore, for any $t \geq 1$

$$\begin{aligned} N_{i,t} &= 82u \leq 1013 \log(\delta_u^{-1}) \Lambda_{ij} \\ &\leq 2026 \min_{j \in \llbracket K \rrbracket} \{\Lambda_{ij}\} \log(216K^2 \delta^{-1} \min_{j \in \llbracket K \rrbracket} \{\Lambda_{ij}\}), \end{aligned}$$

where we used Lemma K.4 with $x = u$ and $c = \delta/(2K^2)$. \square

Proof for Theorem 4.1. Suppose Assumptions 1 and 2 $-\mathcal{B}$ hold. Consider Algorithm 2 with input $\delta \in (0, 1/4)$. Suppose that event (\mathcal{B}_1) holds.

We have by definition of the total number of queries made N :

$$N = \sum_{i \in \llbracket K \rrbracket \setminus \{i^*\}} N_{i,t}.$$

Therefore, Lemma F.2 gives the result. For the second result, consider Algorithm 2 without line (12) (we stop sampling arms directly after their elimination from S_t). Lemma E.1 guarantees with probability at least $1 - \delta$, the optimal arm i^* always belongs to S_t . Therefore, Lemma E.5 guarantees that after at most $\frac{25}{2} \log(\delta_t^{-1}) \Lambda_{ii^*}$ rounds, we have $\hat{\Delta}_{ii^*}(t, \delta) > 0$, which leads to the elimination of i .

G Proof of Theorem 5.1

For any $i \in \llbracket K \rrbracket \setminus \{i^*\}$, let us define Υ'_i by

$$\Upsilon'_i := \text{Arg Min}_{j \in \llbracket K \rrbracket} \left\{ \Lambda'_{ij} \vee \frac{1}{4} \right\}. \quad (38)$$

Lemma G.1. Consider Algorithm 2 with inputs $\delta \in (0, 1)$. If (\mathcal{A}_1) defined in (13) holds, then for any $i \in \llbracket K \rrbracket \setminus \{i^*\}$ and $t \geq 1$:

If $i \in S_t$, then $\Upsilon'_i \cap C_t \neq \emptyset$, where Υ'_i is defined in (38).

Proof. Suppose that (\mathcal{A}_1) holds. Let $t \geq 2$, $i \in \llbracket K \rrbracket \setminus \{i^*\}$. Proceeding by proof via contradiction, suppose that $\Upsilon'_i \cap C_t = \emptyset$. This implies in particular that all elements in Υ'_i were eliminated prior to t . Let j denote the element of Υ'_i with the largest mean:

$$j \in \text{Arg Max}_{l \in \Upsilon'_i} \{\mu_l\}.$$

Let s denote the round where j has failed the test (i.e. $\exists k \in C_s, \hat{\Delta}'_{jk}(s, \delta) > 0$).

Hence, using Lemma E.4, we have

$$\frac{1}{2} \Lambda'_{jk} \log(1/\delta_s) \leq s - 1. \quad (39)$$

We consider two cases:

Case 1: $\Lambda'_{ij} \geq \frac{1}{5}$. Observe that j was kept for testing up to round $35s$ (i.e. $j \in C_{82s}$) and $82s < t$ (since $j \notin C_t$). At round $82s$ we necessarily had $\hat{\Delta}'_{ij}(82s, \delta) \leq 0$.

Therefore, using Lemma E.4

$$82s - 1 \leq 25 \Lambda'_{ij} \log(1/\delta_{35s}). \quad (40)$$

Combining (39) and (40) gives

$$\frac{82}{2} \Lambda'_{jk} \log(1/\delta_s) \leq 25 \Lambda'_{ij} \log(1/\delta_{82s}).$$

Therefore

$$\Lambda'_{jk} \leq \frac{50}{82} \left(1 + \frac{\log(82)}{\log(\delta_s^{-1})} \right) \Lambda'_{ij} \leq \Lambda'_{ij}. \quad (41)$$

Using Lemma E.6, we have

$$\Lambda'_{ik} \leq \max\{\Lambda'_{ij}, \Lambda'_{jk}\}. \quad (42)$$

We plug the bound $\Lambda'_{jk} \leq \Lambda'_{ij}$ from (41) into (42) and obtain $\Lambda'_{ik} \leq \Lambda'_{ij}$. Therefore $k \in \Upsilon'_i$.

To conclude, recall that k eliminates j , hence $\mu_k > \mu_j$. The contradiction arises from $k \in \Upsilon'_i$ and the definition of j as the element with largest mean in Υ'_i .

Case 2: $\Lambda'_{ij} < \frac{1}{5}$. As in the previous case, observe that j was kept for testing up to round $82s$ (i.e. $j \in C_{82s}$) and $82s < t$ (since $j \notin C_t$). At round $82s$ we necessarily had $\hat{\Delta}'_{ij}(82s, \delta) \leq 0$.

Therefore, using Lemma E.4

$$82s - 1 \leq \frac{3}{2 \log(1/(3\sqrt{2}\Lambda'_{ij}))} \log(1/\delta_{82s}). \quad (43)$$

Combining (39) and (43) gives

$$\frac{82}{2} \Lambda'_{jk} \log(1/\delta_s) \leq \frac{3 \log(1/\delta_{82s})}{2 \log(\sqrt{2}/(5\Lambda'_{ij}))}.$$

Therefore

$$\begin{aligned} \Lambda'_{jk} &\leq \frac{3}{82} \left(1 + \frac{\log(82)}{\log(1/\delta_s)} \right) \frac{1}{\log(\sqrt{2}/(5\Lambda'_{ij}))} \\ &\leq \frac{6}{82 \log(1/(3\sqrt{2}\Lambda'_{ij}))}. \end{aligned}$$

Recall that using Lemma E.6, we have: $\Lambda'_{ik} \leq \max\{\Lambda'_{ij}, \Lambda'_{jk}\}$. This implies necessarily that $\Lambda'_{ik} \leq \Lambda'_{jk}$. Otherwise if the maximum of the l.h.s is Λ'_{ij} , we get $\Lambda'_{ij} \leq \Lambda'_{ij}$ then $k \in \Upsilon'_i$, therefore by definition of j as the largest mean of Υ'_i : $\mu_k \leq \mu_j$, which contradicts the fact that k eliminated j . We conclude that (using $\Lambda'_{ij} < 1/5$):

$$\Lambda'_{ik} \leq \frac{6}{82 \log(\sqrt{2}/(5\Lambda'_{ij}))} < \frac{1}{4}.$$

Therefore $k \in \text{Arg Min}_{j \in \llbracket K \rrbracket} \{\Lambda'_{ij} \vee \frac{1}{4}\}$, which similarly to the case above, leads to a contradiction with the definition of j . □

We introduce the following notation. For $i \in \llbracket K \rrbracket$ and $t \geq 1$ let $N_{i,t}$ denote the number of queries made for arm i up to round t

$$N_{i,t} := \sum_{s=1}^t \mathbf{1}(i \in C_s). \quad (44)$$

Lemma G.2. Consider Algorithm 2 with inputs $\delta \in (0, 1)$. If (\mathcal{A}_1) defined in (22) holds, then we have for each $i \in \llbracket K \rrbracket \setminus \{i^*\}$:

$$\forall t \geq 1 : \quad N_{i,t} \leq 4100 \log(216K\delta^{-1} \min_{j \in \llbracket K \rrbracket} \{\Lambda'_{ij} \vee 1\}) \min_{j \in \llbracket K \rrbracket} \left\{ \Lambda'_{ij} \vee \frac{1}{4} \right\},$$

where $N_{i,t}$ is defined in (37).

Proof. Suppose (\mathcal{A}_1) holds. Let $i \in \llbracket K \rrbracket \setminus \{i^*\}$ and $t \geq 1$. Let u denote the last round such that $i \in S_u$. Lemma G.1 states that $\Upsilon'_i \cap C_u \neq \emptyset$, where Υ'_i is defined in (45). Let $j \in \Upsilon'_i \cap C_u$, since $i \in S_u$, we necessarily have

$$\hat{\Delta}'_{ij}(u-1, \delta) \leq 0.$$

Using Lemma E.4, we have

$$u - 1 \leq \frac{25}{\log(1 + 1/\Lambda'_{ij})} \log(\delta_{u-1}^{-1}).$$

Recall that u is the last round such that $i \in S_u$, hence $i \notin C_{35u+1}$. Therefore, for any $t \geq 1$

$$\begin{aligned}
N_{i,t} &= 82u \leq \frac{2050}{\log(1 + 1/\Lambda'_{ij})} \log(\delta^{-1}) \\
&\leq 4100 \log(216K(\Lambda'_{ij} + 1)\delta^{-1}) \frac{1}{\log(1 + 1/\Lambda'_{ij})} \\
&\leq 4100 \log(216K(\Lambda'_{ij} + 1)\delta^{-1}) \frac{1}{\log(1 + 1/\Lambda'_{ij})} \\
&\leq 4100 \log(216K(\Lambda'_{ij} + 1)\delta^{-1}) \frac{1}{\log(1 + 1/(\Lambda'_{ij} \vee 0.25))} \\
&\leq 4100 \log(216K(\Lambda'_{ij} + 1)\delta^{-1}) \frac{8}{7} (\Lambda'_{ij} \vee 0.25) \\
&\leq 4100 \log(216K(\Lambda'_{ij} + 1)\delta^{-1}) \min_{j \in \llbracket K \rrbracket} (\Lambda'_{ij} \vee 0.25),
\end{aligned}$$

$j \in \Upsilon_i$ and Lemma K.4 with $x = u$ and $c = \delta/(2K^2)$. \square

Proof for Theorem 5.1 Suppose Assumptions 1 and 2 $-\mathcal{G}$. Consider Algorithm 2 with input $\delta \in (0, 1/4)$. Suppose that event (\mathcal{B}_1) holds.

We have by definition of the total number of queries made N :

$$N = \sum_{i \in \llbracket K \rrbracket \setminus \{i^*\}} N_{i,t}.$$

Therefore, Lemma F.2 gives the result. For the second result, consider Algorithm 2 without line (12) (we stop sampling arms directly after their elimination from S_t). Lemma E.1 guarantees with probability at least $1 - \delta$, the optimal arm i^* always belongs to S_t . Therefore, Lemma E.5 guarantees that after at most $\frac{25}{2} \log(\delta^{-1}) \frac{1}{\log(1 + 1/\Lambda'_{ii^*})}$ rounds, we have $\hat{\Delta}'_{ii^*}(t, \delta) > 0$, which leads to the elimination of i .

H Proof of Theorem B.1

We provide the same type of guarantees for Algorithm 3.

For any $u, v \in \mathcal{B}_1^+$, we overload the notation $\Xi_i(u)$ into

$$\Xi_u(v) := \begin{cases} +\infty & \text{if } \langle u, \boldsymbol{\mu} \rangle \leq \langle v, \boldsymbol{\mu} \rangle \\ \max \left\{ \frac{\text{Var}(\langle \mathbf{X}, u \rangle - \langle \mathbf{X}, v \rangle)}{(\langle u, \boldsymbol{\mu} \rangle - \langle v, \boldsymbol{\mu} \rangle)^2}, \frac{3\|u-v\|_1}{\langle v, \boldsymbol{\mu} \rangle - \langle u, \boldsymbol{\mu} \rangle} \right\} & \text{otherwise} \end{cases}$$

In particular we have $\Xi_{e_i}(w) = \Xi_i(w)$, where $(e_i)_{i \in \llbracket K \rrbracket}$ is the canonical basis of \mathbb{R}^K . We say that an arm $i \in \llbracket K \rrbracket$ has failed the Γ -test at round t , if

$$\sup_{w \in \mathcal{B}_1^+(C_t \setminus \{i\})} \hat{\Gamma}_i(w, t, \delta) > 0.$$

Lemma H.1. *Let $i \in \llbracket K \rrbracket$, $u, v \in \mathcal{B}_1^+(\llbracket K \rrbracket \setminus \{i\})$, we have*

$$\Xi_i(v) \leq \max\{\Xi_i(u), \Xi_u(v)\}.$$

Proof. Let $i \in \llbracket K \rrbracket$ and $u, v \in \mathcal{B}_1^+(\llbracket K \rrbracket \setminus \{i\})$. Suppose that $\mu_i < \langle u, \boldsymbol{\mu} \rangle < \langle v, \boldsymbol{\mu} \rangle$. We have

$$\begin{aligned}
\frac{\sqrt{\text{Var}(X_i - \langle v, \mathbf{X} \rangle)}}{\langle v, \boldsymbol{\mu} \rangle - \mu_i} &\leq \frac{\sqrt{\text{Var}(X_i - \langle u, \mathbf{X} \rangle)} + \sqrt{\text{Var}(\langle u, \mathbf{X} \rangle - \langle v, \mathbf{X} \rangle)}}{(\langle v, \boldsymbol{\mu} \rangle - \langle u, \boldsymbol{\mu} \rangle) + (\langle u, \boldsymbol{\mu} \rangle - \mu_i)} \\
&\leq \max \left\{ \frac{\sqrt{\text{Var}(X_i - \langle u, \mathbf{X} \rangle)}}{\langle v, \boldsymbol{\mu} \rangle - \langle u, \boldsymbol{\mu} \rangle}, \frac{\sqrt{\text{Var}(\langle u, \mathbf{X} \rangle - \langle v, \mathbf{X} \rangle)}}{\langle u, \boldsymbol{\mu} \rangle - \mu_i} \right\},
\end{aligned}$$

where the first line follows by the triangle inequality and the second is a consequence of the inequality $\frac{a+b}{c+d} \leq \max\{\frac{a}{c}, \frac{b}{d}\}$ for positive numbers (Lemma K.3). Moreover we have

$$\begin{aligned} \frac{\|v - e_i\|_1}{\langle v - e_i, \boldsymbol{\mu} \rangle} &\leq \frac{\|v - u\|_1 + \|u - e_i\|_1}{\langle v - u, \boldsymbol{\mu} \rangle + \langle u - e_i, \boldsymbol{\mu} \rangle} \\ &\leq \max\left\{ \frac{\|v - u\|_1}{\langle v - u, \boldsymbol{\mu} \rangle}, \frac{\|u - e_i\|_1}{\langle u - e_i, \boldsymbol{\mu} \rangle} \right\}, \end{aligned}$$

where we used the inequality $\frac{a+b}{c+d} \leq \max\{\frac{a}{c}, \frac{b}{d}\}$ for positive numbers (Lemma K.3).

Combining the previous bounds, we obtain the result.

If $\mu_i \geq \langle u, \boldsymbol{\mu} \rangle$ or $\langle u, \boldsymbol{\mu} \rangle \geq \langle v, \boldsymbol{\mu} \rangle$. We have

$$\max\{\Xi_i(u); \Xi_u(v)\} = +\infty,$$

which proves the result. \square

For any $i \in \llbracket K \rrbracket \setminus \{i^*\}$, let us define Ψ_i by

$$\Psi_i := \underset{w \in \mathbf{B}_1^+(\llbracket K \rrbracket \setminus \{i\})}{\text{Arg Min}} \Xi_i(w). \quad (45)$$

Lemma H.2. Consider Algorithm 3 with input $\delta \in (0, 1)$. If (\mathcal{B}_2) defined in (13) holds, then for any $i \in \llbracket K \rrbracket \setminus \{i^*\}$, $t \geq 1$:

If $i \in S_t$, then there exists a vector $w^* \in \Psi_i$ such that: $\text{supp}(w^*) \subseteq C_t$.

Proof. Let $t \geq 1$, $i \in \llbracket K \rrbracket \setminus \{i^*\}$. We take w^* to be one of the vectors from the set Ψ_i , such that its support was jointly queried the most up to round t . More formally:

$$w^* \in \underset{w \in \Psi_i}{\text{Arg Max}} \{\langle w, \boldsymbol{\mu} \rangle\}.$$

Proceeding by proof via contradiction, we suppose that $\text{supp}(w^*) \not\subseteq C_t$. Then, we will build a vector $w' \in \mathbf{B}_1^+$, such that $\langle w^*, \boldsymbol{\mu} \rangle < \langle w', \boldsymbol{\mu} \rangle$, the contradiction follows from the definition of w^* . Let j be the first eliminated element in $\text{supp}(w^*)$. Let s denote the round where j has failed the Γ -test (i.e. $\exists \tilde{w} \in \mathbf{B}_1^+, \hat{\Gamma}_j(\tilde{w}, s, \delta) > 0$).

Let us define $w' \in \mathbb{R}^K$ as follows: $w'_j = w_j^* \tilde{w}_j$ and for $k \in \llbracket K \rrbracket \setminus \{j\}$, $w'_k = w_k^* + w_j^* \tilde{w}_k$. Recall that

$$\begin{aligned} \|w'\|_1 &= \sum_{k \in \llbracket K \rrbracket \setminus \{j\}} w_k^* + \sum_{k \in \llbracket K \rrbracket} w_j^* \tilde{w}_k \\ &= 1 - w_j^* + w_j^* \|\tilde{w}\|_1 \\ &= 1, \end{aligned}$$

where we used the fact that $\|\tilde{w}\|_1 = 1$. We conclude that $w' \in \mathbf{B}_1^+$.

Let us show that $w' \in \Psi_i$. We have

$$\begin{aligned} \langle w^* - w', \boldsymbol{\mu} \rangle &= w_j^* (1 - \tilde{w}_j) \mu_j + \sum_{k \in \llbracket K \rrbracket \setminus \{j\}} (w_k^* - w_k^* - w_j^* \tilde{w}_k) \mu_k \\ &= w_j^* \mu_j - w_j^* \tilde{w}_j \mu_j - w_j^* \sum_{k \in \llbracket K \rrbracket \setminus \{j\}} \tilde{w}_k \mu_k \\ &= w_j^* (\mu_j - \langle \tilde{w}, \boldsymbol{\mu} \rangle). \end{aligned} \quad (46)$$

Moreover

$$\begin{aligned} \|w^* - w'\|_1 &= \sum_{k=1}^K |w_k^* - w'_k| \\ &= w_j^* |1 - \tilde{w}_j| + w_j^* \sum_{k \in \llbracket K \rrbracket \setminus \{j\}} \tilde{w}_k \\ &= w_j^* \|\tilde{w} - e_j\|_1. \end{aligned}$$

Using (46) we have

$$\begin{aligned}\Xi_{w^*}(w') &= \max \left\{ \frac{\text{Var}(\langle w^* - w', \mathbf{X} \rangle)}{(\langle w^* - w', \boldsymbol{\mu} \rangle)^2}, \frac{\|w^* - w'\|_1}{\langle w^* - w', \boldsymbol{\mu} \rangle} \right\} \\ &= \max \left\{ \frac{\text{Var}(w_j(X_j - \langle \tilde{w}, \mathbf{X} \rangle))}{w_j^2(\mu_j - \langle \tilde{w}, \mathbf{u} \rangle)^2}, \frac{\|\tilde{w} - e_j\|_1}{\langle \tilde{w}, \boldsymbol{\mu} \rangle - \mu_j} \right\} \\ &= \Xi_j(\tilde{w}).\end{aligned}$$

Therefore, using Lemma H.1

$$\begin{aligned}\Xi_i(w') &\leq \max\{\Xi_i(w^*); \Xi_{w^*}(w')\} \\ &= \max\{\Xi_i(w^*); \Xi_j(\tilde{w})\}.\end{aligned}\tag{47}$$

Recall that $\hat{\Gamma}_j(\tilde{w}, s, \delta) > 0$. Hence using Lemma E.5, we have

$$K \log(2K^2 \delta_s^{-1}) \Xi_j(\tilde{w}) \leq s.\tag{48}$$

Moreover, since j failed the Γ -test at round s , we have by construction of Algorithm 3 $j \in C_{98s}$. Recall that j is the first element of the support of w^* that was eliminated, then we necessarily have $\text{supp}(w^*) \subset C_{98s}$. Since we assumed that $\text{supp}(w^*) \not\subset C_t$, we have $98s < t$, hence $i \in C_{98s}$ and $\hat{\Gamma}_i(w^*, 98s, \delta) \leq 0$. Using Lemma E.5

$$98s \leq 36K \log(2K^2 \delta_{98s}^{-1}) \Xi_i(w^*).\tag{49}$$

Combining inequalities (48) and (49), we have

$$98K \log(2K^2 \delta_s^{-1}) \Xi_j(\tilde{w}) < 36K \log(2K^2 \delta_{98s}^{-1}) \Xi_i(w^*).$$

Therefore

$$\begin{aligned}\Xi_j(\tilde{w}) &\leq \frac{36 \log(2K^2 \delta_{98s}^{-1})}{98 \log(2K^2 \delta_s^{-1})} \Xi_i(w^*) \\ &\leq \frac{36}{98} \left(1 + \frac{\log(98)}{\log(\delta_s^{-1})} \right) \Xi_i(w^*) \\ &\leq \Xi_i(w^*).\end{aligned}$$

Combining the bound above with (47), we conclude that $\Xi_i(w') \leq \Xi_i(w^*)$. Hence $w' \in \Psi_i$.

Finally, recall that by (46) $\langle w', \boldsymbol{\mu} \rangle > \langle w^*, \boldsymbol{\mu} \rangle$ (since \tilde{w} eliminated j : $\langle \tilde{w}, \boldsymbol{\mu}_j \rangle - \mu_j > 0$). The conclusion follows from $w' \in \Psi_i$ and the definition of w^* . \square

Lemma H.3. Consider Algorithm 3 with input $\delta \in (0, 1)$. If (\mathcal{B}_2) defined in (13) holds, then we have for each $i \in \llbracket K \rrbracket$, $t \geq 1$:

$$N_{i,t} \leq 10540 \log(1296K^2 \Xi_i(w^*) \delta^{-1}) K \Xi_i(w^*).$$

Proof. Suppose (\mathcal{A}_2) holds. Let $i \in \llbracket K \rrbracket \setminus \{i^*\}$ and $t \geq 1$. Let u denote the last round such that $i \in S_u$. Lemma H.2 states that there exists $w^* \in \Psi_i$ such that $\text{supp}(w^*) \subset C_u$, where Ψ_i is defined in (32). Since $i \in S_u$, we necessarily have:

$$\hat{\Gamma}_i(w^*, u-1, \delta) \leq 0.$$

Using Lemma E.3, we have

$$u-1 \leq 108K \log(2K^2 \delta_{u-1}^{-1}) \Xi_i(w^*).$$

Recall that u is the last round such that $i \in S_u$, therefore $i \notin C_{98u+1}$. Hence, for any $t \geq 1$

$$\begin{aligned}N_{i,t} &= 98u \leq 10540K \log(2K^2 \delta_u^{-1}) \Xi_i(w^*) \\ &\leq 10540 \log(1296K^2 \Xi_i(w^*) \delta^{-1}) K \Xi_i(w^*),\end{aligned}$$

where we used Lemma K.4 with $x = u$ and $c = \delta/(2K^2)$. \square

The conclusion for the proof of Theorem B.1 is similar to the conclusion of the proof of Theorem 4.1.

I Proof of Theorem 6.1

Without loss of generality assume that $\mu_1 > \mu_2 \geq \dots \geq \mu_K$. For any Bernoulli variables $(X_i)_i$ with means $(\mu_i)_i$. Fix a sequence of positive numbers $(V_{i1})_{i \geq 2}$ such that for each $i \geq 2$:

$$(\mu_1 - \mu_i) - (\mu_1 - \mu_i)^2 \leq \text{Var}(X_1 - X_i) \leq \min(2 - (\mu_1 + \mu_i); \mu_1 + \mu_i) - (\mu_1 - \mu_i)^2. \quad (50)$$

Below we build K Bernoulli variables $(X_i)_{i \in \llbracket K \rrbracket}$ such that $\mathbb{E}[X_i] = \mu_i$, $\text{Var}(X_i - X_1) \leq V_{i1}$ and $1/4 \leq \mu_i \leq 3/4$, for all $i \in \llbracket K \rrbracket$. Recall that the set of Bernoulli variables satisfying the last constraints is denoted $\mathbb{B}_K(\boldsymbol{\mu}, \mathbf{V})$.

Building a distribution in $\mathbb{B}_K(\boldsymbol{\mu}, \mathbf{V})$: Let $(U_t)_{t \geq 1}$ and $(W_{i,t})_{i \in \llbracket K \rrbracket, t \geq 1}$ denote sequences of independent variables following the uniform distribution on the interval $[0, 1]$. Let $(a_i, b_i)_{i \in \llbracket K \rrbracket}$ denote a sequence of numbers in $[0, 1]$ to be specified later. We define the arms variables $(X_i)_{i \in \llbracket K \rrbracket}$ as follows:

- $X_{1,t} = \mathbb{1}(U_t \leq \mu_1)$ for each $t \geq 1$.
- For $i \geq 2, t \geq 1$. We consider two cases:
 - If $V_{i1} \leq \mu_1 + \mu_i - \mu_1^2 - \mu_i^2$: Let $X_{i,t} = \mathbb{1}(\{U_t \leq a_i\} \text{ or } \{W_{i,t} \leq b_i\})$.
 - If $V_{i1} \geq \mu_1 + \mu_i - \mu_1^2 - \mu_i^2$: Let $X_{i,t} = \mathbb{1}(W_{i,t} \leq \mu_i)$.

Now let us specify our choice for the sequences (a_i) and (b_i) .

If $V_{i1} \leq \mu_1 + \mu_i - \mu_1^2 - \mu_i^2$: The first constraint is with respect to the means of (X_i) . We need to have for each $i \in \llbracket K \rrbracket$: $\mathbb{E}[X_i] = \mu_i$, this implies the following for each $i \geq 2$:

$$\begin{aligned} \mathbb{E}[X_i] &= \mathbb{E}[\mathbb{1}(\{U_t \leq a_i\} \text{ or } \{W_{i,t} \leq b_i\})] \\ &= a_i + b_i - a_i b_i. \end{aligned}$$

We therefore have

$$a_i + b_i - a_i b_i = \mu_i \quad (51)$$

The second constraint is with respect to the variance of the variable $(X_1 - X_i)$, we set $\text{Var}(X_1 - X_i) = V_{i1}$. This implies the following

$$\begin{aligned} \text{Var}(X_1 - X_i) &= \mathbb{E}[(\mathbb{1}(U_t \leq \mu_1) - \mathbb{1}(\{U_t \leq a_i\} \text{ or } \{W_{i,t} \leq b_i\}))^2] - (\mu_1 - \mu_i)^2 \\ &= \mu_1 + \mu_i - 2(a_i + (\mu_1 - a_i)b_i) - (\mu_1 - \mu_i)^2. \end{aligned}$$

We therefore have that a_i and b_i satisfy

$$\mu_1 + \mu_i - 2(a_i + (\mu_1 - a_i)b_i) - (\mu_1 - \mu_i)^2 = V_{i1}. \quad (52)$$

Solving in a_i and b_i for the system: (51) and (52) is:

$$\begin{aligned} b_i &= \frac{V_{i1} - ((\mu_1 - \mu_i) - (\mu_1 - \mu_i)^2)}{2(1 - \mu_1)} \\ a_i &= \frac{\mu_i - b_i}{1 - b_i} \end{aligned}$$

Observe that when $V_{i1} \leq \mu_1 + \mu_i - \mu_1^2 - \mu_i^2$, we have :

$$\begin{aligned} b_i &= \frac{V_{i1} - (\mu_1 - \mu_i) + (\mu_1 - \mu_i)^2}{2(1 - \mu_1)} \\ &\leq \frac{2\mu_i - 2\mu_i\mu_1}{2(1 - \mu_1)} \\ &= \mu_i. \end{aligned}$$

Therefore $b_i \leq \mu_i$. Moreover, using Lemma K.8, we have $b_i \geq 0$. We conclude that $b_i \in [0, \mu_i]$, which implies that $a_i \in [0, 1]$.

If $V_{i1} \geq \mu_1 + \mu_i - \mu_1^2 - \mu_i^2$: Let us prove in case that we have $\text{Var}(X_1 - X_i) \leq V_{i1}$. Recall that X_i and X_1 are independent. Therefore:

$$\begin{aligned} \text{Var}(X_1 - X_i) &= \mu_1 + \mu_i - 2\mu_1\mu_i - (\mu_1 - \mu_i)^2 \\ &\leq V_{i1} + \mu_1^2 + \mu_i^2 - 2\mu_1\mu_i - (\mu_1 - \mu_i)^2 \\ &= V_{i1}. \end{aligned}$$

As a conclusion we have in both cases: the distribution belongs to $\mathbb{B}_K(\boldsymbol{\mu}, \boldsymbol{\sigma})$.

Developing the lower bound for T_i in the case $V_{i1} \leq \mu_1 + \mu_i - \mu_1^2 - \mu_i^2$: Given a joint distribution of Bernoulli variables in $\mathbb{B}_K(\boldsymbol{\mu}, \mathbf{V})$, let us develop the corresponding lower bound.

Let \mathcal{A} be a δ sound strategy. Denote $\mathbb{P}^{(1)}$ the joint distribution of X_i defined above. For $i \in \{2, \dots, K\}$, denote by $\mathbb{P}^{(i)}$ the alternative probability distribution where only arm i is modified as follows:

$$X_{i,t} = \mathbb{1}\left(\{U_t \leq a_i\} \text{ or } \left\{W'_{i,t} \leq \frac{\mu_1 + \epsilon - \mu_i}{1 - a_i} + b_i\right\}\right),$$

where $\epsilon > 0$ and $(W'_{i,t})_{t \geq 1}$ is a sequence of independent variables following the uniform distribution in $[0, 1]$. Observe that $\mathbb{E}^{(i)}[X_i] = \mu_1 + \epsilon$, therefore under $\mathbb{P}^{(i)}$ the optimal arm is arm i .

Fix $i \geq 2$, since \mathcal{A} is δ -sound,

$$\mathbb{P}^{(1)}(\psi \neq 1) + \mathbb{P}^{(i)}(\psi \neq i) \leq 2\delta. \quad (53)$$

Using Theorem 2.2 of [29], we have

$$\mathbb{P}^{(1)}(\psi \neq 1) + \mathbb{P}^{(i)}(\psi = 1) \geq 1 - \text{TV}(\mathbb{P}^{(1)}, \mathbb{P}^{(i)}),$$

where $\text{TV}(\mathbb{P}, \mathbb{Q})$ denotes the total variation distance between \mathbb{P} and \mathbb{Q} . Using (58), we conclude that

$$\text{TV}(\mathbb{P}^{(1)}, \mathbb{P}^{(i)}) \geq 1 - 2\delta. \quad (54)$$

For a subset $A \subset \llbracket K \rrbracket$, denote N_A the total number of rounds where the jointly queried arms are the elements of A , and let $\mathbb{P}_A^{(\cdot)}$ denote the joint distribution of $(X_j)_{j \in A}$ under $\mathbb{P}^{(\cdot)}$. Under Protocol 1, the learner has to choose a subset $C_t \subset \llbracket K \rrbracket$ in each round t , and observes only the rewards of arms in C_t . Hence, we can apply Lemma K.5 which bounds the total variation distance in terms of Kullback-Leibler discrepancy to this case where arms correspond to subsets of $\llbracket K \rrbracket$. This gives

$$\begin{aligned} \text{TV}(\mathbb{P}^{(1)}, \mathbb{P}^{(i)}) &\leq 1 - \frac{1}{2} \exp\left\{-\sum_{A \subset \llbracket K \rrbracket} \mathbb{E}^{(1)}[N_A] \text{KL}(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)})\right\} \\ &= 1 - \frac{1}{2} \exp\left\{-\sum_{\substack{A \subset \llbracket K \rrbracket: \\ i \in A}} \mathbb{E}^{(1)}[N_A] \text{KL}(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)})\right\}. \end{aligned} \quad (55)$$

Now fix $A \subset \llbracket K \rrbracket$, such that $i \in A$. Let us calculate $\text{KL}(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)})$.

Denote $\mathbb{Q}_A^{(i)}$ for $i \in \llbracket K \rrbracket$, the joint distribution: $((X_{j,t})_{j \in A}, U_t)$. Using the data processing inequality

$$\text{KL}(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}) \leq \text{KL}(\mathbb{Q}_A^{(1)}, \mathbb{Q}_A^{(i)}).$$

Observe that X_1 follows the same distribution under $\mathbb{P}^{(1)}$ and $\mathbb{P}^{(i)}$. Therefore

$$\text{KL}(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}) \leq \mathbb{E}_{U_t}^{(1)}\left[\text{KL}(\mathbb{P}_A^{(1)}((X_j)_{j \in A} | U_t), \mathbb{P}_A^{(i)}((X_j)_{j \in A} | U_t))\right],$$

where $\mathbb{E}_{U_t}^{(\cdot)}$ denotes the conditional expectation under $\mathbb{P}^{(\cdot)}$ with respect to U_t .

Observe that conditionally to U_t , the variables $(X_j)_{j \neq 1}$ are independent, both under $\mathbb{P}_A^{(1)}$ and $\mathbb{P}_A^{(i)}$. Moreover X_j for $j \in A \setminus \{1, i\}$ have the same probability distributions under both alternatives. Therefore

$$\begin{aligned} \text{KL}\left(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}\right) &\leq \mathbb{E}_{U_t}^{(1)} \left[\text{KL}\left(\mathbb{P}_A^{(1)}((X_j)_{j \in A} | U_t), \mathbb{P}_A^{(i)}((X_j)_{j \in A} | U_t)\right) \right] \\ &= \mathbb{E}_{U_t}^{(1)} \left[\text{KL}\left(\mathbb{P}_A^{(1)}(X_i | U_t), \mathbb{P}_A^{(i)}(X_i | U_t)\right) \right]. \end{aligned}$$

Therefore, we have:

$$\text{KL}\left(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}\right) \leq (1 - a_i) \text{KL}\left(\mathbb{1}(W_{i,t} \leq b_i), \mathbb{1}\left(W'_{i,t} \leq \frac{\mu_1 + \epsilon - \mu_i}{1 - a_i} + b_i\right)\right).$$

Using Lemma K.9, we have

$$\begin{aligned} \text{KL}\left(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}\right) &\leq (1 - a_i) \frac{\left(\frac{\mu_1 + \epsilon - \mu_i}{1 - a_i}\right)^2}{\left(\frac{\mu_1 + \epsilon - \mu_i}{1 - a_i} + b_i\right) \left(1 - \frac{\mu_1 + \epsilon - \mu_i}{1 - a_i} - b_i\right)} \\ &= \frac{(1 - a_i)(\mu_1 + \epsilon - \mu_i)^2}{(\mu_1 + \epsilon - \mu_i + b_i - a_i b_i)(1 - (\mu_1 + \epsilon - \mu_i) - a_i - b_i + a_i b_i)}. \end{aligned}$$

Using $a_i + b_i - a_i b_i = \mu_i$ and $a_i = \frac{\mu_i - b_i}{1 - b_i}$, we have

$$\begin{aligned} \text{KL}\left(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}\right) &\leq \frac{(1 - \mu_i)(\mu_1 + \epsilon - \mu_i)^2}{((\mu_1 + \epsilon)(1 - b_i) - \mu_i + b_i)(1 - \mu_1 - \epsilon)} \\ &= \frac{(1 - \mu_i)(\mu_1 + \epsilon - \mu_i)^2}{(\mu_1 + \epsilon - \mu_i + b_i(1 - \mu_1 - \epsilon))(1 - \mu_1 - \epsilon)} \end{aligned}$$

Next, we plug the inequality above into inequality (55) and obtain:

$$\text{TV}\left(\mathbb{P}^{(1)}, \mathbb{P}^{(i)}\right) \leq 1 - \frac{1}{2} \exp\left\{-\text{KL}\left(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}\right) \sum_{\substack{A \subset \llbracket K \rrbracket: \\ i \in A}} \mathbb{E}^{(1)}[N_A]\right\}.$$

For $j \in \llbracket K \rrbracket$, denote by T_j the total number of rounds where arm j was queried:

$$T_j := \sum_{\substack{A \subset \llbracket K \rrbracket: \\ j \in A}} N_A.$$

Therefore

$$\text{TV}\left(\mathbb{P}^{(1)}, \mathbb{P}^{(i)}\right) \leq 1 - \frac{1}{2} \exp\left\{-\text{KL}\left(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}\right) \mathbb{E}^{(1)}[T_i]\right\}.$$

Combining the inequality above with (54) and taking $\epsilon \rightarrow 0$, we obtain

$$\mathbb{E}^{(1)}[T_i] \geq \frac{(\mu_1 + \epsilon - \mu_i + b_i(1 - \mu_1 - \epsilon))(1 - \mu_1 - \epsilon)}{(1 - \mu_i)(\mu_1 + \epsilon - \mu_i)^2} \log\left(\frac{1}{4\delta}\right).$$

Taking $\epsilon \rightarrow 0$, using the expression of b_i and $\mu_i, \mu_1 \leq 3/4$ yields

$$\begin{aligned} \mathbb{E}^{(1)}[T_i] &\geq \frac{(\mu_1 - \mu_i) + V_{i1} + (\mu_1 - \mu_i)^2}{2(\mu_1 - \mu_i)^2} \cdot \frac{1 - \mu_1}{1 - \mu_i} \log\left(\frac{1}{4\delta}\right) \\ &\geq \frac{V_{i1} + (\mu_1 - \mu_i)}{8(\mu_1 - \mu_i)^2} \log\left(\frac{1}{4\delta}\right). \end{aligned} \tag{56}$$

Developing a lower bound for T_i in the case $V_{i1} \geq \mu_1 + \mu_i - \mu_1^2 - \mu_i^2$. In this case we introduce the following alternative distribution $\mathbb{P}^{(i)}$, where we only change the variable X_i into: $\forall t \geq 1 : X_{i,t} = \mathbb{1}(W_{i,t} \leq \mu_1 + \epsilon)$, for some small positive constant ϵ .

Similarly to the previous case, since all variables are independent: for any $A \subset \llbracket K \rrbracket$:

Therefore

$$\begin{aligned} \text{KL}\left(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}\right) &= \text{KL}(\mathbb{1}(W_{i,t} \leq \mu_i), \mathbb{1}(W_{i,t} \leq \mu_1 + \epsilon)) \\ &\leq \frac{(\mu_1 + \epsilon - \mu_i)^2}{(\mu_1 + \epsilon)(1 - \mu_1 - \epsilon)} \end{aligned}$$

The remainder of the calculation is similar to the previous case and leads to (using $\mu_i \in [1/3, 3/4]$):

$$\begin{aligned} \mathbb{E}^{(1)}[T_i] &\geq \frac{\mu_1(1 - \mu_1)}{(\mu_1 - \mu_i)^2} \log\left(\frac{1}{4\delta}\right) \\ &\geq \frac{3}{16(\mu_1 - \mu_i)^2} \log\left(\frac{1}{4\delta}\right) \\ &\geq \frac{3(V_{i1} + \mu_1 - \mu_i)}{16(\mu_1 - \mu_i)^2} \log\left(\frac{1}{4\delta}\right), \end{aligned} \tag{57}$$

where we used in the last line the fact that: $V_{i1} + \mu_1 - \mu_i \leq 5/4$

Conclusion. Using (56) and (57), we have for any $i \in \{2, \dots, K\}$:

$$\mathbb{E}^{(1)}[T_i] \geq \frac{V_{i1} + (\mu_1 - \mu_i)}{8(\mu_1 - \mu_i)^2} \log\left(\frac{1}{4\delta}\right).$$

Recall that T_i represents the number of rounds where arm i is queried. Hence the total number of queries N satisfies: $N \geq \sum_{i=2}^K T_i$. We conclude that

$$\mathbb{E}^{(1)}[N] \geq \log\left(\frac{1}{4\delta}\right) \sum_{i=2}^K \frac{V_{i1} + (\mu_1 - \mu_i)}{8(\mu_1 - \mu_i)^2}.$$

J Proof of Theorem 6.2

Without loss of generality assume that $\mu_1 > \mu_2 \geq \dots \geq \mu_K$, hence $i^* = 1$. Let \mathcal{A} be a δ sound strategy. Let $\mathbb{P}^{(1)}$ denote the joint distribution of the arms defined as follows: Let Z denote a random variable distributed following $\mathcal{N}(\mu_1, 1)$.

- For $i \in \{2, \dots, K\}$, we set $X_i = Z + Y_i$, where Y_i is a random variable independent of Z and $(Y_i)_{i \geq 2}$, and distributed following $\mathcal{N}(\mu_i - \mu_1, V_{i1})$.
- The optimal arm is given by $X_1 = Z$, let $\sigma_1 := 1$.

Recall that in the configuration above we have for each $i \in \llbracket K \rrbracket$: $\mathbb{E}[X_i] = \mu_i$, $\mathbb{E}[(X_1 - X_i)^2] = V_{i1}$, and $\text{Var}(X_i) \geq 1$, therefore $\mathbb{P}^{(0)} \in \mathbb{G}_K(\boldsymbol{\mu}, \mathbf{V})$.

For $i \in \{2, \dots, K\}$, denote by $\mathbb{P}^{(i)}$ the alternative probability distribution where the only arm modified is arm i . We set: $X_i = Z + Y'$, where $Y' \sim \mathcal{N}(\epsilon, V_{i1} + (\mu_1 - \mu_i)^2)$, where $\epsilon > 0$.

Observe that, under $\mathbb{P}^{(i)}$, the arm i is optimal. Fix $i \geq 2$. Since \mathcal{A} is δ -sound, we have

$$\mathbb{P}^{(1)}(\psi \neq 1) + \mathbb{P}^{(i)}(\psi \neq i) \leq 2\delta. \tag{58}$$

Using Theorem 2.2 of [29], we have

$$\mathbb{P}^{(1)}(\psi \neq 1) + \mathbb{P}^{(i)}(\psi = 1) \geq 1 - \text{TV}\left(\mathbb{P}^{(1)}, \mathbb{P}^{(i)}\right),$$

where $\text{TV}(\mathbb{P}, \mathbb{Q})$ denotes the total variation distance between \mathbb{P} and \mathbb{Q} . Using (58), we deduce that

$$\text{TV}(\mathbb{P}^{(1)}, \mathbb{P}^{(i)}) \geq 1 - 2\delta. \quad (59)$$

For a subset $A \subset \llbracket K \rrbracket$, denote N_A the total number of rounds where the jointly queried arms are the elements of A , and let $\mathbb{P}_A^{(\cdot)}$ denote the joint distribution of $(X_j)_{j \in A}$ under $\mathbb{P}^{(\cdot)}$. Under Protocol 1, the learner has to choose a subset $C_t \subset \llbracket K \rrbracket$ in each round t , and observes only the rewards of arms in C_t . Hence, we can apply Lemma K.5 which bounds the total variation distance in terms of Kullback-Leibler discrepancy to this case where arms correspond to subsets of $\llbracket K \rrbracket$. This gives

$$\begin{aligned} \text{TV}(\mathbb{P}^{(1)}, \mathbb{P}^{(i)}) &\leq 1 - \frac{1}{2} \exp \left\{ - \sum_{A \subset \llbracket K \rrbracket} \mathbb{E}^{(1)}[N_A] \text{KL}(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}) \right\} \\ &= 1 - \frac{1}{2} \exp \left\{ - \sum_{\substack{A \subset \llbracket K \rrbracket: \\ i \in A}} \mathbb{E}^{(1)}[N_A] \text{KL}(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}) \right\}. \end{aligned} \quad (60)$$

Now fix $A \subset \llbracket K \rrbracket$, such that $i \in A$. Let us calculate $\text{KL}(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)})$.

Using the data processing inequality, we deduce that, for any $A \subset \llbracket K \rrbracket$,

$$\text{KL}(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}) \leq \text{KL}(\mathbb{P}_{A \cup \{1\}}^{(1)}, \mathbb{P}_{A \cup \{1\}}^{(i)}).$$

Observe that X_1 follows the same distribution under $\mathbb{P}^{(1)}$ and $\mathbb{P}^{(i)}$. Therefore

$$\text{KL}(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}) \leq \mathbb{E}_{X_1}^{(1)} \left[\text{KL}(\mathbb{P}_A^{(1)}((X_j)_{j \in A} | X_1), \mathbb{P}_A^{(i)}((X_j)_{j \in A} | X_1)) \right],$$

where $\mathbb{E}_{X_1}[\cdot]$ refers to the conditional expectation with respect to X_1 . Recall that conditionally to X_1 , the variables $(X_j)_{j \neq 1}$ are independent, both under $\mathbb{P}_A^{(1)}$ and $\mathbb{P}_A^{(i)}$. Moreover X_j for $j \in A \setminus \{1, i\}$ have the same probability distributions under both alternatives. Therefore, we can the above Kullback divergence in terms of conditional Kullback divergence.

$$\begin{aligned} \text{KL}(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}) &\leq \mathbb{E}_{X_1}^{(1)} \left[\text{KL}(\mathbb{P}_A^{(1)}((X_j)_{j \in A} | X_1), \mathbb{P}_A^{(i)}((X_j)_{j \in A} | X_1)) \right] \\ &= \mathbb{E}_{X_1}^{(1)} \left[\text{KL}(\mathbb{P}_A^{(1)}(X_i | X_1), \mathbb{P}_A^{(i)}(X_i | X_1)) \right]. \end{aligned}$$

Moreover, we have

$$\begin{aligned} \mathbb{E}_A^{(1)}[X_i | X_1] &= \mu_i + X_1 - \mu_1 \\ \mathbb{E}_A^{(i)}[X_i | X_1] &= \mu_1 + \epsilon + X_1 - \mu_1 \\ \text{Var}_A^{(1)}(X_i | X_1) &= 1 + V_{i1} \\ \text{Var}_A^{(i)}(X_i | X_1) &= 1 + V_{i1} + (\mu_1 - \mu_i)^2. \end{aligned}$$

We deduce that

$$\begin{aligned} \text{KL}(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}) &\leq \frac{1}{2} \log \left(\frac{1 + V_{i1} + (\mu_1 - \mu_i)^2}{1 + V_{i1}} \right) + \frac{1 + V_{i1} + (\mu_i - \mu_1 - \epsilon)^2}{2(1 + V_{i1})} - \frac{1}{2} \\ &\leq \frac{1}{2} \log \left(\frac{V_{i1} + (\mu_1 - \mu_i)^2}{V_{i1}} \right) + \frac{(\mu_i - \mu_1 - \epsilon)^2 - (\mu_i - \mu_1)^2}{2(1 + V_{i1} + (\mu_1 - \mu_i)^2)}. \end{aligned}$$

As a conclusion, we have

$$\text{KL}(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}) \leq \frac{1}{2} \log \left(\frac{V_{i1} + (\mu_1 - \mu_i)^2}{V_{i1}} \right) + \frac{(\mu_i - \mu_1 - \epsilon)^2 - (\mu_i - \mu_1)^2}{2(1 + V_{i1} + (\mu_1 - \mu_i)^2)}. \quad (61)$$

Next, we plug the inequality above into (60) and obtain:

$$\mathrm{TV}\left(\mathbb{P}^{(1)}, \mathbb{P}^{(i)}\right) \leq 1 - \frac{1}{2} \exp\left\{-\mathrm{KL}\left(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}\right) \sum_{\substack{A \subset \llbracket K \rrbracket: \\ i \in A}} \mathbb{E}^{(1)}[N_A]\right\}.$$

For $j \in \llbracket K \rrbracket$, denote by T_j the total number of rounds where arm j was queried:

$$T_j := \sum_{\substack{A \subset \llbracket K \rrbracket: \\ j \in A}} N_A.$$

Therefore

$$\mathrm{TV}\left(\mathbb{P}^{(1)}, \mathbb{P}^{(i)}\right) \leq 1 - \frac{1}{2} \exp\left\{-\mathrm{KL}\left(\mathbb{P}_A^{(1)}, \mathbb{P}_A^{(i)}\right) \mathbb{E}^{(1)}[T_i]\right\}.$$

Combining the inequality above with (59) and taking $\epsilon \rightarrow 0$, we obtain

$$\mathbb{E}^{(1)}[T_i] \geq \frac{2}{\log\left(\frac{V_{i1} + (\mu_1 - \mu_i)^2}{V_{i1}}\right)} \log\left(\frac{1}{4\delta}\right).$$

Taking $\epsilon \rightarrow 0$ and using the definition of σ_i for $i \geq 2$, we have

$$\mathbb{E}^{(1)}\left[\sum_{i=2}^K T_i\right] \geq 2 \log\left(\frac{1}{4\delta}\right) \sum_{i=2}^K \frac{1}{\log\left(\frac{V_{i1} + (\mu_1 - \mu_i)^2}{V_{i1}}\right)}.$$

Recall that T_i represents the number of rounds where arm i is queried. Hence the total number of queries N satisfies: $N \geq \sum_{i=2}^K T_i$. We conclude that

$$\mathbb{E}^{(1)}[N] \geq 2 \log\left(\frac{1}{4\delta}\right) \sum_{i=2}^K \frac{1}{\log\left(1 + \frac{(\mu_1 - \mu_i)^2}{V_{i1}}\right)}.$$

K Some technical results

We state below a version of the empirical Bernstein's inequality presented in [3].

Theorem K.1. *Let X_1, \dots, X_t be i.i.d random variables taking their values in $[0, b]$. Let $\mu = \mathbb{E}[X_1]$ be their common expected value. Consider the empirical expectation \bar{X}_t and variance V_t defined respectively by*

$$\bar{X}_t = \frac{\sum_{i=1}^t X_i}{t} \quad \text{and} \quad V_t = \frac{\sum_{i=1}^t (X_i - \bar{X}_t)^2}{t}.$$

Then for any $t \in \mathbb{N}$ and $x > 0$, with probability at least $1 - 3e^{-x}$

$$|\bar{X}_t - \mu| \leq \sqrt{\frac{2V_t x}{t}} + \frac{3bx}{t}.$$

Theorem below corresponds to Theorem 10 in [24].

Theorem K.2. *Let $n \geq 2$ and $\mathbf{X} = (X_1, \dots, X_n)$ be a vector of independent random variables with values in $[0, 1]$. Then for $\delta > 0$ we have, writing $\mathbb{E}V_n$ for $\mathbb{E}_{\mathbf{X}} V_n(\mathbf{X})$,*

$$\begin{aligned} \mathbb{P}\left\{\sqrt{\mathbb{E}V_n} > \sqrt{V_n(\mathbf{X})} + \sqrt{\frac{2 \log(1/\delta)}{n-1}}\right\} &\leq \delta \\ \mathbb{P}\left\{\sqrt{V_n(\mathbf{X})} > \sqrt{\mathbb{E}V_n} + \sqrt{\frac{2 \log(1/\delta)}{n-1}}\right\} &\leq \delta. \end{aligned}$$

The following lemma is technical, it will be used in the proof of Lemma E.6.

Lemma K.3. Let a, b, c and $d > 0$, we have

$$\frac{a+b}{c+d} \leq \max\left\{\frac{a}{c}, \frac{b}{d}\right\}.$$

Proof. Let $\rho = \frac{c}{c+d} \in (0, 1)$. Observe that

$$\frac{a+b}{c+d} = \rho \frac{a}{c} + (1-\rho) \frac{b}{d},$$

and $1-\rho = \frac{d}{c+d} \in (0, 1)$. Taking the maximum of the convex combination above gives the result. \square

Lemma K.4. Let $x \geq 1, c \in (0, 1)$ and $y > 0$ such that:

$$\frac{\log(x/c)}{x} > y. \tag{62}$$

Then:

$$x < \frac{2 \log\left(\frac{1}{cy}\right)}{y}.$$

Proof. Inequality (62) implies

$$x < \frac{\log(x/c)}{y},$$

and further

$$\log(x/c) < \log(1/yc) + \log \log(x/c) \leq \log(1/yc) + \frac{1}{2} \log(x/c),$$

since it can be easily checked that $\log(t) \leq t/2$ for all $t > 0$. Solving and plugging back into the previous display leads to the claim. \square

Lemma K.5 (20, with slight modification). Let ν and ν' be two collections of d probability distributions on \mathbb{R} , such that for all $a \in \llbracket d \rrbracket$, the distributions ν_a and ν'_a are mutually absolutely continuous. For any almost-surely finite stopping time τ with respect to (\mathcal{F}_t) ,

$$\sup_{\mathcal{E} \in \mathcal{F}_\tau} |\mathbb{P}_\nu(\mathcal{E}) - \mathbb{P}_{\nu'}(\mathcal{E})| \leq 1 - \frac{1}{2} \exp\left\{-\sum_{a=1}^d \mathbb{E}_\nu[N_a(\tau)] \text{KL}(\nu_a, \nu'_a)\right\}.$$

Lemma K.6. [30] Let (Y_1, \dots, Y_n) be i.i.d Gaussian variables, with mean 0 and variance 1. Let $Z = \sum_{i=1}^n Y_i^2$. For any number $0 < x < 1$,

$$\begin{aligned} \mathbb{P}\left(Z \geq n - 1 + 2\sqrt{(n-1)\log(1/x)} + 2\log(1/x)\right) &\leq x, \\ \mathbb{P}\left(Z \leq n - 1 - 2\sqrt{(n-1)\log(1/\delta)}\right) &\leq x. \end{aligned}$$

For any positive number $0 < x < 1$

$$\mathbb{P}\left(Z \leq (n-1)Cx^{2/(n-1)}\right) \leq x,$$

where the constant $C = \exp(-1)$.

Concentration bound for the Gaussian variance sample. Let $\mathbf{X} = (X_1, \dots, X_n)$ be a vector of independent standard normal variables. Define the sample variance by

$$V_n(\mathbf{X}) := \frac{1}{n(n-1)} \sum_{1 \leq i < j \leq n} (X_i - X_j)^2. \tag{63}$$

Observe that

$$\begin{aligned}
V_n(\mathbf{X}) &= \frac{1}{2n(n-1)} \sum_{i \neq j} (X_i - X_j)^2 \\
&= \frac{1}{n(n-1)} \left((n-1) \sum_{i=1}^n X_i^2 - \sum_{i \neq j} X_i X_j \right) \\
&= \frac{1}{n(n-1)} \mathbf{X}^\top A \mathbf{X},
\end{aligned}$$

where A is the matrix such that off-diagonal entries are equal to -1 and diagonal entries are equal to $n-1$. Let us compute the eigenvalue of the matrix A : Observe that $A = n\mathbf{I}_n - \mathbb{1}_n \mathbb{1}_n^\top$, hence the eigenvalues of A are n with multiplicity $n-1$ and 0 . Hence, we have

$$\mathbf{X}^\top A \mathbf{X} = n \sum_{i=1}^{n-1} Y_i^2,$$

where (Y_i) are independent and follow the standard normal distribution. Finally using Lemma K.6 we obtain:

Lemma K.7. *Let $\mathbf{X} = (X_1, \dots, X_n)$ be a vector of independent standard normal variables. Let $V_n(\mathbf{X})$ denote the variance sample defined in (63). Let $\delta \in (0, 1/3)$, we have with probability at least $1 - 3\delta$:*

$$\begin{aligned}
V_n(\mathbf{X}) &\geq \max \left\{ 1 - 2\sqrt{\frac{\log(1/\delta)}{n-1}}; C\delta^{2/(n-1)} \right\} \\
V_n(\mathbf{X}) &\leq 1 + 2\sqrt{\frac{\log(1/\delta)}{n-1}} + 2\frac{\log(1/\delta)}{n-1}.
\end{aligned}$$

Lemma K.8. *Let X and Y be two Bernoulli variables with means x and y respectively. We have*

$$\max(x+y-1, 0) \leq \mathbb{E}[XY] \leq \min(x, y).$$

Moreover:

$$x - y - (x - y)^2 \leq \text{Var}(X - Y) \leq \min(2 - (x + y); x + y) - (x - y)^2.$$

Proof. Without loss of generality suppose that $x \leq y$. We have

$$\begin{aligned}
\mathbb{E}[XY] &= \mathbb{P}(XY = 1) \\
&= \mathbb{P}(X = 1 \text{ and } Y = 1) \\
&= \mathbb{P}(X = 1) + \mathbb{P}(Y = 1) - \mathbb{P}(X = 1 \text{ or } Y = 1) \\
&= x + y - \mathbb{P}(X = 1 \text{ or } Y = 1).
\end{aligned}$$

The conclusion follows by using $y \leq \mathbb{P}(X = 1 \text{ or } Y = 1) \leq 1$, and $XY \geq 0$.

Moreover, we have:

$$\begin{aligned}
\text{Var}(X - Y) &= \mathbb{E}[(X - Y)^2] - (x - y)^2 \\
&= \mathbb{E}[X] + \mathbb{E}[Y] - 2\mathbb{E}[XY] - (x - y)^2 \\
&= x + y - (x - y)^2 - 2\mathbb{E}[XY].
\end{aligned}$$

We plug-in the previous bounds on $\mathbb{E}[XY]$ and obtain the result. \square

Lemma K.9. *Let X and Y denote two Bernoulli variables with parameters $x \in (0, 1)$ and $y \in (0, 1)$ respectively. We have*

$$KL(X, Y) \leq \frac{(x - y)^2}{y(1 - y)}.$$

Proof. Using $\log u \leq u - 1$ we get

$$KL(X, Y) = x \log \frac{x}{y} + (1 - x) \log \frac{1 - x}{1 - y} \leq \frac{x^2}{y} + \frac{(1 - x)^2}{1 - y} - 1 = \frac{(x - y)^2}{y(1 - y)}.$$

\square