
Bandits with Ranking Feedback

Davide Maran*
Politecnico di Milano
davide.maran@polimi.it

Francesco Bacchiocchi*
Politecnico di Milano
francesco.bacchiocchi@polimi.it

Francesco Emanuele Stradi*
Politecnico di Milano
francescoemanuele.stradi@polimi.it

Matteo Castiglioni
Politecnico di Milano
matteo.castiglioni@polimi.it

Nicola Gatti
Politecnico di Milano
nicola.gatti@polimi.it

Marcello Restelli
Politecnico di Milano
marcello.restelli@polimi.it

Abstract

In this paper, we introduce a novel variation of multi-armed bandits called *bandits with ranking feedback*. Unlike traditional bandits, this variation provides feedback to the learner that allows them to rank the arms based on previous pulls, without quantifying numerically the difference in performance. This type of feedback is well-suited for scenarios where the arms' values cannot be precisely measured using metrics such as monetary scores, probabilities, or occurrences. Common examples include human preferences in matchmaking problems. Furthermore, its investigation answers the theoretical question on how numerical rewards are crucial in bandit settings. In particular, we study the problem of designing *no-regret* algorithms with ranking feedback both in the *stochastic* and *adversarial* settings. We show that, with stochastic rewards, differently from what happens with non-ranking feedback, no algorithm can suffer a logarithmic regret in the time horizon T in the instance-dependent case. Furthermore, we provide two algorithms. The first, namely DREE, guarantees a superlogarithmic regret in T in the instance-dependent case thus matching our lower bound, while the second, namely R-LPE, guarantees a regret of $\tilde{O}(\sqrt{T})$ in the instance-independent case. Remarkably, we show that no algorithm can have an optimal regret bound in both instance-dependent and instance-independent cases. Finally, we prove that no algorithm can achieve a sublinear regret when the rewards are adversarial.

1 Introduction

Multi-armed bandits are well-known sequential decision-making problems where a learner is given a number of arms whose reward is unknown [Lattimore and Szepesvari, 2017]. At every round, the learner can pull an arm and observe a realization of the reward associated with that arm, which can be generated *stochastically* [Auer et al., 2002] or *adversarially* [Auer et al., 1995]. The central

*Equal Contribution.

question in multi-armed bandits concerns how to address the *exploration/exploitation* tradeoff to minimize the *regret* between the reward provided by the *learning policy* and the optimal *clairvoyant* algorithm. Interestingly, multi-armed bandits come with several flavors capturing a wide range of different applications, *e.g.*, with delayed feedback [Vernade et al., 2017, 2020], combinatorial constraints [Combes et al., 2015], and a continuous set of arms [Kleinberg et al., 2019].

In this paper, we introduce a novel variation of multi-armed bandits that, to the best of our knowledge, is unexplored so far. We name the model as *bandits with ranking feedback*. This feedback provides the learner with a partial observation over the rewards given by the arms. More precisely, the learner can rank the arms based on the previous pulls they experienced, but they cannot quantify numerically the difference in performance. Thus, the learner is not allowed to assess how much an arm is better or worse than another. This type of feedback is well-suited for scenarios where the arms’ values cannot be precisely measured using metrics such as monetary scores, probabilities, or occurrences, and naturally applies to various settings, *e.g.*, when dealing with human preferences such as in matchmaking settings among humans and when the scores cannot be revealed for privacy or security reasons. This latter case can be found, *e.g.*, in online advertising platforms offering automatic bidding services as they have no information on the actual revenue of the advertising campaigns since the advertisers prefer not to reveal these values being sensible data for the companies. Notice that a platform can observe the number of clicks received by an advertising campaign, but it cannot observe the revenue associated with that campaign. Remarkably, our model poses the interesting theoretical question whether the lack of numerical scores precludes the design of sublinear regret algorithms or worsens the regret bounds that are achievable when numerical scores are available.

Original contributions. We study the problem of designing *no-regret* algorithms for bandits with ranking feedback in both *stochastic* and *adversarial* settings. In the case of adversarial rewards, we prove that no algorithm can achieve sublinear regret. In contrast, with stochastic rewards, we show that ranking feedback does not preclude such a possibility. However, it worsens the guarantees in terms of regret bound compared to the non-ranking case. In particular, in the instance-dependent case, we show that no algorithm can achieve logarithmic regret in the time horizon (unlike the non-ranking case), and we provide an algorithm, namely DREE (Dynamical Ranking Exploration-Exploitation), guaranteeing superlogarithmic regret that matches such a lower bound. In the instance-independent case, a crucial question is whether there exists an algorithm providing better regret bound compared to the one achieved by the well-known Explore-then-Commit algorithm, which trivially guarantees $\tilde{O}(T^{2/3})$. We positively answer this question by designing an algorithm, namely R-LPE (Ranking Logarithmic Phased Elimination), which guarantees a regret of $\tilde{O}(\sqrt{T})$ in the instance-independent case if the rewards are Gaussian. To achieve such a result, we derive several non-standard results that allow us to discretize Brownian motions, which are of independent interest. These two different approaches leave open the problem of whether there exists an algorithm achieving optimal performance in both the instance-dependent and instance-independent case. We negatively answer this question by showing that no algorithm can achieve an optimal regret bound in both the two cases, confirming the need of designing two distinct algorithms for the two cases. Finally, we numerically evaluate our DREE and R-LPE algorithms in a testbed, and we compare their performance with some baselines from the literature in different settings. We show that our algorithms dramatically outperform the baselines in terms of empirical regret.

Related works. The field most related to bandits with ranking is *preference learning*, which aims at learning the preferences of one or more agents from some observations [Fürnkranz and Hüllermeier, 2010]. Let us remark that preference learning has recently gained a lot of attention from the scientific community, as it enables the design of AI artifacts capable of interacting with human-in-the-loop (HTL) environments. Indeed, human feedback may be quite misleading when it is asked to report numerical values, while humans are far more effective at reporting ranking preferences. The preference learning literature mainly focuses on two kinds of preference observations: pairwise preferences and ranking. In the first case, the data observed by the learner involves preferences between two objects, *i.e.*, a partial preference is given to the learner. In the latter, a complete ranking of the available data is given as feedback. Our work belongs to the latter branch. Preference learning has been widely investigated by the online learning community, see, *e.g.*, [Bengs et al., 2018].

Precisely, our work presents several similarities with the *dueling bandits* settings [Yue et al., 2012, Saha and Gaillard, 2022, Lekang and Lamperski, 2019], where, in each round, the learner pulls two arms and observes a ranking over them. Nevertheless, although dueling bandits share similarities

to our setting, they present substantial differences. Specifically, in our model, the learner observes a ranking depending on the arms they have pulled so far. In dueling bandits, the learner observes an instantaneous comparison between the arms they have just pulled; thus, the outcome of such a comparison does not depend on the arms previously selected, as is the case of bandits with ranking feedback. As a consequence, while in bandits with ranking feedback the goal of the learner is to exploit the arm with the highest mean, in dueling bandits the goal of the learner is to select the arm winning with the highest probability. Furthermore, while we adopt the classical notion of regret used in the bandit literature to assess the theoretical properties of our algorithms, in dueling bandits, the algorithms are often evaluated with a suitable notion of regret, which differs from the classical one.

Dueling bandits have their reinforcement learning (RL) counterpart in the *preference-based reinforcement learning* (PbRL), see, e.g., [Novoseller et al., 2019] and [Wirth et al., 2017]. Interestingly, PbRL techniques differ from the standard RL approaches in that they allow an algorithm to learn from non-numerical rewards; this is particularly useful when the environment encompasses human-like entities [Chen et al., 2022]. Furthermore, *preference-based reinforcement learning* provides a bundle of results, ranging from theory [Xu et al., 2020] to practice [Christiano et al., 2017, Lee et al., 2021]. In PbRL, preferences may concern both states and actions; contrariwise, our framework is stateless since the rewards gained depend only on the action taken during the learning dynamic. Moreover, the differences outlined between dueling bandits and bandits with ranking feedback still hold for preference-based reinforcement learning, as preferences are considered between observations instead of the empirical mean of the accumulated rewards.

Paper structure. The paper is structured as follows. In Section 2, we report the problem formulation, the setting and the necessary notation. In Section 3, we study the stochastic setting, that is, when the rewards are sampled from fixed distributions. In Section 3.1, we present the lower bound for the instance-dependent setting. Thus, in Section 3.2, we propose our algorithm and we show that it achieves a tight regret bound. In Section 3.3, we show a trade-off between the regret upper-bound achieved in the instance-dependent and instance-independent case. Finally, in Section 3.4, we present our algorithm for the instance-independent case, which achieves optimal regret bound when the rewards are Gaussian. In Section 4, we study the adversarial setting, that is, when no statistical assumptions are made on the rewards. Thus, we present our impossibility result, namely, no algorithm can achieve sublinear regret in adversarial bandits with ranking feedback. In Appendix A and Appendix B, we report the omitted proofs for the stochastic setting. In Appendix C, we report the omitted proof for the adversarial setting. Finally, in Appendix D, we report the empirical evaluation of our algorithms.

2 Problem formulation

In this section, we formally state the model of bandits with ranking feedback and discuss the learner-environment interaction. Subsequently, we define policies and the notion of regret both in the *stochastic* and in the *adversarial* settings.

Setting and interaction. Differently from standard bandits—see, e.g., [Lattimore and Szepesvari, 2017]—in which the learner observes the *reward* associated with the pulled arm, in bandits with ranking feedback the learner can only observe a *ranking* over the arms based on the previous pulls. Formally, we assume the learner-environment interaction to unfold as follows.²

- (i) At every round $t \in [T]$, where T is the time horizon, the learner chooses an arm $i_t \in \mathcal{A} := [n]$, where \mathcal{A} is the set of available arms and $n = |\mathcal{A}| < +\infty$.
- (ii) We study both stochastic and adversarial rewards. In the stochastic setting, the environment draws the reward $r_t(i_t)$ associated with arm i_t from a probability distribution ν_{i_t} , i.e., $r_t(i_t) \sim \nu_{i_t}$, whereas, in the adversarial setting, $r_t(i_t)$ is chosen adversarially by an opponent from a bounded set of reward functions.
- (iii) There is a bandit feedback on the reward of the arm $i_t \in \mathcal{A}$ pulled at round t leading to the estimate of the empirical mean of i_t as follows:

$$\hat{r}_t(i) := \frac{\sum_{j \in \mathcal{W}_t(i)} r_j(i)}{Z_i(t)},$$

²Given $n \in \mathbb{N}_{>0}$ we denote with $[n] := \{1, \dots, n\}$.

where $\mathcal{W}_t(i) := \{\tau \in [t] \mid i_\tau = i\}$ and $Z_i(t) := |\mathcal{W}_t(i)|$.³ However, the learner observes the rank over the empirical means $\{\hat{r}_t(i)\}_{i \in \mathcal{A}}$. We denote with $\mathcal{S}_{\mathcal{A}}$ the set containing all the possible permutations of the elements of set \mathcal{A} . Formally, we assume that the ranking $\mathcal{R}_t \in \mathcal{S}_{\mathcal{A}}$ observed by the learner at round t is such that:

$$\hat{r}_t(\mathcal{R}_{t,i}) \geq \hat{r}_t(\mathcal{R}_{t,j}) \quad \forall t \in [T] \quad \forall i, j \in [n] \text{ s.t. } i \geq j,$$

where $\mathcal{R}_{t,i} \in \mathcal{A}$ denotes the i -th element in the ranking \mathcal{R}_t at round $t \in [T]$.

For the sake of clarity, we provide an example to illustrate bandits with ranking feedback and the corresponding learner-environment interaction.

Example. We consider an environment with two arms, *i.e.*, $\mathcal{A} = \{1, 2\}$, in which the learner plays the first action at rounds $t = 1$ and $t = 3$ and the second action at round $t = 2$, so that $\mathcal{W}_3(1) = \{1, 3\}$ and $\mathcal{W}_3(2) = \{2\}$. Let $r_1(1) = 1$ and $r_3(1) = 5$ be the rewards when playing the first arm at rounds $t = 1$ and $t = 3$, respectively, while let $r_2(2) = 5$ be the reward when playing the second arm at round $t = 2$. The empirical means of the two arms and resulting rankings at every round $t \in [3]$ are given by:

$$\begin{cases} \hat{r}_t(1) = 1, \hat{r}_t(2) = 0 & \mathcal{R}_t = \langle 1, 2 \rangle & t = 1 \\ \hat{r}_t(1) = 1, \hat{r}_t(2) = 5 & \mathcal{R}_t = \langle 2, 1 \rangle & t = 2 \\ \hat{r}_t(1) = 3, \hat{r}_t(2) = 5 & \mathcal{R}_t = \langle 2, 1 \rangle & t = 3 \end{cases}$$

Policies and regret. At every round t , the action played by the learner is prescribed by a policy π . In both the stochastic and adversarial settings, we let the policy π be a randomized map from the history of the interaction $H_{t-1} = (\mathcal{R}_1, i_1, \mathcal{R}_2, i_2, \dots, \mathcal{R}_{t-1}, i_{t-1})$ to the set of all the probability distributions with support \mathcal{A} . Formally, we let $\pi : H_{t-1} \rightarrow \Delta(\mathcal{A})$, for $t \in [T]$, such that $i_t \sim \pi(H_{t-1})$. As it is customary in bandits, the learner's goal is to design a policy π minimizing the cumulative expected regret, whose formal definition is as follows:

$$R_T(\pi) = \mathbb{E} \left[\sum_{t=1}^T r_t(i^*) - r_t(i_t) \right],$$

where the expectation is over the randomness of both the policy and environment in the stochastic setting, and we let $i^* \in \arg \max_{i \in \mathcal{A}} \mu_i$ with $\mu_i = \mathbb{E}[v_i]$, whereas the expectation is over the randomness of the policy in the adversarial setting and we let $i^* \in \arg \max_{i \in \mathcal{A}} \sum_{t=1}^T r_t(i)$. For the sake of simplicity, from here on, we omit the dependence on π , referring to $R_T(\pi)$ as R_T . The impossibility of observing the reward realizations raises several technical difficulties when designing no-regret algorithms since the approaches adopted for standard (non-ranking) bandits do not generalize to our case. In the following sections, we discuss how the lack of this information degrades the performance of the algorithms when the feedback is ranking.

3 Analysis in the stochastic setting

Initially, we observe that approaches based on *optimism-vs.-uncertainty*, such as UCB1, might be challenging to apply within our framework. This is because the learner lacks the information to estimate the reward associated with an arm, making it difficult to infer a confidence bound. Therefore, the most popular class of algorithms one can employ in bandits with ranking feedback is that of *explore-then-commit* (EC) algorithms (see, *e.g.*, Auer et al. [2002]), where the learner either exploits a single arm or explores the others according to a deterministic or randomized exploration strategy.

In the following, we distinguish the instance-dependent case from the instance-independent one. In particular, we provide two algorithms, each guaranteeing a sublinear regret in one of the two cases.

3.1 Instance-dependent lower bound

It is well-known that standard bandits admit algorithms guaranteeing a regret that is logarithmic in time horizon T in the instance-dependent case. We show in this section that such a result does not

³Note that the latter definition is well-posed as long as $|\mathcal{W}_t(i)| > 0$. For each $i \in \mathcal{A}$ and $t \in [T]$ such that $|\mathcal{W}_t(i)| = 0$, we let $\hat{r}_t(i) = 0$.

hold when the feedback is provided as a ranking. More precisely, our result rules out the possibility of having a logarithmic regret. However, in the next section, we prove that we can get a regret whose dependence on T is arbitrarily close to a logarithm, thus showing that the extra cost one has to pay in the instance-dependent case to deal with ranking feedback is asymptotically negligible in T .

Our impossibility result exploits a connection between random walks and arms' cumulative rewards. Formally, we define an (asymmetric) random walk as follows.

Definition 1. A random walk is a stochastic process $\{G_t\}_{t \in \mathbb{N}}$ such that:

$$G_t = \begin{cases} 0 & t = 1 \\ G_{t-1} + \epsilon_t & t > 1 \end{cases},$$

where $\{\epsilon_t\}_{t \in \mathbb{N}}$ is an i.i.d. sequence of random variables, and $\mathbb{E}[\epsilon_t]$ is the drift of the random walk.

We model the cumulative reward collected by a specific arm during the learning process as a random walk, where the drift represents the expected reward associated with that arm. Let us notice that, in bandits where the feedback is not given as a ranking, the learner can completely observe the evolution of the random walks, being able to observe the realizations of the reward associated with each pulled arm. Such observations allow the learner to estimate the difference between the performance of each pair of arms. For instance, the learner can observe whether two arms perform similarly or, instead, whether the gap between their performances is significant. Differently, in our case, the learner only observes the rank without quantify numerically the performance.

This loss of information raises several technical issues that are crucial, especially when the random walks never switch. Intuitively, in bandits with ranking feedback, we can observe how close the expected rewards of two arms are only by observing subsequent switches of their positions in the ranking. However, there is a strictly positive probability that two random walks never switch (thus leading to no intersection) when they have a different drift $\mathbb{E}[\epsilon_t]$ and therefore we may not evaluate how two arms are close. This is shown in the following lemma.

Lemma 1 (Separation lemma). Let G_t, G'_t be two independent random walks defined as:

$$G_{t+1} = G_t + \epsilon_t \quad \text{and} \quad G'_{t+1} = G'_t + \epsilon'_t,$$

where $G_0 = G'_0 = 0$ and the drifts satisfy $\mathbb{E}[\epsilon_t] = p > q = \mathbb{E}[\epsilon'_t]$. Then:

$$\mathbb{P}\left(\forall t, t' \in \mathbb{N}^* \quad G_t/t \geq G'_{t'}/t'\right) \geq c(p, q) > 0.$$

The exact value of the constant $c(p, q)$ depends only on the two means p and q , as well as the probability distribution defining such drifts. In the simpler case of Bernoulli distributions, this quantity can be derived in closed form, as shown in Lemma 4 in the appendix. The rationale behind the above lemma is that when considering two random walks with different drifts, there exists a *separating line* between them with a strictly positive probability. Therefore, with a non-negligible probability, the empirical mean corresponding to the process with the higher drift upper bounds the empirical mean of the process with the lower drift for the entire duration of the learner-environment interaction. In bandits with ranking feedback, such a separation lemma shows that the problem of distinguishing two different instances is harder than in the standard, non-ranking feedback case. Before stating our result, as is customary in bandit literature, we denote with $\Delta_i := \mu_i^* - \mu_i$, where we let $i^* \in \arg \max_{i \in \mathcal{A}} \mu_i$ and $\mu_i := \mathbb{E}[\nu_i]$. Now, we can state the following result for the instance-dependent case.

Theorem 1 (Instance-dependent lower bound). Let π be any policy for the bandits with ranking feedback, then, for any $C(\cdot) : [0, +\infty) \rightarrow [0, +\infty)$, there is $\{\Delta_i\}_{i \in [n]}$ and a time horizon $T > 0$ such that $R_T > \sum_{i=1}^n C(\Delta_i) \log(T)$.

To prove Theorem 1, we consider two stochastic bandit instances, each with two available arms. In both instances, the first arms have the same distribution, while the distributions of the second arms are slightly different. Specifically, the first arm is optimal in the first instance, and the second arm is optimal in the second instance. Thus, if the learner employs a policy that achieves $\mathcal{O}(\log(T))$ in the first instance, such a policy may result in super-logarithmic regret in the second one as a consequence of Lemma 1.

3.2 Instance-dependent upper bound

We introduce the Dynamical Ranking Exploration-Exploitation algorithm (DREE). The pseudo-code is provided in Algorithm 1. As usual in bandit algorithms, in the first n rounds, a pull for each arm is performed (Lines 2–4). At every subsequent round $t > n$, the exploitation/exploration tradeoff is addressed by playing the best arm according to the received feedback unless there is at least one arm whose number of pulls at t is smaller than a superlogarithmic function $f(t) : (0, \infty) \rightarrow \mathbb{R}_+$.⁴ More precisely, the algorithm plays an arm i at round t if it has been pulled less than $f(t)$ times (Lines 5–6), where ties due to multiple arms pulled less than $f(t)$ times are broken arbitrarily. Instead, if all arms have been pulled at least $f(t)$ times, the arm in the highest position of the last ranking feedback is pulled (Lines 7–9). Each round ends once the learner receives the feedback in terms of ranking over the arms (Line 10). Let us observe that the exploration strategy of Algorithm 1 is deterministic, and the only source of randomness concerns the realization of the arms' rewards.

We state the following result, providing the upper regret bound of Algorithm 1 as a function of f .

Theorem 2 (Instance-dependent upper bound). *Assume that the reward distribution of every arm is 1-subgaussian. Let $f : (0, \infty) \rightarrow \mathbb{R}$ be a superlogarithmic function in t , then there is a term $C(f, \Delta_i)$ for each sub-optimal arm $i \in [n]$ which does not depend on T , such that Algorithm 1 satisfies $R_T \leq f(T) \sum_{i=1}^n \Delta_i + \log(T) \sum_{i=1}^n C(f, \Delta_i)$.*

To minimize the asymptotic dependence in T of the cumulative regret suffered by the algorithm, we can choose, e.g., $f : (0, \infty) \rightarrow \mathbb{R}$ as $f(t) = \log(t)^{1+\delta}$, where parameter $\delta > 0$ is as small as possible. However, the minimization of δ comes at the cost of increasing the terms $C(f, \Delta_i)$ as they grow exponentially as $\delta > 0$ goes to zero as long as $\Delta_i < 1$. In particular, the terms $C(f, \Delta_i)$ are defined as stated in the following corollary.

Corollary 3. *Let $\delta > 0$ and $f(t) = \log(t)^{1+\delta}$ be the superlogarithmic function used in Algorithm 1, then we have:*

$$C(f, \Delta_i) = \frac{2\Delta_i \left(e^{((2/\Delta_i^2)^{1/\delta})} + 1 \right)}{1 - e^{-\Delta_i^2/2}}$$

We remark that the term $C(f, \Delta_i)$ depends exponentially on Δ_i , suggesting that $C(f, \Delta_i)$ may be large even when adopting values of δ that are not arbitrarily close to zero. At this point, a natural question arises. Is such an exponential dependence with respect to the gap Δ_i unavoidable, or is it a consequence of the kind of feedback the learner receives? In the next section, we answer this question.

Furthermore, let us observe that Algorithm 1 satisfies important properties in the instance-dependent stochastic setting. More precisely, (i) it matches the instance-dependent regret lower-bound, since $f(\cdot)$ can be chosen arbitrarily close to $\log(t)$, (ii) it works without requiring the knowledge of the time horizon T , thus being an *any-time algorithm*.

Algorithm 1 Dynamical Ranking Exploration-Exploitation (DREE)

- 1: **for** $t \in [T]$ **do**
 - 2: **if** $t \leq n$ **then**
 - 3: play arm i_t
 - 4: **end if**
 - 5: **if** There is an arm i played less than $f(t)$ times **then**
 - 6: Play $i_t = i$
 - 7: **else**
 - 8: Play $i_t = \mathcal{R}_{t-1,1}$
 - 9: **end if**
 - 10: Receive updated ranking \mathcal{R}_t
 - 11: **end for**
-

3.3 Instance dependent/independent trade-off

In this section, we provide a negative result, showing that *no algorithm* can perform well in both the instance-dependent and instance-independent cases, thus suggesting that the two cases need to be studied separately. Initially, we state the following result that relates to the upper regret bounds in the two (instance-dependent/independent) cases.

Theorem 4 (Instance Dependent/Independent Trade-off). *Let π be any policy for the bandits with ranking feedback problem. If π satisfies the following properties:*

⁴A function $f(t)$ is superlogarithmic when $\lim_{t \rightarrow \infty} \frac{f(t)}{\log(t)} = +\infty$.

- (instance-dependent upper regret bound) $R_T \leq \sum_{i=1}^n C(\Delta_i)T^\alpha$
- (instance-independent upper regret bound) $R_T \leq nCT^\beta$

then, $2\alpha + \beta \geq 1$, where $\alpha, \beta \geq 0$.

Proof. (Sketch) Let $p_1 = 0.5, p_2 = 0.5 - \varepsilon, p_2^* = 0.5 + \varepsilon$ and consider two instances:

$$P : \begin{cases} \nu_1 = Be(p_1) \\ \nu_2 = Be(p_2) \end{cases} \quad P^* : \begin{cases} \nu_1 = Be(p_1) \\ \nu_2 = Be(p_2^*) \end{cases}$$

Clearly, the optimal arm is arm 1 in instance P , while arm 2 is optimal in P^* . We then define the event:

$$E_t = \bigcap_{\tau=1}^t \{\mathcal{R}_\tau = \langle 1, 2 \rangle\}.$$

In the first instance, E_t corresponds to the optimal arm being ranked in the first position for the first t time steps, while the opposite holds for the second instance. The main question at this point reduces to: how many times should a policy select the arm ranked in the last position to distinguish between the two instances? To ensure that the instance-independent regret is sublinear in the first instance, we have to guarantee that $\mathbb{E}_P[Z_2(t)|E_t] \leq Ct^{\alpha+\eta}$ for quantity η specified in the proof. Indeed, if in the first instance we pull the arm 2 too many times under the event E_t , we may suffer a potentially large instance-dependent regret bound. Then, we prove that if the previous consideration holds, in the instance P^* , we have:

$$R_t \geq \Omega(t^{1-2\rho\alpha}),$$

for a constant ρ close to one. This lower bound on the instance independent regret entails that $\beta \geq 1 - 2\alpha$, or, equivalently $\beta + 2\alpha \geq 1$. \square

From Theorem 4, we can easily infer the following impossibility result.

Corollary 5. *There is no algorithm for bandits with ranking feedback achieving both subpolynomial regret in the instance-dependent case, i.e., $\forall \alpha > 0, \exists C(\cdot) : R_T \leq \sum_{i=1}^n C(\Delta_i)T^\alpha$, and sublinear regret in the instance-independent case.*

To ease the interpretation of Corollary 5, we discuss the performance of Algorithm 1 in the instance-independent case in the following result.

Corollary 6. *For every choice of $\delta > 0$ in $f(t) = \log(t)^{1+\delta}$, there is no value of $\eta > 0$ for which Algorithm 1 achieves an instance-independent regret bound of the form $R_T \leq \mathcal{O}(T^{1-\eta})$.*

The above result shows that Algorithm 1 suffers from linear regret in T in the instance-independent case except for logarithmic terms.

Moreover, the following corollary of Theorem 4 answers the question raised in the previous section. We can prove that the terrible dependence on the gaps Δ_i is not a feature of Algorithm 1; instead, it cannot be avoided until the instance-dependent regret has a good order in T .

Corollary 7. *Let π be any policy for the bandits with ranking feedback problem that satisfies and instance-dependent upper regret bound of the form $R_T \leq \sum_{i=1}^n C(\Delta_i)f(T)$, where $f(T)$ is sub-polynomial. Then, $C(\Delta)$ is super-polynomial in $1/\Delta$.*

Proof. We prove the opposite implication, namely that if $C(\Delta)$ is polynomial in $1/\Delta$, then $f(T)$ cannot be sub-polynomial in T . By assumption, in case of two arms with just one gap Δ ,

$$R_T \leq \sum_{\ell=1}^p C_\ell \Delta^{-\ell} f(T),$$

which implied that the instance independent regret can be bounded in the following way,

$$\begin{aligned} R_T &\leq \sup_{\Delta > 0} \min \left\{ \sum_{\ell=1}^p C_\ell \Delta^{-\ell} f(T), \Delta T \right\} \\ &\leq \sum_{\ell=1}^p \sup_{\Delta > 0} \min \{ C_\ell \Delta^{-\ell} f(T), \Delta T \}. \end{aligned}$$

Notice that, for $\Delta \geq T^{-1/(\ell+1)}$ the first term is less than $CT^{\frac{\ell}{\alpha+1}} f(T)$, while for $\Delta \leq T^{-1/(\ell+1)}$, the second one is less than $T^{\frac{\ell}{\alpha+1}}$. Therefore, the full instance independent regret is bounded by:

$$R_T \leq \sum_{\ell=1}^p C_\ell T^{\frac{\ell}{\alpha+1}} f(T),$$

which is polynomial in T . If, by contradiction, $f(T)$ were sub-polynomial, this bound would be sublinear (growing roughly as $\mathcal{O}(T^{\frac{p}{p-1}})$), but this contradicts the result of Corollary 5. \square

3.4 Instance-independent upper bound

The impossibility result stated by Corollary 5 pushes for the need for an algorithm guaranteeing a sublinear regret in the instance-independent case. Initially, we observe that the standard Explore-then-Commit (EC) algorithm [Auer et al., 2002] can be applied, achieving a regret bound $\mathcal{O}(T^{2/3})$ in the instance-independent case. Let us briefly summarize the functioning of the EC algorithm. It divides the time horizon into two phases as follows: (i) *exploration phase*: the arms are pulled uniformly for the first $m \cdot n$ rounds, where m is a parameter of the algorithm one can tune to minimize the regret; (ii) *commitment phase*: the arm maximizing the estimated reward is pulled. In the case of bandits with ranking feedback, the EC algorithm explores the arms in the first $m \cdot n$ rounds and subsequently pulls the arm in the first position of the ranking feedback received at round $t = m \cdot n$. As is customary in standard (non-ranking) bandits, the best regret bound can be achieved by setting $m = \lceil T^{2/3} \rceil$, thus obtaining $\mathcal{O}(T^{2/3})$. We show that we can get a regret bound better than that of the EC algorithm. In particular, we provide the Ranking Logarithmic Phased Elimination (R-PLE) algorithm, which breaks the barrier of $\mathcal{O}(T^{2/3})$ guaranteeing a regret $\tilde{\mathcal{O}}(\sqrt{T})$ when neglecting logarithmic terms. Due to the mathematical instruments involved, the proof of this regret bound only holds for the case of Gaussian noise, as the ones presented in a similar setting by Garivier et al. [2016]. The pseudocode of R-PLE is reported in Algorithm 2.

R-LPE algorithm. In order to properly analyze the algorithm, we need to introduce the two following definitions. Initially, we introduce the definition of the loggrid set as follows,

Definition 2 (Loggrid). *Given two real numbers a, b s.t. $a < b$ and a constant value T , we define*

$$LG(a, b, T) := \left\{ \lfloor T^{\lambda_j b + (1-\lambda_j)a} \rfloor : \lambda_j = \frac{j}{\lfloor \log(T) \rfloor}, \forall j = 0, \dots, \lfloor \log(T) \rfloor \right\}.$$

Next, we give the notion of active set, which the algorithm employs to cancel out sub-optimal arms.

Definition 3 (Filtering condition). *Let S be the active set of the algorithm, at a certain timestep. We define a timestep t fair if all active arms have been pulled the same number of times times. In any fair timestep t , we define the active set $\mathcal{F}_t(\zeta)$ the set of arms*

$$\mathcal{F}_t(\zeta) := \left\{ i \in S : \forall j \in S \quad \sum_{\tau=1 \text{ s.t. } \tau \text{ fair}}^t \{\mathcal{R}_\tau(i) > \mathcal{R}_\tau(j)\} \geq \zeta \right\}.$$

*This condition will be called **filtering condition**.*

Initially, we observe that R-LPE differs from Algorithm 1, as it takes into account the whole history of the process and not only the last ranking \mathcal{R}_t received. It also requires the knowledge of T . Set S denotes the active set of arms used by the algorithm. Initially, set S comprises all the possible arms available in the problem (Line 1). Furthermore, the set which drives the update of the decision space S , namely \mathcal{L} , is initialized as the loggrid built on parameters $1/2, 1, T$ (Line 2). At every round $t \in [T]$, R-LPE chooses the arm from active set S with the minimum number of pulls, namely i s.t. $Z_i(t)$ is minimized (Line 4); ties are broken by index order. Next, the number of times arm i_t has been pulled, namely $Z_i(t)$, is updated accordingly (Line 5). The peculiarity of the algorithm is that set S changes every time the condition $\min_i Z_i(t) \in \mathcal{L}$ is satisfied (Line 7). When the aforementioned condition is met, the set of active arms S is filtered to avoid the exploration on sub-optimal arms. Precisely, S is filtered given the time dependent parameter α (Line 8-9).

Regret bound. We state the following theorem providing a regret bound to Algorithm 2 in the instance-independent case.

Theorem 8. *In the stochastic bandits with ranking feedback setting, when the noise is Gaussian, Algorithm 2 achieves $R_T \leq 62n^4 \log(T)^2 T^{1/2}$.*

Proof. (Sketch) To prove the theorem, we define, for every pair of indices $i, j \in [n]$, the event:

$$E_{ij}^\psi := \begin{cases} E_{ij} & \mu_j - \mu_i > \psi \\ \emptyset & \text{else,} \end{cases} \quad (1)$$

where the E_{ij} corresponds to the event in which arm i eliminates arm j at some point in the process, while ψ is a constant defined in the following. The probability that at least one of these events holds is bounded (by Lemma 7 and employing a Union Bound) as $\mathbb{P}(\Psi) := \mathbb{P}\left(\bigcup_{i \neq j, i, j \in [n]} E_{ij}^\psi\right) \leq \mathcal{O}(n^2 \log(T)^2 T^{-1/2} \psi^{-1})$, as their number is at most $n(n-1)/2$.

Therefore, if the complement of the event Ψ holds, there exists an arm i^* with a gap (w.r.t. the first arm) less than $(n-1)\psi$, which is not eliminated until the last round. This is because at most $n-1$ eliminations can happen, and if the complement of the event Ψ holds such eliminations concern pairs of arms with a difference in mean of at most ψ .

Since the arm $i^* \in [n]$ is not eliminated until the last round under the event Ψ^C , the probability of event $E_{ii^*}^*$, corresponding to the event in which the suboptimal arm $i \in [n]$ survives for more than $\mathcal{O}(\log(T)T^{1/2}\Delta_{ii^*}^{-1})$ pulls, is bounded by $2T^{-1/2}$ thanks to Lemma 8. As a result, employing a union bound over all possible values of i^* , we can say that the probability that any event $E_{ii^*}^*$ with $i^* \in [n]$ happens is at most $2(n-1)T^{-1/2}$. Thus, fixing $\psi = \Delta_i/(2(n-1))$ ensures that $\Delta_{ii^*}^{-1} \geq \Delta_i/n$, which entails $\mathbb{P}(Z_i(T) \geq \mathcal{O}(\log(T)T^{1/2}\Delta_i^{-1})) \leq \mathcal{O}(n^3 \log(T)^2 T^{-1/2} \Delta_i^{-1})$, and thus, it holds $\mathbb{E}[\Delta_i Z_i(T)] = \mathcal{O}(n^3 \log(T)^2 T^{-1/2})$. Finally, using the Regret Decomposition Lemma [Lattimore and Szepesvari, 2017] we can conclude the proof. Proving the two lemmas, however, is not trivial since it requires to see the whole process as a discretization of a biased Brownian motion, and then applying results for this kind of stochastic processes. \square

At first glance, the result presented in Theorem 8 may seem unsurprising. Indeed, there are several elimination algorithms achieving $\mathcal{O}(\sqrt{T})$ regret bounds in different bandit settings (see, for example, [Auer and Ortner, 2010, Lattimore et al., 2020, Li and Scarlett, 2022]). Nevertheless, our setting poses several additional challenges compared to existing ones. For instance, in our framework, it is not possible to rely on concentration bounds, as the current feedback is heavily correlated with the past ones. This is precisely the reason for the anomalous growth of the regret in terms of the number of arms n : due to the extremely correlated feedback, two union bounds are necessary through the last proof, and this leads to an increase in the dependence of the order of n . In the impossibility of using concentration inequalities, our analysis employs novel arguments, drawing from recent results in the theory of Brownian Motions, which allow to properly model the particular ranking feedback.

4 Analysis in the adversarial setting

We focus on bandits with ranking feedback in adversarial settings. In particular, we show that no algorithm provides sublinear regret without statistical assumptions on the rewards.

Theorem 9. *In adversarial bandits with ranking feedback, no algorithm achieves $o(T)$ regret with respect to the best arm in hindsight with a probability of $1 - \epsilon$ for any $\epsilon > 0$.*

Proof. (Sketch) The proof introduces three instances in an adversarial setting in a way that no algorithm can achieve sublinear regret in all the three. The main reason behind such a negative

Algorithm 2 Ranking Logarithmic Phased Elimination (R-LPE)

- 1: Initialize $S = [n]$
 - 2: Initialize $\mathcal{L} = LG(1/2, 1, T)$
 - 3: **for** $t \in [T]$ **do**
 - 4: Play $i_t \in \arg \min_{i \in S} Z_i(t)$
 - 5: Update $Z_i(t)$ number of times i_t has been pulled
 - 6: Observe ranking \mathcal{R}_t
 - 7: **if** $\min_{i \in S} Z_i(t) \in \mathcal{L}$ **then**
 - 8: $\alpha = \frac{\log(\min_{i \in S} Z_i(t))}{\log(T)} - \frac{1}{2}$
 - 9: $S = \mathcal{F}_t(T^{2\alpha})$
 - 10: **end if**
 - 11: **end for**
-

result is that ranking feedback obfuscates the value of the rewards so as not to allow the algorithm to distinguish two or more instances where the rewards are non-stationary. The three instances employed in the proof are divided into three phases such that the instances are similar in terms of rewards for the first two phases, while they are extremely different in the third phase. In summary, if the learner receives the same ranking when playing in two instances with different best arms in hindsight, it is not possible to achieve a small regret in both of them. \square

Acknowledgments

This paper is supported by the Italian MIUR PRIN 2022 Project “Targeted Learning Dynamics: Computing Efficient and Fair Equilibria through No-Regret Algorithms”, by the FAIR (Future Artificial Intelligence Research) project, funded by the NextGenerationEU program within the PNRR-PE-AI scheme (M4C2, Investment 1.3, Line on Artificial Intelligence), and by the EU Horizon project ELIAS (European Lighthouse of AI for Sustainability, No. 101120237).

References

- P. Auer, N. Cesa-Bianchi, Y. Freund, and R.E. Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, pages 322–331, 1995. doi: 10.1109/SFCS.1995.492488.
- Peter Auer and Ronald Ortner. Ucb revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61(1-2):55–65, 2010.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002.
- Paolo Baldi. Stochastic calculus. In *Stochastic Calculus*, pages 215–254. Springer, 2017.
- Viktor Bengs, Robert Busa-Fekete, Adil El Mesaoudi-Paul, and Eyke Hüllermeier. Preference-based online learning with dueling bandits: A survey. 2018. doi: 10.48550/ARXIV.1807.11398. URL <https://arxiv.org/abs/1807.11398>.
- Xiaoyu Chen, Han Zhong, Zhuoran Yang, Zhaoran Wang, and Liwei Wang. Human-in-the-loop: Provably efficient preference-based reinforcement learning with general function approximation. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 3773–3793. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/chen22ag.html>.
- Paul Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences, 2017. URL <https://arxiv.org/abs/1706.03741>.
- Richard Combes, M. Sadegh Talebi, Alexandre Proutiere, and Marc Lelarge. Combinatorial bandits revisited, 2015. URL <https://arxiv.org/abs/1502.03475>.
- Johannes Fürnkranz and Eyke Hüllermeier, editors. *Preference Learning*. Springer, 2010. ISBN 978-3-642-14124-9. doi: 10.1007/978-3-642-14125-6. URL <https://doi.org/10.1007/978-3-642-14125-6>.
- Aurélien Garivier, Tor Lattimore, and Emilie Kaufmann. On explore-then-commit strategies. *Advances in Neural Information Processing Systems*, 29, 2016.
- Robert Kleinberg, Aleksandrs Slivkins, and Eli Upfal. Bandits and experts in metric spaces. *J. ACM*, 66(4), may 2019. ISSN 0004-5411. doi: 10.1145/3299873. URL <https://doi.org/10.1145/3299873>.
- Tor Lattimore and Csaba Szepesvari. Bandit algorithms. 2017. URL <https://tor-lattimore.com/downloads/book/book.pdf>.

- Tor Lattimore, Csaba Szepesvari, and Gellert Weisz. Learning with good feature representations in bandits and in rl with a generative model. In *International Conference on Machine Learning*, pages 5662–5670. PMLR, 2020.
- Kimin Lee, Laura Smith, Anca Dragan, and Pieter Abbeel. B-pref: Benchmarking preference-based reinforcement learning, 2021. URL <https://arxiv.org/abs/2111.03026>.
- Tyler Lekang and Andrew Lampserski. Simple algorithms for dueling bandits, 2019. URL <https://arxiv.org/abs/1906.07611>.
- Zihan Li and Jonathan Scarlett. Gaussian process bandit optimization with few batches. In *International Conference on Artificial Intelligence and Statistics*, pages 92–107. PMLR, 2022.
- Ellen R. Novoseller, Yibing Wei, Yanan Sui, Yisong Yue, and Joel W. Burdick. Dueling posterior sampling for preference-based reinforcement learning, 2019. URL <https://arxiv.org/abs/1908.01289>.
- Aadirupa Saha and Pierre Gaillard. Versatile dueling bandits: Best-of-both-world analyses for online learning from preferences, 2022. URL <https://arxiv.org/abs/2202.06694>.
- Lajos Takács. On a generalization of the arc-sine law. *The Annals of Applied Probability*, 6(3): 1035–1040, 1996.
- Claire Vernade, Olivier Cappé, and Vianney Perchet. Stochastic bandit models for delayed conversions. *ArXiv*, abs/1706.09186, 2017.
- Claire Vernade, Andras Gyorgy, and Timothy Mann. Non-stationary delayed bandits with intermediate observations. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 9722–9732. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/vernade20b.html>.
- Christian Wirth, Riad Akrou, Gerhard Neumann, and Johannes Fürnkranz. A survey of preference-based reinforcement learning methods. *J. Mach. Learn. Res.*, 18(1):4945–4990, jan 2017. ISSN 1532-4435.
- Yichong Xu, Ruosong Wang, Lin F. Yang, Aarti Singh, and Artur Dubrawski. Preference-based reinforcement learning with finite-time guarantees, 2020. URL <https://arxiv.org/abs/2006.08910>.
- Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. The k-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012. ISSN 0022-0000. doi: <https://doi.org/10.1016/j.jcss.2011.12.028>. URL <https://www.sciencedirect.com/science/article/pii/S0022000012000281>. JCSS Special Issue: Cloud Computing 2011.

A Proofs of instance dependent stochastic analysis

A.1 Proof of instance dependent lower bound and lemmas

Lemma 2 (Separation lemma). *Let G_t, G'_t be two independent random walks defined as:*

$$G_{t+1} = G_t + \epsilon_t \quad \text{and} \quad G'_{t+1} = G'_t + \epsilon'_t,$$

where $G_0 = G'_0 = 0$ and the drifts satisfy $\mathbb{E}[\epsilon_t] = p > q = \mathbb{E}[\epsilon'_t]$. Then:

$$\mathbb{P}\left(\forall t, t' \in \mathbb{N}^* \quad G_t/t \geq G'_{t'}/t'\right) \geq c(p, q) > 0.$$

Proof. Let us consider the random walk

$$\tilde{G}_{t+1} = \tilde{G}_t + \epsilon_t - \frac{p+q}{2}.$$

Being $\mathbb{E}[\epsilon_t - \frac{p+q}{2}] > 0$, from the well-known fact that a random walk with drift is transient, we know that there is a strictly positive probability $c_1(p, q)$ that $\{\tilde{G}_t > 0, \forall t > 0\}$.

On the opposite side, we can see that the random walk

$$\tilde{G}'_{t+1} = \tilde{G}'_t + \epsilon'_t - \frac{p+q}{2}$$

satisfies the opposite inequality, $\mathbb{E}[\epsilon'_t - \frac{p+q}{2}] < 0$, so that, for the same reason, the event $\{\tilde{G}'_t < 0, \forall t > 0\}$ has a strictly positive probability $c_2(p, q)$.

Therefore, being the two processes independent, one has

$$\mathbb{P}\left(\bigcap_{t, t'=1}^{\infty} \{\tilde{G}_{t+1} > 0, \tilde{G}'_{t'+1} < 0\}\right) > 0,$$

which entails:

$$\begin{aligned} c_1(p, q)c_2(p, q) &= \mathbb{P}\left(\bigcap_{t, t'=1}^{\infty} \{\tilde{G}_t > 0, \tilde{G}'_{t'} < 0\}\right) \\ &= \mathbb{P}\left(\bigcap_{t, t'=1}^{\infty} \left\{G_t - t\frac{p+q}{2} > 0, G'_{t'} - t'\frac{p+q}{2} < 0\right\}\right) \\ &= \mathbb{P}\left(\bigcap_{t, t'=1}^{\infty} \left\{G_t/t > \frac{p+q}{2}, G'_{t'}/t' < \frac{p+q}{2}\right\}\right) \\ &\leq \mathbb{P}\left(\bigcap_{t, t'=1}^{\infty} \{G_t/t > G'_{t'}/t'\}\right). \end{aligned}$$

which can be reformulated as in the statement. □

In order to prove the lower bound, we will need the following Lemma.

Lemma 3. *Let $\{X_n\}_n$ be a sequence of i.i.d. Bernoulli random variables. Then, for every event $E \in \mathcal{F}_n$, where \mathcal{F}_n is the filtration generated by X_1, \dots, X_n ,*

$$\mathbb{P}_{X_1, \dots, X_n \sim Be(p)}(E) \leq \mathbb{P}_{X_1, \dots, X_n \sim Be(p')}(E) \max\left(\frac{p}{p'}, \frac{1-p}{1-p'}\right)^n$$

Proof. Without loss of generality, let us assume that $p' < p$.

$$\begin{aligned}
\mathbb{P}_{X_1, \dots, X_n \sim \text{Be}(p)}(E) &= \int_{\{0,1\}^n} \mathbf{1}_E(\underline{x}) \prod_{i=1}^n p^{x_i} (1-p)^{1-x_i} d\underline{x} \\
&\leq \int_{\{0,1\}^n} \mathbf{1}_E(\underline{x}) \prod_{i=1}^n (p/p')^{x_i} p'^{x_i} (1-p')^{1-x_i} d\underline{x} \quad (2) \\
&\leq \left(\frac{p}{p'}\right)^n \int_{\{0,1\}^n} \mathbf{1}_E(\underline{x}) \prod_{i=1}^n p'^{x_i} (1-p')^{1-x_i} d\underline{x} \\
&= \left(\frac{p}{p'}\right)^n \mathbb{P}_{X_1, \dots, X_n \sim \text{Be}(p')}(E).
\end{aligned}$$

where Inequality (2) follows from the assumption $p' < p$. The other way round can be proved substituting p and p' . \square

We can now prove the following instance dependent lower bound.

Theorem 1 (Instance-dependent lower bound). *Let π be any policy for the bandits with ranking feedback, then, for any $C(\cdot) : [0, +\infty) \rightarrow [0, +\infty)$, there is $\{\Delta_i\}_{i \in [n]}$ and a time horizon $T > 0$ such that $R_T > \sum_{i=1}^n C(\Delta_i) \log(T)$.*

Proof. Let $p_1 = 0.5$, $p_2 = 0.5 - \varepsilon$, $p_2^* = 0.5 + \varepsilon$. Let us consider two problems:

$$P : \begin{cases} \nu_1 = \text{Be}(p_1) \\ \nu_2 = \text{Be}(p_2) \end{cases} \quad P^* : \begin{cases} \nu_1 = \text{Be}(p_1) \\ \nu_2 = \text{Be}(p_2^*) \end{cases}$$

Clearly, the optimal arm is 1 for P and 2 for P^* . Let us now define the following event on the rankings received:

$$E_t = \bigcap_{\tau=1}^t \{\mathcal{R}_\tau = \langle 1, 2 \rangle\}.$$

The event E_t can be interpreted as "up to time t , we have always observed the ranking $\langle 1, 2 \rangle$ ".

Let π be any policy. Then, at least one of the following is true:

- the policy is "light tail":

$$\liminf_{t \rightarrow \infty} \frac{\sum_{\tau=1}^t \pi(\text{pull } 2 | E_\tau)}{\log(t)} = c_\pi < +\infty.$$

- the policy is "heavy tailed":

$$\limsup_{t \rightarrow \infty} \frac{\sum_{\tau=1}^t \pi(\text{pull } 2 | E_\tau)}{\log(t)} = +\infty.$$

In this latter case, it is obvious that the regret cannot be logarithmic in t , so we will focus on light-tailed policies.

Considering the first case, there is a sequence of times t_k such that

$$\lim_k \frac{\mathbb{E}_P[Z_2(t_k) | E_{t_k}]}{\log(t_k)} = c_\pi.$$

Where the expectation is over the random variables of the arms following the distribution given in (P). Therefore, by Markov's inequality, we have, for sufficiently large k ,

$$\forall h > 0 \quad \mathbb{P}_P(Z_2(t_k) < h | E_{t_k}) \geq 1 - \frac{2c_\pi \log(t_k)}{h}.$$

Now, notice that the event $Z_2(t_k) < h$ is contained in the σ -algebra generated by the first h pulls of 2 (and all the pulls of arm 1, but this is irrelevant since ν_1 corresponds to the same distribution in the two problems). Therefore, from Lemma 3 we have:

$$\begin{aligned} \forall h > 0 \quad \mathbb{P}_{P^*}(Z_2(t_k) < h) &\geq \left(\frac{0.5 - \varepsilon}{0.5 + \varepsilon}\right)^h \mathbb{P}_P(Z_2(t_k) < h) \\ &\geq \left(\frac{0.5 - \varepsilon}{0.5 + \varepsilon}\right)^h \mathbb{P}_P(Z_2(t_k) < h, E_{t_k}) \\ &= \left(\frac{0.5 - \varepsilon}{0.5 + \varepsilon}\right)^h \mathbb{P}_P(Z_2(t_k) < h | E_{t_k}) \mathbb{P}_P(E_{t_k}). \end{aligned}$$

Thus, we have

1. From the previous steps,

$$\forall h > 0 \quad \mathbb{P}_P(Z_2(t_k) < h | E_{t_k}) \geq 1 - \frac{2c_\pi \log(t_k)}{h}.$$

2. $\inf_{t>0} \mathbb{P}_P(E_t) = q > 0$ thanks to Lemma 2.

The two point together entail that

$$\forall h > 0 \quad \mathbb{P}_{P^*}(Z_2(t_k) < h) \geq q \left(1 - \frac{2c_\pi \log(t_k)}{h}\right) \left(\frac{0.5 - \varepsilon}{0.5 + \varepsilon}\right)^h.$$

Here, for every $\varepsilon > 0$ the inequality $\frac{0.5 - \varepsilon}{0.5 + \varepsilon} \geq 1 - 4\varepsilon$ holds, so that

$$\forall h > 0 \quad \left(\frac{0.5 - \varepsilon}{0.5 + \varepsilon}\right)^h \geq (1 - 4\varepsilon)^h,$$

therefore, taking $h = 4c_\pi \log(t_k)$, we have:

$$\begin{aligned} \mathbb{P}_{P^*}(Z_2(t_k) < 2c_\pi \log(t_k)) &\geq \frac{q}{2} e^{4c_\pi \log(t_k) \log(1 - 4\varepsilon)} \\ &\geq \frac{q}{2} e^{-4c_\pi \log(t_k) 4\varepsilon} \\ &= \frac{q}{2} t_k^{-4c_\pi 4\varepsilon}. \end{aligned}$$

If we then put $\varepsilon < \frac{1}{17c_\pi}$, we have this lower bound on the regret in case of instance P^* :

$$\begin{aligned} R_{t_k} |_{\sim P^*} &\geq \varepsilon \mathbb{E}_{P^*}[(t_k - Z_2(t_k))] \\ &\geq \varepsilon (t_k - 2c_\pi \log(t_k)) \mathbb{P}_{P^*}(Z_2(t_k) < 2c_\pi \log(t_k)) \\ &\geq \frac{1}{2} \varepsilon t_k \mathbb{P}_{P^*}(Z_2(t_k) < 2c_\pi \log(t_k)) \\ &\geq \frac{q}{4} \varepsilon t_k^{-16/17} t_k = \frac{q}{4} \varepsilon t_k^{1/17}. \end{aligned}$$

which grows polynomially with time. Therefore, whichever the value of $C(\Delta_i)$ (which is $C(\varepsilon)$ in this case), we can always find T in the sequence t_k such that

$$R_t |_{\sim P^*} \geq \frac{q}{4} \varepsilon t^{1/17} > C(\varepsilon) \log(t).$$

□

A.2 Proof of instance dependent upper bound

Theorem 2 (Instance-dependent upper bound). *Assume that the reward distribution of every arm is 1-subgaussian. Let $f : (0, \infty) \rightarrow \mathbb{R}$ be a superlogarithmic function in t , then there is a term $C(f, \Delta_i)$ for each sub-optimal arm $i \in [n]$ which does not depend on T , such that Algorithm 1 satisfies $R_T \leq f(T) \sum_{i=1}^n \Delta_i + \log(T) \sum_{i=1}^n C(f, \Delta_i)$.*

Proof. Let $i^* \in [n]$ be the optimal arm, and let $Z_i(t)$ the number of pulls of arm i up to time t . For any sub-optimal arm a_i , we have

$$\mathbb{E}[Z_i(t)] = \sum_{\tau=1}^t \mathbb{P}(I_\tau = i) \quad (3)$$

$$= \sum_{\tau=1}^t \mathbb{P}(I_\tau = i, Z_i(\tau-1) < f(\tau)) + \sum_{\tau=1}^t \mathbb{P}(I_\tau = i, Z_i(\tau-1) \geq f(\tau)), \quad (4)$$

where we let I_τ be the arm pulled at time $\tau \in [T]$. We split the proof in two parts, providing a bound for each term defining Equation (4).

Claim 1: The first term of Equation (4) is bounded by $f(t)$. indeed, notice that if:

$$\sum_{\tau=1}^t \mathbf{1}\{I_\tau = i, Z_i(\tau-1) < f(\tau)\} \geq f(t),$$

then, there is $t_0 \leq t$ such that $Z_i(t_0) = f(t) - 1$. Thus, we could rewrite the latter term as follows:

$$\sum_{\tau=1}^{t_0} \mathbf{1}\{I_\tau = i, Z_i(\tau-1) < f(\tau)\} + \sum_{\tau=t_0+1}^t \mathbf{1}\{I_\tau = i, Z_i(\tau-1) < f(\tau)\}$$

By definition, the first sum is bounded by $f(t)$, while the second one is bounded by

$$\sum_{\tau=t_0+1}^t \mathbf{1}\{I_\tau = i, Z_i(\tau-1) < f(t)\} = 0.$$

since, for $\tau > t_0$, $Z_i(\tau-1) \geq f(t)$.

Claim 2: The second term is bounded by $C(\Delta_i) \log(t)$ for some $C(\Delta_i)$. We know, by design of the algorithm, that the arm a_i can be pulled only if:

1. It has the highest empirical mean.
2. Every other arm has been pulled at least $f(t)$ times, including arm i^* .

In particular, defining the event:

$$E_{i,t} := \{Z_i(t) \geq f(t)\}$$

we have:

$$\mathbb{P}(I_\tau = i, Z_i(\tau-1) \geq f(\tau)) \leq \mathbb{P}(\bar{X}_{i,\tau} > \bar{X}_{i^*,\tau}, E_{i,t}, E_{i^*,t})$$

which can be true only if at least one of the following holds:

1. $\bar{X}_{i,\tau} > \mu_i + \Delta_i/2$, which, intersected with $E_{i,\tau}$, by Hoeffding's inequality is true with probability at most:

$$\begin{aligned} \mathbb{P}(\bar{X}_{i,\tau} > \mu_i + \Delta_i/2, E_{i,\tau}) &\leq \sum_{y=f(\tau)}^{\infty} \mathbb{P}(\bar{X}_{i,\tau} > \mu_i + \Delta_i/2, Z_i(\tau) = y) \\ &\leq \sum_{y=f(\tau)}^{\infty} e^{-\frac{y\Delta_i^2}{2}} = \frac{e^{-\frac{f(\tau)\Delta_i^2}{2}}}{1 - e^{-\Delta_i^2/2}} \end{aligned}$$

2. $\bar{X}_{i^*,\tau} < \mu_* + \Delta_i/2$, which, intersected to $E_{i^*,\tau}$, by Hoeffding's inequality is also true with the same probability of before.

Therefore, we have proved that:

$$\sum_{\tau=1}^t \mathbb{P}(I_\tau = i, Z_i(\tau-1) \geq f(\tau)) \leq 2(1 - e^{-\Delta_i^2/2})^{-1} \sum_{\tau=1}^t e^{-\frac{f(\tau)\Delta_i^2}{2}},$$

which grows slower than $\log(t)$. Indeed, being $f(\cdot)$ superlogathmic, we have:

$$\lim_{t \rightarrow \infty} t e^{-\frac{f(t)\Delta_i^2}{2}} = \lim_{t \rightarrow \infty} t e^{-\log(t) \frac{f(t)\Delta_i^2}{2 \log(t)}} = \lim_{t \rightarrow \infty} t \left(\frac{1}{t} \right)^{\frac{f(t)\Delta_i^2}{2 \log(t)}} = 0.$$

Thus, for every $c > 0$, we can find t_0 satisfying $e^{-\frac{f(t)\Delta_i^2}{2}} \leq \frac{c}{t} \forall t \geq t_0$, so that:

$$\lim_{t \rightarrow \infty} \frac{\sum_{\tau=1}^t e^{-\frac{f(\tau)\Delta_i^2}{2}}}{c \log(t)} \leq \lim_{t \rightarrow \infty} \frac{\sum_{\tau=1}^t e^{-\frac{f(\tau)\Delta_i^2}{2}}}{c \sum_{\tau=1}^t \frac{1}{\tau}} \leq \lim_{t \rightarrow \infty} \underbrace{\frac{\sum_{\tau=1}^{t_0} e^{-\frac{f(\tau)\Delta_i^2}{2}}}{c \sum_{\tau=1}^t \frac{1}{\tau}}}_{\rightarrow 0} + \underbrace{\frac{\sum_{\tau=t_0}^t e^{-\frac{f(\tau)\Delta_i^2}{2}}}{c \sum_{\tau=t_0}^t \frac{1}{\tau}}}_{\leq 1} \leq 1.$$

This fact allows us to state (since a convergent sequence is always bounded) that:

$$C_0(\Delta_i) = 2(1 - e^{-\Delta_i^2/2})^{-1} \sup_{t > 1} \frac{\sum_{\tau=1}^t e^{-\frac{f(\tau)\Delta_i^2}{2}}}{\log(t)} < +\infty,$$

proving that for every suboptimal arm i

$$\mathbb{E}[Z_i(t)] \leq f(t) + C_0(\Delta_i) \log(t)$$

To conclude the proof, it is sufficient to redefine $C(\Delta_i) := \Delta_i C_0(\Delta_i)$ and see that:

$$R_t = \sum_{i=1}^N \Delta_i \mathbb{E}[Z_i(t)].$$

□

Corollary 10. Let $\delta > 0$ and $f(t) = \log(t)^{1+\delta}$ be the superlogarithmic function used in Algorithm 1, then we have:

$$C(f, \Delta_i) = \frac{2\Delta_i \left(e^{((2/\Delta_i^2)^{1/\delta})} + 1 \right)}{1 - e^{-\Delta_i^2/2}}$$

Proof. Let $t_0 \in \mathbb{N}$ be smallest integer such that:

$$e^{-\frac{1}{2} \log(t)^{1+\delta} \Delta_i^2} \leq \frac{1}{t}, \quad \forall t \geq t_0 > 1.$$

By rearranging the latter inequality we have that:

$$t_0 = \lceil e^{((2/\Delta_i^2)^{1/\delta})} \rceil.$$

From the previous proof, it was known that $C(\Delta_i, f) = \Delta_i C_0(\Delta_i)$, where

$$C_0(\Delta_i) = 2(1 - e^{-\Delta_i^2/2})^{-1} \sup_{t > 1} \frac{\sum_{\tau=1}^t e^{-\frac{f(\tau)\Delta_i^2}{2}}}{\log(t)}.$$

Therefore, we have for this specific choice of f ,

$$\begin{aligned} C_0(\Delta_i) \log(t) &\leq 2(1 - e^{-\Delta_i^2/2})^{-1} \sum_{\tau=1}^t e^{-\frac{f(\tau)\Delta_i^2}{2}} \\ &\leq 2(1 - e^{-\Delta_i^2/2})^{-1} \frac{\sum_{\tau=1}^t e^{-\frac{f(\tau)\Delta_i^2}{2}}}{\log(t)} \log(t) \\ &\leq 2(1 - e^{-\Delta_i^2/2})^{-1} \left(\frac{\sum_{\tau=1}^{t_0} e^{-\frac{f(\tau)\Delta_i^2}{2}}}{\log(t)} + \frac{\sum_{\tau=t_0}^t e^{-\frac{f(\tau)\Delta_i^2}{2}}}{\sum_{\tau=t_0}^t \frac{1}{\tau}} \right) \log(t) \end{aligned}$$

Where the last step is due to the fact that for $t_0 > 2$ we have $\sum_{\tau=t_0}^t \frac{1}{\tau} \leq \log(t)$. From this point, we can note that in the fraction $\frac{\sum_{\tau=t_0}^t e^{-\frac{f(\tau)\Delta_i^2}{2}}}{\sum_{\tau=t_0}^t \frac{1}{\tau}}$ for fixed τ , each term of the upper sum $e^{-\frac{f(\tau)\Delta_i^2}{2}}$ is less or equal than each term of the lower $\frac{1}{\tau}$. Therefore, we have

$$\frac{\sum_{\tau=t_0}^t e^{-\frac{f(\tau)\Delta_i^2}{2}}}{\sum_{\tau=t_0}^t \frac{1}{\tau}} \leq 1.$$

With this consideration, we are able to conclude:

$$\begin{aligned} C_0(\Delta_i) &\leq 2(1 - e^{-\Delta_i^2/2})^{-1} \left(\frac{\sum_{\tau=1}^{t_0} e^{-\frac{f(\tau)\Delta_i^2}{2}}}{\log(t)} + \frac{\sum_{\tau=t_0}^t e^{-\frac{f(\tau)\Delta_i^2}{2}}}{\sum_{\tau=t_0}^t \frac{1}{\tau}} \right) \\ &\leq 2(1 - e^{-\Delta_i^2/2})^{-1} \left(\frac{t_0}{\log(t)} + 1 \right) \\ &\leq 2(1 - e^{-\Delta_i^2/2})^{-1} + 2(1 - e^{-\Delta_i^2/2})^{-1} e^{((2/\Delta_i^2)^{1/\delta})}. \end{aligned}$$

Where in the last step we have used the fact that $\log(t) > 1$ for $t > 2$. Recollecting all the terms we have that:

$$C(\Delta_i, \log(t)^{1+\delta}) = \Delta_i C_0(\Delta_i, \log(t)^{1+\delta}) \leq 2\Delta_i (1 - e^{-\Delta_i^2/2})^{-1} (e^{((2/\Delta_i^2)^{1/\delta})} + 1),$$

concluding the proof. \square

B Proofs in the instance independent stochastic analysis

B.1 Instance dependent/independent trade-off

Lemma 4. *Let us define a random walk*

$$G_{t+1} = G_t + \epsilon_t \quad \epsilon_t = \begin{cases} 1 & p \\ -1 & 1-p \end{cases}.$$

where $G_0 = 1$, with $p = 1/2 + \Delta/2 > 0.5$ (for $\Delta \in (0, 1)$). Then, we have

$$\mathbb{P} \left(\bigcup_{t=1}^{\infty} \{G_t \leq 0\} \right) = \left(\frac{1-\Delta}{1+\Delta} \right).$$

Proof. Define

$$f_n = \mathbb{P}(G_0 = n, \exists t : G_t = 0)$$

which satisfies, for $n \geq 0$:

$$f_n = pf_{n-1} + (1-p)f_{n+1}$$

with $f_n = 1$ for $n \leq 0$. The equation corresponding to the aforementioned dynamical system is:

$$(1-p)\lambda^2 - \lambda + p = 0$$

which has two solutions:

$$\lambda = \frac{1 \pm \sqrt{1 - 4p(1-p)}}{2(1-p)}$$

Thus, we obtain, $\forall n > 0$

$$f_n = A \left(\frac{1 + \sqrt{1 - 4p(1-p)}}{2(1-p)} \right)^n + B \left(\frac{1 - \sqrt{1 - 4p(1-p)}}{2(1-p)} \right)^n$$

where $A = 0$ (otherwise, the equation does not define a probability) and $B = 1$ (since $f_1 \rightarrow 1$ for $\Delta \rightarrow 0$). Therefore, from the definition of p :

$$\begin{aligned} f_n &= \left(\frac{1 - \sqrt{1 - 4p(1-p)}}{2(1-p)} \right)^n \\ &= \left(\frac{1 - \sqrt{1 - 4(1/2 - \Delta/2)(1/2 + \Delta/2)}}{1 + \Delta} \right)^n \\ &= \left(\frac{1 - \Delta}{1 + \Delta} \right)^n. \end{aligned}$$

for $n = 1$ we have the result. \square

Theorem 4 (Instance Dependent/Independent Trade-off). *Let π be any policy for the bandits with ranking feedback problem. If π satisfies the following properties:*

- (instance-dependent upper regret bound) $R_T \leq \sum_{i=1}^n C(\Delta_i)T^\alpha$
- (instance-independent upper regret bound) $R_T \leq nCT^\beta$

then, $2\alpha + \beta \geq 1$, where $\alpha, \beta \geq 0$.

Proof. Let $p_1 = 0.5$, $p_2 = 0.5 - \varepsilon$, $p_2^* = 0.5 + \varepsilon$. Let us consider two problems:

$$P : \begin{cases} \nu_1 = Be(p_1) \\ \nu_2 = Be(p_2) \end{cases} \quad P^* : \begin{cases} \nu_1 = Be(p_1) \\ \nu_2 = Be(p_2^*) \end{cases}$$

Clearly, the optimal arm is 1 for P and 2 for P^* . Let us now define the event:

$$E_t = \bigcap_{\tau=1}^t \{\mathcal{R}_\tau = \langle 1, 2 \rangle\}.$$

By assumption, policy π has a sub- t^α instance-dependent regret, therefore,

$$\forall \eta > 0, \limsup_{t \rightarrow \infty} \frac{\sum_{\tau=1}^t \pi(\text{pull } 2 | E_\tau)}{t^{\alpha+\eta}} = 0,$$

otherwise we would have an instance dependent regret of order $t^{\alpha+\eta}$ in the simple case of $p_1 = 1$, $p_2 = 0$. This means that:

$$\limsup_t \frac{\mathbb{E}_P[Z_2(t) | E_t]}{t^{\alpha+\eta}} = 0 \implies \exists C > 0 \forall t > 0 : \quad \mathbb{E}_P[Z_2(t) | E_t] \leq Ct^{\alpha+\eta}.$$

Where \mathbb{E}_P is the expectation over the random variables of the arms following the distribution given in (P) . Therefore, by Markov's inequality, we have,

$$\mathbb{P}_P(Z_2(t) > 2Ct^{\alpha+\eta} | E_t) \leq \frac{\mathbb{E}_P[Z_2(t) | E_t]}{2Ct^{\alpha+\eta}} \leq \frac{Ct^{\alpha+\eta}}{2Ct^{\alpha+\eta}} \leq \frac{1}{2}.$$

Now, note that for every $h > 0$ the event $Z_2(t) < h$ is contained in the σ -algebra generated by the first h pulls of arm 2 (and all the pulls of arm 1, but this is irrelevant since arm 1 corresponds to the same distribution in the two problems). Therefore, thanks to Lemma 3 we have:

$$\begin{aligned} \forall h > 0 \quad \mathbb{P}_{P^*}(Z_2(t) \leq h) &\geq \left(\frac{0.5 - \varepsilon}{0.5 + \varepsilon} \right)^h \mathbb{P}_P(Z_2(t) \leq h) \\ &\geq \left(\frac{0.5 - \varepsilon}{0.5 + \varepsilon} \right)^h \mathbb{P}_P(T_2(t) \leq h, E_t) \\ &= \left(\frac{0.5 - \varepsilon}{0.5 + \varepsilon} \right)^h \mathbb{P}_P(T_2(t) \leq h | E_t) \mathbb{P}_P(E_t). \end{aligned}$$

By the previous step, we have $\mathbb{P}_P(Z_2(t) \leq 2Ct^{\alpha+\eta}|E_t) \geq 1/2$, so that:

$$\mathbb{P}_{P^*}(Z_2(t) \leq 2Ct^{\alpha+\eta}) \geq \frac{1}{2} \left(\frac{0.5 - \varepsilon}{0.5 + \varepsilon} \right)^{2Ct^{\alpha+\eta}} \mathbb{P}_P(E_t),$$

while, thanks to Lemma 4, we have $\mathbb{P}_P(E_t) \geq 1 - \frac{1-\varepsilon}{1+\varepsilon} \geq 2\varepsilon$, meaning that:

$$\mathbb{P}_{P^*}(Z_2(t) \leq 2Ct^{\alpha+\eta}) \geq \frac{1}{2} \left(\frac{0.5 - \varepsilon}{0.5 + \varepsilon} \right)^{2Ct^{\alpha+\eta}} \frac{2\varepsilon}{1 + \varepsilon}.$$

At this point we are using this result to provide a lower bound for the regret in the instance independent case. Analyzing the instance independent regret, by definition we have to fix t as time horizon and let the arm gap ε depend on t .

Let us now fix $\rho > 1$. With the choice $\varepsilon = t^{-\rho\alpha}$, we have

$$\begin{aligned} \mathbb{P}_{P^*}(Z_2(t) \leq 2Ct^{\alpha+\eta}) &\geq \frac{1}{2} \left(\frac{0.5 - t^{-\rho\alpha}}{0.5 + t^{-\rho\alpha}} \right)^{2Ct^{\alpha+\eta}} 2t^{-\rho\alpha} \\ &\geq \frac{1}{2} \left(1 - 4t^{-\rho\alpha} \right)^{2Ct^{\alpha+\eta}} \frac{2t^{-\rho\alpha}}{1 + t^{-\rho\alpha}} \\ &\geq \left(1 - 4t^{-\rho\alpha} \right)^{2Ct^{\alpha+\eta}} \frac{t^{-\rho\alpha}}{2}. \end{aligned} \tag{5}$$

At this point if we choose $\eta = \alpha(\rho - 1)/2$, the following fact holds:

$$\begin{aligned} \lim_{t \rightarrow \infty} \left(1 - 4t^{-\rho\alpha} \right)^{2Ct^{\alpha+\eta}} &= \lim_{y \rightarrow 0^+} \left(1 - 4y \right)^{Cy^{\frac{-(\alpha+\eta)}{\rho\alpha}}} \\ &= \lim_{y \rightarrow 0^+} \left(1 - 4y \right)^{Cy^{\frac{-\alpha(1/2+\rho/2)}{\rho\alpha}}} \\ &= \lim_{y \rightarrow 0^+} \left(1 - 4y \right)^{Cy^{\frac{-(1/2+\rho/2)}{\rho}}} \end{aligned}$$

where in the first equality we substituted $y = t^{-\rho\alpha}$. Here, $y^{\frac{-(1/2+\rho/2)}{\rho}} = y^{-1} \cdot y^{\frac{\rho-1}{2\rho}}$, where the second exponent is strictly positive.

$$\begin{aligned} \lim_{t \rightarrow \infty} \left(1 - 4t^{-\rho\alpha} \right)^{2Ct^{\alpha+\eta}} &= \lim_{y \rightarrow 0^+} \left(\left(1 - 4y \right)^{-y} \right)^{Cy^{\frac{\rho-1}{2\rho}}} \\ &= \lim_{y \rightarrow 0^+} (1/e^4)^{Cy^{\frac{\rho-1}{2\rho}}} = 1. \end{aligned}$$

This limit shows that there is $c_\rho > 0$ such that:

$$\left(1 - 4t^{-\rho\alpha} \right)^{2Ct^{\alpha+\eta}} \geq c_\rho \quad \forall t \text{ sufficiently big.}$$

Substituting this property in the previously found Equation (5), we get:

$$\mathbb{P}_{P^*}(Z_2(t) \leq 2Ct^{\alpha(1/2+\rho/2)}) \geq \frac{c_\rho}{2} t^{-\rho\alpha},$$

holding for every $\rho > 0$ and sufficiently big t .

This proves that, with $\varepsilon = t^{-\rho\alpha}$:

$$\begin{aligned} R_t|_{\sim P^*} &\geq \varepsilon(t - 2Ct^{\alpha(1/2+\rho/2)}) \mathbb{P}_{P^*}(Z_2(t) \leq 2Ct^{\alpha(1/2+\rho/2)}) \\ &\geq \frac{1}{2} t \cdot \frac{c_\rho}{2} t^{-2\rho\alpha} = \frac{c_\rho}{4} t^{1-2\rho\alpha} \quad \forall t \text{ sufficiently big.} \end{aligned}$$

Therefore, for $\beta \leq 1 - 2\rho\alpha$ it is not possible to have an upper bound for the instance independent regret. Since this is valid for every $\rho > 1$, we can also extend the result to any $\beta < 1 - 2\alpha$, which leads to the conclusion that the necessary condition to satisfy both:

- (Instance Dependent regret bound)

$$R_t \leq \sum_{i=1}^n C(\Delta_i) t^\alpha \quad \forall t > 0$$

- (Instance Independent regret bound)

$$R_t \leq nCt^\beta \quad \forall t > 0$$

for the same policy π is

$$2\alpha + \beta \geq 1.$$

□

B.2 Proofs of instance independent regret upper bound

To understand the tractation of the instance independent regret analysis, we will need some results from the theory of stochastic processes. We devote the following subsections to develop all the results required to prove the regret bound on the algorithm.

B.2.1 Discretizing the Brownian motion

In this section, we prove some results about the relationship between Random Walk G_i and Brownian Motion B_t , that will be crucial in the proof of the regret bound. For this scope, we will introduce this quantity

$$|B_t + t\mu_0 < \eta| = \int_0^1 \mathbf{1}_{(-\infty, \eta)}(\tau) d\tau,$$

corresponding to the Lebesgue measure of the set $\{t \in [0, 1] : B_t + t\mu_0 < \eta\}$.

We start with a lemma that bounds the increments in a standard brownian Motion.

Lemma 5. *Let $(B_t)_{t \in [0,1]}$ be a standard Brownian motion. Define*

$$\forall i \in \{0, \dots, n-1\}, \quad I_i := [i/n, (i+1)/n].$$

Then, for every $\eta \geq 0$,

$$\mathbb{P} \left(\sup_{i \in \{0, \dots, n-1\}} \left(\sup_{t \in I_i} B_t - B_{i/n} \right) \geq \eta \right) = \mathbb{P} \left(\inf_{i \in \{0, \dots, n-1\}} \left(\inf_{t \in I_i} B_t - B_{i/n} \right) \leq -\eta \right) \leq \frac{2\sqrt{n} \exp\left(-\frac{\eta^2 n}{2\sigma^2}\right)}{\eta/\sigma\sqrt{2\pi}}.$$

Proof. Indeed, the Brownian motion satisfies:

$$\begin{aligned} \mathbb{P} \left(\sup_{i \in \{0, \dots, n-1\}} \left(\sup_{t \in I_i} B_t - B_{i/n} \right) \geq \eta \right) &= \mathbb{P} \left(\bigcup_{i=0}^{n-1} \sup_{t \in I_i} B_t - B_{i/n} > \eta \right) \\ &\leq \sum_{i=0}^{n-1} \mathbb{P} \left(\sup_{t \in I_i} B_t - B_{i/n} > \eta \right) \\ &= \sum_{i=0}^{n-1} 2\mathbb{P}(B_{(i+1)/n} - B_{i/n} > \eta) \\ &= \sum_{i=0}^{n-1} 2\mathbb{P}(\mathcal{N}(0, \sigma^2/n) > \eta) \\ &\leq \sum_{i=0}^{n-1} \frac{2 \exp\left(-\frac{\eta^2 n}{2\sigma^2}\right)}{\eta/\sigma\sqrt{2n\pi}} \leq \frac{2\sqrt{n} \exp\left(-\frac{\eta^2 n}{2\sigma^2}\right)}{\eta/\sigma\sqrt{2\pi}}. \end{aligned}$$

where the third equality holds from the reflection principle (see [Baldi, 2017]), and last inequality holds since it is well-known that $\mathbb{P}(\mathcal{N}(0, \beta^2) > y) \leq \frac{\exp(-y^2/2\beta^2)}{y/\beta\sqrt{2\pi}}$ for tail bound on Gaussian distributions.

In the exact same way, we can prove that

$$\mathbb{P}\left(\inf_{i \in \{0, \dots, n-1\}} \left(\inf_{t \in I_i} B_t - B_{i/n}\right) \leq -\eta\right) \leq \frac{2\sqrt{n} \exp\left(-\frac{\eta^2 n}{2\sigma^2}\right)}{\eta/\sigma\sqrt{2\pi}}.$$

Together, the two results imply the thesis. \square

We are now ready to prove a theorem that links the Brownian Motion and a Random Walk in term of the probability that each of them stays in the interval $[0, \infty)$.

Lemma 6 (Discretization lemma). *Let $(G_i)_{i \in \{0, \dots, n-1\}}$ be a Gaussian 0-mean unit variance random walk, and $\mu \in \mathbb{R}$, and $(B_t)_{t \in [0,1]}$ a standard Brownian motion. Then, for every $s \in (0, 1)$ we have,*

$$\mathbb{P}(|B_t + t\mu_0 > \eta| > s) - P(n, \eta) \leq \mathbb{P}\left(\sum_{i=0}^{n-1} \mathbf{1}_{(0, \infty)}(G_i + i\mu) > sn\right) \leq \mathbb{P}(|B_t + t\mu_0 > -\eta| > s) + P(n, \eta)$$

and

$$\mathbb{P}(|B_t + t\mu_0 \leq \eta| \leq s) - P(n, \eta) \leq \mathbb{P}\left(\sum_{i=0}^{n-1} \mathbf{1}_{(-\infty, 0]}(G_i + i\mu) \leq sn\right) \leq \mathbb{P}(|B_t + t\mu_0 \leq -\eta| \leq s) + P(n, \eta)$$

with $P(n, \eta) = \frac{2\sqrt{n} \exp(-\eta^2 n/2)}{\eta\sqrt{2\pi}}$ and $\mu_0 = \sqrt{n}\mu$.

Proof. We only prove the first part, as the second one follows trivially by substituting $s \leftarrow 1 - s$, $\mu \leftarrow -\mu$, $G_i \leftarrow -G_i$, $B_t \leftarrow -B_t$.

Let $(B_t)_{t \in [0,1]}$ be a standard Brownian motion. Define

$$\forall i \in \{0, \dots, n-1\}, \quad I_i := [i/n, (i+1)/n].$$

Let us set

$$\mu_0 = \sqrt{n}\mu.$$

With this definition, we have the following set inclusions, for any $s \in [0, 1]$ and $\eta > 0$:

$$\begin{aligned} \{|B_t + t\mu_0 > \eta| > s\} &= \left\{ \int_0^1 \mathbf{1}_{(\eta, \infty)}(B_\tau + \tau\mu_0) d\tau > s \right\} \\ &= \left\{ \sum_{i=0}^{n-1} \int_{I_i} \mathbf{1}_{(\eta, \infty)}(B_\tau + \tau\mu_0) d\tau > s \right\} \\ &\subseteq \left\{ \sum_{i=0}^{n-1} \sup_{\tau \in I_i} \mathbf{1}_{(\eta, \infty)}(B_\tau + \tau\mu_0) > sn \right\} \\ &\subseteq \left\{ \sum_{i=0}^{n-1} \mathbf{1}_{(0, \infty)}\left(B_{i/n} + \frac{i}{n}\mu_0\right) > sn \right\} \cup \left\{ \sup_{i \in \{0, \dots, n-1\}} \left(\sup_{t \in I_i} B_t - B_{i/n}\right) \geq \eta \right\}. \end{aligned}$$

Moreover, it is also true that, using the same steps

$$\begin{aligned} \{|B_t + t\mu_0 > -\eta| > s\} &= \left\{ \int_0^1 \mathbf{1}_{(-\eta, \infty)}(B_\tau + \tau\mu_0) d\tau > s \right\} \\ &\supseteq \left\{ \sum_{i=0}^{n-1} \mathbf{1}_{(0, \infty)}\left(B_{i/n} + \frac{i}{n}\mu_0\right) > sn \right\} \cap \left\{ \inf_{i \in \{0, \dots, n-1\}} \left(\inf_{t \in I_i} B_t - B_{i/n}\right) \geq -\eta \right\}. \end{aligned}$$

Now, note that the random variable $B_{i/n}$, for $i = 1, \dots, n$ has the same distribution of G_i/\sqrt{n} , so that

$$\begin{aligned} \mathbb{P}\left(\sum_{i=0}^{n-1} \mathbf{1}_{(0,\infty)}\left(B_{i/n} + \frac{i}{n}\mu_0\right) > sn\right) &= \mathbb{P}\left(\sum_{i=0}^{n-1} \mathbf{1}_{(0,\infty)}\left(\sqrt{n}B_{i/n} + \frac{i}{\sqrt{n}}\mu_0\right) > sn\right) \\ &= \mathbb{P}\left(\sum_{i=0}^{n-1} \mathbf{1}_{(0,\infty)}(G_i + i\mu) > sn\right). \end{aligned}$$

Therefore, by union bound:

$$\mathbb{P}(|B_t + t\mu_0 > \eta| > s) \leq \mathbb{P}\left(\sum_{i=0}^{n-1} \mathbf{1}_{(0,\infty)}(G_i + i\mu) > sn\right) + \mathbb{P}\left(\sup_{i \in \{0, \dots, n-1\}} \left(\sup_{t \in I_i} B_t - B_{i/n}\right) \geq \eta\right),$$

and

$$\mathbb{P}(|B_t + t\mu_0 > -\eta| > s) \geq \mathbb{P}\left(\sum_{i=0}^{n-1} \mathbf{1}_{(0,\infty)}(G_i + i\mu) > sn\right) - \mathbb{P}\left(\inf_{i \in \{0, \dots, n-1\}} \left(\inf_{t \in I_i} B_t - B_{i/n}\right) \leq -\eta\right).$$

The proof is completed applying lemma 5 and reordering the terms. \square

For our setting, it is convenient to state a corollary of the previous result that will be used in the next proofs.

Corollary 11. *Let $(G_i)_{i \in \{0, \dots, n-1\}}$ be a Gaussian 0-mean unit variance random walk, and $\mu \in \mathbb{R}$. Then, for every $s \in (0, 1)$ we have,*

$$\mathbb{P}\left(\sum_{i=0}^{n-1} \mathbf{1}_{(-\infty, 0]}(G_i + i\mu) \leq sn\right) \in$$

$$\left[\mathbb{P}\left(\left|B_t + t\mu_0 \leq \frac{2\log(n)}{\sqrt{n}}\right| \leq s\right) - \frac{\sqrt{2}}{\sqrt{\pi n \log(n)}}, \mathbb{P}\left(\left|B_t + t\mu_0 \leq -\frac{2\log(n)}{\sqrt{n}}\right| \leq s\right) + \frac{\sqrt{2}}{\sqrt{\pi n \log(n)}} \right],$$

where $\mu_0 = \sqrt{n}\mu$.

Proof. It is sufficient to make the substitution

$$\eta = \frac{2\log(n)}{\sqrt{n}},$$

in the previous lemma. Indeed, we have

$$\begin{aligned} P(n, \eta) &= \frac{2\sqrt{n} \exp(-\eta^2 n/2)}{\eta\sqrt{2\pi}} \\ &= \frac{2\sqrt{n} \exp(-\log(n)^2)}{\frac{\log(n)}{\sqrt{n}}\sqrt{2\pi}} \\ &= \frac{2n \exp(-\log(n)^2)}{\log(n)\sqrt{2\pi}} = \frac{\sqrt{2}}{\sqrt{\pi n \log(n)}}. \end{aligned}$$

\square

B.2.2 Proofs of filtering inequalities

All the proof of this subsections will be based on the following very powerful result, which studies the time spent by a Brownian Motion with drift in the half-line $[0, \infty)$.

Theorem 12 (Takács [1996]). *Let B_t be a standard Brownian motion on $t \in [0, 1]$, and let us note as $|\cdot|$ the Lebesgue measure of a set. For $\mu_0 \in \mathbb{R}$ and $\eta > 0$, we have*

$$\mathbb{P}(|B_t + t\mu_0 \leq \eta| \leq s) = 2 \int_0^s \left[\frac{\varphi(\mu_0\sqrt{1-\tau})}{\sqrt{1-\tau}} + \mu_0\Phi(\mu_0\sqrt{1-\tau}) \right] \times \left[\frac{\varphi(\eta/\sqrt{\tau} - \mu_0\sqrt{\tau})}{\sqrt{\tau}} - \mu_0 e^{2\mu_0\eta} \Phi(-\eta/\sqrt{\tau} - \mu_0\sqrt{\tau}) \right] d\tau,$$

where

$$\varphi(x) := \frac{1}{\sqrt{2\pi}} e^{-x^2/2} \quad \Phi(x) := \int_{-\infty}^x \varphi(u) du.$$

Thanks to the previous theorem, we can prove the following crucial results.

Theorem 13. *Let T be a sufficiently large constant. Let $(G_i)_{i \in \{0, \dots, n-1\}}$ be a Gaussian 0-mean unit variance random walk, and $\mu \in \mathbb{R}$. If $\mu \geq CT^{-\alpha}$, for some $\alpha \in (0, 1/2)$ and $C = 4 \log(T)$, then setting $n = \lceil T^{1/2+\alpha} \rceil$ we have*

$$\mathbb{P} \left(\sum_{i=0}^{n-1} \mathbf{1}_{(-\infty, 0]}(G_i + i\mu) \leq T^{2\alpha} \right) \geq 1 - 2T^{-1/2}.$$

Proof. In the rest of the proof, we will assume, for ease of notation, that T is such that $T^{1/2+\alpha}$ an integer, so that $n = T^{1/2+\alpha}$. This is done without loss of generality, since substituting n with $n+1$ leads to a negligible difference for T sufficiently big. Applying the discretization corollary 11, we have that for every $s \in (0, 1)$

$$\mathbb{P} \left(\sum_{i=0}^{n-1} \mathbf{1}_{(-\infty, 0]}(G_i + i\mu) \leq sn \right) \geq \mathbb{P} \left(\left| B_t + t\mu_0 \leq \frac{2 \log(n)}{\sqrt{n}} \right| \leq s \right) - \frac{\sqrt{2}}{\sqrt{\pi n \log(n)}}, \quad (6)$$

where $\mu_0 = \sqrt{n}\mu$. Therefore, by assumption,

$$\mu_0 = \sqrt{n}\mu \geq (T^{1/2+\alpha})^{1/2} CT^{-\alpha} = CT^{1/4-\alpha/2}.$$

At this point, we can apply Theorem 12 to have, for any $\eta > 0$,

$$\mathbb{P}(|B_t + t\mu_0 \leq \eta| \leq s) = 2 \int_0^s \left(\frac{\phi(\mu_0\sqrt{1-\tau})}{\sqrt{1-\tau}} + \mu_0\Phi(\mu_0\sqrt{1-\tau}) \right) \times \left(\phi\left(\frac{\eta - \mu_0\tau}{\sqrt{\tau}}\right) \frac{1}{\sqrt{\tau}} - \mu_0 e^{2\mu_0\eta} \Phi\left(\frac{-\eta - \mu_0\tau}{\sqrt{\tau}}\right) \right) d\tau$$

which means that

$$\mathbb{P}(|B_t + t\mu_0 \leq \eta| \leq s) = 1 - 2 \int_s^1 \left(\underbrace{\frac{\phi(\mu_0\sqrt{1-\tau})}{\sqrt{1-\tau}}}_{(1)} + \underbrace{\mu_0\Phi(\mu_0\sqrt{1-\tau})}_{(2)} \right) \times \left(\underbrace{\phi\left(\frac{\eta - \mu_0\tau}{\sqrt{\tau}}\right) \frac{1}{\sqrt{\tau}}}_{(3)} - \underbrace{\mu_0 e^{2\mu_0\eta} \Phi\left(\frac{-\eta - \mu_0\tau}{\sqrt{\tau}}\right)}_{(4)} \right) d\tau.$$

Here, we have to consider that

- $\eta = \frac{2\log(n)}{\sqrt{n}} \leq 2\log(T)T^{-\alpha/2-1/4}$
- $\mu_0 \geq CT^{1/4-\alpha/2}$.

Moreover, to have the thesis, we are interested in a value of s such that $sn = T^{2\alpha}$, corresponding to $T^{-1/2+\alpha}$. Therefore, in the interval $[T^{-1/2+\alpha}, 1]$, we have

1. Consider term (3):

$$\begin{aligned} \phi\left(\frac{\eta - \mu_0\tau}{\sqrt{\tau}}\right) \frac{1}{\sqrt{\tau}} &\leq \phi\left(\frac{\eta - \mu_0T^{-1/2+\alpha}}{T^{-1/4+\alpha/2}}\right) \frac{1}{T^{-1/4+\alpha/2}} \\ &= \phi\left(\eta T^{1/4-\alpha/2} - \mu_0 T^{-1/4+\alpha/2}\right) \frac{1}{T^{-1/4+\alpha/2}}. \end{aligned}$$

Here, since $\eta = \frac{2\log(n)}{\sqrt{n}} \leq 2\log(T)T^{-\alpha/2-1/4}$, the part $\eta T^{1/4-\alpha/2}$ is bounded by $2\log(T)$.

Instead, $\mu_0 T^{-1/4+\alpha/2} \geq CT^{1/4-\alpha/2}T^{-1/4+\alpha/2} = C$.

2. Term (4) is non-negative.

Therefore, for $C = 4\log(T)$, we have that in the interval $[T^{-1/2+\alpha}, 1]$

$$(3) + (4) \leq \phi(2\log(T)) \frac{1}{T^{-1/4+\alpha/2}} = \frac{1}{\sqrt{2\pi}T^{-1/4+\alpha/2}} e^{-2\log(T)^2} \leq \frac{T^{-1}}{\sqrt{2\pi}}.$$

With this inequality, we have

$$\begin{aligned} \mathbb{P}\left(|B_t + t\mu_0 \leq \eta| \leq T^{-1/2+\alpha}\right) &= 1 - 2 \frac{T^{-1}}{\sqrt{2\pi}} \int_{T^{-1/2+\alpha}}^1 \left(\frac{\phi(\mu_0\sqrt{1-\tau})}{\sqrt{1-\tau}} + \mu_0\Phi(\mu_0\sqrt{1-\tau}) \right) d\tau \\ &\geq 1 - 2 \frac{T^{-1}}{\sqrt{2\pi}} \int_{T^{-1/2+\alpha}}^1 \frac{1}{\sqrt{2\pi}(1-\tau)} + |\mu_0| d\tau \\ &\geq 1 - 2 \frac{T^{-1}}{\sqrt{2\pi}} \int_0^1 \frac{1}{\sqrt{2\pi}(1-\tau)} + |\mu_0| d\tau \\ &= 1 - 2 \frac{T^{-1}}{\sqrt{2\pi}} \left(\frac{\sqrt{2}}{\sqrt{\pi}} + \mu_0 \right). \end{aligned}$$

At this point, knowing from the assumptions that $n < T$, we have $\mu_0 \leq \sqrt{T}$, which implies

$$\mathbb{P}\left(|B_t + t\mu_0 \leq \eta| \leq T^{-1/2+\alpha}\right) \geq 1 - \frac{T^{-1/2}}{\pi}.$$

Substituting this result into Equation 6, we get, for $s = T^{-1/2+\alpha}$ and $n \geq T^{1/2+\alpha}$

$$\begin{aligned} \mathbb{P}\left(\sum_{i=0}^{n-1} \mathbf{1}_{(-\infty, 0]}(G_i + i\mu) < T^{2\alpha}\right) &\geq 1 - \frac{T^{-1/2}}{\pi} - \frac{\sqrt{2}}{\sqrt{\pi}T^{1/2+\alpha} \log(T^{1/2+\alpha})} \\ &\geq 1 - 2T^{-1/2}. \end{aligned}$$

□

The second result is the following

Theorem 14. Let T be a sufficiently large constant. Let $(G_i)_{i \in \{0, \dots, n-1\}}$ be a Gaussian 0-mean unit variance random walk, and $\mu \in \mathbb{R}$ such that $\mu \leq -CT^{-\theta}$, for some $\theta \in (0, 1/2)$ and $C = 2\sqrt{\log(T)} + 2$. Then, for any $\alpha \in (0, 1/2)$, setting $n = \lfloor T^{1/2+\alpha} \rfloor$ we have

$$\mathbb{P} \left(\sum_{i=0}^{n-1} \mathbf{1}_{(-\infty, 0]}(G_i + i\mu) \leq T^{2\alpha} \right) \leq 3T^{-1/2+\theta}.$$

Proof. In the rest of the proof, we will assume, for ease of notation, that T is such that $T^{1/2+\alpha}$ an integer, so that $n = T^{1/2+\alpha}$. This is done without loss of generality, since substituting n with $n + 1$ leads to a negligible difference for T sufficiently large. Applying the discretization corollary 11, we have that for every $s \in (0, 1)$

$$\mathbb{P} \left(\sum_{i=0}^{n-1} \mathbf{1}_{(-\infty, 0]}(G_i + i\mu) \leq sn \right) \leq \mathbb{P} \left(\left| B_t + t\mu_0 \leq -\frac{2\log(n)}{\sqrt{n}} \right| \leq s \right) + \frac{\sqrt{2}}{\sqrt{\pi n \log(n)}}, \quad (7)$$

where $\mu_0 = \sqrt{n}\mu$. Therefore, by assumption,

$$\mu_0 = \sqrt{n}\mu \leq -(T^{1/2+\alpha})^{1/2}CT^{-\theta} = -CT^{1/4+\alpha/2-\theta}.$$

Differently from the previous proof, here we cannot directly apply Theorem 12, since $\eta = -\frac{2\log(n)}{\sqrt{n}} < 0$.

Still, we can say that

$$\begin{aligned} \mathbb{P} \left(\left| B_t + t\mu_0 \leq -\frac{2\log(n)}{\sqrt{n}} \right| \leq s \right) &= \mathbb{P} \left(\left| -B_t - t\mu_0 > \frac{2\log(n)}{\sqrt{n}} \right| \leq s \right) \\ &= \mathbb{P} \left(\left| -B_t - t\mu_0 \leq \frac{2\log(n)}{\sqrt{n}} \right| > 1 - s \right). \end{aligned}$$

At this point, we set $\eta = \frac{2\log(n)}{\sqrt{n}}$, $\tilde{\mu}_0 = -\mu_0$ and $B_t = -B_t$ (it is not necessary to rename it since its distribution is symmetric). In this way we can apply Theorem 12 having that the previous probability corresponds to

$$\begin{aligned} \mathbb{P} (|B_t + t\tilde{\mu}_0 \leq \eta| > 1 - s) &= 2 \int_{1-s}^1 \left(\underbrace{\frac{\phi(\tilde{\mu}_0\sqrt{1-\tau})}{\sqrt{1-\tau}}}_{(1)} + \underbrace{\tilde{\mu}_0\Phi(\tilde{\mu}_0\sqrt{1-\tau})}_{(2)} \right) \\ &\quad \times \left(\underbrace{\phi\left(\frac{\eta - \tilde{\mu}_0\tau}{\sqrt{\tau}}\right)}_{(3)} \frac{1}{\sqrt{\tau}} - \underbrace{\tilde{\mu}_0 e^{2\mu_0\eta}\Phi\left(\frac{-\eta - \tilde{\mu}_0\tau}{\sqrt{\tau}}\right)}_{(4)} \right) d\tau. \end{aligned}$$

Here, we have to consider that

- $\eta = \frac{2\log(n)}{\sqrt{n}} \leq 2\log(T)T^{-\alpha/2-1/4}$
- $\tilde{\mu}_0 \geq CT^{1/4+\alpha/2-\theta}$.

Moreover, to have the thesis, we are interested in a value of s such that $sn = T^{2\alpha}$, corresponding to $T^{-1/2+\alpha}$.

Here, it is convenient to divide the proof in two cases, depending on the sign of $1/4 + \alpha/2 - \theta$.

1. Assume $(1/4 + \alpha/2 - \theta > 0)$. Then, considering term (3) we have that for $\tau \in [1/2, 1]$

$$(3) \leq \phi\left(\frac{\eta - \tilde{\mu}_0\tau}{\sqrt{\tau}}\right) \frac{1}{\sqrt{\tau}} \leq \sqrt{2}\phi\left(\sqrt{2}\eta - \tilde{\mu}_0/\sqrt{2}\right).$$

Moreover, since term (4) is nonnegative we also have

$$(3) + (4) \leq \sqrt{2}\phi\left(\sqrt{2}\eta - \tilde{\mu}_0/\sqrt{2}\right) = \frac{1}{\sqrt{\pi}}e^{-(\sqrt{2}\eta - \tilde{\mu}_0/\sqrt{2})^2/2}.$$

Being $1/4 + \alpha/2 - \theta > 0$ and $\eta < 1$, the exponent is less than $-(\sqrt{2} - C/\sqrt{2})^2/2$. This means that for $C = 2\sqrt{\log(T)} + 2$ the full term is bounded by

$$(3) + (4) \leq \frac{1}{\sqrt{\pi}}e^{-(\sqrt{2}-C/\sqrt{2})^2/2} = \frac{1}{\sqrt{\pi}}e^{-(\sqrt{2\log(T)})^2/2} = \frac{T^{-1}}{\sqrt{\pi}}.$$

Substituting this inequality, we get

$$\begin{aligned} \mathbb{P}\left(|B_t + t\tilde{\mu}_0 \leq \eta| > 1 - T^{-1/2+\alpha}\right) &\leq \frac{2T^{-1}}{\sqrt{\pi}} \int_{1-T^{-1/2+\alpha}}^1 \left(\frac{\phi(\tilde{\mu}_0\sqrt{1-\tau})}{\sqrt{1-\tau}} + \tilde{\mu}_0\Phi(\tilde{\mu}_0\sqrt{1-\tau})\right) d\tau \\ &\leq \frac{2T^{-1}}{\sqrt{\pi}} \int_{1-T^{-1/2+\alpha}}^1 \frac{1}{\sqrt{2\pi(1-\tau)}} + |\tilde{\mu}_0| d\tau \\ &\leq \frac{2T^{-1}}{\sqrt{\pi}}(2 + T^{-1/2+\alpha}\tilde{\mu}_0) \leq \frac{6T^{-1}}{\sqrt{\pi}}. \end{aligned}$$

This quantity is of course less than $T^{-\theta}$, since $\theta \in (0, 1/2)$ by assumption

2. Assume $(1/4 + \alpha/2 - \theta < 0)$. In this case, we have, being $\tilde{\mu}_0 \geq 0$, the following inequality

$$\mathbb{P}(|B_t + t\tilde{\mu}_0 \leq \eta| > 1 - s) \leq \mathbb{P}(|B_t \leq \eta| > 1 - s).$$

This simplified form leads to

$$\begin{aligned} \mathbb{P}(|B_t + t\tilde{\mu}_0 \leq \eta| > 1 - s) &\leq 2 \int_{1-s}^1 \frac{\phi(0)}{\sqrt{1-\tau}} \phi\left(\frac{\eta}{\sqrt{\tau}}\right) \frac{1}{\sqrt{\tau}} d\tau \\ &\leq 2 \int_{1-s}^1 \frac{\phi(0)}{\sqrt{1-\tau}} \phi(0) \frac{1}{\sqrt{\tau}} d\tau \\ &= \frac{1}{\pi} \int_{1-s}^1 \frac{1}{\sqrt{\tau(1-\tau)}} d\tau. \end{aligned}$$

Since in our case $s = T^{-1/2+\alpha} < 1/2$, this can be further simplified as

$$\begin{aligned} \mathbb{P}(|B_t + t\tilde{\mu}_0 \leq \eta| > 1 - s) &\leq \frac{1}{\pi} \int_{1-s}^1 \frac{1}{\sqrt{\tau(1-\tau)}} d\tau \\ &= \frac{2}{\pi} \int_{1-s}^1 \frac{1}{\sqrt{1-\tau}} d\tau \\ &\stackrel{y=1-\tau}{=} \frac{2}{\pi} \int_0^s \frac{1}{\sqrt{y}} dy = \frac{4}{\pi} \sqrt{s}. \end{aligned}$$

This leads to

$$\mathbb{P}\left(|B_t + t\tilde{\mu}_0 \leq \eta| > 1 - T^{-1/2+\alpha}\right) \leq \frac{4}{\pi}T^{-1/4+\alpha/2}.$$

By assumption, $1/4 + \alpha/2 - \theta < 0$ the exponent is $-1/4 + \alpha/2 < T^{-1/2+\theta}$. Therefore, we have

$$\mathbb{P}\left(|B_t + t\tilde{\mu}_0 \leq \eta| > 1 - T^{-1/2+\alpha}\right) \leq \frac{4}{\pi}T^{-1/2+\theta}.$$

Thus, we have proved that in both cases

$$\mathbb{P}\left(|B_t + t\tilde{\mu}_0 \leq \eta| > 1 - s\right) \leq \frac{4}{\pi}T^{-1/2+\theta}.$$

Finally, applying Equation (7) and substituting the value of n , we get

$$\mathbb{P}\left(\sum_{i=0}^{n-1} \mathbf{1}_{(-\infty, 0]}(G_i + i\mu) \leq T^{2\alpha}\right) \leq \frac{4}{\pi}T^{-1/2+\theta} + \frac{\sqrt{2}}{\sqrt{\pi}T^{1/2+\alpha} \log(T^{1/2+\alpha})},$$

which implies

$$\mathbb{P}\left(\sum_{i=0}^{n-1} \mathbf{1}_{(-\infty, 0]}(G_i + i\mu) \leq T^{2\alpha}\right) \leq 3T^{-1/2+\theta}.$$

□

B.2.3 Regret bound

Before the actual proof, we are stating a simple proposition about the structure of the loggrid, which will ease the next computations.

Proposition 15. *Let*

$$LG(1/2, 1, T) := \left\{ \lfloor T^{\lambda_j + (1-\lambda_j)/2} \rfloor : \lambda_j = \frac{j}{\lfloor \log(T) \rfloor}, \forall j = 0, \dots, \lfloor \log(T) \rfloor \right\}.$$

The following identities hold

1. $LG(1/2, 1, T)$ can be equivalently defined as

$$LG(1/2, 1, T) := \left\{ \lfloor T^{1/2 + \frac{j}{2\lfloor \log(T) \rfloor}} \rfloor, \forall j = 0, \dots, \lfloor \log(T) \rfloor \right\}.$$

2. Let ℓ_j the j -th element of $LG(1/2, 1, T)$, and $\alpha_j = \frac{\log(\ell_j)}{\log(T)} - 1/2$. Then $\alpha_j = \frac{j}{2\lfloor \log(T) \rfloor} + o(T^{-1/2})$.

3. The ratio of two consecutive values of ℓ_j is $\frac{\ell_{j+1}}{\ell_j} \approx T^{\frac{1}{2\lfloor \log(T) \rfloor}} \in [\sqrt{e}, 2]$ for $T \geq 51$.

Next, we prove the following lemmas, which concern some features of our algorithm.

Lemma 7. *For any arm i , the probability of the event E_{ii_0} , corresponding to i eliminating the another i_0 arm such that their gap is $\Delta_{ii_0} := \mu_{i_0} - \mu_i > 0$, is, at most*

$$\mathbb{P}(E_{ii_0}) \leq 6 \log(T)(4 \log(T) + 2)T^{-1/2} \Delta_{ii_0}^{-1}.$$

Proof. Let us call:

$$\tilde{\Delta}_{ii_0} = \frac{\Delta_{ii_0}}{4 \log(T) + 2}.$$

At this point, there are two possibilities,

1. $\tilde{\Delta}_{ii_0} \leq T^{-1/2}$: in this case, the statement of the lemma is vacuous.
2. $\tilde{\Delta}_{ii_0} > T^{-1/2}$: in this case, by assumption, there are two consecutive $\ell_{j_*}, \ell_{j_*+1} \in \mathcal{L}$ such that

$$\tilde{\Delta}_i \in \left(\frac{T^{1/2}}{\ell_{j_*+1}}, \frac{T^{1/2}}{\ell_{j_*}} \right].$$

This is true due to the fact that the sequence ℓ_j spans from $T^{1/2}$ to T . By Proposition 15, this can be equivalently expressed by saying that

$$\tilde{\Delta}_{ii_0} \in \left(T^{-\frac{j_*+1}{2\lceil \log(T) \rceil}}, T^{-\frac{j_*}{2\lceil \log(T) \rceil}} \right].$$

Let us define the following family of events:

$$E_{ii_0}(t) := \text{arm } i \text{ eliminates arm } i_0 \text{ when both have been pulled } z \text{ times.}$$

The probability of E_{ii_0} is bounded by:

$$\begin{aligned} \mathbb{P}(E_{ii_0}) &= \mathbb{P}\left(\bigcup_{z=1}^T E_{ii_0}(z)\right) = \mathbb{P}\left(\bigcup_{z \in \mathcal{L}} E_{ii_0}(z)\right) \\ &\leq \sum_{j=1}^{|\mathcal{L}|} \mathbb{P}(E_{ii_0}(\ell_j)). \end{aligned}$$

Here, we have applied the fact that, by design of the algorithm, the arms can only be discarded in fair steps for which $z \in \mathcal{L}$ and then the union bound. Here, remember that by definition of the filtering condition, defining t_j as the fair time-step where both arms have been played ℓ_j times, this event can be again rewritten as:

$$\left\{ \sum_{\tau=1:t_j \text{ fair}} \{\mathcal{R}_\tau(i) > \mathcal{R}_\tau(i_0)\} \geq T^{2\alpha_j} \right\},$$

where $\alpha_j = \frac{\log(\ell_j)}{\log(T)} - \frac{1}{2}$. If we call $\hat{\mu}_{\tau, i_0}, \hat{\mu}_{\tau, i}$ the empirical means of arms i_0, i after τ pulls of each, the previous event can be interpreted as the time in which the random walk given by the difference of the rewards of the two arms stays in $(-\infty, 0]$:

$$\begin{aligned} \sum_{\tau=1:t_j \text{ fair}} \{\mathcal{R}_\tau(i) > \mathcal{R}_\tau(i_0)\} &= \sum_{\tau=1}^{\ell_j} \mathbf{1}\{\hat{\mu}_{\tau, i} \geq \hat{\mu}_{\tau, i_0}\} \\ &= \sum_{\tau=1}^{\ell_j} \mathbf{1}_{(-\infty, 0]} \left(\underbrace{\sum_{k=1}^{\tau} r_{i,k} - \sum_{j=1}^{\tau} r_{1,k}}_{G_\tau} \right), \end{aligned}$$

where $\sum_{k=1}^{\tau} r_{i,k}$ is the cumulative reward of arm i and $\sum_{k=1}^{\tau} r_{1,k}$ is the cumulative reward of arm i_0 . Therefore, we have written this quantity as the time spent by the random walk G_τ in the interval $(-\infty, 0]$, for $\tau = 1, \dots, \ell_j$. The drift term for this random walk is given by:

$$\mathbb{E}[r_{i,k} - r_{1,k}] = \mu_{i_0} - \mu_i = -\Delta_{ii_0}.$$

Therefore, we can apply Theorem 14 for the following choice of parameters,

- (a) $\alpha = \alpha_j = \frac{j}{2\lceil \log(T) \rceil} + o(T^{-1/2})$ (Proposition 15), which implies $n = \lceil T^{1/2+\alpha_j} \rceil = \ell_j$.
- (b) $\theta = \frac{j_*+1}{2\lceil \log(T) \rceil}$. (We can use this choice since the drift is

$$-\Delta_{ii_0} = - \underbrace{(4 \log(T) + 2)}_{\geq 2\sqrt{\log(T)+2}} \underbrace{\tilde{\Delta}_{ii_0}}_{\geq T^{-\frac{j_*+1}{2\lceil \log(T) \rceil}}},$$

therefore the assumptions of the theorem are respected.)

Applying the theorem, we have:

$$\mathbb{P}(E_i(2\ell_j)) \leq 3T^{-1/2+\theta} = 3T^{-1/2+\frac{j_*+1}{2\lfloor \log(T) \rfloor}}.$$

Summing over j , we get,

$$\begin{aligned} \mathbb{P}\left(\bigcup_{t=1}^T E_1(t)\right) &\leq \sum_{j=1}^{|\mathcal{L}|} \mathbb{P}(E_1(2\ell_j)) \\ &\leq 3 \log(T) T^{-1/2+\frac{j_*+1}{2\lfloor \log(T) \rfloor}}. \end{aligned}$$

To conclude, consider that, by definition ,

$$\ell_j = T^{1/2+\frac{j}{2\lfloor \log(T) \rfloor}};$$

this means that:

$$\tilde{\Delta}_{ii_0} \in \left(T^{-\frac{j_*+1}{2\lfloor \log(T) \rfloor}}, T^{-\frac{j_*}{2\lfloor \log(T) \rfloor}} \right],$$

so that, in particular,

$$\tilde{\Delta}_{ii_0}^{-1} \geq T^{\frac{j_*}{2\lfloor \log(T) \rfloor}}$$

and

$$\Delta_{ii_0}^{-1} \geq (4 \log(T) + 2) T^{\frac{j_*}{2\lfloor \log(T) \rfloor}}.$$

Substituting in the bound we have just found results in,

$$\begin{aligned} 3 \log(T) T^{-1/2+\frac{j_*+1}{2\lfloor \log(T) \rfloor}} &\leq 3 \log(T) (4 \log(T) + 2) T^{-1/2} T^{\frac{1}{2\lfloor \log(T) \rfloor}} \Delta_{ii_0}^{-1} \\ &\leq 6 \log(T) (4 \log(T) + 2) T^{-1/2} \Delta_{ii_0}^{-1}. \end{aligned}$$

□

Lemma 8. *For any arm i , the probability of the event E_i^* , corresponding to i not being eliminated after $(8 \log(T) + 4) T^{1/2} \Delta_{ii_0}^{-1}$ pulls if there is an active arm i_0 with $\Delta_{ii_0} := \mu_{i_0} - \mu_i > 0$ is, at least,*

$$\mathbb{P}(E_{ii_0}^*) \leq 2T^{-1/2}.$$

Proof. Define ℓ_j, α_j as in the previous lemma, so that $\alpha_j = \frac{\log(\ell_j)}{\log(T)} - \frac{1}{2}$. As in the previous lemma, we define:

$$\tilde{\Delta}_{ii_0} = \frac{\Delta_{ii_0}}{4 \log(T) + 2},$$

and $j_* \in \{1, \dots, \lfloor \log(T) \rfloor\}$ such that:

$$\tilde{\Delta}_{ii_0} \in \left(T^{-\frac{j_*+1}{2\lfloor \log(T) \rfloor}}, T^{-\frac{j_*}{2\lfloor \log(T) \rfloor}} \right].$$

Here, remember that by definition of the filtering condition, defining t_{j_*} as the fair timestep where both arms have been played ℓ_{j_*} times, E_i^* is included in the following event:

$$\left\{ \sum_{\tau=1:t_{j_*}} \mathbf{1}_{\text{fair}} \{\mathcal{R}_\tau(i_0) > \mathcal{R}_\tau(i)\} \geq T^{2\alpha_{j_*}} \right\}.$$

As before, this event can be interpreted as the difference between two random walks being negative, due to the fact that:

$$\begin{aligned} \sum_{\tau=1:t_{j_*}} \mathbf{1}_{\text{fair}} \{\mathcal{R}_\tau(i_0) > \mathcal{R}_\tau(i)\} &= \sum_{\tau=1}^{\ell_{j_*}} \mathbf{1} \{\hat{\mu}_{\tau,1} \geq \hat{\mu}_{\tau,i}\} \\ &= \sum_{\tau=1}^{\ell_{j_*}} \mathbf{1}_{(-\infty, 0]} \left(\underbrace{\sum_{k=1}^{\tau} r_{i_0,k} - \sum_{j=1}^{\tau} r_{i,k}}_{G_\tau} \right). \end{aligned}$$

In this formulation, we have written the quantity of interest for the filtering condition after ℓ_{j_*+1} pulls as the time spent by the random walk G_τ in the interval $(-\infty, 0]$, for $\tau = 1, \dots, \ell_{j_*+1}$. This time, the drift term is:

$$\mathbb{E}[r_{i_0, k} - r_{i, k}] = \mu_{i_0} - \mu_i = \Delta_{ii_0}.$$

Therefore, we can apply theorem 13 for $\alpha = \alpha_{j_*}$, since, by assumption,

$$\Delta_{ii_0} = \underbrace{(4 \log(T) + 2)}_{\geq 4 \log(T)} \underbrace{\tilde{\Delta}_{ii_0}}_{\geq T^{-\frac{j_*+1}{2 \lceil \log(T) \rceil}}}.$$

This theorem leads to:

$$\begin{aligned} \mathbb{P}(E_{ii_0}^*) &\leq 1 - \mathbb{P}\left(\sum_{\tau=1}^{\ell_{j_*+1}} \mathbf{1}_{(-\infty, 0]}(G_\tau) \leq T^{2\alpha_{j_*+1}}\right) \\ &\stackrel{thm.13}{\leq} 2T^{-1/2}. \end{aligned} \quad (8)$$

To get the thesis, is sufficient to reformulate the critical time-step ℓ_{j_*+1} in terms of Δ_{ii_0} . By definition,

$$\tilde{\Delta}_{ii_0} \leq T^{-\frac{j_*}{2 \lceil \log(T) \rceil}} = \frac{T^{1/2}}{\ell_{j_*}} \leq 2 \frac{T^{1/2}}{\ell_{j_*+1}}.$$

Therefore, $\Delta_{ii_0} \leq (8 \log(T) + 4) \frac{T^{1/2}}{\ell_{j_*+1}}$. From this, it immediately follows,

$$\ell_{j_*+1} \leq (8 \log(T) + 4) T^{1/2} \Delta_{ii_0}^{-1}.$$

□

We are finally able to prove our main result about the instance independent regret of the algorithm.

Theorem 8. *In the stochastic bandits with ranking feedback setting, when the noise is Gaussian, Algorithm 2 achieves $R_T \leq 62n^4 \log(T)^2 T^{1/2}$.*

Proof. Fix a sub-optimal arm i with corresponding gap Δ_i with respect to the optimal arm and let ψ a parameter to be chosen later. Define, for every couple of indices i_1, i_0 the event:

$$E_{i_1 i_0}^\psi := \begin{cases} E_{i_1 i_0} & \mu_{i_0} - \mu_{i_1} > \psi \\ \emptyset & \text{else,} \end{cases}$$

where the event $E_{i_1 i_0}$ is defined as arm i_1 eliminating arm i_0 in some point of the process. The probability that at least one of this events verifies is bounded by lemma 7 and union bound with:

$$\mathbb{P}(\Psi) := \mathbb{P}\left(\bigcup_{i_1=1, i_0=0}^n E_{i_1 i_0}^\psi\right) \leq 3n(n-1) \log(T) (4 \log(T) + 2) T^{-1/2} \psi^{-1},$$

as their number is at most $n(n-1)/2$. Therefore, under the complementary of Ψ , as no elimination with gap larger than ψ happens, we are sure that an arm i^* with gap (w.r.t. the first arm) less than $(n-1)\psi$ survives until the last: in fact, at most $n-1$ eliminations may happen, and all between pair of arms with a difference at most ψ .

As the arm i^* is active until the last, the probability of event $E_{ii^*}^*$ that i survives for more than $(8 \log(T) + 4) T^{1/2} \Delta_{ii^*}^{-1}$ pulls is bounded by Lemma 8 with $2T^{-1/2}$. Making the union bound over all possible values of i^* , which is a random variable, we can say that the probability that any of this event happens is at most $2(n-1) T^{-1/2}$. Summarizing, we have the following bound on $Z_i(T)$:

1. $Z_i(T) \leq T$ if either Ψ verifies, which happens with probability,

$$3n(n-1) \log(T) (4 \log(T) + 2) T^{-1/2} \psi^{-1},$$

- or if E_{ii^*} verifies, which happens with probability at most $2(n-1) T^{-1/2}$

2. $Z_i(T) \leq T(8 \log(T) + 4)T^{1/2}(\Delta_i - (n-1)\psi)^{-1}$ otherwise. The last comes just from the fact that $\Delta_{ii^*} \geq \Delta_i - (n-1)\psi$, by definition of i^* .

If we take $\psi = \frac{\Delta}{2(n-1)}$, it results in,

$$\begin{aligned}
\mathbb{E}[\Delta_i Z_i(T)] &\leq \Delta_i T \mathbb{P}(\cup_{i_1=1, i_0=0}^n E_{i_1 i_0}^\psi) + T \Delta_i \mathbb{P}(E_{ii^*}) \\
&\quad + \Delta_i T (8 \log(T) + 4) T^{1/2} (\Delta_i - (n-1)\psi)^{-1} \\
&\leq \Delta_i T \left(3n(n-1) \log(T) (4 \log(T) + 2) T^{-1/2} \psi^{-1} + 2n T^{-1/2} \right) \\
&\quad + \Delta_i T (8 \log(T) + 4) T^{1/2} (\Delta_i - (n-1)\psi)^{-1} \\
&\stackrel{\psi = \frac{\Delta}{2(n-1)}}{\leq} \Delta_i T \left(6n(n-1)^2 \log(T) (4 \log(T) + 2) T^{-1/2} \Delta_i^{-1} + 2n T^{-1/2} \right) \\
&\quad + 2 \Delta_i T (8 \log(T) + 4) T^{1/2} \Delta_i^{-1} \\
&= 6n(n-1)^2 \log(T) (4 \log(T) + 2) T^{1/2} + 2n \Delta_i T^{1/2} + 2T(8 \log(T) + 4) T^{1/2}.
\end{aligned}$$

Being $\Delta_i \in (0, 1)$ and $n \geq 2$, the previous quantity is bounded by,

$$\mathbb{E}[\Delta_i Z_i(T)] \leq 62n^3 \log(T)^2 T^{1/2}.$$

The proof is completed by using the the Regret Decomposition Lemma Lattimore and Szepesvari [2017]. \square

C Proof for adversarial setting

Theorem 9. *In adversarial bandits with ranking feedback, no algorithm achieves $o(T)$ regret with respect to the best arm in hindsight with a probability of $1 - \epsilon$ for any $\epsilon > 0$.*

Proof. This negative result follows from the impossibility to achieve $R_T \leq CT$ regret by any algorithm, with C properly set constant and probability $1 - \bar{\epsilon}$, in all three instances reported next. Please notice that, this result implies that even the No-Regret property cannot be achieved in the Bandit with Ranking Feedback setting.

Without loss of generality we consider rewards function bounded in $[0, 10]$. Consider three instances, with two arms a_0, a_1 for each and the associated rewards, defined as follows:

$$\begin{aligned}
\text{Instance } \textcircled{1} : & \begin{cases} a_0 : \frac{1}{2} & \forall t \in \boxed{1}, & \frac{1}{2} & \forall t \in \boxed{2}, & \frac{1}{2} & \forall t \in \boxed{3} \\ a_1 : 0 & \forall t \in \boxed{1}, & 0 & \forall t \in \boxed{2}, & 0 & \forall t \in \boxed{3} \end{cases} \\
\text{Instance } \textcircled{2} : & \begin{cases} a_0 : \delta & \forall t \in \boxed{1}, & 0 & \forall t \in \boxed{2}, & 0 & \forall t \in \boxed{3} \\ a_1 : 0 & \forall t \in \boxed{1}, & 1 & \forall t \in \boxed{2}, & 1 & \forall t \in \boxed{3} \end{cases} \\
\text{Instance } \textcircled{3} : & \begin{cases} a_0 : \delta & \forall t \in \boxed{1}, & 0 & \forall t \in \boxed{2}, & 10 & \forall t \in \boxed{3} \\ a_1 : 0 & \forall t \in \boxed{1}, & 1 & \forall t \in \boxed{2}, & 0 & \forall t \in \boxed{3} \end{cases}
\end{aligned}$$

where Phase $\boxed{1}$ is made by the first $T/4$ rounds, Phase $\boxed{2}$ is made by the next $T/4$ rounds, Phase $\boxed{3}$ is made by the last $T/2$ rounds and δ is near to 0.

In phase $\boxed{1}$ all the instances have the same ranking feedback, as the first action gives higher rewards with respect to the second one. To make instance $\textcircled{1}$ receive $R_T \leq CT$, it is necessary:

$$\frac{1}{2}T - \frac{1}{2}\mathbb{E}[n_{a_0}] \leq CT \Rightarrow \mathbb{E}[n_{a_0}] \geq (1 - 2C)T \quad (9)$$

where n_{a_0} is the number of times the first arm has been pulled, and the expected value is taken on the randomization of the algorithm. From previous equation we obtain that in all instances:

$$\mathbb{E} \left[n_{a_0}^{\boxed{1}} \right] \geq (1 - 2C)T - \frac{3}{4}T = (1 - C_1)T/4 \quad (10)$$

where $C_1 = 8C$, n_{a_0} is the number of time the first arm has to be pulled in phase $\boxed{1}$ and the inequality is computed considering that a_0 is played in all the next phases.

By reverse Markov inequality:

$$\mathbb{P} \left(n_{a_0}^{\boxed{1}} > (1 - \bar{C}_1)T/4 \right) \geq \frac{\bar{C}_1 - C_1}{C_1} \quad (11)$$

Setting the probability equal to 9/10 we obtain:

$$\bar{C}_1 = 10C_1 \quad (12)$$

from which follow that with probability 9/10 we have:

$$n_{a_0}^{\boxed{1}} > (1 - 10C_1)T/4 \quad (13)$$

and consequently:

$$n_{a_1}^{\boxed{1}} \leq 10C_1T/4. \quad (14)$$

We observe that in the second Phase, Instances $\textcircled{2}$ and $\textcircled{3}$ have the same feedback. Proceeding as done before, to make instance $\textcircled{2}$ receive $R_T \leq CT$ it is necessary:

$$\frac{3}{4}T - \mathbb{E}[n_{a_1}] \leq CT \Rightarrow \mathbb{E}[n_{a_1}] \geq \left(\frac{3}{4} - C \right) T \quad (15)$$

From previous equation we obtain that in instances $\textcircled{2}$ and $\textcircled{3}$:

$$\mathbb{E} \left[n_{a_1}^{\boxed{2}} \right] \geq \left(\frac{3}{4} - C \right) T - T/2 = (1 - C_2)T/4 \quad (16)$$

where the inequality is computed considering that a_1 is played in the next phases and $C_2 = 4C$. By Reverse Markov Inequality, we obtain that, with probability 9/10:

$$n_{a_1}^{\boxed{2}} > (1 - 10C_2)T/4 \quad (17)$$

and consequently:

$$n_{a_0}^{\boxed{2}} \leq 10C_2T/4 \quad (18)$$

We neglect the δ value for now, as it can be chosen to be insignificant with respect to the previous computation.

Now we focus on the third phase, in which instance $\textcircled{2}$ should play:

$$\mathbb{E} \left[n_{a_1}^{\boxed{3}} \right] \geq \left(\frac{3}{4} - C \right) T - T/4 = (1 - C_3)T/2, \quad (19)$$

where $C_3 = 2C$. By Reverse Markov Inequality, we obtain that, with probability 9/10:

$$n_{a_1}^{\boxed{3}} > (1 - 10C_3)T/2 \quad (20)$$

and consequently:

$$n_{a_0}^{\boxed{3}} \leq 10C_3T/2 \quad (21)$$

Now, we compute the number of rounds needed in the third instance to switch the ranking in the third phase, namely q . Notice that, until this switch, the last two instances receive the same feedback.

We compute q in the best-case scenario (that is, when small q value is sufficient to allow the switch) that satisfies the constraints previously shown. Precisely, q is computed so that the empirical mean of

arm a_0 is greater than the arm a_1 one, given that $n_{a_0}^{\boxed{1}} > (1 - 10C_1)T/4$ and $n_{a_1}^{\boxed{2}} > (1 - 10C_2)T/4$. Formally:

$$\frac{0(1 - 10C_1)T/4 + 10q}{q + (1 - 10C_1)T/4} \geq \frac{0C_110T/4 + (1 - 10C_2)T/4 + 0T/2}{10C_1T/4 + (1 - 10C_2)T/4 + T/2} \quad (22)$$

We now show that for proper C value we can lower bound the right side with $\frac{1}{4}$. In particular:

$$\frac{0C_110T/4 + (1 - 10C_2)T/4 + 0T/2}{10C_1T/4 + (1 - 10C_2)T/4 + T/2} > \frac{1}{4} \Rightarrow C < 1/200 \quad (23)$$

which means that, for $C < \frac{1}{200}$, we can substitute the right side of the equation with $\frac{1}{4}$ to simplify the computation. Moreover, notice that gap between $\frac{1}{4}$ and $\frac{0C_110T/4+(1-10C_2)T/4+0T/2}{10C_1T/4+(1-10C_2)T/4+T/2}$ allowed us to neglect the computations with δ . Then:

$$\frac{0(1 - 10C_1)T/4 + 10q}{q + (1 - 10C_1)T/4} \geq 1/4 \Rightarrow q \geq \frac{4}{39} \left(\frac{1}{4} - 20C \right) T/4 \quad (24)$$

To achieve a contradiction, it sufficient to find C so that $q + n_{a_1}^{\boxed{3}} > T/2$; indeed, the previous inequality shows the impossibility to gain enough rewards to make the ranking change and, at the same time, guarantee the minimum rewards to make instance ② No-Regret. Given that the ranking switch is a necessary condition to make instance ③ No-Regret, the result of impossibility follows for:

$$\frac{4}{39} \left(\frac{1}{4} - 20C \right) T/4 + (1 - 20C)T/2 > T/2 \Rightarrow C < \frac{1}{1640} \quad (25)$$

To conclude the proof, we show that the intersection between the events derived by Reverse Markov Inequality (namely E_i with $i \in [3]$) holds with constant probability:

$$\begin{aligned} \mathbb{P} \left(\bigcap_{i \in [3]} E_i \right) &= 1 - \mathbb{P} \left(\bigcup_{i \in [3]} E_i^c \right) \\ &\geq 1 - \sum_{i \in [3]} \mathbb{P}(E_i^c) \\ &= 1 - \frac{3}{10} = \frac{7}{10} \end{aligned}$$

where the inequality holds by Union Bound. Substituting all the previous results in the definition of Regret we obtain, with probability $\frac{7}{10} = 1 - \bar{\epsilon}$ and $C < \frac{1}{1640}$, $R_T \geq CT = \Omega(T)$ which concludes the proof. \square

D Numerical evaluation

This section presents a numerical evaluation of the algorithms proposed in the paper for the *stochastic settings*, namely, DREE and R-LPE. The goal of such a study is to show two crucial results: firstly, the comparison of our algorithms with a well-known bandit baseline, and secondly, the need to develop distinct algorithms tailored for instance-dependent and instance-independent scenarios.

To establish a benchmark for comparison, we consider the EC (Explore-Then-Commit) algorithm, which is one of the most popular algorithms among the explore-then-commit class providing sub-linear regret guarantees. In the following, we evaluate the DREE algorithm with different choices of the δ parameter in the function $f(t) = \log(t)^{1+\delta}$; precisely, we choose $\delta \in \{1.0, 1.5, 2.0\}$. Furthermore, we consider four stochastic instances whose specific parameters are discussed below. In all these instances, we assume the rewards to be drawn from Gaussian random variables with unit variance, *i.e.*, $\sigma^2 = 1$, and we let the time horizon be equal to $T = 2 \cdot 10^5$. Finally, for each algorithm, we evaluate the cumulative regret averaged over 50 runs.

We structure the presentation of the experimental results into two groups. In the first, the instances have a small Δ_{\min} , while in the second, the instances have a large Δ_{\min} .

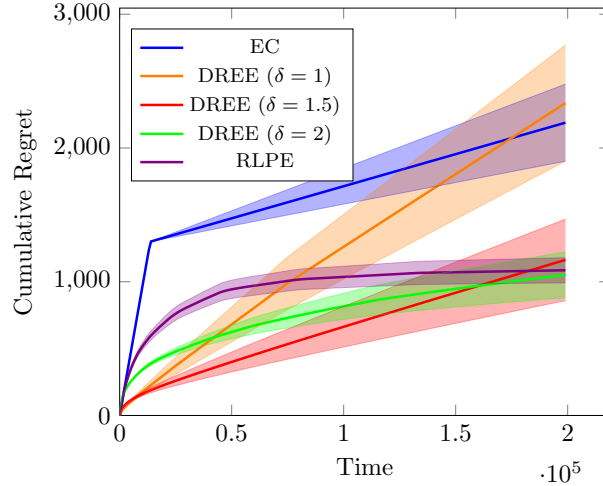


Figure 1: Instance with $\Delta_{\min} = 0.03$ and all the gaps small.




Figure 2: Instance with $\Delta_{\min} = 0.03$ and the other gaps big

Small values of Δ_{\min} We focus on two instances with $\Delta_{\min} < 0.05$. In the first of these two instances, we consider $n = 4$ arms, and a minimum gap of $\Delta_{\min} = 0.03$. In the second instance, we consider $n = 6$ arms, with $\Delta_{\min} = 0.03$. The expected values of the rewards of each arm are reported in Section D.1, while the experimental results in terms of average cumulative regret are reported in Figures 1–2. We observe that in the first instance (see Figure 1) all the DREE algorithms exhibits a linear regret bound, confirming the strong sensitivity of this family of algorithms on the parameter Δ_{\min} in terms of regret bound. In contrast, the R-LPE algorithm exhibits better performances in terms of regret bound, as its theoretical guarantee are independent on the values of Δ_{\min} . Furthermore, Figure 2 shows that the DREE algorithms (with $\delta \in 1.0, 1.5$) achieve a better regret bound when the number of arms is increased. Indeed, these regret bounds are comparable to the ones achieved by the R-LPE algorithm. The previous result is reasonable as the presence of Δ_i -s in the regret bound lowers the dependence on the number of arms. It is worth noticing that all our algorithms outperform the baseline EC.

Large values of Δ_{\min} We focus on two instances with $\Delta_{\min} \geq 0.25$. In the first instance, we consider $n = 4$ arms with a minimum gap of $\Delta_{\min} = 0.5$ among their expected rewards. In the




Figure 3: Instance with $\Delta_{\min} = 0.5$.




Figure 4: Instance with $\Delta_{\min} = 0.25$.

second instance, we instead consider a larger number of arms, specifically $n = 8$, with a minimum gap equal to $\Delta_{\min} = 0.25$. The expected values of the rewards are reported in Section D.1, while the experimental results in terms of average cumulative regret are provided in Figures 3–4. As it clear from both Figures 3–4 when Δ_{\min} is sufficiently large, the DREE algorithms (with $\delta \in \{1.0, 1.5\}$) achieves better performances with respect both the EC and R-PLPE algorithms in terms of cumulative regret. Furthermore, there is empirical evidence that a small δ guarantees better performance, which is reasonable according to theory. Indeed, when δ is small, the function $f(t)$, which drives the exploration, is closer to a logarithm. Also, as shown in Corollary 3, when Δ_{\min} is large enough, the parameter δ affects the dimension of $C(f, \Delta_i)$ more weakly, which results in a better regret bound.

D.1 Experiments

For the sake of clarity, we report in the followings additional details on the four instances presented in Figures 1,2,3,4. Notice that, all the plots present the cumulative regret averaged for 50 runs, with 95% confidence interval. Furthermore:

- *Instance of Figure 1*: time horizon $T = 2 \cdot 10^5$, arms $n = 4$, mean reward vector

$$\boldsymbol{\mu} = [0.9, 1.05, 1.12, 1.15],$$

unitary variance for each arm, $\Delta_{\min} = 0.03$;

- *Instance of Figure 2*: time horizon $T = 2 \cdot 10^5$, arms $n = 6$, mean reward vector

$$\boldsymbol{\mu} = [0.03, 0.07, 0.1, 0.08, 0.97, 1],$$

unitary variance for each arm, $\Delta_{\min} = 0.03$;

- *Instance of Figure 3*: time horizon $T = 2 \cdot 10^5$, arms $n = 4$, mean reward vector

$$\boldsymbol{\mu} = [0.05, 0.25, 0.5, 1.0],$$

unitary variance for each arm, $\Delta_{\min} = 0.5$;

- *Instance of Figure 4*: time horizon $T = 2 \cdot 10^5$, arms $n = 8$, mean reward vector

$$\boldsymbol{\mu} = [0.05, 0.05, 0.1, 0.15, 0.25, 0.5, 0.75, 1.0],$$

unitary variance for each arm, $\Delta_{\min} = 0.25$;

D.2 Detailed explanation of the experiments

In this section, we report all the details of the experiments performed in the paper. These are important to ensure the truthfulness of the results and the claims based on empirical validation.

Training details In the main paper we have presented four experiments, each corresponding to a different environment. Each experiment is performed for fifty random seeds, and the computation is split in 10 parallel processes by the library `joblib`. The overall computational time for one experiment is around 337.92 seconds, that is roughly five minutes and one half.

Compute As stated, the numerical simulations resulted to be very fast. For this reason, it was not necessary to run them on a server, and we used a personal computer with the following specifications:

- CPU: 11th Gen Intel(R) Core(TM) i7-1165G7 2.80 GHz
- RAM: 16,0 GB
- Operating system: Windows 11
- System type: 64 bit

Reproducibility Due to the stochastic nature of the bandit problem, all the simulations have been repeated several times. We have performed all the experiments with 50 different random seeds, corresponding precisely to the first 50 natural numbers. The seed influences the generation of the reward by the environment, while all algorithms proposed, being deterministic, are independent on the seed.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: All we say in the abstract is done in the work.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We clarify that the contribution is theoretical, which is the only aspect that can be seen as a limitation of this work.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We leave the complete and formal proofs in the appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: In the appendix we report every detail of the experiments, even the amount of CPU time used.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide the code.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: These information are reported in the appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Every experiment reports, together with the regret curves, a shaded area indicating the uncertainty.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: These information are reported in the appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: Experiments are just Python simulation, no concerns about real data.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: Only theoretical work.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We did not use data.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: We only use standard Python libraries to make the experiments.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: Answered in previous questions.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We did not do any crowdsourcing.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: We have no study participants.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.