

Complexity Analysis of a Countable-armed Bandit Problem

Anand Kalvit AKALVIT22@GSB.COLUMBIA.EDU and **Assaf Zeevi** ASSAF@GSB.COLUMBIA.EDU
Columbia University

Editors: Shipra Agrawal and Francesco Orabona

Abstract

We consider a stochastic multi-armed bandit (MAB) problem motivated by “large” action spaces, and endowed with a population of arms containing exactly K arm-types, each characterized by a distinct mean reward. The decision maker is oblivious to the statistical properties of reward distributions as well as the population-level distribution of different arm-types, and is precluded also from observing the type of an arm after play. We study the classical problem of minimizing the expected cumulative regret over a horizon of play n , and propose algorithms that achieve a rate-optimal finite-time instance-dependent regret of $\mathcal{O}(\log n)$. We also show that the instance-independent (minimax) regret is $\tilde{\mathcal{O}}(\sqrt{n})$ when $K = 2$. While the order of regret and complexity of the problem suggests a great degree of similarity to the classical MAB problem, properties of the performance bounds and salient aspects of algorithm design are quite distinct from the latter, as are the key primitives that determine complexity along with the analysis tools needed to study them.

Keywords: Infinite-armed bandits, Explore-then-Commit, UCB, Adaptivity to reservoir-distribution.

1. Introduction

Background and motivation. The MAB model is a classical machine learning paradigm capturing the essence of the *exploration-exploitation* trade-offs characteristic of sequential decision making problems under uncertainty. In its simplest formulation, the decision maker (DM) must play at each time instant $t \in \{1, \dots, n\}$ one out of a set of $K \ll n$ possible alternatives (*aka* arms), each characterized by a distribution of rewards. Oblivious to their statistical properties, the DM must play a sequence of n arms so as to maximize her cumulative expected payoffs, an objective often converted to minimizing *regret* relative to an oracle with perfect ex ante knowledge of the best arm.

Since the seminal paper of [Lai and Robbins \(1985\)](#) that laid the main theoretical foundations, there has been a plethora of work developing more advanced MAB models to encapsulate more realistic data-driven decision processes. These include formulations with covariate or contextual information, choice-models, budget constraints, non-stationary rewards, and metric space embeddings, among many others that utilize some structure in the arms, reward distributions, or physics of the problem (see [Slivkins et al. \(2019\)](#); [Lattimore and Szepesvári \(2020\)](#) for a comprehensive survey). In this paper, we are motivated by the *choice overload* phenomenon pervading modern MAB applications with a prohibitively large action space such as those encountered in online marketplaces, matching platforms and the likes.

Modeling choice overload. In several applications of MAB, it is quite common for the number of arms to be “large” to the extent that it may potentially exceed even the horizon of play, i.e. $K \geq n$. For example, the problem faced by recommendation systems in large retail platforms, such

as Amazon, is characterized by a prohibitively large number of arms (products of certain type) and limited “display space,” creating a very challenging combinatorial problem (see, e.g., Agrawal et al. (2019)). Naturally, the canonical MAB model is ill-suited for the study of such settings. Among problems of this nature, a simple yet illuminating abstraction is one where an infinite population of arms is partitioned into K different *arm-types*, each characterized uniquely by some reward statistic (e.g., the mean), and the fraction of each arm-type in the population of arms (or the arm-reservoir) remains fixed over the horizon of play. The motivation to study such settings stems from several contemporary applications, e.g., in a prototypical task-matching problem arising in the online gig economy: the platform must choose upon each task arrival, one agent from a large pool of available agents characterized by unknown (or only partially known) skill proficiencies. Such settings arise naturally in populations endowed with latent *low-dimensional* representations, i.e., an agent can only belong to one of finitely many possible *types*, each characterized distinctly by some attribute. Market segmentation based on types is central also to the operations research literature, see, e.g., Bui et al. (2012); Banerjee et al. (2017); Johari et al. (2021), etc., for examples involving analyses of online service and recommendation systems, among several other areas.

The countable-armed bandit problem. We provide an abstraction of the aforementioned decision problem as a bandit with countably many arms, each queried from an infinite population of arms (*aka* the arm-reservoir). There are K possible arm-types given by $\mathcal{K} := \{1, \dots, K\}$ with K known a priori¹, and the probability vector $\alpha = \{\alpha_i : i \in \mathcal{K}\}$ denotes their corresponding fraction, i.e., relative prevalence in the reservoir, which is unknown.

Intuitively, the statistical complexity of regret minimization in the simplest formulation of the countable-armed bandit (CAB) with $K = 2$ arm-types is governed by three principal primitives: (i) the sub-optimality gap $\Delta := \mu_1 - \mu_2$ between the mean rewards $\mu_1 > \mu_2$ of the optimal and inferior arm-types; (ii) the probability α_1 of sampling from the population an arm of the optimal type; and (iii) the horizon of play n . Absent knowledge of (Δ, α_1) , exploration is challenging owing to the “large” number of arms. In particular, in contrast with the classical two-armed MAB, in the CAB problem, any finite selection of arms may only contain the mean μ_2 . Consequently, any algorithm limited to such a consideration set will suffer a linear regret. Absence of information on α_1 further exacerbates the challenges in the study of the CAB problem.

Contributions. There has been limited technical development in this area and the literature remains sparse. In this work, we resolve several foundational questions pertaining to the complexity of this problem setting and provide a comprehensive understanding of various other aspects thereof. Our theoretical contributions can be projected along the following axes:

(i) Complexity of regret. We establish information-theoretic performance lower bounds that are order-wise tight (in the horizon n) in the instance-dependent setting (Theorem 2). In the instance-independent (minimax) setting, we answer affirmatively an open question on achievability of $\tilde{O}(\sqrt{n})$ regret when $K = 2$ and show that this order is best achievable up to poly-logarithmic factors in n (Theorem 5). In addition, we provide a *uniform* lower bound on achievable performance that is tight in n and explicitly captures the scaling behavior w.r.t. the fraction α_1 of optimal arms, and furthermore, has a novel non-information-theoretic proof based entirely on convex analysis (Theorem 3).

1. It is routine for platforms to run pilot experiments during initial rounds to gather information on key primitives such as the size and stability of clusters, if any exist in the population. One can therefore safely assume in settings where such clusters strongly exist that the number K of possible types is accurately estimated.

(ii) Algorithm design. We design algorithms that achieve aforementioned regret guarantees relying only on knowledge of K and are agnostic to information pertaining to the reward distributions as well as the frequency of occurrence of different types. Our design principles (Algorithms 1 and 2) are functionally distinct from extant work on finite-armed bandits which reflects in a fundamentally different scaling of regret (see Theorems 5 and 7). We also provide resolution to an outstanding design issue in extant literature for $K = 2$ (see Algorithm 3 and Theorem 8).

(iii) Regret behavior and arm-type distribution. In contrast to some observations on related models in the literature that show a higher order than $\log n$ (instance-dependent) regret-behavior w.r.t. α , we establish that when the learner has knowledge of K but not of α , one can still achieve $\mathcal{O}(\log n)$ regret where the dependence on α only manifests as an additive loss (Theorems 7 and 8).

Before proceeding with a formal description of our model, we provide a brief overview of related works below.

Extant literature on bandits with infinitely many arms. These problems involve an infinite population of arms and a fixed *reservoir distribution* over a (typically uncountable) set of arm-types; a common reward statistic (usually the mean) uniquely characterizes each arm-type. The infinite-armed bandit problem traces its roots to [Berry et al. \(1997\)](#) where the problem was studied under Bernoulli rewards and a reservoir distribution of Bernoulli parameters that is Uniform on $[0, 1]$. Subsequent works have considered more general reward and reservoir distributions on $[0, 1]$, see, e.g., [Wang et al. \(2009\)](#); [Bonald and Proutiere \(2013\)](#); [Carpentier and Valko \(2015\)](#); [Chan and Hu \(2018\)](#). In terms of the statistical complexity of regret minimization, an uncountably rich set of arm-types is tantamount to the minimal achievable regret being polynomial in the horizon of play (see aforementioned references). In contrast, the recently studied models in [Kalvit and Zeevi \(2020\)](#); [de Heide et al. \(2021\)](#) that our work is most closely related to, are fundamentally simpler owing to a finite set of arm-types; this is central to the achievability of logarithmic (instance-dependent) regret in this class of problems. These two works are briefly discussed below.

The CAB problem first appeared in [Kalvit and Zeevi \(2020\)](#) together with an online adaptive policy achieving $\mathcal{O}(\log n)$ regret when $K = 2$. This policy is derived from UCB1 ([Auer et al., 2002](#)) and relies on certain newly discovered concentration and convergence properties thereof (see [Kalvit and Zeevi \(2021\)](#) for a detailed discussion of said properties). However, the analysis of this policy cannot be adapted to $K > 2$ types; we will later provide an example with $K > 2$ where the policy will likely run into issues that can be effectively mitigated by the algorithms proposed in this paper. There is also recent literature ([de Heide et al., 2021](#)) on a related setting where the set of inferior arm-types may be arbitrary as long as it is $\underline{\Delta}$ -separated from the optimal mean. However, ex ante knowledge of the proportion α_1 of optimal arms is necessary to achieve logarithmic regret in this setting. This aspect distinguishes their setting from CAB and will be discussed at length later.

Lastly, a formulation of the countable-armed problem based on *pure exploration*, referred to as the “heaviest coin identification problem,” was studied in [Chandrasekaran and Karp \(2014\)](#) for $K = 2$ (see [Jamieson et al. \(2016\)](#) for subsequent developments). In contrast, our problem is based on optimization of cumulative payoff (or regret); as a result, it shares little similarity with cited works.

Outline of the paper. A formal description of the model is provided in §2; §3 discusses lower bounds on achievable performance. We propose our algorithms in §4 and state supporting theoretical guarantees. Numerical experiments, proofs, and auxiliary results are relegated to the appendix.

2. Problem formulation

The set of arm-types is given by $\mathcal{K} = \{1, \dots, K\}$, and the decision maker (DM) only knows the cardinality K of \mathcal{K} . Each type $i \in \mathcal{K}$ is characterized by a unique mean reward μ_i ; the reservoir is thus characterized by the collection $\boldsymbol{\mu} := \{\mu_i : i \in \mathcal{K}\}$ of possible mean rewards. Without loss of generality, we assume $\mu_1 > \dots > \mu_K$ and refer to type 1 as the *optimal* type (we may refer to the others as inferior types). Define $\bar{\Delta} := \mu_1 - \mu_K$ and $\underline{\Delta} := \mu_1 - \mu_2$ as the maximal and minimal sub-optimality gaps respectively, and $\delta := \min_{1 \leq i < j \leq K} (\mu_i - \mu_j)$ as the minimal reward gap. Finally, $\boldsymbol{\alpha} := (\alpha_i : i \in \mathcal{K})$ denotes the vector of reservoir probabilities for each type (*aka* the reservoir distribution), coordinate-wise bounded away from 0. These primitives will be important drivers of the statistical complexity of the regret minimization problem, as we shall later see. The horizon of play is n , and the DM must play one arm at each time $t \in \{1, \dots, n\}$.

The set of arms that have been played up to and including time $t \in \{1, 2, \dots\}$ is denoted by \mathcal{I}_t (and $\mathcal{I}_0 := \emptyset$). The set of actions available to the DM at time t is given by $\mathcal{A}_t = \mathcal{I}_{t-1} \cup \{\text{new}_t\}$; the DM must either play an arm from \mathcal{I}_{t-1} at time t or select the action “new _{t} ” which corresponds to querying (and playing) a new arm from the reservoir. This new arm is optimal-typed with probability α_1 and sub-optimal otherwise. The DM is oblivious to $\boldsymbol{\alpha}$, and furthermore precluded from observing the type of an arm upon query or play. A policy $\pi := (\pi_1, \pi_2, \dots)$ is an adaptive allocation rule that prescribes at time t an action π_t from \mathcal{A}_t (possibly randomized). Each pull (or play) of an arm results in a stochastic reward. The sequence of rewards realized from the first k pulls of an arm labeled i (henceforth called arm i) is denoted by $(X_{i,j})_{1 \leq j \leq k}$; these are mean-preserving in time keeping the arm fixed, independent across arms and time, and take values in $[0, 1]$. The natural filtration at time t , denoted by \mathcal{F}_t and defined w.r.t. the sequence of rewards realized up to and including time t , is given by $\mathcal{F}_t := \sigma \left\{ (\pi_s)_{1 \leq s \leq t}, \left\{ (X_{i,j})_{1 \leq j \leq N_i(t)} : i \in \mathcal{I}_t \right\} \right\}$ (with $\mathcal{F}_0 := \emptyset$), where $N_i(t)$ denotes the number of pulls of arm i up to and including time t . The cumulative *pseudo-regret* of policy π after n plays is given by $R_n^\pi := \sum_{t=1}^n (\mu_1 - \mu_{\mathcal{T}(\pi_t)})$, where $\mathcal{T}(\pi_t) \in \mathcal{K}$ denotes the type of the arm played by π at time t ; note that R_n^π is a sample path-dependent quantity. The DM is interested in the classical problem of minimizing the expectation of the *cumulative regret* $\hat{R}_n^\pi := \sum_{t=1}^n (\mu_1 - X_{\pi_t, N_{\pi_t}(t)})$, given by

$$\inf_{\pi \in \Pi} \mathbb{E} \hat{R}_n^\pi = \inf_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=1}^n (\mu_1 - X_{\pi_t, N_{\pi_t}(t)}) \right] \stackrel{(\dagger)}{=} \inf_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=1}^n (\mu_1 - \mu_{\mathcal{T}(\pi_t)}) \right] = \inf_{\pi \in \Pi} \mathbb{E} R_n^\pi, \quad (1)$$

where (\dagger) follows from the Tower property of expectation, the infimum is over policies satisfying the *non-anticipation* property $\pi_t : \mathcal{F}_{t-1} \rightarrow \mathcal{P}(\mathcal{A}_t)$ for $t \in \{1, 2, \dots\}$; $\mathcal{P}(\mathcal{A}_t)$ denotes the probability simplex on \mathcal{A}_t . Accordingly, the expectation in (1) is w.r.t. all the possible sources of randomness in the problem (rewards, policy, and the arm-reservoir).

3. Lower bounds for natural policy classes

There are three fundamental primitives governing the complexity of achievable regret in this setting, viz., (i) the minimal sub-optimality gap $\underline{\Delta}$; (ii) the proportion α_1 of the optimal arm-type in the reservoir; and (iii) the number of arm-types K . We next characterize lower bounds on achievable performance w.r.t. each of these primitives.

3.1. Achievable performance w.r.t. $\underline{\Delta}$: Information-theoretic lower bounds

The statistical complexity of this problem setting is best illustrated via the paradigmatic case of $K = 2$ and $\alpha_1 \leq 1/2$. In this case, one anticipates the problem to be at least as hard as the classical two-armed bandit with a mean reward gap of $\underline{\Delta}$. Indeed, we establish this in Theorem 2 via information-theoretic reductions adapted to handle a countable number of arms (proof is provided in Appendix C). In what follows, an instance ν of the problem refers to a collection of reward distributions with gap $\underline{\Delta}$; note that we are excluding the reservoir probabilities $\alpha = (\alpha_1, \alpha_2)$ from the definition of an instance. Recall that α_1 denotes the reservoir probability associated with the optimal mean reward, and $\alpha_2 = 1 - \alpha_1$ is the probability of the inferior. We will overload the notation for expected cumulative regret slightly to emphasize its dependence on ν as well as α .

Definition 1 (Admissible policies when $K = 2$) *A policy π is deemed admissible if for any instance ν , reservoir distributions $\alpha = (\alpha_1, 1 - \alpha_1)$, $\alpha' = (\alpha'_1, 1 - \alpha'_1)$, and horizon n , one has that $\mathbb{E}R_n^\pi(\nu, \alpha') \geq \mathbb{E}R_n^\pi(\nu, \alpha)$ whenever $\alpha'_1 \leq \alpha_1$. The set of such policies is denoted by Π_{adm} .*

We remark that the aforementioned definition is not restrictive in our problem setting since it is only natural that any reasonable policy should incur a larger cumulative regret (in expectation) in problems where the reservoir holds fewer optimal arms (in proportion).

Theorem 2 (Information-theoretic lower bounds when $K = 2$) *There exists an absolute constant $C > 0$ such that the following holds under any $\pi \in \Pi_{adm}$ and any α with $\alpha_1 \leq 1/2$:*

1. *For any $\underline{\Delta} > 0$, there exists a problem instance ν such that $\mathbb{E}R_n^\pi(\nu, \alpha) \geq C \log n / \underline{\Delta}$ for large enough n , where the “large enough n ” depends exclusively on $\underline{\Delta}$.*
2. *For any $n \in \mathbb{N}$, there exists a problem instance ν such that $\mathbb{E}R_n^\pi(\nu, \alpha) \geq C \sqrt{n}$.*

Distinction from classical MAB. Although the above result bears resemblance to classical information-theoretic lower bounds for finite-armed bandits, it is imperative to note that the setting has a fundamentally greater complexity that requires a more nuanced analysis vis-à-vis the finite-armed problem. Traditional proofs, as a result, cannot be tailored to our setting in a translational manner. To see this, note that when α_1 is high, a query of the reservoir is very likely to return an arm of the optimal type; in the limit as $\alpha_1 \rightarrow 1$, the problem becomes degenerate as all policies incur zero expected regret. Clearly, the problem cannot be harder than a two-armed bandit with gap $\underline{\Delta}$ uniformly over all values of α_1 . While we conjecture $\alpha_1 < 1$ to be a sufficient condition for the existence of $\Omega(\log n / \underline{\Delta})$ instance-dependent and $\Omega(\sqrt{n})$ instance-independent (minimax) lower bounds, there are technical challenges due to probabilistic type associations over countably many arms. The restriction to $\alpha_1 \leq 1/2$ and *admissible* policies (Definition 1) is then necessary for tractability of the proof and it remains unclear if this can be generalized further. Furthermore, when $K > 2$, even defining an appropriate notion of admissibility à la Definition 1 is non-trivial and will likely involve dependencies on μ in addition to α ; pursuits in this direction are currently left to future work.

3.2. Achievable performance w.r.t. α_1 : A uniform lower bound for front-loaded policies

Though the bounds in Theorem 2 are tight in n as we shall later see, they fail to provide any actionable insights w.r.t. α . A natural question in the CAB setting is whether and in what manner does the presence of countably many arms affect achievable regret. In particular, how does the difficulty associated with finding optimal arms from the reservoir (and the dependence on the distribution α)

come into play. Below, we propose a lower bound that explicitly captures this dependence, albeit with respect to a somewhat restricted policy class.

Theorem 3 (α -dependent lower bound for any $K \geq 2$) *Denote by $\tilde{\Pi}$ the class of policies under which the decision to query the arm-reservoir at any time $s \in \{1, 2, \dots\}$ is independent of \mathcal{F}_{s-1} . Then, for all problem instances ν with a minimal sub-optimality gap of at least $\underline{\Delta} > 0$, one has*

$$\liminf_{n \rightarrow \infty} \inf_{\pi \in \tilde{\Pi}} \frac{\mathbb{E} R_n^\pi(\nu, \alpha)}{\log n} \geq \frac{(1 - \alpha_1)^2 \underline{\Delta}}{\alpha_1}.$$

Discussion. The proof is located in Appendix D. Several comments are in order. **(i)** The class $\tilde{\Pi}$, in particular, includes policies that front-load queries, i.e., query the reservoir upfront for a pre-specified number of arms and then deploy a regret minimizing routine on the queried set until the end of the playing horizon, see, e.g., the `SAMPLING-UCB` algorithm due to de Heide et al. (2021). **(ii)** The cited paper also derives an information-theoretic $\Omega(\log n / \alpha_1 \underline{\Delta})$ lower bound based on a standard reduction to a hypothesis testing problem, although notably their setting is non-trivially distinct from ours (this reflects starkly different sensitivities of achievable regret to $1/\alpha_1$ -scale, as we shall later see). Importantly though, akin to Theorem 2, their bound too, establishes the *existence* of an instance with logarithmic regret. On the other hand, the foremost noticeable aspect of Theorem 3 that differs from aforementioned results is that it provides a uniform lower bound *over all instances* that are at least $\underline{\Delta}$ -separated in the mean reward, as opposed to merely establishing their existence. **(iii)** The presence of $\underline{\Delta}$ in the numerator (unlike traditional bounds where $\underline{\Delta}$ resides in the denominator) suggests that while this bound may be vacuous if $\underline{\Delta}$ is “small,” it certainly becomes most relevant when $\underline{\Delta}$ is “well-separated.” At the same time, it should be noted that Theorem 3 does not contradict the $\mathcal{O}(\log n / \alpha_1 \underline{\Delta})$ upper bound (up to logarithmic factors in $\underline{\Delta}$) of `SAMPLING-UCB`; it merely provides a tool to separate regimes of $\underline{\Delta}$ where one bound captures the dominant effects vis-à-vis the other. **(iv)** A novelty of Theorem 3 lies in its proof, which differs from classical lower bound proofs in that it is based entirely on ideas from convex analysis as opposed to the information-theoretic and change-of-measure techniques hitherto used in the literature.

Remarks. **(i)** It is not impossible to avoid the $1/\alpha_1$ -scaling of the instance-dependent logarithmic regret. We will later show via an upper bound for one of our algorithms (`ALG2`(n)) that the α_1 -dependence can, in fact, be relegated to constant order terms (`ALG2`(n) queries arms *adaptively* based on sample-history and therefore does not belong to $\tilde{\Pi}$). Importantly, this will somewhat surprisingly establish that the instance-dependent logarithmic bound in Theorem 2 is optimal w.r.t. to its dependence on α_1 (the scaling w.r.t. $\underline{\Delta}$, however, may not be best possible as forthcoming upper bounds suggest). **(ii)** Theorem 3 holds also for any arm-reservoir where the optimal type is at least $\underline{\Delta}$ -separated from the rest, the nature of types (countable or uncountable) notwithstanding.

3.3. Achievable performance w.r.t. K : The Bandit and the Coupon-collector

In the classical K -armed bandit problem, the (instance-dependent) regret scales linearly with the number of arms. We will next show that the K -typed countable-armed setting studied in this paper differs from its K -armed counterpart on account of a fundamentally distinct scaling of regret w.r.t. K . We will illustrate this by pivoting to a full information setting with one-sample learning, i.e., a setting where the decision maker observes the mean reward of an arm immediately upon pulling it, but does not learn whether it is optimal. Under such a setting, the optimal policy π^* for the K -armed

problem will pull each of the K arms once and subsequently commit the residual budget of play to the optimal arm, thus incurring a *lifetime regret* of $\mathbb{E}R_\infty^{\pi^*} = \sum_{i=2}^K (\mu_1 - \mu_i) = \Theta(K)$. The optimal policy for the K -typed countable-armed setting will, analogously, keep querying new arms from the reservoir until it has collected one of each of the K types, and will subsequently commit to the arm within said collection that has the best mean reward. In this case, regret will only accrue until the decision maker has pulled one arm of each type.

Theorem 4 (Regret scaling w.r.t. K) *In the full information setting, the lifetime regret of any policy under reservoir distribution α and mean reward vector μ is at least $\sum_{i=2}^K \alpha_i (\mu_1 - \mu_i) K \log K$.*

If the reservoir distribution remains non-degenerate w.r.t. the optimal type, i.e., the fraction α_1 of optimal arms in the reservoir remains bounded away from 1 as K increases, it is ensured that $\sum_{i=2}^K \alpha_i (\mu_1 - \mu_i)$ remains non-vanishing in K . Consequently, the lower bound in Theorem 4 grows as $\Omega(K \Delta \log K)$.

This result establishes a fundamentally distinct scaling of regret w.r.t. K in the countable-armed setting vis-à-vis the K -armed one (in the full information setting). When the true type of an arm is not immediately observable, one only expects the $\Omega(K \log K)$ scaling to exacerbate. In fact, when the learning horizon is n , we conjecture that regret grows at least as $\Omega(K \log K \log n)$, where the $\Omega(\cdot)$ is modulo gap-dependent constants. Characterizing the information-theoretic optimal rate, however, remains a challenging open problem. The proof of Theorem 4 is provided in Appendix E.

4. Gap and reservoir adaptive policies

As discussed, our goal here is to investigate regret achievable under μ -adaptive algorithms that are agnostic also to ex ante information on the distribution α of possible arm types. We propose two algorithms; $\text{ALG1}(n)$ and $\text{ALG2}(n)$, that are both predicated on ex ante knowledge of the horizon of play n . §4.1 discusses the first of these, $\text{ALG1}(n)$, which uses knowledge of n to calibrate the duration of its exploration phases. $\text{ALG1}(n)$ serves as an insightful basal motif for algorithm design in that it satisfies the desiderata of an $\mathcal{O}(\log n)$ instance-dependent regret for general $K \geq 2$ as well as an $\tilde{\mathcal{O}}(\sqrt{n})$ instance-independent (minimax) regret when $K = 2$; the latter property settling an open problem in the literature. However, its regret has a sub-optimal dependence on α . We leverage structural insights from the analysis of $\text{ALG1}(n)$ to explore another design in $\text{ALG2}(n)$ in §4.2, which determines its exploration phase lengths *adaptively*, as opposed to pre-specifying them upfront. This new design guarantees the best possible dependence of regret on α . Finally in §4.3, we discuss a fully sequential adaptive algorithm from extant literature for $K = 2$, and propose a simple modification to rid it of a certain fragile assumption pertaining to ex ante knowledge of the support of reward distributions. We also provide new sharper bounds for the modified algorithm and discuss potential issues with its generalization to $K \geq 2$ vis-à-vis $\text{ALG1}(n)$ and $\text{ALG2}(n)$.

4.1. Explore-then-commit with a pre-specified exploration schedule

In what follows (and all subsequent algorithms), a *new* arm is one that is freshly queried from the reservoir (an arm without a history of previous pulls). This arm belongs to type i with probability α_i independent of the problem history thus far (collection of arms and types queried and the corresponding reward realizations until the current time).

Algorithm 1 ALG1(n) (Fixed design ETC)

- 1: **Input:** Horizon of play n .
 - 2: Set budget $T = n$; set epoch counter $k = 1$.
 - 3: **Initialize new epoch:** Query K new arms; call it consideration set $\mathcal{A} = \{1, \dots, K\}$.
 - 4: Set exploration duration $L = \lceil e^{2\sqrt{k}} \log n \rceil$.
 - 5: $m \leftarrow \min(L, \lfloor T/K \rfloor)$.
 - 6: Play each arm in \mathcal{A} m times; observe rewards $\{(X_{1,j}, \dots, X_{K,j}) : j = 1, \dots, m\}$.
 - 7: Update budget: $T \leftarrow T - Km$.
 - 8: **if** $\exists a, b \in \mathcal{A}, a < b$ s.t. $\left| \sum_{j=1}^m (X_{a,j} - X_{b,j}) \right| < 2me^{-\sqrt{k}}$ **then**
 - 9: Permanently discard \mathcal{A} .
 - 10: $k \leftarrow k + 1$.
 - 11: Repeat from step (3).
 - 12: **else**
 - 13: Permanently commit to arm $a^* \in \arg \max_{a \in \mathcal{A}} \left\{ \sum_{j=1}^m X_{a,j} \right\}$.
-

Discussion of Algorithm 1. The foremost noticeable feature of this algorithm is the (nearly) exponentially increasing exploration schedule. Specifically, in the k^{th} epoch, each of the K arms in the consideration set is played $\lceil e^{2\sqrt{k}} \log n \rceil$ times. It suffices to cap the size of the consideration set at K since the decision maker is a priori aware of the existence of exactly K arm-types in the reservoir. Upon completion of the k^{th} epoch, the cumulative-difference-of-reward statistic for each of the $\binom{K}{2}$ arm-pairs is compared against a threshold of $2e^{-\sqrt{k}}m$, where $2e^{-\sqrt{k}}$ should be imagined as a proxy for a lower bound on the minimal reward gap δ . If said statistic is small relative to the threshold for some pair, the pair is likely to contain arms of the same type (equal means), in which case, the algorithm discards the *entire* consideration set and ushers in a new epoch with a larger exploration phase. This is done to avoid the possibility of incurring linear regret should an optimal arm be missing from the consideration set (e.g., when all K arms belong to type 2). On the other hand, if all arm-pairs are sufficiently separated, the algorithm simply commits permanently to the empirically best arm. The intuition behind the (nearly) exponential schedule is that as k grows, $2e^{-\sqrt{k}}$ will eventually provide a lower bound on δ , and one may hope to achieve appropriate levels of error control using window sizes in the k^{th} epoch. Full proof is provided in Appendix F.

Theorem 5 (Upper bound on the regret of ALG1(n)) For a horizon of play $n \geq K$, the expected cumulative regret of the policy π given by ALG1(n) is bounded as

$$\mathbb{E}R_n^\pi \leq \frac{\tilde{C}_\alpha \bar{\Delta} \log n}{\delta^2} \log^2 \left(\frac{4}{\delta} \right) + 2K\bar{\Delta},$$

where \tilde{C}_α is some constant that depends only on α ; an exact expression is provided in (13). In particular, $\tilde{C}_\alpha \rightarrow \infty$ as $\prod_{i=1}^K \alpha_i$ approaches 0.

Discussion. The dependence on the *minimal reward gap* δ in Theorem 5 is not an artifact of our analysis but, in fact, reflective of the operating principle of the algorithm. ALG1(n) keeps querying new consideration sets of size K until it determines with high enough confidence that the queried arms are all distinct-typed; this is the genesis of δ in the upper bound. Importantly, equipped just

with knowledge of K , it remains unclear if there exists an alternative sampling strategy that does not rely on assessing pairwise similarities between the queried arms, without necessitating any additional information on α . Furthermore, while $\text{ALG1}(n)$ is evidently rate-optimal in n (in the instance-dependent sense), the scaling of its upper bound w.r.t. α is far from optimal. In particular, the α -dependent factor in the leading term is attributable to a naive pre-determined exploration schedule. This dependence can, in fact, be relegated to $\mathcal{O}(1)$ terms using a more sophisticated policy that operates based on an adaptive determination of stopping and re-initialization times.

Remarks. (i) When $K = 2$, the upper bound in Theorem 5 reduces to $\left(\tilde{C}_\alpha/\underline{\Delta}\right) \log^2(4/\underline{\Delta}) \log n + 4\underline{\Delta}$, leading to a worst-case regret (w.r.t. $\underline{\Delta}$) of $\tilde{\mathcal{O}}\left(\tilde{C}_\alpha\sqrt{n}\right)$, where the big-Oh only hides poly-logarithmic factors in n . This settles an open question concerning the best achievable minimax regret in the countable-armed problem with two types.² (ii) While specifying the exploration schedule, the choice of the exponent in k can be fairly general as long as it is coercive and grows sufficiently fast but sub-linearly; the square-root function is chosen for technical convenience. Instead, if one were to use a linear function of k in the exponent, the algorithm's performance would become fragile w.r.t. ex ante knowledge of α ; an ill-calibrated $\text{ALG1}(n)$ can potentially incur linear regret.

4.2. Explore-then-commit with an adaptive stopping time

Algorithm 2 $\text{ALG2}(n)$ (ETC with adaptive stopping times)

- 1: **Input:** Horizon of play n .
 - 2: Set budget $T = n$; Burn-in samples (per arm) s_n .
 - 3: **Initialize new epoch:** Query K new arms; call it consideration set $\mathcal{A} = \{1, \dots, K\}$.
 - 4: Play each arm in \mathcal{A} for s_n periods; observe rewards $\{X_{a,j} : a \in \mathcal{A}, j = 1, \dots, s_n\}$.
 - 5: Set per-arm sample count $m = s_n$.
 - 6: Update budget: $T \leftarrow T - s_n K$.
 - 7: Generate $\binom{K}{2}$ independent standard Gaussian random variables $\{\mathcal{Z}_{a,b} : a, b \in \mathcal{A}, a < b\}$.
 - 8: **while** $T \geq K$ **do**
 - 9: **if** $\exists a, b \in \mathcal{A}, a < b$ s.t. $\left|\mathcal{Z}_{a,b} + \sum_{j=1}^m (X_{a,j} - X_{b,j})\right| < 4\sqrt{m \log m}$ **then**
 - 10: Permanently discard \mathcal{A} and repeat from step (3).
 - 11: **else**
 - 12: **if** $\left|\sum_{j=1}^m (X_{a,j} - X_{b,j})\right| \geq 4\sqrt{m \log n} \forall a, b \in \mathcal{A}, a < b$ **then**
 - 13: Permanently commit to arm $a^* \in \arg \max_{a \in \mathcal{A}} \left\{\sum_{j=1}^m X_{a,j}\right\}$.
 - 14: **else**
 - 15: Play each arm in \mathcal{A} once; observe rewards $\{X_{a,m+1} : a \in \mathcal{A}\}$.
 - 16: $m \leftarrow m + 1$.
 - 17: $T \leftarrow T - K$.
-

Discussion of Algorithm 2. At any time, the algorithm computes two thresholds; $\mathcal{O}(\sqrt{m \log m})$ and $\mathcal{O}(\sqrt{m \log n})$ for the $\binom{K}{2}$ pairwise difference-of-reward statistics, m being the per-arm sample count. If said statistic is dominated by the former threshold for some arm-pair, it is likely to contain arms of the same type (equal means). The explanation stems from the Law of the Iterated Logarithm

² Minimax guarantees in previous work were polynomially bounded away from \sqrt{n} ; see [Kalvit and Zeevi \(2020\)](#).

(see [Durrett \(2019\)](#), Theorem 8.5.2): a zero-drift length- m random walk has its envelope bounded by $\mathcal{O}(\sqrt{m \log \log m})$. In the aforementioned scenario, the algorithm discards the *entire* consideration set and ushers in a new one. This is done to avoid the possibility of incurring linear regret should an optimal arm be missing from the consideration set. In the other scenario that the difference-of-reward statistic dominates the larger threshold for all arm-pairs, the consideration set is likely to contain arms of distinct types (no two have equal means) and the algorithm simply commits to the empirically best arm. Lastly, if difference-of-reward lies between the two thresholds (signifying insufficient learning), the sample count for each arm is advanced by one, and the entire process repeats.

Reason for introducing zero-mean corruptions supported on \mathbb{R} . Centered Gaussian noise is added to the difference-of-reward statistic in step (9) of $\text{ALG2}(n)$ to avoid the possibility of incurring linear regret should the support of the reward distributions be a “very small” subset of $[0, 1]$. To illustrate this point, suppose that $K = 2$, $s_n = 1$, and the rewards associated with the two types are deterministic with $\underline{\Delta} < 2\sqrt{2 \log 2}$. Then, as soon as the algorithm queries a heterogeneous consideration set (one arm optimal and the other inferior) and the per-arm sample count reaches 2, the difference-of-reward statistic will satisfy $\left| \sum_{j=1}^2 (X_{1,j} - X_{2,j}) \right| = 2\underline{\Delta} < 4\sqrt{2 \log 2}$, resulting in the consideration set getting discarded. On the other hand, if the consideration set is homogeneous (both arms simultaneously optimal or inferior), the algorithm will still re-initialize as soon as the per-arm count reaches 2.³ This will force the algorithm to keep querying new arms from the reservoir at rate that is linear in time, which is tantamount to incurring linear regret in the horizon. The addition of centered Gaussian noise hedges against this risk by guaranteeing that the difference-of-reward process essentially has an infinite support at all times even when the reward distributions might be degenerate. This rids the regret performance of its fragility w.r.t. the support of reward distributions. The next proposition crystallizes this discussion; proof is provided in [Appendix G](#).

Proposition 6 (Persistence of heterogeneous consideration sets) *Suppose $\{X_{a,j} : j = 1, 2, \dots\}$ is a collection of independent samples from an arm of type $a \in \{1, \dots, K\} =: \mathcal{A}$. Let $\{\mathcal{Z}_{a,b} : a, b \in \mathcal{A}, a < b\}$ be a collection of $\binom{K}{2}$ independently generated standard Normal random variables. Then,*

$$\mathbb{P} \left(\bigcap_{m \geq 1} \bigcap_{a, b \in \mathcal{A}, a < b} \left\{ \left| \mathcal{Z}_{a,b} + \sum_{j=1}^m (X_{a,j} - X_{b,j}) \right| \geq 4\sqrt{m \log m} \right\} \right) > \frac{\bar{\Phi}(f(T_0))}{2} =: \beta_{\delta, K} > 0, \quad (2)$$

where $\bar{\Phi}(\cdot)$ is the right tail of the standard Normal CDF, and $T_0 := \max \left(\lceil (64/\delta^2) \log^2 \left(\frac{64}{\delta^2} \right) \rceil, \Lambda_K \right)$ with $\Lambda_K := \inf \left\{ p \in \mathbb{N} : \sum_{m=p}^{\infty} \frac{1}{m^8} \leq \frac{1}{2K^2} \right\}$. Lastly, $f(x) := x + 4\sqrt{x \log x}$ for all $x \geq 1$.

Interpretation of $\beta_{\delta, K}$. First of all, note that $\beta_{\delta, K}$ admits a closed-form characterization in terms of standard functions and satisfies $\beta_{\delta, K} > 0$ for $\delta > 0$ with $\lim_{\delta \rightarrow 0} \beta_{\delta, K} = 0$. Secondly, $\beta_{\delta, K}$ depends exclusively on δ and K , and represents a lower bound on the probability that $\text{ALG2}(n)$ will never discard a consideration set containing arms of distinct types. This meta-result will be key to the upper bound on the regret of $\text{ALG2}(n)$ stated next in [Theorem 7](#).

3. $\left| \sum_{j=1}^m (X_{1,j} - X_{2,j}) \right| = 0$ identically in this case for any $m \in \mathbb{N}$ while $4\sqrt{m \log m} > 0$ only for $m \geq 2$.

Theorem 7 (Upper bound on the regret of $\text{ALG2}(n)$) For a horizon of play $n \geq K$ and per-arm burn-in phase of $1 \leq s_n \leq n/K$ samples, the expected cumulative regret of the policy π given by $\text{ALG2}(n)$ is bounded as

$$\mathbb{E}R_n^\pi \leq \frac{CK^3\bar{\Delta}}{\gamma(s_n)} \left(\frac{\log n}{\delta^2} + \frac{s_n}{K! \prod_{i=1}^K \alpha_i} \right),$$

where C is some absolute constant, and $\gamma(s_n)$ is as defined in (14). In particular, $\gamma(t)$ is monotone increasing in t with $\gamma(t) \rightarrow 1$ as $t \rightarrow \infty$, and $\gamma(1) = \beta_{\delta,K}$, where the latter is as defined in (2).

Remarks. The dependence on δ in Theorem 7 is not incidental and has the same genesis as discussed in the context of Theorem 5. However, there is a prominent distinction from Theorem 5 in that the dependence on α is captured exclusively through the constant term (as opposed to the logarithmic term). This should be viewed in light of the lower bound in Theorem 3; by allowing for policies that query the arm-reservoir *adaptively*, one can potentially make the regret performance robust w.r.t. α . Absence of α from the leading term also leads to the somewhat remarkable conclusion that the lower bound in Theorem 2 is optimal w.r.t. dependence on α . The proof is provided in Appendix H.

More on the inverse scaling w.r.t. $\gamma(s_n)$. This multiplicative factor is likely a consequence of the countable nature of arms (as opposed to finite). When $K = 2$, $\alpha_1 \leq 1/2$, and the burn-in phase s_n has a fixed duration independent of n , the upper bound in Theorem 7 reduces to $\mathcal{O}\left(\beta_{\underline{\Delta},2}^{-1}(\log n/\underline{\Delta} + \underline{\Delta}/\alpha_1)\right)$, where the big-Oh only hides absolute constants. Evidently, there is an inflation by $\beta_{\underline{\Delta},2}^{-1}$ relative to the optimal $\mathcal{O}(\log n/\underline{\Delta})$ rate achievable in the paradigmatic two-armed bandit with gap $\underline{\Delta}$. By setting s_n as a coercive sub-logarithmic function of the horizon n (e.g., $s_n = \sqrt{\log n}$), one can shave off the $\beta_{\underline{\Delta},2}^{-1}$ factor to achieve $\mathcal{O}(\log n/\underline{\Delta})$ regret. This establishes tightness of the instance-dependent lower bound in Theorem 2 when $K = 2$. On the other hand, owing to the dependence of $\gamma(s_n)$ (and $\beta_{\underline{\Delta},2}$) on $\underline{\Delta}$, the worst-case (instance-independent) upper bound of $\text{ALG2}(n)$ can be observed from Theorem 7 to be bounded away from $\Omega(\sqrt{n})$. However, recall that Theorem 5 already settles the issue of characterizing the optimal minimax rate when $K = 2$ (up to logarithmic factors in n). Thus, we provide a complete characterization of the complexity of this problem when $K = 2$, thereby answering all the open problems in Kalvit and Zeevi (2020). For $K > 2$, Theorem 7 guarantees an upper bound of $\mathcal{O}(K^3\bar{\Delta}/\delta^2 \log n)$ under a coercive sub-logarithmic burn-in phase. In this case, characterizing the optimal dependence on $\bar{\Delta}$ and δ remains an open problem. The scaling w.r.t. K , however, cannot be improved to $\mathcal{O}(K)$ as suggested by the $\Omega(K \log K)$ lower bound in Theorem 4. A full characterization of the complexity of the general setting with $K > 2$ arm-types remains challenging and is left to future work.

4.3. Towards fully sequential adaptive strategies: Optimism in exploration

In this section, we revisit the UCB-based adaptive policy proposed in Kalvit and Zeevi (2020) for $K = 2$. The policy is restated as ALG3 below after suitable modifications for reasons discussed next. The original policy (Algorithm 2 in cited paper) achieves an instance-dependent regret of $\mathcal{O}(\log n)$. Additionally, this algorithm enjoys the benefit of being anytime in n owing to the use of UCB1 (Auer et al., 2002) as a subroutine. However, it suffers a major limitation through its dependence on ex ante knowledge of the support of reward distributions. In particular, the algorithm requires the reward

distributions associated with the arm-types to have “full support” on $[0, 1]$, e.g., only distributions such as Bernoulli(\cdot), Uniform on $[0, 1]$, Beta(\cdot, \cdot), etc., are amenable to its performance guarantees; Uniform on $[0, 0.5]$, on the other hand, is not. We identify a simple fix to this issue: Drawing inspiration from the design of $\text{ALG2}(n)$, we propose adding a centered Gaussian noise term to the difference-of-reward statistic (see step (6) of Algorithm 3) to essentially create an unbounded support. This rids the algorithm of fragility w.r.t. the reward support while also preserving $\mathcal{O}(\log n)$ regret guarantees (see Theorem 8).

Remark. The original Algorithm 2 in cited paper uses a threshold that is distinct from $4\sqrt{m \log m}$ (see step (6) of Algorithm 3); the choice of $4\sqrt{m \log m}$ here aims to unify the technical presentation with $\text{ALG2}(n)$, and facilitate a more transparent comparison between the corresponding upper bounds.

Algorithm 3 ALG3 (Nested UCB1 for $K = 2$ types)

- 1: **Initialize new epoch** (resets clock $t \leftarrow 0$): Query two *new* arms; call it set $\mathcal{A} = \{1, 2\}$.
 - 2: Play each arm in \mathcal{A} once; observe rewards $\{X_{a,1} : a \in \mathcal{A}\}$.
 - 3: Minimum per-arm sample count $m \leftarrow 1$.
 - 4: Generate a standard Gaussian random variable \mathcal{Z} .
 - 5: **for** $t \in \{3, 4, \dots\}$ **do**
 - 6: **if** $\left| \mathcal{Z} + \sum_{j=1}^m (X_{1,j} - X_{2,j}) \right| < 4\sqrt{m \log m}$ **then**
 - 7: Permanently discard \mathcal{A} and repeat from step (1).
 - 8: **else**
 - 9: Play arm $a_t \in \arg \max_{a \in \mathcal{A}} \left(\frac{\sum_{j=1}^{N_a(t-1)} X_{a,j}}{N_a(t-1)} + \sqrt{\frac{2 \log(t-1)}{N_a(t-1)}} \right)$.
 - 10: Observe reward $X_{a_t, N_{a_t}(t)}$.
 - 11: **if** $m < \min_{a \in \mathcal{A}} N_a(t)$ **then**
 - 12: $m \leftarrow m + 1$.
-

Discussion of Algorithm 3. Similar to $\text{ALG2}(n)$, ALG3 also has an episodic dynamic with exactly one pair of arms played per episode. The distinction, however, resides in the fact that ALG3 plays arms according to UCB1 in every episode as opposed to playing them equally often until committing to the empirically superior one. Secondly, unlike $\text{ALG2}(n)$, ALG3 never “commits” to an arm (or a consideration set). The implication is that the algorithm will keep querying new consideration sets throughout the playing horizon; this property is at the core of its anytime nature. Despite these differences, the performance guarantees of the two algorithms are essentially identical when $K = 2$, as the next result illustrates. The proof is provided in Appendix J.

Theorem 8 (Upper bound on the regret of ALG3 when $K = 2$) *The expected cumulative regret of the policy π given by ALG3 after any number of pulls $n \geq 2$ is bounded as*

$$\mathbb{E}R_n^\pi \leq \frac{C}{\beta_{\underline{\Delta}, 2}} \left(\frac{\log n}{\underline{\Delta}} + \frac{\underline{\Delta}}{\alpha_1} \right),$$

where $\beta_{\underline{\Delta}, 2}$ is as defined in (2) with $\delta \leftarrow \underline{\Delta}$ and $K \leftarrow 2$, and C is some absolute constant.

Remark. It is possible to shave off the $\beta_{\underline{\Delta}, 2}$ factor by introducing in ALG3 a horizon-dependent burn-in phase à la $\text{ALG2}(n)$. This may be achieved at the expense of ALG3 's anytime property.

Limitation of ALG3. The performance stated in Theorem 8 together with its anytime property might appear to give an edge to ALG3 over ALG2(n). However, the former is theoretically disadvantaged in that its logarithmic upper bound is not currently amenable to extensions to the general K -typed setting. The issue traces its roots to the use of UCB1 as a subroutine. The concentration behavior of UCB1 leveraged towards the analysis of ALG3 when $K = 2$ fails to hold when $K > 2$, rendering proofs intractable. This is illustrated via a simple example with $K = 3$ types discussed below.

Technical issues with generalizing ALG3 to K types. When $K = 2$, there are only two possibilities for what a consideration set could be; arms can have means that are either (i) distinct, or (ii) equal. In the former case, an optimal arm is guaranteed to exist in the consideration set and UCB1 will spend the bulk of its sampling effort on it, which is good for regret performance. In the latter scenario, since arms have equal means, UCB1 will split samples approximately equally between the two with high probability (see Theorem 4(i) in Kalvit and Zeevi (2020)); subsequently the consideration set will be discarded within a finite number of samples in expectation (see steps (6) and (7) of ALG3). Contrast this with an alternative setting with $K = 3$ and mean rewards $\mu_1 > \mu_2 > \mu_3$. A natural generalization of ALG3 (see ALG4 in Appendix B) will query consideration sets of size 3. Thus, a query can potentially return one arm with mean μ_2 and two with mean μ_3 . Since an optimal arm (mean μ_1) is missing, the algorithm will incur linear regret on this set; it is therefore imperative to discard it at the earliest. Unfortunately though, UCB1 will invest an overwhelming majority of its sampling effort in the “locally optimal” arm (mean μ_2) and allocate logarithmically fewer samples among the other two. This logarithmic rate of sampling arms with mean μ_3 is proof-inhibiting (vis-à-vis the $K = 2$ case where the rate is linear as previously discussed), making it difficult to theoretically answer if ALG4 might still be able to discard the arms within, say, logarithmically many pulls of the horizon. This is an open research question and at the moment, an $\mathcal{O}(\log n)$ bound exists only for $K = 2$; we could only establish asymptotic-optimality ($o(n)$ regret) when $K > 2$ (see Theorem 9). Among other things, identifying the optimal (instance-dependent) scaling factors w.r.t. (μ, α) and the optimal order of minimax regret when $K > 2$ remain open problems.

5. Concluding remarks and open problems

This paper provides a first-order characterization of the complexity of the K -typed countable-armed bandit problem with matching lower and upper bounds for $K = 2$. For $K > 2$, we establish an instance-dependent upper bound of $\mathcal{O}(K^3 \bar{\Delta} / \delta^2 \log n)$ and show that the scaling w.r.t. K cannot beat $\Omega(K \log K)$; the latter property differentiates this setting fundamentally from the classical K -armed problem. Another key takeaway from our work is that achievable regret in this setting only has a *second-order* dependence on the reservoir distribution, i.e., dependence on α only manifests through sub-logarithmic terms (see Theorem 7 and 8). Although this work is predicated on countably many arms, our algorithms can easily be adapted to settings with a large but finite number of arms. For example, the result on second-order dependence w.r.t. α has profound implications for the N -armed bandit problem with K arm-types, where each type is characterized by a unique mean reward. A naive implementation of standard MAB algorithms in this setting will result in a regret that scales linearly with N . Instead, one can simulate a K -typed reservoir over the N arms and deploy ALG2(n) to achieve an $\mathcal{O}(K^3)$ scaling of the leading term; if $K \ll N$, performance improvement can be substantial vis-à-vis naive MAB algorithms. Another important direction concerns adaptivity to K : This paper provides algorithms that adapt to α assuming perfect knowledge of K ; performance characterization given only an approximation thereof remains an open problem.

Acknowledgments

We thank the reviewers for their feedback and suggestions.

References

- Shipra Agrawal, Vashist Avadhanula, Vineet Goyal, and Assaf Zeevi. Mnl-bandit: A dynamic learning approach to assortment selection. *Operations Research*, 67(5):1453–1485, 2019.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- Siddhartha Banerjee, Sreenivas Gollapudi, Kostas Kollias, and Kamesh Munagala. Segmenting two-sided markets. In *Proceedings of the 26th International Conference on World Wide Web*, pages 63–72, 2017.
- Donald A Berry, Robert W Chen, Alan Zame, David C Heath, Larry A Shepp, et al. Bandit problems with infinitely many arms. *The Annals of Statistics*, 25(5):2103–2116, 1997.
- Thomas Bonald and Alexandre Proutiere. Two-target algorithms for infinite-armed bandits with bernoulli rewards. In *Advances in Neural Information Processing Systems*, pages 2184–2192, 2013.
- Loc Bui, Ramesh Johari, and Shie Mannor. Clustered bandits. *arXiv preprint arXiv:1206.4169*, 2012.
- Alexandra Carpentier and Michal Valko. Simple regret for infinitely many armed bandits. In *International Conference on Machine Learning*, pages 1133–1141, 2015.
- Hock Peng Chan and Shouri Hu. Infinite arms bandit: Optimality via confidence bounds. *arXiv preprint arXiv:1805.11793*, 2018.
- Karthekeyan Chandrasekaran and Richard Karp. Finding a most biased coin with fewest flips. In *Conference on Learning Theory*, pages 394–407. PMLR, 2014.
- Arghya Roy Chaudhuri and Shivaram Kalyanakrishnan. Quantile-regret minimisation in infinitely many-armed bandits. In *UAI*, pages 425–434, 2018.
- Rianne de Heide, James Cheshire, Pierre Ménard, and Alexandra Carpentier. Bandits with many optimal arms. In *Advances in Neural Information Processing Systems*, volume 34, pages 22457–22469, 2021.
- Rick Durrett. *Probability: theory and examples*, volume 49. Cambridge university press, 2019.
- Philippe Flajolet, Daniele Gardy, and Loÿs Thimonier. Birthday paradox, coupon collectors, caching algorithms and self-organizing search. *Discrete Applied Mathematics*, 39(3):207–229, 1992.
- W Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

- Kevin Jamieson, Daniel Haas, and Ben Recht. On the detection of mixture distributions with applications to the most biased coin problem. *arXiv preprint arXiv:1603.08037*, 2016.
- Ramesh Johari, Vijay Kamble, and Yash Kanoria. Matching while learning. *Operations Research*, 69(2):655–681, 2021.
- Anand Kalvit and Assaf Zeevi. From finite to countable-armed bandits. In *Advances in Neural Information Processing Systems*, volume 33, pages 8259–8269, 2020.
- Anand Kalvit and Assaf Zeevi. A closer look at the worst-case behavior of multi-armed bandit algorithms. In *Advances in Neural Information Processing Systems*, volume 34, pages 8807–8819, 2021.
- Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Aleksandrs Slivkins et al. Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning*, 12(1-2):1–286, 2019.
- Yizao Wang, Jean-Yves Audibert, and Rémi Munos. Algorithms for infinitely many-armed bandits. In *Advances in Neural Information Processing Systems*, pages 1729–1736, 2009.
- Yinglun Zhu and Robert Nowak. On regret with multiple best arms. *Advances in Neural Information Processing Systems*, 33:9050–9060, 2020.

Supplementary material: General organization

1. Appendix [A](#) provides numerical experiments.
2. Appendix [B](#) generalizes `ALG3` to $K > 2$ and states performance guarantees thereof.
3. Appendix [C](#) provides the proof of [Theorem 2](#).
4. Appendix [D](#) provides the proof of [Theorem 3](#).
5. Appendix [E](#) provides the proof of [Theorem 4](#).
6. Appendix [F](#) provides the proof of [Theorem 5](#).
7. Appendix [G](#) provides the proof of [Proposition 6](#).
8. Appendix [H](#) provides the proof of [Theorem 7](#).
9. Appendix [I](#) provides auxiliary results used in the analysis of `ALG3`.
10. Appendix [J](#) provides the proof of [Theorem 8](#).
11. Appendix [K](#) provides auxiliary results used in the analysis of `ALG4`.
12. Appendix [L](#) provides the proof of [Theorem 9](#).

Appendix A. Numerical experiments

We evaluate the empirical performance of our algorithms for $K = 2$ and $K = 3$ on synthetic data.

Experiments. In what follows, the graphs show the performance of different algorithms simulated on synthetic data. The horizon is capped at $n = 10^5$ for $K = 2$ and at $n = 10^4$ for $K = 3$. Each regret trajectory is averaged over at least 100 independent experiments (sample-paths). The shaded regions indicate standard 95% confidence intervals. For horizon-dependent algorithms, regret is plotted for discrete values of the horizon n indicated by “*” and interpolated; for anytime algorithms, regret accrued until each $t \in \{1, \dots, n\}$ is plotted.

Baseline policies. We will benchmark the performance of our algorithms against two policies: (i) `Sampling-UCB` (de Heide et al., 2021), and (ii) `ETC- ∞ (2)` (Kalvit and Zeevi, 2020). The former is a UCB-styled policy based on front-loading exploration of *new* arms (Theorem 3 thus applies to this policy). It is, however, noteworthy that `Sampling-UCB` is predicated on ex ante knowledge of (a lower bound on) the probability α_1 of sampling an optimal arm from the reservoir; we reemphasize that this is not the setting of interest in our paper. Furthermore, its regret scales as $\tilde{O}(\log n / (\alpha_1 \underline{\Delta}))$ (up to poly-logarithmic factors in $1/\underline{\Delta}$), which is inferior in terms of its dependence on α_1 relative to `ALG2(n)` and `ALG3` (see Theorem 7 and 8 respectively). There exist other algorithms as well (see, e.g., Chaudhuri and Kalyanakrishnan (2018); Zhu and Nowak (2020)) developed for formulations with prohibitively large number of arms. However, these are either sensitive to certain parametric assumptions on the probability of sampling an optimal arm, or focus on a different notion of regret altogether; both directions remain outside the ambit of our setting.

The second policy `ETC- ∞ (2)` is a non-adaptive explore-then-commit-styled algorithm for reservoirs with $K = 2$ types; this policy requires ex ante knowledge of a lower bound on the difference between the two mean rewards. Although `ETC- ∞ (2)` was originally proposed only for $K = 2$, it is easily generalizable and we present in Algorithm 5 a version (`ETC- ∞ (K)`) that is adapted to K types.

Setup 1 [Figures 1, 2 and 3]. In this setting, we consider $K = 2$ with $\alpha_1 = 0.5$, i.e., two equiprobable arm-types, characterized by Bernoulli(0.6) and Bernoulli(0.4) rewards. Via this setup, we intend to illustrate the difference between the empirical performance achievable in the countable-armed setting vis-à-vis its traditional two-armed counterpart. Refer to Figure 1. The red curve indicates the empirical performance of `ALG3` in this setting. For reference, the blue one shows the empirical performance of `UCB1` (Auer et al., 2002) in a two-armed bandit with Bernoulli(0.6) and Bernoulli(0.4) rewards; the green curve indicates the best achievable instance-dependent regret (Lai and Robbins, 1985) in said two-armed configuration. As expected, the regret of `ALG3` is inflated relative to `UCB1`. This is owing to the $\beta_{\delta,2} \leq 1$ factor present in the denominator of `ALG3`’s upper bound; characterization of the sharpest lower bound on the probability in (2) (see Proposition 6) is challenging owing to the limited theoretical tools available to this end and we leave it as an open problem at the moment. Figure 2 shows the empirical performance of the algorithms proposed in this paper as well as `Sampling-UCB` initialized with $\alpha_1 = 1/2$ and `ETC- ∞ (2)` initialized with $\underline{\delta} = \delta/2 = 0.1$. Evidently, the (adaptive) explore-then-commit approach in `ALG2(n)` outperforms the pre-specified exploration schedule-based approach of `ALG1(n)`, and performs almost as good as the *gap-aware* approach in `ETC- ∞ (2)`. While `Sampling-UCB` outmatches all explore-then-

commit styled approaches, the best performing algorithm is ALG3. Surprisingly, this is despite the fact that the theoretical performance bounds for $ALG2(n)$ and $ALG3$ are identical (modulo numerical multiplicative constants) when $K = 2$ and $\alpha_1 \leq 0.5$ (see Theorem 7 and 8). A similar hierarchy in performances is also observable in Figure 3, which corresponds to a slightly “easier” instance with $\delta = 0.4$ (as opposed to 0.2) and equiprobable Bernoulli(0.9) and Bernoulli(0.5) rewards.




Figure 1: $K = 2$ and $\alpha = (1/2, 1/2)$: Achievable regret in 2-CAB vis-à-vis 2-MAB.




Figure 2: $K = 2$ and $\alpha = (1/2, 1/2)$: An instance with Bernoulli 0.6, 0.4 rewards.

Setup 2 [Figure 4]. Here, we consider a setting with $K = 3$ arm-types characterized by Bernoulli rewards with means 0.9, 0.5, 0.1, each occurring with probability $1/3$. We compare the performance of $ALG1(n)$, $ALG2(n)$ and $ALG4$ (generalization of $ALG3$ to $K \geq 2$; see Appendix B) with $ETC-\infty(3)$ initialized with $\underline{\delta} = \delta/2 = 0.2$, and $Sampling-UCB$ initialized with $\alpha_1 = 1/3$. It is noteworthy that despite $ALG4$'s significantly superior empirical performance relative to aforementioned algorithms, only a weak $o(n)$ theoretical guarantee on its regret is currently available (see Theorem 9) due to reasons discussed earlier in the paper. Investigating best achievable rates under $ALG4$ is an area of active research at the moment.




Figure 3: $K = 2$ and $\alpha = (1/2, 1/2)$: An instance with Bernoulli 0.9, 0.5 rewards.




Figure 4: $K = 3$ and $\alpha = (1/3, 1/3, 1/3)$: Instance with Bernoulli 0.9, 0.5, 0.1 rewards.

Appendix B. More algorithms for reservoirs with $K \geq 2$ arm-types

Algorithm 4 ALG4 (Nested UCB1 for K types)

- 1: **Initialize new epoch** (resets clock $t \leftarrow 0$): Query K new arms; call it set $\mathcal{A} = \{1, \dots, K\}$.
 - 2: Play each arm in \mathcal{A} once; observe rewards $\{X_{a,1} : a \in \mathcal{A}\}$.
 - 3: Minimum per-arm sample count $m \leftarrow 1$.
 - 4: Generate $\binom{K}{2}$ independent standard Gaussian random variables $\{\mathcal{Z}_{a,b} : a, b \in \mathcal{A}, a < b\}$.
 - 5: **for** $t \in \{K+1, K+2, \dots\}$ **do**
 - 6: **if** $\exists a, b \in \mathcal{A}, a < b$ s.t. $|\mathcal{Z}_{a,b} + \sum_{j=1}^m (X_{a,j} - X_{b,j})| < 4\sqrt{m \log m}$ **then**
 - 7: Permanently discard \mathcal{A} and repeat from step (1).
 - 8: **else**
 - 9: Play arm $a_t \in \arg \max_{a \in \mathcal{A}} \left(\frac{\sum_{j=1}^{N_a(t-1)} X_{a,j}}{N_a(t-1)} + \sqrt{\frac{2 \log(t-1)}{N_a(t-1)}} \right)$.
 - 10: Observe reward $X_{a_t, N_{a_t}(t)}$.
 - 11: **if** $m < \min_{a \in \mathcal{A}} N_a(t)$ **then**
 - 12: $m \leftarrow m + 1$.
-

Theorem 9 (Upper bound on the regret of ALG4) *The expected cumulative regret of the policy π given by ALG4 after any number $n \geq 1$ of plays is bounded as*

$$\mathbb{E}R_n^\pi \leq \frac{CK}{\beta_{\delta,K}} \left(\frac{\log n}{\Delta} + \bar{\Delta} \right) + o \left(\frac{\bar{\Delta} n}{\beta_{\delta,K} \prod_{i=1}^K \alpha_i} \right),$$

where C is some absolute constant, $\beta_{\delta,K}$ is as defined in (2), and the little-Oh is asymptotic in n and only hides multiplicative factors in K .

Proof is provided in Appendix L.

Algorithm 5 ETC- ∞ (K)

- 1: **Input:** (i) Horizon of play n , (ii) A lower bound $\underline{\delta} \in (0, \delta]$ on the minimal reward gap δ .
 - 2: Set budget $T = n$.
 - 3: **Initialize new epoch:** Query K new arms; call it consideration set $\mathcal{A} = \{1, \dots, K\}$.
 - 4: Set exploration duration $L = \lceil 2\underline{\delta}^{-2} \log n \rceil$.
 - 5: $m \leftarrow \min(L, \lfloor T/K \rfloor)$.
 - 6: Play each arm in \mathcal{A} m times; observe rewards $\{(X_{1,j}, \dots, X_{K,j}) : j = 1, \dots, m\}$.
 - 7: Update budget: $T \leftarrow T - Km$.
 - 8: **if** $\exists a, b \in \mathcal{A}, a < b$ s.t. $|\sum_{j=1}^m (X_{a,j} - X_{b,j})| < \underline{\delta}m$ **then**
 - 9: Permanently discard \mathcal{A} , and repeat from step (3).
 - 10: **else**
 - 11: Permanently commit to arm $a^* \in \arg \max_{a \in \mathcal{A}} \left\{ \sum_{j=1}^m X_{a,j} \right\}$.
-

Appendix C. Proof of Theorem 2

Notation. For each $i \in \{1, 2\}$, let $\mathcal{G}_i(x)$ be an arbitrary collection of distributions with mean $x \in \mathbb{R}$. The tuple $(\mathcal{G}_1(x), \mathcal{G}_2(y))$ will be referred to as an *instance*.

Since the horizon of play is fixed at n , the decision maker may play at most n distinct arms. Therefore, it suffices to focus only on the sequence of the first n arms that may be played. A *realization* of an instance $\nu = (\mathcal{G}_1(\mu_1), \mathcal{G}_2(\mu_2))$ is defined as the n -tuple $r \equiv (r_i)_{1 \leq i \leq n}$, where $r_i \in \mathcal{G}_1(\mu_1) \cup \mathcal{G}_2(\mu_2)$ denotes the reward distribution of arm $i \in \{1, \dots, n\}$. It must be noted that the decision maker need not play every arm in r . Let $i^* := \arg \max_{i \in \{1, 2\}} \mu_i$. Suppose that the distribution over possible realizations of $\nu = (\mathcal{G}_1(\mu_1), \mathcal{G}_2(\mu_2))$ in $\{r : r_i \in \mathcal{G}_1(\mu_1) \cup \mathcal{G}_2(\mu_2), 1 \leq i \leq n\}$ satisfies $\mathbb{P}(r_i \in \mathcal{G}_{i^*}(\mu_{i^*})) = \alpha^*$ (where $\alpha^* \in (0, 1)$ is arbitrary) for all $i \in \{1, \dots, n\}$, i.e., optimal arms occur in the reservoir with probability α^* .

Recall that the cumulative pseudo-regret after n plays of a policy π on $\nu = (\mathcal{G}_1(\mu_1), \mathcal{G}_2(\mu_2))$ is given by $R_n^\pi(\nu) = \sum_{t=1}^n (\mu_{i^*} - \mu_{\mathcal{T}(\pi_t)})$, where $\mathcal{T}(\pi_t) \in \{1, 2\}$ indicates the type of the arm played by π at time t . Our goal is to lower bound $\mathbb{E}R_n^\pi(\nu)$, where the expectation is w.r.t. the randomness in π as well as the distribution over possible realizations of ν . To this end, we define the notion of expected cumulative regret of π on a realization r of $\nu = (\mathcal{G}_1(\mu_1), \mathcal{G}_2(\mu_2))$ by

$$S_n^\pi(\nu, r) := \mathbb{E}^\pi \left[\sum_{t=1}^n (\mu_{i^*} - \mu_{\mathcal{T}(\pi_t)}) \right],$$

where the expectation \mathbb{E}^π is w.r.t. the randomness in π . Note that $\mathbb{E}R_n^\pi(\nu) = \mathbb{E}^\nu S_n^\pi(\nu, r)$, where the expectation \mathbb{E}^ν is w.r.t. the distribution over possible realizations of ν . We define our problem class \mathcal{N}_Δ as the collection of Δ -separated instances given by

$$\mathcal{N}_\Delta := \{(\mathcal{G}_1(\mu_1), \mathcal{G}_2(\mu_2)) : \mu_1 - \mu_2 = \Delta, (\mu_1, \mu_2) \in \mathbb{R}^2\}.$$

Fix an arbitrary $\Delta > 0$ and consider an instance $\nu = (\{Q_1\}, \{Q_2\}) \in \mathcal{N}_\Delta$, where (Q_1, Q_2) are unit-variance Gaussian distributions with means (μ_1, μ_2) respectively. Consider an arbitrary realization $r \in \{Q_1, Q_2\}^n$ of ν and let $\mathcal{I} \subseteq \{1, \dots, n\}$ denote the set of inferior arms in r (arms with reward distribution Q_2). Consider another instance $\nu' \in \mathcal{N}_\Delta$ given by $\nu' = (\{\tilde{Q}_1\}, \{Q_1\})$, where \tilde{Q}_1 is another unit variance Gaussian with mean $\mu_1 + \Delta$. Now consider a realization $r' \in \{\tilde{Q}_1, Q_1\}^n$ of ν' that is such that the arms at positions in \mathcal{I} have distribution \tilde{Q}_1 while those at positions in $\{1, \dots, n\} \setminus \mathcal{I}$ have distribution Q_1 . Notice that \mathcal{I} is the set of optimal arms in r' (arms with reward distribution \tilde{Q}_1). Then, the following always holds:

$$S_n^\pi(\nu, r) + S_n^\pi(\nu', r') \geq \left(\frac{\Delta n}{2}\right) \left(\mathbb{P}_{\nu, r}^\pi \left(\sum_{i \in \mathcal{I}} N_i(n) > \frac{n}{2} \right) + \mathbb{P}_{\nu', r'}^\pi \left(\sum_{i \in \mathcal{I}} N_i(n) \leq \frac{n}{2} \right) \right),$$

where $\mathbb{P}_{\nu, r}^\pi(\cdot)$ and $\mathbb{P}_{\nu', r'}^\pi(\cdot)$ denote the probability measures w.r.t. the instance-realization pairs (ν, r) and (ν', r') respectively, and $N_i(n)$ denotes the number of plays up to and including time n of arm $i \in \{1, \dots, n\}$. Using the Bretagnolle-Huber inequality (Theorem 14.2 of [Lattimore and Szepesvári \(2020\)](#)), we obtain

$$S_n^\pi(\nu, r) + S_n^\pi(\nu', r') \geq \left(\frac{\Delta n}{4}\right) \exp(-D(\mathbb{P}_{\nu, r}^\pi, \mathbb{P}_{\nu', r'}^\pi)),$$

where $D\left(\mathbb{P}_{\nu,r}^\pi, \mathbb{P}_{\nu',r'}^\pi\right)$ denotes the KL-Divergence between $\mathbb{P}_{\nu,r}^\pi$ and $\mathbb{P}_{\nu',r'}^\pi$. Using Divergence decomposition (Lemma 15.1 of [Lattimore and Szepesvári \(2020\)](#)), we further obtain

$$S_n^\pi(\nu, r) + S_n^\pi(\nu', r') \geq \left(\frac{\underline{\Delta}n}{4}\right) \exp\left(-\left(\frac{D(Q_2, \tilde{Q}_1)}{\underline{\Delta}}\right) S_n^\pi(\nu, r)\right) = \left(\frac{\underline{\Delta}n}{4}\right) \exp(-2\underline{\Delta}S_n^\pi(\nu, r)),$$

where the equality follows since \tilde{Q}_1 and Q_2 are unit variance Gaussian distributions with means separated by $2\underline{\Delta}$. Next, taking the expectation \mathbb{E}^ν on both sides followed by a direct application of Jensen's inequality yields

$$\mathbb{E}R_n^\pi(\nu) + \mathbb{E}^\nu S_n^\pi(\nu', r') \geq \left(\frac{\underline{\Delta}n}{4}\right) \exp(-2\underline{\Delta}\mathbb{E}R_n^\pi(\nu)). \quad (3)$$

Consider the $\mathbb{E}^\nu S_n^\pi(\nu', r')$ term in (3). Using a simple change-of-measure argument, we obtain

$$\mathbb{E}^\nu S_n^\pi(\nu', r') = \mathbb{E}^{\nu'} \left[S_n^\pi(\nu', r') \left(\frac{1 - \alpha^*}{\alpha^*}\right)^{2(\Lambda(r') - n/2)} \right],$$

where $\Lambda(r')$ is the number of optimal arms in realization r' .

Since α^* is arbitrary, we fix $\alpha^* = 1/2$ to obtain

$$\mathbb{E}^\nu S_n^\pi(\nu', r') = \mathbb{E}^{\nu'} S_n^\pi(\nu', r') = \mathbb{E}R_n^\pi(\nu'), \quad (4)$$

Now, from (3) and (4), we have that for $\alpha^* = 1/2$,

$$\begin{aligned} \mathbb{E}R_n^\pi(\nu) + \mathbb{E}R_n^\pi(\nu') &\geq \frac{\underline{\Delta}n}{4} \exp(-2\underline{\Delta}\mathbb{E}R_n^\pi(\nu)) \\ \implies \tilde{R}_n &\geq \frac{\underline{\Delta}n}{8} \exp(-2\underline{\Delta}\tilde{R}_n), \end{aligned} \quad (5)$$

where $\tilde{R}_n := \max(\mathbb{E}R_n^\pi(\nu), \mathbb{E}R_n^\pi(\nu'))$.

Instance-dependent lower bound

The assertion of the theorem follows from the fact that the inequality (5) is fulfilled only if for any $\varepsilon \in (0, 1)$, \tilde{R}_n satisfies for all n large enough $\tilde{R}_n \geq (1 - \varepsilon) \log n / (2\underline{\Delta})$. Therefore, there exists an instance ν with gap $\underline{\Delta}$ such that $\mathbb{E}R_n^\pi(\nu) \geq C \log n / \underline{\Delta}$ for some absolute constant C and n large enough, whenever $\alpha^* = 1/2$. In fact, said statement holds for all $\alpha^* \leq 1/2$ since the policy π satisfies Definition 1.

Instance-independent (minimax) lower bound

Since $\tilde{R}_n \leq \underline{\Delta}n$, it follows from (5) that

$$\tilde{R}_n \geq \frac{\underline{\Delta}n}{8} \exp(-2\underline{\Delta}^2n).$$

Setting $\underline{\Delta} = 1/\sqrt{n}$, and noting that the inequality, in fact, holds for all $\alpha^* \leq 1/2$ (owing to the admissibility of π ; see Definition 1), proves the stated assertion. \square

Appendix D. Proof of Theorem 3

Note that this result is stated for general $K \geq 2$ and is not specific to $K = 2$. In fact, the nature of the set of possible sub-optimal types is inconsequential to the proof that follows as long as said set is at least $\underline{\Delta}$ -separated from the optimal mean reward. Consider an arbitrary policy $\pi \in \tilde{\Pi}$. Denote by A_n^π the number of *distinct* arms played by π until time n . Consider an arbitrary $k \in \{1, \dots, n\}$. Then, conditioned on $A_n^\pi = k$, the expected cumulative regret incurred by π is at least

$$\mathbb{E}[R_n^\pi | A_n^\pi = k] \geq (1 - \alpha_1)\underline{\Delta}k + (1 - \alpha_1)^k \underline{\Delta}(n - k) =: f(k). \quad (6)$$

Intuition behind (6). Each of the k arms played during the horizon has at least one pull associated with it. Consider a clairvoyant policy coupled to π that learns the best among the A_n^π arms played by π as soon as each has been pulled exactly once, i.e., after a total of A_n^π pulls. Clearly, the regret incurred by said clairvoyant policy lower bounds $\mathbb{E}R_n^\pi$. Further, since A_n^π is independent of the sample-history of arms, it follows that the A_n^π arms are statistically identical. Thus, conditioned on $A_n^\pi = k$, the expected regret from the first k pulls of the clairvoyant policy is at least $(1 - \alpha_1)\underline{\Delta}k$. Also, the probability that each of the k arms is inferior-typed is $(1 - \alpha_1)^k$; the clairvoyant policy thus incurs a regret of at least $(1 - \alpha_1)^k \underline{\Delta}(n - k)$ going forward. This explains the lower bound in (6). Therefore, for any $k \in \{1, 2, \dots, n\}$, we have

$$\mathbb{E}[R_n^\pi | A_n^\pi = k] \geq \min_{k \in \{1, 2, \dots, n\}} f(k) \geq \min_{x \in [0, n]} f(x).$$

We will show that $f(x)$ is strictly convex over $[0, n]$ with $f'(0) < 0$ and $f'(n) > 0$. Then, it would follow that $f(\cdot)$ admits a unique minimizer $x_n^* \in (0, n)$ given by the solution to $f'(x) = 0$. The minimum $f(x_n^*)$ will turn out to be logarithmic in n . Observe that

$$\begin{aligned} f'(x) &= (1 - \alpha_1)\underline{\Delta} + (1 - \alpha_1)^x \underline{\Delta} [(n - x) \log(1 - \alpha_1) - 1], \\ f''(x) &= -(1 - \alpha_1)^x \underline{\Delta} [2 - (n - x) \log(1 - \alpha_1)] \log(1 - \alpha_1). \end{aligned}$$

Since $\underline{\Delta} > 0$, it follows that $f''(x) > 0$ over $[0, n]$. Further, note that

$$\begin{aligned} f'(0) &= -\alpha_1 \underline{\Delta} + \underline{\Delta} n \log(1 - \alpha_1) < 0, \\ f'(n) &= (1 - \alpha_1)\underline{\Delta} - (1 - \alpha_1)^n \underline{\Delta} > 0. \end{aligned}$$

Solving $f'(x_n^*) = 0$ for the unique minimizer x_n^* , we obtain

$$\begin{aligned} \left(\frac{1}{1 - \alpha_1}\right)^{x_n^* - 1} - 1 &= (n - x_n^*) \log\left(\frac{1}{1 - \alpha_1}\right) \\ \implies \left(\frac{1}{1 - \alpha_1}\right)^{x_n^*} + x_n^* \log\left(\frac{1}{1 - \alpha_1}\right) &> n \log\left(\frac{1}{1 - \alpha_1}\right) \\ \implies 2 \left(\frac{1}{1 - \alpha_1}\right)^{x_n^*} &> n \log\left(\frac{1}{1 - \alpha_1}\right), \end{aligned}$$

where the last inequality follows using $y > \log y$. Therefore, we have

$$\begin{aligned} \left(\frac{1}{1-\alpha_1}\right)^{x_n^*} &> \frac{n}{2} \log\left(\frac{1}{1-\alpha_1}\right) \\ \implies x_n^* &> \frac{\log n + \log \log\left(\frac{1}{1-\alpha_1}\right) - \log 2}{\log\left(\frac{1}{1-\alpha_1}\right)}. \end{aligned}$$

Thus, for any $k \in \{1, \dots, n\}$,

$$\begin{aligned} \mathbb{E}[R_n^\pi | A_n^\pi = k] &\geq f(x_n^*) > (1-\alpha_1)\underline{\Delta}x_n^* > (1-\alpha_1) \left(\frac{\log n + \log \log\left(\frac{1}{1-\alpha_1}\right) - \log 2}{\log\left(\frac{1}{1-\alpha_1}\right)} \right) \underline{\Delta} \\ \implies \mathbb{E}R_n^\pi &\geq (1-\alpha_1) \left(\frac{\log n + \log \log\left(\frac{1}{1-\alpha_1}\right) - \log 2}{\log\left(\frac{1}{1-\alpha_1}\right)} \right) \underline{\Delta} \\ \implies \inf_{\pi \in \tilde{\Pi}} \frac{\mathbb{E}R_n^\pi}{\log n} &\geq (1-\alpha_1) \left(\frac{1}{\log\left(\frac{1}{1-\alpha_1}\right)} + \frac{\log \log\left(\frac{1}{1-\alpha_1}\right) - \log 2}{(\log n) \log\left(\frac{1}{1-\alpha_1}\right)} \right) \underline{\Delta} \\ \implies \inf_{\pi \in \tilde{\Pi}} \frac{\mathbb{E}R_n^\pi}{\log n} &\stackrel{(\dagger)}{\geq} (1-\alpha_1) \left(\frac{1-\alpha_1}{\alpha_1} + \frac{\log \log\left(\frac{1}{1-\alpha_1}\right) - \log 2}{(\log n) \log\left(\frac{1}{1-\alpha_1}\right)} \right) \underline{\Delta} \\ \implies \inf_{\pi \in \tilde{\Pi}} \frac{\mathbb{E}R_n^\pi}{\log n} &\geq \frac{(1-\alpha_1)^2 \underline{\Delta}}{\alpha_1} + (1-\alpha_1) \left(\frac{\log \log\left(\frac{1}{1-\alpha_1}\right) - \log 2}{(\log n) \log\left(\frac{1}{1-\alpha_1}\right)} \right) \underline{\Delta}, \end{aligned}$$

where (\dagger) follows using $\log y \leq y - 1$. Taking the appropriate limit now proves the assertion. \square

Appendix E. Proof of Theorem 4

The reservoir distribution is given by $\alpha = (\alpha_1, \dots, \alpha_K)$. In the full information setting, the decision maker observes the true mean reward of an arm immediately upon pulling it. Let $\pi = (\pi_t : t = 1, 2, \dots)$ be the policy that pulls a new arm from the reservoir in each period. Let N denote the first time at which one arm of each of the K types is collected under π . Then, it follows from classical results (see Theorem 4.1 in Flajolet et al. (1992)) for the *Coupon-collector problem* that

$$\begin{aligned} \mathbb{E}N &= \int_0^\infty \left(1 - \prod_{j=1}^K (1 - \exp(-\alpha_j y)) \right) dy \stackrel{(\dagger)}{\geq} \int_0^\infty \left(1 - \prod_{j=1}^K \left(1 - \exp\left(-\frac{y}{K}\right) \right) \right) dy \\ &= \sum_{j=1}^K \frac{K}{j} \stackrel{(\ddagger)}{\geq} K \log K, \end{aligned} \tag{7}$$

where (†) follows as $\prod_{j=1}^K (1 - \exp(-\alpha_j y))$ is maximized when α is the Uniform distribution; (‡) is a classical result (see previous reference). The optimal policy π^* follows π until time N , and subsequently commits to the arm with the highest mean among the first N arms. The lifetime regret of π^* is then given by

$$\begin{aligned} \mathbb{E}R_\infty^{\pi^*} &= \mathbb{E} \left[\sum_{t=1}^N \sum_{i=2}^K (\mu_1 - \mu_i) \mathbb{1} \{ \mathcal{T}(\pi_t) = i \} \right] \\ &= \mathbb{E} \left[\sum_{t=1}^{\infty} \sum_{i=2}^K (\mu_1 - \mu_i) \mathbb{1} \{ \mathcal{T}(\pi_t) = i, t \leq N \} \right] \\ &= \sum_{t=1}^{\infty} \sum_{i=2}^K (\mu_1 - \mu_i) \mathbb{P}(\mathcal{T}(\pi_t) = i, t \leq N), \end{aligned} \quad (8)$$

where the last equality follows from Tonelli's Theorem. Note that

$$\begin{aligned} \mathbb{P}(\mathcal{T}(\pi_t) = i, t \leq N) &= \mathbb{P}(\mathcal{T}(\pi_t) = i) - \mathbb{P}(\mathcal{T}(\pi_t) = i, t > N) \\ &= \alpha_i - \mathbb{P}(\mathcal{T}(\pi_t) = i \mid t > N) \mathbb{P}(t > N) \\ &= \alpha_i - \mathbb{P}(\mathcal{T}(\pi_t) = i) \mathbb{P}(t > N) \\ &\stackrel{(\star)}{=} \alpha_i - \alpha_i \mathbb{P}(t > N) \\ &= \alpha_i \mathbb{P}(N \geq t), \end{aligned} \quad (9)$$

where (★) follows since $\mathcal{T}(\pi_t)$ is i.i.d. in time t . Therefore, from (8) and (9), we have

$$\begin{aligned} \mathbb{E}R_\infty^{\pi^*} &= \sum_{t=1}^{\infty} \sum_{i=2}^K \alpha_i (\mu_1 - \mu_i) \mathbb{P}(N \geq t) \\ &= \sum_{i=2}^K \alpha_i (\mu_1 - \mu_i) \sum_{t=1}^{\infty} \mathbb{P}(N \geq t) \\ &= \sum_{i=2}^K \alpha_i (\mu_1 - \mu_i) \mathbb{E}N. \end{aligned} \quad (10)$$

Finally, from (10) and (7), one obtains

$$\mathbb{E}R_\infty^{\pi^*} \geq \sum_{i=2}^K \alpha_i (\mu_1 - \mu_i) K \log K.$$

□

Appendix F. Proof of Theorem 5

Let $\left\{ \left(X_{1,j}^l, \dots, X_{K,j}^l \right) : j = 1, \dots, m_l \right\}$ be the reward sequence associated with the K arms played in the l^{th} epoch, where $m_l = \lceil e^{2\sqrt{l}} \log n \rceil$. Let $\mathcal{A} := \{1, \dots, K\}$ and define

$$I := \inf \left\{ l \in \mathbb{N} : \left| \sum_{j=1}^{m_l} \left(X_{a,j}^l - X_{b,j}^l \right) \right| \geq 2m_l e^{-\sqrt{l}} \forall a, b \in \mathcal{A}, a < b \right\}.$$

Then,

$$\begin{aligned}
 \mathbb{P}(I \geq k) &= \mathbb{P}\left(\bigcap_{l=1}^{k-1} \bigcup_{a,b \in \mathcal{A}, a < b} \left\{ \left| \sum_{j=1}^{m_l} (X_{a,j}^l - X_{b,j}^l) \right| < 2m_l e^{-\sqrt{l}} \right\}\right) \\
 &= \prod_{l=1}^{k-1} \mathbb{P}\left(\bigcup_{a,b \in \mathcal{A}, a < b} \left\{ \left| \sum_{j=1}^{m_l} (X_{a,j}^l - X_{b,j}^l) \right| < 2m_l e^{-\sqrt{l}} \right\}\right) \\
 &\leq \prod_{l=1}^{k-1} \left[1 - \mathbb{P}(\mathbb{D}) + \mathbb{P}(\mathbb{D}) \mathbb{P}_{\mathbb{D}}\left(\bigcup_{a,b \in \mathcal{A}, a < b} \left\{ \left| \sum_{j=1}^{m_l} (X_{a,j}^l - X_{b,j}^l) \right| < 2m_l e^{-\sqrt{l}} \right\}\right)\right],
 \end{aligned}$$

where \mathbb{D} denotes the event that the K arms played in epoch l are ‘‘all-distinct,’’ i.e., no two arms belong to the same type, $\mathbb{P}_{\mathbb{D}}(\cdot) := \mathbb{P}(\cdot | \mathbb{D})$ denotes the corresponding conditional measure, and $\mathbb{P}(\mathbb{D}) = K! \prod_{i=1}^K \alpha_i$. Using the Union bound, we obtain

$$\mathbb{P}(I \geq k) \leq \prod_{l=1}^{k-1} \left[1 - \mathbb{P}(\mathbb{D}) + \mathbb{P}(\mathbb{D}) \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}_{\mathbb{D}}\left(\left| \sum_{j=1}^{m_l} (X_{a,j}^l - X_{b,j}^l) \right| < 2m_l e^{-\sqrt{l}}\right) \right]. \quad (11)$$

Define $\tau := K \sum_{l=1}^I m_l$. Consider the following events:

$\mathbf{E}_1 := \{\text{None of the arms played in epoch } I \text{ belongs to the optimal type}\}$,

$\mathbf{E}_2 := \{\text{At least one optimal-typed arm is played in epoch } I, \text{ and the empirically best arm is \underline{not} \text{ optimal-typed}\}$.

Recall that $\bar{\Delta} = \mu_1 - \mu_K$ denotes the maximal sub-optimality gap. Then, the cumulative pseudo-regret R_n (superscript π suppressed for notational convenience) of $\text{ALG1}(n)$ is bounded as

$$\begin{aligned}
 R_n &\leq \mathbb{1}\{\tau \leq n\} [\bar{\Delta}\tau + \mathbb{1}\{\mathbf{E}_1 \cup \mathbf{E}_2\} \bar{\Delta}n] + \mathbb{1}\{\tau > n\} \bar{\Delta}n \\
 &\leq \bar{\Delta}\tau + (\mathbb{1}\{\tau \leq n, \mathbf{E}_1\} + \mathbb{1}\{\tau \leq n, \mathbf{E}_2\}) \bar{\Delta}n + \mathbb{1}\{\tau > n\} \bar{\Delta}n.
 \end{aligned}$$

Taking expectations,

$$\begin{aligned}
 \mathbb{E}R_n &\leq \bar{\Delta}\mathbb{E}\tau + [\mathbb{P}(\tau \leq n, \mathbf{E}_1) + \mathbb{P}(\tau \leq n, \mathbf{E}_2)] \bar{\Delta}n + \mathbb{P}(\tau > n) \bar{\Delta}n \\
 &\leq 2\bar{\Delta}\mathbb{E}\tau + [\mathbb{P}(\tau \leq n, \mathbf{E}_1) + \mathbb{P}(\tau \leq n, \mathbf{E}_2)] \bar{\Delta}n,
 \end{aligned}$$

where the last step uses Markov’s inequality. Therefore,

$$\begin{aligned}
 \mathbb{E}R_n &\leq 2K\bar{\Delta}\mathbb{E}[Im_I] + [\mathbb{P}(\tau \leq n, \mathbf{E}_1) + \mathbb{P}(\tau \leq n, \mathbf{E}_2)] \bar{\Delta}n \\
 &\leq 4K\bar{\Delta}\mathbb{E}\left[Ie^{2\sqrt{I}} \right] \log n + [\mathbb{P}(\tau \leq n, \mathbf{E}_1) + \mathbb{P}(\tau \leq n, \mathbf{E}_2)] \bar{\Delta}n.
 \end{aligned}$$

F.1. Upper bounding $\mathbb{E}\left[Ie^{2\sqrt{I}} \right]$

Recall that $\delta = \min_{1 \leq i < j \leq K} (\mu_i - \mu_j)$ denotes the smallest gap between any two distinct mean rewards. Then, on the event \mathbb{D} , for any $a, b \in \mathcal{A}, a < b$, we either have $\mathbb{E}\left[X_{a,j}^l - X_{b,j}^l \right] \geq \delta$ or

$\mathbb{E} [X_{a,j}^l - X_{b,j}^l] \leq -\delta$. Without loss of generality, suppose that $\mathbb{E} [X_{a,j}^l - X_{b,j}^l] \geq \delta$. Then,

$$\begin{aligned} \mathbb{P}_{\mathcal{D}} \left(\left| \sum_{j=1}^{m_l} (X_{a,j}^l - X_{b,j}^l) \right| < 2m_l e^{-\sqrt{l}} \right) &\leq \mathbb{P}_{\mathcal{D}} \left(\sum_{j=1}^{m_l} (X_{a,j}^l - X_{b,j}^l) < 2m_l e^{-\sqrt{l}} \right) \\ &= \mathbb{P}_{\mathcal{D}} \left(\sum_{j=1}^{m_l} (X_{a,j}^l - X_{b,j}^l - \delta) < -m_l (\delta - 2e^{-\sqrt{l}}) \right). \end{aligned}$$

Then, for $l > \lceil \log^2(4/\delta) \rceil =: k^*$, one has that

$$\mathbb{P}_{\mathcal{D}} \left(\left| \sum_{j=1}^{m_l} (X_{a,j}^l - X_{b,j}^l) \right| < 2m_l e^{-\sqrt{l}} \right) \leq \mathbb{P}_{\mathcal{D}} \left(\sum_{j=1}^{m_l} (X_{a,j}^l - X_{b,j}^l - \delta) < -2m_l e^{-\sqrt{l}} \right) \leq n^{-2}, \quad (12)$$

where the final inequality follows using the Chernoff-Hoeffding bound (Hoeffding, 1963), together with the fact that $-1 \leq X_{a,j}^l - X_{b,j}^l \leq 1$ and $m_l = \lceil e^{2\sqrt{l}} \log n \rceil$. Using (11) and (12), we obtain for $k > k^* + 1$ that

$$\begin{aligned} \mathbb{P}(I \geq k) &\leq \prod_{l=1}^{k-1} \left[1 - \mathbb{P}(\mathcal{D}) + \mathbb{P}(\mathcal{D}) \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}_{\mathcal{D}} \left(\left| \sum_{j=1}^{m_l} (X_{a,j}^l - X_{b,j}^l) \right| < 2m_l e^{-\sqrt{l}} \right) \right] \\ &\leq \prod_{k^* < l \leq k-1} \left[1 - \mathbb{P}(\mathcal{D}) + \frac{K^2 \mathbb{P}(\mathcal{D})}{n^2} \right] \\ &\leq \left[1 - \frac{3\mathbb{P}(\mathcal{D})}{4} \right]^{k-k^*-1}, \end{aligned}$$

where the last inequality holds for $n \geq 2K$.⁴ Thus, for any $k \geq 1$ and $n \geq 2K$, we have

$$\mathbb{P}(I \geq k^* + k) \leq \left[1 - \frac{3\mathbb{P}(\mathcal{D})}{4} \right]^{k-1}.$$

4. We will ensure that all guarantees hold for $n \geq K$ by offsetting regret by $2K\bar{\Delta}$ in the end.

Now,

$$\begin{aligned}
 \mathbb{E} \left[I e^{2\sqrt{I}} \right] &\leq k^* e^{2\sqrt{k^*}} + \sum_{k=1}^{\infty} (k^* + k) e^{2\sqrt{k^*+k}} \mathbb{P}(I = k^* + k) \\
 &\stackrel{(\dagger)}{\leq} k^* e^{2\sqrt{k^*}} + \sum_{k=1}^{\infty} (k^* + k) e^{2\sqrt{k^*}} e^{2\sqrt{k}} \mathbb{P}(I = k^* + k) \\
 &\stackrel{(\ddagger)}{\leq} k^* e^{2\sqrt{k^*}} + 2 \sum_{k=1}^{\infty} k^* k e^{2\sqrt{k^*}} e^{2\sqrt{k}} \mathbb{P}(I = k^* + k) \\
 &\leq k^* e^{2\sqrt{k^*}} + 2 \sum_{k=1}^{\infty} k^* k e^{2\sqrt{k^*}} e^{2\sqrt{k}} \mathbb{P}(I \geq k^* + k) \\
 &\leq k^* e^{2\sqrt{k^*}} + 2 \sum_{k=1}^{\infty} k^* k e^{2\sqrt{k^*}} e^{2\sqrt{k}} \left[1 - \frac{3\mathbb{P}(\mathcal{D})}{4} \right]^{k-1} \\
 &= k^* e^{2\sqrt{k^*}} \left[1 + 2 \sum_{k=1}^{\infty} k e^{2\sqrt{k}} \left[1 - \frac{3\mathbb{P}(\mathcal{D})}{4} \right]^{k-1} \right] \\
 &\leq 4k^* e^{2\sqrt{k^*}} \sum_{k=1}^{\infty} k e^{2\sqrt{k}} \left[1 - \frac{3\mathbb{P}(\mathcal{D})}{4} \right]^{k-1},
 \end{aligned}$$

where (\dagger) follows using $\sqrt{k^* + k} \leq \sqrt{k^*} + \sqrt{k}$, and (\ddagger) using $k^* + k \leq 2k^*k$ (since $k^*, k \in \mathbb{N}$). Define

$$C_{\alpha} := \sum_{k=1}^{\infty} k e^{2\sqrt{k}} \left[1 - \frac{3\mathbb{P}(\mathcal{D})}{4} \right]^{k-1}.$$

Since $\mathbb{P}(\mathcal{D}) = K! \prod_{i=1}^K \alpha_i$, note that the infinite summation is finite since $\alpha = (\alpha_i : i = 1, \dots, K)$ is coordinate-wise bounded away from 0. Therefore,

$$\mathbb{E} \left[I e^{2\sqrt{I}} \right] \leq 4C_{\alpha} k^* e^{2\sqrt{k^*}}.$$

Note that

$$e^{2\sqrt{k^*}} = e^{2\sqrt{\lceil \log^2(4/\delta) \rceil}} \leq e^{2\sqrt{\log^2(4/\delta)+1}} \leq e^{2(\sqrt{\log^2(4/\delta)+1})} \leq \frac{16e^2}{\delta^2}.$$

Therefore, in conclusion,

$$\mathbb{E} \left[I e^{2\sqrt{I}} \right] \leq \frac{128e^2 C_{\alpha}}{\delta^2} \log^2 \left(\frac{4}{\delta} \right).$$

F.2. Upper bounding $\mathbb{P}(\tau \leq n, \mathbf{E}_1)$

Note that on the event $\{\tau \leq n\}$, the duration of epoch I is Km_I . On the event \mathbf{E}_1 , the consideration set contains at least two arms that belong to the same type. Without loss of generality, suppose that

these arms are indexed 1 and 2. Then,

$$\begin{aligned}
 \mathbb{P}(\tau \leq n, \mathbf{E}_1) &\leq \mathbb{P}\left(\left|\sum_{j=1}^{m_I} (X_{1,j}^I - X_{2,j}^I)\right| \geq 2m_I e^{-\sqrt{I}}, \tau \leq n\right) \\
 &\leq \mathbb{P}\left(\left|\sum_{j=1}^{m_I} (X_{1,j}^I - X_{2,j}^I)\right| \geq 2m_I e^{-\sqrt{I}}, I \leq n\right) \\
 &\leq \sum_{l=1}^n \mathbb{P}\left(\left|\sum_{j=1}^{m_l} (X_{1,j}^l - X_{2,j}^l)\right| \geq 2m_l e^{-\sqrt{l}}\right) \\
 &\leq \sum_{l=1}^n \frac{2}{n^2} \\
 &= \frac{2}{n},
 \end{aligned}$$

where the last inequality follows using the Chernoff-Hoeffding bound (Hoeffding, 1963).

F.3. Upper bounding $\mathbb{P}(\tau \leq n, \mathbf{E}_2)$

Note that on the event $\{\tau \leq n\}$, the duration of epoch I is Km_I . On the event \mathbf{E}_2 , the consideration set \mathcal{A} contains at least one arm of the optimal type, and the empirically best arm belongs to an inferior type. Without loss of generality, suppose that arm 1 belongs to the optimal type and $\mathcal{I} \subset \mathcal{A}$ denotes the set of inferior-typed arms. We then have

$$\begin{aligned}
 \mathbb{P}(\tau \leq n, \mathbf{E}_2) &\leq \mathbb{P}\left(\bigcup_{a \in \mathcal{I}} \left\{ \sum_{j=1}^{m_I} (X_{a,j}^I - X_{1,j}^I) \geq 2m_I e^{-\sqrt{I}}, \tau \leq n \right\}\right) \\
 &\leq \sum_{a \in \mathcal{I}} \mathbb{P}\left(\sum_{j=1}^{m_I} (X_{a,j}^I - X_{1,j}^I) \geq 2m_I e^{-\sqrt{I}}, I \leq n\right) \\
 &\leq \sum_{a \in \mathcal{I}} \sum_{l=1}^n \mathbb{P}\left(\sum_{j=1}^{m_l} (X_{a,j}^l - X_{1,j}^l) \geq 2m_l e^{-\sqrt{l}}\right) \\
 &\leq \sum_{a \in \mathcal{I}} \sum_{l=1}^n \frac{1}{n^2} \\
 &\leq \frac{K}{n},
 \end{aligned}$$

where the second-to-last inequality follows using the Chernoff-Hoeffding bound (Hoeffding, 1963)

since $\mathbb{E}[X_{a,j}^l - X_{1,j}^l] \leq -\underline{\Delta} < 0 \forall a \in \mathcal{I}$.

F.4. Putting everything together

In conclusion, the expected cumulative regret of the policy π given by $\text{ALG1}(n)$ is bounded for any $n \geq K$ as

$$\mathbb{E}R_n^\pi \leq \frac{\tilde{C}_\alpha K \bar{\Delta} \log n}{\delta^2} \log^2 \left(\frac{4}{\delta} \right) + 4K \bar{\Delta},$$

where \tilde{C}_α is a finite constant that depends only on $\alpha = (\alpha_i : i = 1, \dots, K)$. In particular, \tilde{C}_α is given by the following infinite summation:

$$\tilde{C}_\alpha := 512e^2 C_\alpha = 512e^2 \sum_{k=1}^{\infty} k e^{2\sqrt{k}} \left[1 - \frac{3K! \prod_{i=1}^K \alpha_i}{4} \right]^{k-1}. \quad (13)$$

□

Appendix G. Proof of Proposition 6

Consider the following stopping time:

$$\tau := \inf \left\{ m \in \mathbb{N} : \exists a, b \in \mathcal{A}, a < b \text{ s.t. } \mathcal{Z}_{a,b} + \sum_{j=1}^m (X_{a,j} - X_{b,j}) < 4\sqrt{m \log m} \right\}.$$

Since $\mathbb{P} \left(\bigcap_{m \geq 1} \bigcap_{a,b \in \mathcal{A}, a < b} \left\{ \left| \mathcal{Z}_{a,b} + \sum_{j=1}^m (X_{a,j} - X_{b,j}) \right| \geq 4\sqrt{m \log m} \right\} \right) \geq \mathbb{P}(\tau = \infty)$, it suffices to show that $\mathbb{P}(\tau = \infty)$ is bounded away from 0. To this end, define the following entities:

$$\begin{aligned} \Lambda_K &:= \inf \left\{ p \in \mathbb{N} : \sum_{m=p}^{\infty} \frac{1}{m^8} \leq \frac{1}{2K^2} \right\}, \\ T_0 &:= \max \left(\left\lceil \left(\frac{64}{\delta^2} \right) \log^2 \left(\frac{64}{\delta^2} \right) \right\rceil, \Lambda_K \right), \\ f(x) &:= x + 4\sqrt{x \log x} \quad \text{for } x \geq 1. \end{aligned}$$

Lemma 10 *For any $a, b \in \mathcal{A}$ s.t. $a < b$, it is the case that*

$$\{\mathcal{Z}_{a,b} > f(T_0)\} \subseteq \bigcap_{m=1}^{T_0} \left\{ \mathcal{Z} + \sum_{j=1}^m (X_{a,j} - X_{b,j}) \geq 4\sqrt{m \log m} \right\}.$$

Proof of Lemma 10. Note that

$$\begin{aligned} \mathcal{Z}_{a,b} &> f(T_0) \\ &= T_0 + 4\sqrt{T_0 \log T_0} \\ &\geq m + 4\sqrt{m \log m} \quad \forall 1 \leq m \leq T_0 \\ &\stackrel{(a)}{\geq} \sum_{j=1}^m (X_{b,j} - X_{a,j}) + 4\sqrt{m \log m} \quad \forall 1 \leq m \leq T_0 \\ \implies \mathcal{Z}_{a,b} + \sum_{j=1}^m (X_{a,j} - X_{b,j}) &\geq 4\sqrt{m \log m} \quad \forall 1 \leq m \leq T_0, \end{aligned}$$

where (a) follows since the rewards are bounded in $[0, 1]$, i.e., $|X_{a,j} - X_{b,j}| \leq 1$. \square

Lemma 11 *For $m \geq T_0$, it is the case that*

$$\delta \geq 8\sqrt{\frac{\log m}{m}}.$$

Proof of Lemma 11. First of all, note that $T_0 \geq (64/\delta^2) \log^2 (64/\delta^2) \geq 64$ (since $\delta \leq 1$). For $s = (64/\delta^2) \log^2 (64/\delta^2)$, one has

$$\delta^2 = \frac{64 \log^2 \left(\frac{64}{\delta^2}\right)}{s} \underset{(b)}{\geq} \frac{64 \left[\log \left(\frac{64}{\delta^2}\right) + 2 \log \log \left(\frac{64}{\delta^2}\right)\right]}{s} = \frac{64 \log s}{s},$$

where (b) follows since the function $g(x) := x^2 - x - 2 \log x$ is monotone increasing for $x \geq \log 64$ (think of $\log (64/\delta^2)$ as x), and therefore attains its minimum at $x = \log 64$; one can verify that this minimum is strictly positive. Furthermore, since $\log s/s$ is monotone decreasing for $s \geq 64$, it follows that for any $m \geq T_0$,

$$\delta^2 \geq \frac{64 \log m}{m}.$$

\square

Now coming back to the proof of Proposition 6, consider an arbitrary $l \in \mathbb{N}$ such that $l > T_0$. Then,

$$\begin{aligned} \mathbb{P}(\tau \leq l) &= \mathbb{P}(\tau \leq l, \mathcal{Z} > f(T_0)) + \mathbb{P}(\tau \leq l, \mathcal{Z} \leq f(T_0)) \\ &\leq \mathbb{P}(\tau \leq l, \mathcal{Z} > f(T_0)) + \Phi(f(T_0)). \end{aligned}$$

Now,

$$\begin{aligned}
 & \mathbb{P}(\tau \leq l, \mathcal{Z} > f(T_0)) \\
 &= \mathbb{P}\left(\bigcup_{m=1}^l \bigcup_{a,b \in \mathcal{A}, a < b} \left\{ \mathcal{Z}_{a,b} + \sum_{j=1}^m (X_{a,j} - X_{b,j}) < 4\sqrt{m \log m}, \mathcal{Z}_{a,b} > f(T_0) \right\}\right) \\
 &\stackrel{(\dagger)}{=} \mathbb{P}\left(\bigcup_{m=T_0}^l \bigcup_{a,b \in \mathcal{A}, a < b} \left\{ \mathcal{Z}_{a,b} + \sum_{j=1}^m (X_{a,j} - X_{b,j}) < 4\sqrt{m \log m}, \mathcal{Z}_{a,b} > f(T_0) \right\}\right) \\
 &\leq \sum_{m=T_0}^l \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}\left(\mathcal{Z}_{a,b} + \sum_{j=1}^m (X_{a,j} - X_{b,j}) < 4\sqrt{m \log m}, \mathcal{Z}_{a,b} > f(T_0)\right) \\
 &= \sum_{m=T_0}^l \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}\left(\mathcal{Z}_{a,b} + \sum_{j=1}^m (X_{a,j} - X_{b,j} - \delta) < -m \left(\delta - 4\sqrt{\frac{\log m}{m}}\right), \mathcal{Z}_{a,b} > f(T_0)\right) \\
 &\leq \sum_{m=T_0}^l \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}\left(\sum_{j=1}^m (X_{a,j} - X_{b,j} - \delta) < -m \left(\delta - 4\sqrt{\frac{\log m}{m}}\right), \mathcal{Z}_{a,b} > f(T_0)\right) \\
 &\stackrel{(\ddagger)}{\leq} \sum_{m=T_0}^l \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}\left(\sum_{j=1}^m (X_{a,j} - X_{b,j} - \delta) < -4\sqrt{m \log m}, \mathcal{Z}_{a,b} > f(T_0)\right) \\
 &= \bar{\Phi}(f(T_0)) \sum_{m=T_0}^l \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}\left(\sum_{j=1}^m (X_{a,j} - X_{b,j} - \delta) < -4\sqrt{m \log m}\right) \\
 &\stackrel{(*)}{\leq} \bar{\Phi}(f(T_0)) \sum_{m=T_0}^l \sum_{a,b \in \mathcal{A}, a < b} \frac{1}{m^8} \\
 &\leq \bar{\Phi}(f(T_0)) K^2 \sum_{m=T_0}^{\infty} \frac{1}{m^8} \\
 &\stackrel{(*)}{\leq} \frac{\bar{\Phi}(f(T_0))}{2},
 \end{aligned}$$

where (\dagger) follows from Lemma 10, (\ddagger) from Lemma 11, $(*)$ follows using the Chernoff-Hoeffding bound

and finally, $(*)$ follows from the definition of T_0 . Therefore, we have

$$\begin{aligned}
 \mathbb{P}(\tau \leq l) &\leq \frac{\bar{\Phi}(f(T_0))}{2} + \Phi(f(T_0)) = 1 - \frac{\bar{\Phi}(f(T_0))}{2} \\
 \implies \mathbb{P}(\tau > l) &\geq \frac{\bar{\Phi}(f(T_0))}{2}.
 \end{aligned}$$

Taking the limit $l \rightarrow \infty$ and appealing to the continuity of probability, we obtain

$$\begin{aligned} \mathbb{P}(\tau = \infty) &\geq \frac{\bar{\Phi}(f(T_0))}{2} \\ \implies \mathbb{P}\left(\bigcap_{m \geq 1} \bigcap_{a, b \in \mathcal{A}, a < b} \left\{ \left| \mathcal{Z}_{a,b} + \sum_{j=1}^m (X_{a,j} - X_{b,j}) \right| \geq 4\sqrt{m \log m} \right\}\right) &\geq \frac{\bar{\Phi}(f(T_0))}{2}. \end{aligned}$$

□

Appendix H. Proof of Theorem 7

We will initially assume $\delta > 8\sqrt{\log n/n}$ for technical convenience. In the final step leading up to the asserted bound, we will relax this assumption by offsetting regret appropriately.

Let $\mathcal{A} := \{1, \dots, K\}$. Define the following stopping times:

$$\begin{aligned} \tau_1(s_n) &:= \inf \left\{ m \geq s_n : \exists a, b \in \mathcal{A}, a < b \text{ s.t. } \left| \mathcal{Z}_{a,b} + \sum_{j=1}^m (X_{a,j} - X_{b,j}) \right| < 4\sqrt{m \log m} \right\}, \\ \tau_2(s_n) &:= \inf \left\{ m \geq s_n : \left| \sum_{j=1}^m (X_{a,j} - X_{b,j}) \right| \geq 4\sqrt{m \log n} \forall a, b \in \mathcal{A}, a < b \right\}. \end{aligned}$$

To keep notations simple, we will suppress the argument and denote $\tau_1(s_n)$ and $\tau_2(s_n)$ by τ_1 and τ_2 respectively (the dependence on s_n will be implicit going forward). Let R_t denote the cumulative pseudo-regret of $\text{ALG2}(n)$ after $t \leq n$ pulls. Let D denote the event that the first batch of K arms queried from the reservoir is ‘‘all-distinct,’’ i.e., no two arms in this batch belong to the same type; let D^c be the complement of this event. Let CI denote the event that the algorithm commits to an inferior-typed arm. Let $\tilde{R}_\cdot, \bar{R}_\cdot$ be independently drawn from the same distribution as R_\cdot . Let $x^+ := \max(x, 0)$ for $x \in \mathbb{R}$. Then, R_n evolves according to the following stochastic recursion:

$$\begin{aligned} &R_n \\ &\leq \mathbb{1}\{\mathsf{D}\} \left[\mathbb{1}\{\tau_1 < \tau_2\} \left(\bar{\Delta} \min(K\tau_1, n) + \tilde{R}_{(n-K\tau_1)^+} \right) + \mathbb{1}\{\tau_1 \geq \tau_2\} \left(\bar{\Delta} \min(K\tau_2, n) + \mathbb{1}\{\mathsf{CI}\} \bar{\Delta} (n - K\tau_2)^+ \right) \right] \\ &\quad + \mathbb{1}\{\mathsf{D}^c\} \left[\mathbb{1}\{\tau_1 < \tau_2\} \left(\bar{\Delta} \min(K\tau_1, n) + \bar{R}_{(n-K\tau_1)^+} \right) + \mathbb{1}\{\tau_1 \geq \tau_2\} \bar{\Delta} n \right] \\ &\leq \mathbb{1}\{\mathsf{D}\} \left[\mathbb{1}\{\tau_1 < \tau_2\} \left(\bar{\Delta} \min(K\tau_2, n) + \tilde{R}_n \right) + \mathbb{1}\{\tau_1 \geq \tau_2\} \left(\bar{\Delta} \min(K\tau_2, n) + \mathbb{1}\{\mathsf{CI}\} \bar{\Delta} n \right) \right] \\ &\quad + \mathbb{1}\{\mathsf{D}^c\} \left[\mathbb{1}\{\tau_1 < \tau_2\} \left(\bar{\Delta} \min(K\tau_1, n) + \bar{R}_n \right) + \mathbb{1}\{\tau_1 \geq \tau_2\} \bar{\Delta} n \right] \\ &\leq \mathbb{1}\{\mathsf{D}\} \left[2\bar{\Delta} \min(K\tau_2, n) + \mathbb{1}\{\tau_1 < \tau_2\} \tilde{R}_n + \mathbb{1}\{\mathsf{CI}\} \bar{\Delta} n \right] + \mathbb{1}\{\mathsf{D}^c\} \left[\bar{\Delta} K\tau_1 + \bar{R}_n + \mathbb{1}\{\tau_1 \geq \tau_2\} \bar{\Delta} n \right]. \end{aligned}$$

Taking expectations on both sides, one recovers using the independence of \tilde{R}_n, \bar{R}_n that

$$\begin{aligned} &\mathbb{E}R_n \\ &\leq \frac{\bar{\Delta}}{\mathbb{P}(\tau_1 \geq \tau_2 | \mathsf{D})} \left[2\mathbb{E}[\min(K\tau_2, n) | \mathsf{D}] + \left(\frac{\mathbb{P}(\mathsf{D}^c)}{\mathbb{P}(\mathsf{D})} \right) \mathbb{E}[K\tau_1 | \mathsf{D}^c] + \left(\mathbb{P}(\mathsf{CI} | \mathsf{D}) + \left(\frac{\mathbb{P}(\mathsf{D}^c)}{\mathbb{P}(\mathsf{D})} \right) \mathbb{P}(\tau_1 \geq \tau_2 | \mathsf{D}^c) \right) n \right], \end{aligned}$$

where $\mathbb{P}(\mathsf{D}) = K! \prod_{i=1}^K \alpha_i$.

H.1. Lower bounding $\mathbb{P}(\tau_1 \geq \tau_2 | \mathbf{D})$

Define the following:

$$\gamma(s_n) := \mathbb{P}(\tau_1(s_n) = \infty | \mathbf{D}), \quad (14)$$

where $\tau_1(s_n)$ and \mathbf{D} are as defined before. We will suppress the dependence on s_n to keep notations minimal. Note that

$$\begin{aligned} \mathbb{P}(\tau_1 < \tau_2 | \mathbf{D}) &= \mathbb{P}(\tau_1 < \tau_2, \tau_2 = \infty | \mathbf{D}) + \mathbb{P}(\tau_1 < \tau_2, \tau_2 < \infty | \mathbf{D}) \\ &\leq \mathbb{P}(\tau_2 = \infty | \mathbf{D}) + \mathbb{P}(\tau_1 < \infty | \mathbf{D}) \\ &= \mathbb{P}(\tau_1 < \infty | \mathbf{D}) \\ &= 1 - \gamma(s_n), \end{aligned}$$

where the equality in the third step follows since τ_2 is almost surely finite on the event \mathbf{D} (proved in §H.1.1 below), and the final equality is due to (14). Thus, $\mathbb{P}(\tau_1 \geq \tau_2 | \mathbf{D}) \geq \gamma(s_n)$.

H.1.1. PROOF THAT τ_2 IS ALMOST SURELY FINITE ON \mathbf{D}

Let $\mathbb{P}_{\mathbf{D}}(\cdot) := \mathbb{P}(\cdot | \mathbf{D})$ be the conditional measure w.r.t. the event \mathbf{D} . Let $\mathcal{A} := \{1, \dots, K\}$. Then, by continuity of probability, we have

$$\begin{aligned} \mathbb{P}_{\mathbf{D}}(\tau_2 = \infty) &= \lim_{l \rightarrow \infty} \mathbb{P}_{\mathbf{D}}(\tau_2 > l) \\ &= \lim_{l \rightarrow \infty} \mathbb{P}_{\mathbf{D}} \left(\bigcap_{m=s_n}^l \bigcup_{a,b \in \mathcal{A}, a < b} \left\{ \left| \sum_{j=1}^m (X_{a,j} - X_{b,j}) \right| < 4\sqrt{m \log n} \right\} \right) \\ &\leq \lim_{l \rightarrow \infty} \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}_{\mathbf{D}} \left(\left| \sum_{j=1}^l (X_{a,j} - X_{b,j}) \right| < 4\sqrt{l \log n} \right). \end{aligned}$$

On \mathbf{D} , it must be that $|\mathbb{E}[X_{a,j} - X_{b,j}]| \geq \delta$. Without loss of generality, assume that $\mathbb{E}[X_{a,j} - X_{b,j}] \geq \delta$. Then,

$$\begin{aligned} \mathbb{P}_{\mathbf{D}}(\tau_2 = \infty) &\leq \lim_{l \rightarrow \infty} \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}_{\mathbf{D}} \left(\sum_{j=1}^l (X_{a,j} - X_{b,j}) < 4\sqrt{l \log n} \right) \\ &= \lim_{l \rightarrow \infty} \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}_{\mathbf{D}} \left(\sum_{j=1}^l (X_{a,j} - X_{b,j} - \delta) < -l \left(\delta - 4\sqrt{\frac{\log n}{l}} \right) \right) \\ &\leq \lim_{l \rightarrow \infty} \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}_{\mathbf{D}} \left(\sum_{j=1}^l (X_{a,j} - X_{b,j} - \delta) < -4l\sqrt{\log n} \left(\frac{2}{\sqrt{n}} - \frac{1}{\sqrt{l}} \right) \right), \end{aligned}$$

where the last inequality follows since $\delta > 8\sqrt{\log n/n}$ (by assumption). Now, using the Chernoff-Hoeffding bound (Hoeffding, 1963) together with the fact that $-1 \leq X_{a,j} - X_{b,j} \leq 1$, we obtain for $l > n$ and any $a, b \in \mathcal{A}, a < b$ that

$$\begin{aligned} \mathbb{P}_{\mathcal{D}} \left(\sum_{j=1}^l (X_{a,j} - X_{b,j} - \delta) < -4l\sqrt{\log n} \left(\frac{2}{\sqrt{n}} - \frac{1}{\sqrt{l}} \right) \right) &\leq \exp \left[-8l \left(\frac{2}{\sqrt{n}} - \frac{1}{\sqrt{l}} \right)^2 \log n \right] \\ &= \exp \left[-8 \left(\frac{4l}{n} - 4\sqrt{\frac{l}{n}} + 1 \right)^2 \log n \right]. \end{aligned}$$

Summing over $a, b \in \mathcal{A}$, $a < b$ and taking the limit $l \rightarrow \infty$ proves the stated assertion. \square

H.2. Upper bounding $\mathbb{E}[\min(K\tau_2, n) | \mathcal{D}]$

Let $\mathbb{P}_{\mathcal{D}}(\cdot) := \mathbb{P}(\cdot | \mathcal{D})$ be the conditional measure w.r.t. the event \mathcal{D} . Let $\mathcal{A} := \{1, \dots, K\}$. Then,

$$\begin{aligned} \mathbb{E}[\min(K\tau_2, n) | \mathcal{D}] &= K\mathbb{E} \left[\min \left(\tau_2, \frac{n}{K} \right) | \mathcal{D} \right] \\ &\leq K\mathbb{E}[\min(\tau_2, n) | \mathcal{D}] \\ &\leq Ks_n + K \sum_{k=s_n+1}^n \mathbb{P}_{\mathcal{D}}(\tau_2 \geq k) \\ &\leq Ks_n + K \sum_{k=s_n}^n \mathbb{P}_{\mathcal{D}}(\tau_2 \geq k+1) \\ &\leq Ks_n + K \sum_{k=1}^n \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}_{\mathcal{D}} \left(\left| \sum_{j=1}^k (X_{a,j} - X_{b,j}) \right| < 4\sqrt{k \log n} \right). \end{aligned}$$

On \mathcal{D} , it must be that $|\mathbb{E}[X_{a,j} - X_{b,j}]| \geq \delta$. Without loss of generality, assume that $\mathbb{E}[X_{a,j} - X_{b,j}] \geq \delta$. Then,

$$\begin{aligned} \mathbb{E}[\min(K\tau_2, n) | \mathcal{D}] &\leq Ks_n + K \sum_{k=1}^n \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}_{\mathcal{D}} \left(\sum_{j=1}^k (X_{a,j} - X_{b,j}) < 4\sqrt{k \log n} \right) \\ &= Ks_n + K \sum_{k=1}^n \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}_{\mathcal{D}} \left(\sum_{j=1}^k (X_{a,j} - X_{b,j} - \delta) < -k \left(\delta - 4\sqrt{\frac{\log n}{k}} \right) \right) \\ &\leq Ks_n + \frac{32K^3 \log n}{\delta^2} + K \sum_{k=\lceil \frac{64 \log n}{\delta^2} \rceil}^n \sum_{a,b \in \mathcal{A}, a < b} \mathbb{P}_{\mathcal{D}} \left(\sum_{j=1}^k (X_{a,j} - X_{b,j} - \delta) < \frac{-k\delta}{2} \right), \end{aligned}$$

where the last step follows since $\delta > 8\sqrt{\log n/n}$ (by assumption) implies $n > 64 \log n / \delta^2$, and $k \geq 64 \log n / \delta^2$ implies $\delta - 4\sqrt{\log n/k} \geq \delta/2$. Finally, using the Chernoff-Hoeffding inequality (Hoeffding, 1963) together with the fact that $|X_{a,j} - X_{b,j}| \leq 1$, one obtains

$$\mathbb{E}[\min(K\tau_2, n) | \mathcal{D}] \leq Ks_n + \frac{32K^3 \log n}{\delta^2} + \frac{K^3}{2} \sum_{k=\lceil \frac{64 \log n}{\delta^2} \rceil}^n \exp \left(\frac{-\delta^2 k}{8} \right) \leq Ks_n + \frac{64K^3 \log n}{\delta^2}.$$

H.3. Upper bounding $\mathbb{E}[K\tau_1 | \mathcal{D}^c]$

The event \mathcal{D}^c will be implicitly assumed and we will drop the conditional argument for notational simplicity. Let $\mathcal{A} := \{1, \dots, K\}$. Without loss of generality, suppose that arm 1 and 2 belong to the same type. Then,

$$\begin{aligned}
 \mathbb{E}[K\tau_1 | \mathcal{D}^c] &= Ks_n + K \sum_{k \geq s_n+1} \mathbb{P}(\tau_1 \geq k) \\
 &= Ks_n + K \sum_{k \geq s_n} \mathbb{P}(\tau_1 \geq k+1) \\
 &= Ks_n + K \sum_{k \geq s_n} \mathbb{P} \left(\bigcap_{m=s_n}^k \bigcap_{a,b \in \mathcal{A}, a < b} \left\{ \left| \mathcal{Z}_{a,b} + \sum_{j=1}^m (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{m \log m} \right\} \right) \\
 &\leq Ks_n + K \sum_{k \geq 1} \mathbb{P} \left(\left| \mathcal{Z}_{1,2} + \sum_{j=1}^k (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{k \log k} \right).
 \end{aligned}$$

Since $\mathcal{Z}_{a,b}$ is a standard Gaussian, and the increments $X_{1,j} - X_{2,j}$ are zero-mean sub-Gaussian with variance proxy 1, it follows from the Chernoff-Hoeffding concentration bound (Hoeffding, 1963) that

$$\mathbb{E}[K\tau_1 | \mathcal{D}^c] \leq Ks_n + 2K \sum_{k \geq 1} \frac{1}{k^4} = \left(s_n + \frac{\pi^4}{45} \right) K < (s_n + 3) K.$$

H.4. Upper bounding $\mathbb{P}(\text{CI} | \mathcal{D})$

Let $\mathbb{P}_{\mathcal{D}}(\cdot) := \mathbb{P}(\cdot | \mathcal{D})$ be the conditional measure w.r.t. the event \mathcal{D} . Let $\mathcal{A} := \{1, \dots, K\}$ and without loss of generality, suppose that arm 1 is optimal (mean μ_1). Then,

$$\begin{aligned}
 \mathbb{P}(\text{CI} | \mathcal{D}) &\leq \mathbb{P}_{\mathcal{D}} \left(\bigcup_{b=2}^K \left\{ \sum_{j=1}^{\tau_2} (X_{1,j} - X_{b,j}) \leq -4\sqrt{\tau_2 \log n} \right\} \right) \\
 &\leq \sum_{b=2}^K \sum_{k=s_n}^n \mathbb{P}_{\mathcal{D}} \left(\sum_{j=1}^k (X_{1,j} - X_{b,j}) \leq -4\sqrt{k \log n} \right) + \sum_{b=2}^K \mathbb{P}_{\mathcal{D}}(\tau_2 > n) \\
 &\leq \sum_{b=2}^K \sum_{k=1}^n \frac{1}{n^8} + \frac{K^3}{n^8} \\
 &\leq \frac{K}{n^7} + \frac{K^3}{n^8},
 \end{aligned}$$

where the second-to-last step follows using the Chernoff-Hoeffding inequality (Hoeffding, 1963).

H.5. Upper bounding $\mathbb{P}(\tau_1 \geq \tau_2 | \mathcal{D}^c)$

Let $\mathbb{P}_{\mathcal{D}^c}(\cdot) := \mathbb{P}(\cdot | \mathcal{D}^c)$ be the conditional measure w.r.t. the event \mathcal{D}^c . Let $\mathcal{A} := \{1, \dots, K\}$. On \mathcal{D}^c , there exist 2 arms in \mathcal{A} that belong to the same type; without loss of generality suppose that these arms are indexed by 1, 2. Then,

$$\begin{aligned}
 \mathbb{P}(\tau_1 \geq \tau_2 | \mathcal{D}^c) &\leq \mathbb{P}(\tau_1 > n | \mathcal{D}^c) + \mathbb{P}(\tau_2 \leq n | \mathcal{D}^c) \\
 &\leq \frac{2}{n^4} + \mathbb{P}_{\mathcal{D}^c}(\tau_2 \leq n) \\
 &= \frac{2}{n^4} + \mathbb{P}_{\mathcal{D}^c} \left(\bigcup_{m=s_n}^n \bigcap_{a,b \in \mathcal{A}, a < b} \left\{ \left| \sum_{j=1}^m (X_{a,j} - X_{b,j}) \right| \geq 4\sqrt{m \log n} \right\} \right) \\
 &\leq \frac{2}{n^4} + \sum_{m=1}^n \mathbb{P}_{\mathcal{D}^c} \left(\left| \sum_{j=1}^m (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{m \log n} \right) \\
 &\leq \frac{2}{n^4} + \frac{2}{n^7}, \tag{15}
 \end{aligned}$$

where the last step follows using the Chernoff-Hoeffding bound ([Hoeffding, 1963](#)).

H.6. Putting everything together

Combining everything, one finally obtains that when $\delta > 8\sqrt{\log n/n}$,

$$\mathbb{E}R_n \leq \frac{CK^3\bar{\Delta}}{\gamma(s_n)} \left(\frac{\log n}{\delta^2} + \frac{s_n}{\mathbb{P}(\mathcal{D})} \right),$$

where $\gamma(s_n)$ is as defined in (14), $\mathbb{P}(\mathcal{D}) = K! \prod_{i=1}^K \alpha_i$, and C is some absolute constant. When $\delta \leq 8\sqrt{\log n/n}$, regret is at most $\bar{\Delta}n \leq 64\bar{\Delta}/\delta^2 \log n$. Thus, the aforementioned bound, in fact, holds generally for some large enough absolute constant C . \square

Appendix I. Auxiliary results used in the analysis of ALG3

Lemma 12 (Persistence of heterogeneous consideration sets) *Consider a two-armed bandit with rewards bounded in $[0, 1]$, means $\mu_1 > \mu_2$, and gap $\underline{\Delta} = \mu_1 - \mu_2$. Let $\{X_{i,j} : j = 1, 2, \dots\}$ denote the sequence of rewards collected from arm $i \in \{1, 2\}$ by UCB1 ([Auer et al., 2002](#)). Let \mathcal{Z} be an independently generated standard Gaussian random variable. Let $(N_1(n), N_2(n))$ be the per-arm sample counts under UCB1 up to and including time n . Define*

$$\begin{aligned}
 M_n &:= \min(N_1(n), N_2(n)), \\
 \tau &:= \inf \left\{ n \in \mathbb{N} : \left| \mathcal{Z} + \sum_{j=1}^{M_n} (X_{1,j} - X_{2,j}) \right| < 4\sqrt{M_n \log M_n} \right\}.
 \end{aligned}$$

Then, $\mathbb{P}(\tau = \infty) > \beta_{\underline{\Delta}, 2}$, where $\beta_{\underline{\Delta}, 2}$ is as defined in (2) with $\delta \leftarrow \underline{\Delta}$ and $K \leftarrow 2$.

Lemma 13 (Fast rejection of homogeneous consideration sets) *Consider a two-armed bandit where both arms have equal means. Let $\{X_{i,j} : j = 1, 2, \dots\}$ denote the sequence of rewards collected from arm $i \in \{1, 2\}$ by UCB1 (Auer et al., 2002). Let \mathcal{Z} be an independently generated standard Gaussian random variable. Let $(N_1(n), N_2(n))$ be the per-arm sample counts under UCB1 up to and including time n . Define*

$$M_n := \min(N_1(n), N_2(n)),$$

$$\tau := \inf \left\{ n \in \mathbb{N} : \left| \mathcal{Z} + \sum_{j=1}^{M_n} (X_{1,j} - X_{2,j}) \right| < 4\sqrt{M_n \log M_n} \right\}.$$

Then, there exists an absolute constant C such that $\mathbb{E}\tau \leq C$.

I.1. Proof of Lemma 12

Since the rewards are uniformly bounded in $[0, 1]$, it follows that $N_i(n) \rightarrow \infty$ for each arm $i \in \{1, 2\}$ as $n \rightarrow \infty$ on every sample-path. This is due to the structure of the upper confidence bounds used by UCB1. Consequently, $M_n = \min(N_1(n), N_2(n)) \rightarrow \infty$ as $n \rightarrow \infty$ on every sample-path. Also note that M_n is a weakly increasing integer-valued process (starting from 1) with unit increments, wherever they exist. Thus, it follows on every sample-path that τ , in fact, weakly dominates the stopping time τ' defined below

$$\tau' := \inf \left\{ m \in \mathbb{N} : \left| \mathcal{Z} + \sum_{j=1}^m (X_{1,j} - X_{2,j}) \right| < 4\sqrt{m \log m} \right\}. \quad (16)$$

Therefore, $\mathbb{P}(\tau = \infty) \geq \mathbb{P}(\tau' = \infty) > \beta_{\underline{\Delta}, 2}$, where the last inequality follows from Proposition 6 with $\delta \leftarrow \underline{\Delta}$ and $K \leftarrow 2$. \square

I.2. Proof of Lemma 13

Note that

$$\begin{aligned}
 \mathbb{E}\tau &= 1 + \sum_{k \geq 2} \mathbb{P}(\tau \geq k) \\
 &= 1 + \sum_{k \geq 2} \mathbb{P} \left(\bigcap_{n=1}^{k-1} \left\{ \left| \mathcal{Z} + \sum_{j=1}^{M_n} (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{M_n \log M_n} \right\} \right) \\
 &\leq 1 + \sum_{k \geq 1} \mathbb{P} \left(\left| \mathcal{Z} + \sum_{j=1}^{M_k} (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{M_k \log M_k} \right) \\
 &= 1 + \sum_{k \geq 1} \sum_{m=1}^k \mathbb{P} \left(\left| \mathcal{Z} + \sum_{j=1}^{M_k} (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{M_k \log M_k}, N_1(k) = m \right) \\
 &= 1 + \sum_{k \geq 1} \sum_{1 \leq m \leq k/2} \mathbb{P} \left(\left| \mathcal{Z} + \sum_{j=1}^m (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{m \log m}, N_1(k) = m \right) \\
 &\quad + \sum_{k \geq 1} \sum_{k/2 < m \leq k} \mathbb{P} \left(\left| \mathcal{Z} + \sum_{j=1}^{(k-m)} (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{(k-m) \log(k-m)}, N_1(k) = m \right) \\
 &= 1 + \sum_{k \geq 1} \sum_{1 \leq m \leq k/2} \mathbb{P} \left(\left| \mathcal{Z} + \sum_{j=1}^m (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{m \log m}, N_1(k) = m \right) \\
 &\quad + \sum_{k \geq 1} \sum_{1 \leq m < k/2} \mathbb{P} \left(\left| \mathcal{Z} + \sum_{j=1}^m (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{m \log m}, N_2(k) = m \right) \\
 &\leq 1 + 2 \sum_{k \geq 1} \sum_{\theta k \leq m \leq k/2} \mathbb{P} \left(\left| \mathcal{Z} + \sum_{j=1}^m (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{m \log m} \right) \\
 &\quad + \sum_{k \geq 1} [\mathbb{P}(N_1(k) \leq \theta k) + \mathbb{P}(N_2(k) \leq \theta k)],
 \end{aligned}$$

where $\theta = 1/2 - \sqrt{15}/8$. Using Theorem 4(i) of [Kalvit and Zeevi \(2020\)](#) with $\epsilon = \sqrt{15}/8$, one obtains

$$\begin{aligned}
 \mathbb{E}\tau &\leq 1 + 2 \sum_{k \geq 1} \sum_{\theta k \leq m \leq k/2} \mathbb{P} \left(\left| \mathcal{Z} + \sum_{j=1}^m (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{m \log m} \right) + 16 \sum_{k \geq 1} \frac{1}{k^2} \\
 &\leq 1 + 4 \sum_{k \geq 1} \sum_{\theta k \leq m \leq k/2} \frac{1}{m^4} + 16 \sum_{k \geq 1} \frac{1}{k^2} \\
 &\leq 1 + \frac{4}{\theta^4} \sum_{k \geq 1} \frac{1}{k^3} + 16 \sum_{k \geq 1} \frac{1}{k^2}.
 \end{aligned}$$

□

Appendix J. Proof of Theorem 8

Consider the first epoch and define the following:

$$M_n := \min(N_1(n), N_2(n)),$$

$$\tau := \inf \left\{ n \in \mathbb{N} : \left| \mathcal{Z} + \sum_{j=1}^{M_n} (X_{1,j} - X_{2,j}) \right| < 4\sqrt{M_n \log M_n} \right\},$$

where $(N_1(n), N_2(n))$ are the per-arm sample counts under UCB1 up to and including time n . Note that τ marks the termination of epoch 1.

Let R_n denote the cumulative pseudo-regret of ALG3 after n pulls (superscript π suppressed for notational convenience). Let S_n denote the cumulative pseudo-regret of UCB1 after n pulls in a two-armed bandit with gap $\underline{\Delta}$. Let D and I respectively denote the events that the two arms queried in epoch 1 have distinct and identical types. Similarly, let OPT and INF respectively denote the events that the two arms have ‘‘optimal’’ and ‘‘inferior’’ types (Note that $\text{I} = \text{OPT} \cup \text{INF}$). Let $\tilde{R}_n, \bar{R}_n, \hat{R}_n$ be independently drawn from the same distribution as R_n . Then, note that R_n admits the following stochastic evolution:

$$\begin{aligned} R_n &= \mathbb{1}\{\text{D}\} \left[S_{\min(\tau, n)} + \tilde{R}_{(n-\tau)^+} \right] + \mathbb{1}\{\text{INF}\} \left[\underline{\Delta} \min(\tau, n) + \bar{R}_{(n-\tau)^+} \right] + \mathbb{1}\{\text{OPT}\} \hat{R}_{(n-\tau)^+} \\ &\leq \mathbb{1}\{\text{D}\} \left[S_n + \mathbb{1}\{\tau < n\} \tilde{R}_n \right] + \mathbb{1}\{\text{INF}\} \left[\underline{\Delta} \tau + \bar{R}_n \right] + \mathbb{1}\{\text{OPT}\} \hat{R}_n \\ &\leq \mathbb{1}\{\text{D}\} \left[S_n + \mathbb{1}\{\tau < \infty\} \tilde{R}_n \right] + \mathbb{1}\{\text{INF}\} \left[\underline{\Delta} \tau + \bar{R}_n \right] + \mathbb{1}\{\text{OPT}\} \hat{R}_n, \end{aligned}$$

where the first inequality follows since ALG3 is agnostic to n , and hence the pseudo-regret R_n is weakly increasing in n . Taking expectations on both sides, one recovers using the independence of $\tilde{R}_n, \bar{R}_n, \hat{R}_n$ that

$$\begin{aligned} \mathbb{E}R_n &\leq \frac{1}{\mathbb{P}(\tau = \infty | \text{D})} \left[\mathbb{E}S_n + \left(\frac{1 - \alpha_1}{2\alpha_1} \right) \underline{\Delta} \mathbb{E}[\tau | \text{INF}] \right] \\ &\leq \frac{1}{\beta_{\underline{\Delta}, 2}} \left[\mathbb{E}S_n + C \left(\frac{1 - \alpha_1}{2\alpha_1} \right) \underline{\Delta} \right], \end{aligned}$$

where $\beta_{\underline{\Delta}, 2}$ is as defined in (2) with $\delta \leftarrow \underline{\Delta}$ and $K \leftarrow 2$, and C is some absolute constant; the last inequality follows using Lemma 12 and 13. The stated assertion now follows since $\mathbb{E}S_n \leq C'(\log n / \underline{\Delta} + \underline{\Delta})$ for some absolute constant C' (Auer et al., 2002). \square

Appendix K. Auxiliary results used in the analysis of ALG4

Lemma 14 (Persistence of heterogeneous consideration sets) *Consider a K -armed bandit with rewards bounded in $[0, 1]$ and means $\mu_1 > \dots > \mu_K$. Let $\{X_{a,j} : j = 1, 2, \dots\}$ denote the rewards collected from arm $a \in \{1, \dots, K\} =: \mathcal{A}$ by UCB1 (Auer et al., 2002). Let $\{\mathcal{Z}_{a,b} : a, b \in \mathcal{A}, a < b\}$*

be a collection of $\binom{K}{2}$ independent standard Gaussian random variables. Let $N_a(n)$ be the sample count of arm a under UCB1 until time n . Define

$$M_l := \min_{a \in \mathcal{A}} N_a(l),$$

$$\tau := \inf \left\{ l \geq K : \exists a, b \in \mathcal{A}, a < b \text{ s.t. } \left| \mathcal{Z}_{a,b} + \sum_{j=1}^{M_l} (X_{a,j} - X_{b,j}) \right| < 4\sqrt{M_l \log M_l} \right\}.$$

Then, $\mathbb{P}(\tau = \infty) > \beta_{\delta, K}$, where $\beta_{\delta, K}$ is as defined in (2).

Lemma 15 (Path-wise lower bound on the arm-sampling rate of UCB1) Consider a K -armed bandit with rewards bounded in $[0, 1]$. Let $N_a(n)$ be the sample count of arm $a \in \{1, \dots, K\} =: \mathcal{A}$ under UCB1 (Auer et al., 2002) until time n . Then, for all $n \geq K$,

$$M_n := \min_{a \in \mathcal{A}} N_a(n) \geq f(n),$$

where $(f(n) : n = K, K + 1, \dots)$ is some deterministic monotone non-decreasing integer-valued sequence satisfying $f(K) = 1$ and $f(n) \rightarrow \infty$ as $n \rightarrow \infty$.

K.1. Proof of Lemma 14

Suppose that there exists a sample-path on which some non-empty subset of arms $\mathfrak{A} \subset \mathcal{A}$ receives a bounded number of pulls asymptotically in the horizon of play. Also suppose that \mathfrak{A} is the maximal such subset, i.e., each arm in $\mathcal{A} \setminus \mathfrak{A}$ is played infinitely often asymptotically on said sample-path. This implies that the UCB score of any arm in $\mathcal{A} \setminus \mathfrak{A}$ is at most $1 + o(\sqrt{\log t})$ at time t (since the empirical mean term remains bounded in $[0, 1]$). At the same time, the boundedness hypothesis implies that the UCB score of any arm in \mathfrak{A} is at least $\Omega(\sqrt{\log t})$. Thus, for t large enough, UCB scores of arms in \mathfrak{A} will start to dominate those in $\mathcal{A} \setminus \mathfrak{A}$ and the algorithm will end up playing an arm from \mathfrak{A} at some point, thus increasing the cumulative sample-count of arms in \mathfrak{A} by 1. As t grows further, one can replicate the preceding argument an arbitrary number of times to conclude that \mathfrak{A} receives an unbounded number of pulls on the sample-path under consideration, thereby contradicting the boundedness hypothesis. Therefore, it must be the case that each arm in \mathcal{A} is played infinitely often on every sample-path. Consequently, $M_n = \min_{a \in \mathcal{A}} N_a(n) \rightarrow \infty$ as $n \rightarrow \infty$ on every sample-path.

Now since $(M_n : n = K, K + 1, \dots)$ is an integer-valued process (starting from $M_K = 1$) with unit increments (wherever they exist), it follows that on every sample-path, τ , in fact, weakly dominates the stopping time τ' given by

$$\tau' := \inf \left\{ m \in \mathbb{N} : \exists a, b \in \mathcal{A}, a < b \text{ s.t. } \left| \mathcal{Z}_{a,b} + \sum_{j=1}^m (X_{a,j} - X_{b,j}) \right| < 4\sqrt{m \log m} \right\}. \quad (17)$$

Therefore, $\mathbb{P}(\tau = \infty) \geq \mathbb{P}(\tau' = \infty) > \beta_{\delta, K}$; the last inequality follows from Proposition 6. \square

K.2. Proof of Lemma 15

Suppose that $\mathcal{S}_n = \{(N_a(n) : a \in \mathcal{A})\}$ denotes the set of possible sample-count realizations under UCB1 when the horizon of play is n . Define $f(n) := \min_{(N_a(n) : a \in \mathcal{A}) \in \mathcal{S}_n} \min_{a \in \mathcal{A}} N_a(n)$.

Since \mathcal{S}_n is finite, aforementioned minimum is attained at some $(N_a^*(n) : a \in \mathcal{A}) \in \mathcal{S}_n$. Note that $(N_a^*(n) : a \in \mathcal{A})$ is not a random vector as it corresponds to a specific set of sample-paths (possibly non-unique) on which $\min_{a \in \mathcal{A}} N_a(n)$ is minimized. Therefore, $f(n) = \min_{a \in \mathcal{A}} N_a^*(n)$ is deterministic. We have already established in the proof of Lemma 14 that for each $a \in \mathcal{A}$, $N_a(n) \rightarrow \infty$ as $n \rightarrow \infty$ on every sample-path. In particular, this also implies $N_a^*(n) \rightarrow \infty$ as $n \rightarrow \infty$. Thus, we have established the existence of a sequence $f(n)$ satisfying the assertions of the lemma. \square

Appendix L. Proof of Theorem 9

Let $\mathcal{A} := \{1, \dots, K\}$ be the collection of K arms queried during the first epoch. Consider an arbitrary $l \in \mathbb{N}$ s.t. $l \geq K$ and define the following:

$$M_l := \min_{a \in \mathcal{A}} N_a(l),$$

$$\tau := \inf \left\{ l \geq K : \exists a, b \in \mathcal{A}, a < b \text{ s.t. } \left| Z_{a,b} + \sum_{j=1}^{M_l} (X_{a,j} - X_{b,j}) \right| < 4\sqrt{M_l \log M_l} \right\},$$

where $N_a(l)$ denotes the sample count from arm a under UCB1 until time l . Note that τ marks the termination of epoch 1.

Let R_n denote the cumulative pseudo-regret of ALG4 after n pulls (superscript π suppressed for notational convenience). Let S_n denote the cumulative pseudo-regret of UCB1 after n pulls in a K -armed bandit with means $\mu_1 > \mu_2 > \dots > \mu_K$. Let \mathbb{D} denote the event that the K arms queried in epoch 1 have distinct types (no two belong to the same type). Let \tilde{R}_n, \bar{R}_n be independently drawn from the same distribution as R_n . Then, the evolution of R_n satisfies

$$\begin{aligned} R_n &\leq \mathbb{1}\{\mathbb{D}\} \left[S_{\min(\tau, n)} + \tilde{R}_{(n-\tau)^+} \right] + \mathbb{1}\{\mathbb{D}^c\} \left[\bar{\Delta} \min(\tau, n) + \bar{R}_{(n-\tau)^+} \right] \\ &\stackrel{(\dagger)}{\leq} \mathbb{1}\{\mathbb{D}\} \left[S_n + \mathbb{1}\{\tau < n\} \tilde{R}_n \right] + \mathbb{1}\{\mathbb{D}^c\} \left[\bar{\Delta} \min(\tau, n) + \bar{R}_n \right] \\ &\leq \mathbb{1}\{\mathbb{D}\} \left[S_n + \mathbb{1}\{\tau < \infty\} \tilde{R}_n \right] + \mathbb{1}\{\mathbb{D}^c\} \left[\bar{\Delta} \min(\tau, n) + \bar{R}_n \right], \end{aligned}$$

where (\dagger) follows since ALG4 is agnostic to n , and hence the pseudo-regret R_n is weakly increasing in n . Taking expectations on both sides, one recovers using the independence of \tilde{R}_n, \bar{R}_n that

$$\mathbb{E}R_n \leq \frac{1}{\mathbb{P}(\tau = \infty | \mathbb{D})} \left[\mathbb{E}S_n + \left(\frac{\bar{\Delta} \mathbb{E}[\min(\tau, n) | \mathbb{D}^c]}{\mathbb{P}(\mathbb{D})} \right) \right] \leq \frac{1}{\beta_{\delta, K}} \left[\mathbb{E}S_n + \left(\frac{\bar{\Delta} \mathbb{E}[\min(\tau, n) | \mathbb{D}^c]}{\mathbb{P}(\mathbb{D})} \right) \right],$$

where $\mathbb{P}(\mathbb{D}) = K! \prod_{i=1}^K \alpha_i$, and the last inequality follows using Lemma 14 with $\beta_{\delta, K}$ as defined in (2). We know that $\mathbb{E}S_n \leq CK (\log n / \underline{\Delta} + \bar{\Delta})$ for some absolute constant C (Auer et al., 2002). The rest of the proof is geared towards showing that $\mathbb{E}[\min(\tau, n) | \mathbb{D}^c] = o(n)$.

L.1. Proof of $\mathbb{E}[\min(\tau, n) | \mathcal{D}^c] = o(n)$

Let $\mathbb{P}_{\mathcal{D}^c}(\cdot) := \mathbb{P}(\cdot | \mathcal{D}^c)$ be the conditional measure w.r.t. the event \mathcal{D}^c . On \mathcal{D}^c , there exist two arms in the consideration set \mathcal{A} that belong to the same type. Without loss of generality, suppose that these are indexed by 1, 2. Then,

$$\begin{aligned}
 \mathbb{E}[\min(\tau, n) | \mathcal{D}^c] &\leq K + \sum_{k=K+1}^n \mathbb{P}_{\mathcal{D}^c}(\tau \geq k) \\
 &\leq K + \sum_{k=K}^n \mathbb{P}_{\mathcal{D}^c}(\tau \geq k+1) \\
 &= K + \sum_{k=K}^n \mathbb{P}_{\mathcal{D}^c} \left(\bigcap_{l=1}^k \bigcap_{a, b \in \mathcal{A}, a < b} \left\{ \left| \mathcal{Z}_{a,b} + \sum_{j=1}^{M_l} (X_{a,j} - X_{b,j}) \right| \geq 4\sqrt{M_l \log M_l} \right\} \right) \\
 &\leq K + \sum_{k=K}^n \mathbb{P}_{\mathcal{D}^c} \left(\left| \mathcal{Z}_{1,2} + \sum_{j=1}^{M_k} (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{M_k \log M_k} \right) \\
 &= K + \sum_{k=K}^n \sum_{m=1}^k \mathbb{P}_{\mathcal{D}^c} \left(\left| \mathcal{Z}_{1,2} + \sum_{j=1}^{M_k} (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{M_k \log M_k}, M_k = m \right) \\
 &\stackrel{(\dagger)}{=} K + \sum_{k=K}^n \sum_{m=f(k)}^k \mathbb{P}_{\mathcal{D}^c} \left(\left| \mathcal{Z}_{1,2} + \sum_{j=1}^{M_k} (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{M_k \log M_k}, M_k = m \right) \\
 &\leq K + \sum_{k=K}^n \sum_{m=f(k)}^k \mathbb{P}_{\mathcal{D}^c} \left(\left| \mathcal{Z}_{1,2} + \sum_{j=1}^m (X_{1,j} - X_{2,j}) \right| \geq 4\sqrt{m \log m} \right) \\
 &\stackrel{(\ddagger)}{\leq} K + 2 \sum_{k=K}^n \sum_{m=f(k)}^k \frac{1}{m^4} \\
 &= K + 2 \sum_{k=K}^n \left(\frac{1}{(f(k))^4} + \sum_{m=f(k)+1}^k \frac{1}{m^4} \right) \\
 &\leq K + 2 \sum_{k=K}^n \left(\frac{1}{(f(k))^4} + \frac{1}{3(f(k))^3} \right),
 \end{aligned}$$

where (\dagger) follows from Lemma 15, and (\ddagger) using the Chernoff-Hoeffding bound (Hoeffding, 1963). Since $f(k)$ is monotone non-decreasing and coercive in k , it follows that $\mathbb{E}[\min(\tau, n) | \mathcal{D}^c] = o(n)$, where the little-Oh only hides dependence on K . \square