
A Closer Look at the Worst-case Behavior of Multi-armed Bandit Algorithms

Anand Kalvit¹ and Assaf Zeevi²
Columbia University
New York, USA
{¹akalvit22,²assaf}@gsb.columbia.edu

Abstract

One of the key drivers of complexity in the classical (stochastic) multi-armed bandit (MAB) problem is the difference between mean rewards in the top two arms, also known as the instance gap. The celebrated Upper Confidence Bound (UCB) policy is among the simplest optimism-based MAB algorithms that naturally adapts to this gap: for a horizon of play n , it achieves optimal $\mathcal{O}(\log n)$ regret in instances with “large” gaps, and a near-optimal $\mathcal{O}(\sqrt{n \log n})$ minimax regret when the gap can be arbitrarily “small.” This paper provides new results on the *arm-sampling* behavior of UCB, leading to several important insights. Among these, it is shown that arm-sampling rates under UCB are asymptotically deterministic, *regardless* of the problem complexity. This discovery facilitates new sharp asymptotics and a novel alternative proof for the $\mathcal{O}(\sqrt{n \log n})$ minimax regret of UCB. Furthermore, the paper also provides the first complete process-level characterization of the MAB problem under UCB in the conventional *diffusion scaling*. Among other things, the “small” gap worst-case lens adopted in this paper also reveals profound distinctions between the behavior of UCB and Thompson Sampling, such as an *incomplete learning* phenomenon characteristic of the latter.

1 Introduction

Background and motivation. The MAB paradigm provides a succinct abstraction of the quintessential *exploration* vs. *exploitation* trade-offs inherent in many sequential decision making problems. This has origins in clinical trial studies dating back to 1933 [30] which gave rise to the earliest known MAB heuristic, Thompson Sampling [1]. Today, the MAB problem manifests itself in various forms with applications ranging from dynamic pricing and online auctions to packet routing, scheduling, e-commerce and matching markets among others (see [9] for a comprehensive survey of different formulations). In the canonical stochastic MAB problem, a decision maker (DM) pulls one of K arms sequentially at each time $t \in \{1, 2, \dots\}$, and receives a random payoff drawn according to an arm-dependent distribution. The DM, oblivious to the statistical properties of the arms, must balance exploring new arms and exploiting the best arm played thus far in order to maximize her cumulative payoff over the horizon of play. This objective is equivalent to minimizing the *regret* relative to an oracle with perfect ex ante knowledge of the optimal arm (the one with the highest mean reward). The classical stochastic MAB problem is fully specified by the tuple $((\mathcal{P}_i)_{1 \leq i \leq K}, n)$, where \mathcal{P}_i denotes the distribution of rewards associated with the i^{th} arm, and n the horizon of play.

The statistical complexity of regret minimization in the stochastic MAB problem is governed by a key primitive called the *gap*, denoted by Δ , which accounts for the difference between the top two arm mean rewards in the problem. For a “well-separated” or “large gap” instance, i.e., a fixed Δ bounded away from 0, the seminal paper [22] showed that the order of the smallest achievable regret

is logarithmic in the horizon. There has been a plethora of subsequent work involving algorithms which can be fine-tuned to achieve a regret arbitrarily close to the optimal rate discovered in [22] (see [6, 14, 13, 2, 4], etc., for a few notable examples). On the other hand, no algorithm can achieve an expected regret smaller than $C\sqrt{n}$ for a fixed n (the constant hides dependence on the number of arms) *uniformly* over all problem instances (also called *minimax* regret); see, e.g., [23], Chapter 15. The *saddle-point* in this minimax formulation occurs at a gap that satisfies $\Delta \asymp 1/\sqrt{n}$. This has a natural interpretation: approximately $1/\Delta^2$ samples are required to distinguish between two distributions with means separated by Δ ; at the $1/\sqrt{n}$ -scale, it becomes statistically impossible to distinguish between samples from the top two arms within n rounds of play. If the gap is smaller, despite the increased difficulty in the hypothesis test, the problem becomes “easier” from a regret perspective. Thus, $\Delta \asymp 1/\sqrt{n}$ is the statistically “hardest” scale for regret minimization. A number of popular algorithms achieve the \sqrt{n} minimax-optimal rate (modulo constants), see, e.g., [4, 2], and many more do this within poly-logarithmic factors in n . Many of these are variations of the celebrated upper confidence bound algorithms, e.g., UCB1 [7], that achieve a minimax regret of $\mathcal{O}(\sqrt{n \log n})$, and at the same time also deliver the logarithmic regret achievable in the instance-dependent setting of [22].

A major driver of the regret performance of an algorithm is its arm-sampling characteristics. For example, in the instance-dependent (large gap) setting, optimal regret guarantees imply that the fraction of time the optimal arm(s) are played approaches 1 in probability, as n grows large. However, this fails to provide any meaningful insights as to the distribution of arm-pulls for smaller gaps, e.g., the $\Delta \asymp 1/\sqrt{n}$ “small gap” that governs the “worst-case” instance-independent setting.

An illustrative numerical example involving “small gap.” Consider an A/B testing problem (e.g., a vaccine clinical trial) where the experimenter is faced with two competing objectives: first, to estimate the efficacy of each alternative with the best possible precision given a budget of samples, and second, keeping the overall cost of the experiment low. This is a fundamentally hard task and algorithms incurring a low cumulative cost typically spend little time exploring sub-optimal alternatives, resulting in a degraded estimation precision (see, e.g., [5]). In other words, algorithms tailored for (cumulative) regret minimization may lack *statistical power* [31]. While this trade-off is unavoidable in “well-separated” instances, numerical evidence suggests a plausible resolution in instances with “small” gaps as illustrated below. For example, such a situation might arise in trials conducted using two similarly efficacious vaccines (abstracted away as $\Delta \approx 0$). To illustrate the point more vividly, consider the case where Δ is exactly 0 (of course, this information is not known to the experimenter). This setting is numerically illustrated in Figure 1, which shows the empirical distribution of $N_1(n)/n$ (the fraction of time arm 1 is played until time n) in a two-armed bandit with $\Delta = 0$, under two different algorithms (UCB and Thompson Sampling), and reward configurations.

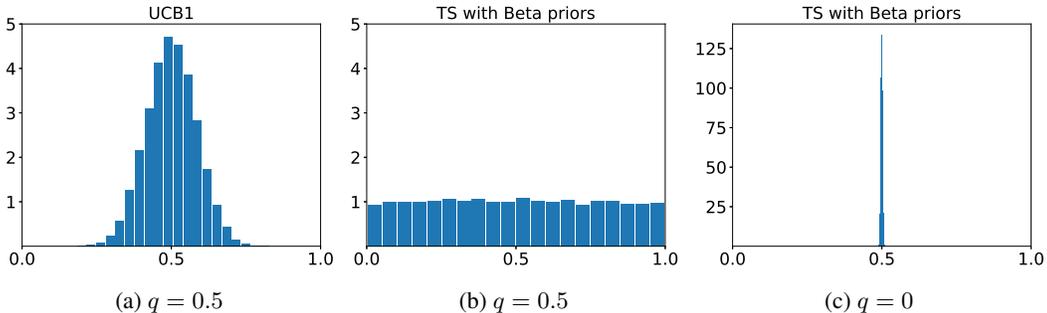


Figure 1: **Incomplete learning under Thompson Sampling.** A two-armed bandit with arms having $\text{Bernoulli}(q)$ rewards each: Histograms show the empirical (probability) distribution of $N_1(n)/n$ for $n = 10,000$ pulls, plotted using 20,000 experiments. Algorithms: UCB1 [7], TS with Beta priors [1].

A desirable property of the outcome in this setting is to have a *linear allocation* of the sampling budget *per arm* on almost every sample-path of the algorithm, as this leads to “complete learning:” an algorithm’s ability to discern statistical indistinguishability of the arm-means, and induce a “balanced” allocation in that event. However, despite the simplicity of the zero-gap scenario, it is far from obvious whether the aforementioned property may be satisfied for standard bandit algorithms such as UCB and Thompson Sampling. Indeed, Figure 1 exhibits a striking difference between the two. The concentration around 1/2 observable in Figure 1(a) indicates that UCB results in an approximately “balanced” sample-split, i.e., the allocation is roughly $n/2$ per arm for large n

(and this is observed for “most” sample-paths). In fact, we will later see that the “bell curve” in Figure 1(a) eventually collapses into the Dirac measure at $1/2$ (Theorem 1). On the other hand, under Thompson Sampling, the allocation of samples across arms *may* be arbitrarily “imbalanced” despite the arms being statistically identical, as seen in Figure 1(b) (see, for contrast, Figure 1(c), where the allocation is perfectly “balanced”). Namely, the distribution of the posterior *may* be such that arm 1 is allocated anywhere from almost no sampling effort all the way to receiving almost the entire sampling budget, as Figure 1(b) suggests. Non-degeneracy of arm-sampling rates is observable also under the more widely used version of the algorithm that is based on Gaussian priors and Gaussian likelihoods (Algorithm 2 in [2]); see Figure 2(a). Such behavior can be detrimental for ex post causal inference in the general A/B testing context, and the vaccine testing problem referenced earlier. This is demonstrated via an instructional example of a two-armed bandit with one deterministic reference arm (aka the “one-armed” bandit paradigm), illustrated in Figure 2, and discussed below.

A numerical example illustrating inference implications. Consider a model where arm 1 returns a constant reward of 0.5, while arm 2 yields rewards distributed as $\text{Bernoulli}(0.5)$. In this setup, the estimate of the gap Δ (*average treatment effect* in causal inference parlance) after n rounds of play is given by $\hat{\Delta} = \bar{X}_2(n) - 0.5$, where $\bar{X}_2(n)$ denotes the empirical mean reward of arm 2 at time n . The \mathcal{Z} statistic associated with this gap estimator is given by $\mathcal{Z} = 2\sqrt{N_2(n)}\hat{\Delta}$, where $N_2(n)$ is the visitation count of arm 2 at time n . In the absence of any sample-adaptivity in the arm 2 data, results from classical statistics such as the Central Limit Theorem (CLT) would posit an asymptotically Normal distribution for \mathcal{Z} . However, since the algorithms that play the arms are adaptive in nature, e.g., UCB and Thompson Sampling, asymptotic-normality may no longer be guaranteed. Indeed, the numerical evidence in Figure 2(b) strongly points to a significant departure from asymptotic-normality of the \mathcal{Z} statistic associated with the gap estimator under Thompson Sampling (TS). Non-normality of the \mathcal{Z} statistic can be problematic for inferential tasks, e.g., it can lead to statistically unsupported inferences in the binary hypothesis test $\mathcal{H}_0 : \Delta = 0$ vs. $\mathcal{H}_1 : \Delta \neq 0$ performed using confidence intervals constructed as per the conventional CLT approximation. In sharp contrast, our work shows that UCB satisfies a certain “balanced” sampling property (such as that in Figure 1(a)) in instances with “small” gaps, formally stated as Theorem 1, that drives the \mathcal{Z} statistic towards asymptotic-normality in the aforementioned binary hypothesis testing example (asymptotic-normality being a consequence of Theorem 4). Furthermore, since the \sqrt{n} -normalized “stochastic” regret (defined in (1) in §2) equals $-\left(\sqrt{N_2(n)}/(4n)\right)\mathcal{Z}$, it follows that this too, satisfies asymptotic-normality under UCB (due to Theorem 4, in conjunction with Theorem 1). These properties are evident in Figure 2(c) below, and signal reliability of ex post causal inference (under classical assumptions like validity of CLT) from “small gap” data collected by UCB vis-à-vis Thompson Sampling (TS). The reliability of inference under TS may be doubtful even in the limit of infinite data, as Figure 2(b) suggests.

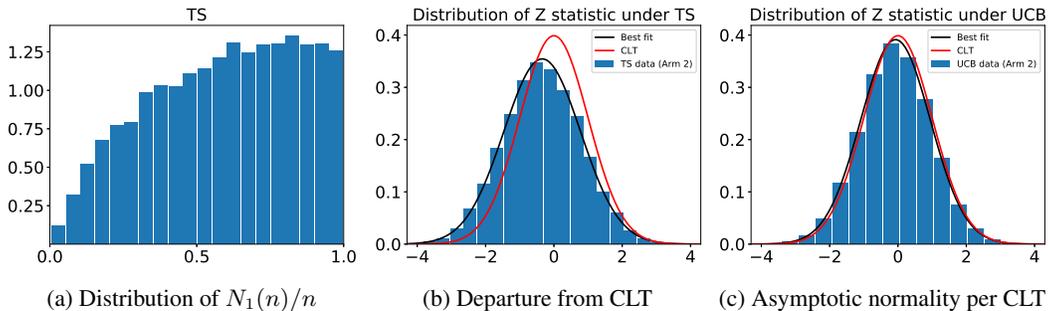


Figure 2: **Failure of CLT under TS.** A two-armed bandit with $\Delta = 0$: Arm 1 returns a constant reward of 0.5, and arm 2 yields rewards distributed as $\text{Bernoulli}(0.5)$. In (a), the histogram shows the empirical (probability) distribution of $N_1(n)/n$. Algorithms: TS (Algorithm 2 in [2]) and UCB (UCB1 in [7]). All histograms have $n = 10,000$ pulls, and are plotted using $\aleph = 20,000$ experiments.

While traditional literature on stochastic bandits is dedicated primarily to the regret minimization problem, there has been significant recent interest also in finer-grain properties of popular “adaptive” MAB algorithms such as UCB and Thompson Sampling. For example, a recent line of work ([27, 28, 29]) investigates the “bias” of optimistic algorithms like UCB. The focus of our work is on understanding the distribution of arm-pulls, which as discussed earlier, has significant bearings on

ex post causal inference from data collected adaptively by bandit algorithms (see, e.g., [33, 16, 32], etc., and references therein for recent developments), algorithmic fairness in the broader context of fairness in machine learning (see [25] for a survey), as well as on novel formulations of the MAB problem such as [20]. Below, we discuss extant literature relevant to our line of work.

Previous work. The study of “well-separated” instances, or the *large gap* regime, is supported by rich literature. For example, [6] provides high-probability bounds on arm-sampling rates under a parametric family of UCB algorithms. However, as the gap diminishes, leading to the so called *small gap* regime, the aforementioned bounds become vacuous. The understanding of arm-sampling behavior remains relatively under-studied here even for popular algorithms such as UCB and Thompson Sampling. This regime is of special interest in that it also covers the classical *diffusion scaling*¹, where $\Delta \asymp 1/\sqrt{n}$, which as discussed earlier, corresponds to instances that statistically constitute the “worst-case” for hypothesis testing and regret minimization. Recently, a partial diffusion-limit characterization of the arm-sampling distribution under a version of Thompson Sampling with horizon-dependent prior variances², was provided in [32] as a solution to a certain stochastic differential equation (SDE). The numerical solution to said SDE was observed to have a non-degenerate distribution on $[0, 1]$. Similar numerical observations on non-degeneracy of the arm-sampling distribution also under standard versions of Thompson Sampling were reported in [10, 20], among others, albeit limited only to the special case of $\Delta = 0$, and absent a theoretical explanation for the aforementioned observations. Thus, outside of the so called “easy” problems, where Δ is bounded away from 0 by an absolute constant, theoretical understanding of the sampling behavior of bandit algorithms remains an open area of research.

Contributions. In this paper, we provide the first complete asymptotic characterization of arm-sampling distributions under canonical UCB (Algorithm 1) as a function of the gap Δ (Theorem 1). This gives rise to a fundamental insight: arm-sampling rates are asymptotically deterministic under UCB regardless of the hardness of the instance. We also provide the first theoretical explanation for an “incomplete learning” phenomenon under Thompson Sampling (Algorithm 2) alluded to in Figure 1, as well as a sharp dichotomy between Thompson Sampling and UCB evident therein (Theorem 2). This result earmarks an “instability” of Thompson Sampling in terms of the limiting arm-sampling distribution. As a sequel to Theorem 1, we provide the first *complete* characterization of the worst-case performance of canonical UCB (Theorem 3). One consequence is that the $\mathcal{O}(\sqrt{n \log n})$ minimax regret of UCB is strictly unimprovable in a precise sense. Moreover, our work also leads to the first process-level characterization of the two-armed bandit problem under canonical UCB in the classical *diffusion limit*, according to which a suitably normalized cumulative reward process converges in law to a Brownian motion with fully characterized drift and infinitesimal variance (Theorem 4). To the best of our knowledge, this is the first such characterization of UCB-type algorithms. Theorem 4 facilitates a complete distribution-level characterization of UCB’s diffusion-limit regret, thereby providing sharp insights as to the problem’s minimax complexity. Such distribution-level information may also be useful for a variety of inferential tasks, e.g., construction of confidence intervals (see the binary hypothesis testing example referenced in Figure 2(c)), among others. We believe our results may also present new design considerations, in particular, how to achieve, loosely speaking, the “best of both worlds” for Thompson Sampling, by addressing its “small gap” instability. Lastly, we note that our proof techniques are markedly different from the conventional methodology adopted in MAB literature ([6, 9, 2]), and may be of independent interest in the study of related learning algorithms.

Organization of the paper. A formal description of the model and the canonical UCB algorithm is provided in §2. All theoretical propositions are stated in §3, along with a high-level overview of their scope and proof sketch; detailed proofs and ancillary results are relegated to the supplementary material. Finally, concluding remarks and open problems are presented in §4.

2 The model and notation

The technical development in this paper will focus on the two-armed problem purely for expositional reasons; we provide extensions to the general K -armed setting in Appendix B. The restriction to two-armed bandits has precedence also in the literature due to its tractability for sharp asymptotic

¹This is a standard technique for performance evaluation of stochastic systems, commonly used in the operations research and mathematics literature, see, e.g., [15].

²Assumed to be vanishing in n ; standard versions of the algorithm involve fixed (positive) prior variances.

analyses, see, e.g., [21]. This setting encapsulates the core statistical complexity of the MAB problem in the “small gap” regime, as well as concisely highlighting the key novelties in our approach. Before describing the model formally, we introduce the following asymptotic conventions.

Notation. We say $f(n) = o(g(n))$ or $g(n) = \omega(f(n))$ if $\lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} = 0$. Similarly, $f(n) = \mathcal{O}(g(n))$ or $g(n) = \Omega(f(n))$ if $\limsup_{n \rightarrow \infty} \left| \frac{f(n)}{g(n)} \right| \leq C$ for some constant C . If $f(n) = \mathcal{O}(g(n))$ and $f(n) = \Omega(g(n))$ hold simultaneously, we say $f(n) = \Theta(g(n))$, or $f(n) \asymp g(n)$, and we write $f(n) \sim g(n)$ in the special case where $\lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} = 1$. If either sequence $f(n)$ or $g(n)$ is random, and one of the aforementioned ratio conditions holds in probability, we use the subscript p with the corresponding Landau symbol. For example, $f(n) = o_p(g(n))$ if $f(n)/g(n) \xrightarrow{p} 0$ as $n \rightarrow \infty$. Lastly, the notation ‘ \Rightarrow ’ will be used for weak convergence.

The model. The arms are indexed by $\{1, 2\}$. Each arm $i \in \{1, 2\}$ is characterized by a reward distribution \mathcal{P}_i supported on $[0, 1]$ with mean μ_i . The difference between the two mean rewards, aka the *gap*, is given by $\Delta = |\mu_1 - \mu_2|$; as discussed earlier, this captures the *hardness* of an instance. The sequence of rewards associated with the *first* m pulls of arm i is denoted by $(X_{i,j})_{1 \leq j \leq m}$. The rewards are assumed to be i.i.d. in time, and independent across arms.³ The number of pulls of arm i up to (and including) time t is denoted by $N_i(t)$. A policy $\pi := (\pi_t)_{t \in \mathbb{N}}$ is an adapted sequence that prescribes pulling an arm $\pi_t \in \mathcal{S}$ at time t , where \mathcal{S} denotes the probability simplex on $\{1, 2\}$. The natural filtration at time t is given by $\mathcal{F}_t := \sigma \left\{ (\pi_s)_{s \leq t}, \left((X_{i,j})_{j \leq N_i(t)} : i = 1, 2 \right) \right\}$. The stochastic regret of policy π after n plays, denoted by R_n^π , is given by

$$R_n^\pi := \sum_{t=1}^n [\max(\mu_1, \mu_2) - X_{\pi_t, N_{\pi_t}(t)}]. \quad (1)$$

The decision maker is interested in the problem of minimizing the *expected regret*, given by

$$\inf_{\pi \in \Pi} \mathbb{E} R_n^\pi,$$

where Π is the set of policies satisfying the non-anticipation property $\pi_t : \mathcal{F}_{t-1} \rightarrow \mathcal{S}$, $1 \leq t \leq n$, and the expectation is w.r.t. the randomness in reward realizations as well as possible randomness in the policy π . In this paper, we will focus primarily on the canonical UCB policy given by Algorithm 1 below. This policy is parameterized by an exploration coefficient ρ , which controls its arm-exploring rate. The standard UCB1 policy [7] corresponds to Algorithm 1 with $\rho = 2$; the effect of ρ on the expected and high-probability regret bounds of the algorithm is well-documented in [6] for problems with a “large gap.” In what follows, $\bar{X}_i(t-1)$ denotes the empirical mean reward from arm $i \in \{1, 2\}$ at time $t-1$, i.e., $\bar{X}_i(t-1) := \frac{\sum_{j=1}^{N_i(t-1)} X_{i,j}}{N_i(t-1)}$.

Algorithm 1 The canonical UCB policy for two-armed bandits.

- 1: **Input:** Exploration coefficient $\rho \in \mathbb{R}_+$.
 - 2: At $t = 1, 2$, play each arm $i \in \{1, 2\}$ once.
 - 3: **for** $t \in \{3, 4, \dots\}$ **do**
 - 4: Play arm $\pi_t \in \arg \max_{i \in \{1, 2\}} \left(\bar{X}_i(t-1) + \sqrt{\frac{\rho \log(t-1)}{N_i(t-1)}} \right)$.
-

3 Main results

Algorithm 1 is known to achieve $\mathbb{E} R_n^\pi = \mathcal{O}(\log n)$ in the instance-dependent setting, and $\mathbb{E} R_n^\pi = \mathcal{O}(\sqrt{n \log n})$ in the “small gap” minimax setting. The primary focus of this paper is on the distribution of arm-sampling rates, i.e., $N_i(n)/n$, $i \in \{1, 2\}$. Our main results are split across two sub-sections; §3.1 examines the behavior of UCB (Algorithm 1) as well as another popular bandit algorithm, Thompson Sampling (specified in Algorithm 2). §3.2 is dedicated to results on the (stochastic) regret of Algorithm 1 under the $\Delta \asymp \sqrt{(\log n)/n}$ “worst-case” gap and the $\Delta \asymp 1/\sqrt{n}$ “diffusion-scaled” gap.

³These assumptions can be relaxed in the spirit of [7]; our results also extend to sub-Gaussian rewards.

3.1 Asymptotics of arm-sampling rates

Theorem 1 (Arm-sampling rates under UCB) *Let $i^* \in \arg \max \{\mu_i : i = 1, 2\}$ with ties broken arbitrarily. Then, the following results hold for arm i^* as $n \rightarrow \infty$ under Algorithm 1 with $\rho > 1$:*

(I) “Large gap:” *If $\Delta = \omega\left(\sqrt{\frac{\log n}{n}}\right)$, then $N_{i^*}(n)/n \xrightarrow{P} 1$.*

(II) “Small gap:” *If $\Delta = o\left(\sqrt{\frac{\log n}{n}}\right)$, then $N_{i^*}(n)/n \xrightarrow{P} 1/2$.*

(III) “Moderate gap:” *If $\Delta \sim \sqrt{\frac{\theta \log n}{n}}$ for some fixed $\theta \geq 0$, then $N_{i^*}(n)/n \xrightarrow{P} \lambda_\rho^*(\theta)$, where the limit is the unique solution (in λ) to*

$$\frac{1}{\sqrt{1-\lambda}} - \frac{1}{\sqrt{\lambda}} = \sqrt{\frac{\theta}{\rho}}, \quad (2)$$

and is monotone increasing in θ , with $\lambda_\rho^(0) = 1/2$ and $\lambda_\rho^*(\theta) \rightarrow 1$ as $\theta \rightarrow \infty$.*

Remark 1 (Permissible values of ρ in Algorithm 1) *For $\rho > 1$, the expected regret of the policy π given by Algorithm 1 is bounded as $\mathbb{E}R_n^\pi \leq C\rho\left(\frac{\log n}{\Delta} + \frac{\Delta}{\rho-1}\right)$ for some absolute constant $C > 0$; the upper bound becomes vacuous for $\rho \leq 1$ (see [6], Theorem 7). We therefore restrict Theorem 1 to $\rho > 1$ to ensure that $\mathbb{E}R_n^\pi$ remains non-trivially bounded for all Δ .*

Discussion and intuition. Theorem 1 essentially asserts that the sampling rates $N_i(n)/n$, $i \in \{1, 2\}$ are asymptotically deterministic in probability under canonical UCB; Δ only serves to determine the value of the limiting constant. The “moderate” gap regime offers a continuous interpolation from instances with zero gaps to instances with “large” gaps as θ sweeps over \mathbb{R}_+ in that $\lambda_\rho^*(\theta)$ increases monotonically from $1/2$ at $\theta = 0$ to 1 at $\theta = \infty$, consistent with intuition. The special case of $\theta = 0$ is numerically illustrated in Figure 1(a). The tails of $N_{i^*}(n)/n$ decay polynomially fast near the end points of the interval $[0, 1]$ with the best possible rate approaching $\mathcal{O}(n^{-3})$, occurring for $\theta = 0$. However, as $N_{i^*}(n)/n$ approaches its limit, convergence becomes slower and is dominated by fatter $\Theta\left(\sqrt{\frac{\log \log n}{\log n}}\right)$ tails. The behavior in this regime is regulated by the $\mathcal{O}(\sqrt{n \log \log n})$ envelope of the zero-drift random walk process driving the algorithm’s regret (see proof of Theorem 1 for details).

Proof sketch. To provide the most intuitive explanation, we pivot to the special case where the arms have *identical* reward distributions, and in particular, $\Delta = 0$. The natural candidate then for the limit of the empirical sampling rate is $1/2$. On a high level, the proof relies on polynomially decaying bounds in n for ϵ -deviations of the form $\mathbb{P}\left(\left|\frac{N_1(n)}{n} - \frac{1}{2}\right| \geq \epsilon\right)$ derived using the standard trick for bounding the number of pulls of any arm on a given sample-path, to wit, for any $u, n \in \mathbb{N}$, $N_1(n)$ can be bounded above by $u + \sum_{t=u+1}^n \mathbb{1}\{\pi_t = 1, N_1(t-1) \geq u\}$, *path-wise*. Setting $u = \lceil (1/2 + \epsilon)n \rceil$ in this expression, one can subsequently show via an analysis involving careful use of the policy structure together with appropriate Chernoff bounds that with high probability (approaching 1 as $n \rightarrow \infty$), $N_1(n)/n \leq 1/2 + \epsilon_\rho$ for some $\epsilon_\rho \in (0, 1/2)$ that depends only on ρ . An identical result would naturally hold also for the other arm by symmetry arguments, and therefore we arrive at a meta-conclusion that $N_i(n)/n \geq 1/2 - \epsilon_\rho > 0$ for both arms $i \in \{1, 2\}$ with high probability (approaching 1 as $n \rightarrow \infty$). It is noteworthy that said conclusion cannot be arrived at for an arbitrary $\epsilon > 0$ (in place of ϵ_ρ) since the polynomial upper bounds on $\mathbb{P}\left(\left|\frac{N_1(n)}{n} - \frac{1}{2}\right| \geq \epsilon\right)$ derived using the aforementioned path-wise upper bound on $N_1(n)$, become vacuous if u is set “too close” to $n/2$, i.e., if ϵ is “near” 0. Extension to the full generality of $\epsilon > 0$ is achieved via a refined analysis that uses the Law of the Iterated Logarithm (see [11], Theorem 8.5.2), together with the previous meta-conclusion, to obtain fatter $\mathcal{O}\left(\sqrt{\frac{\log \log n}{\log n}}\right)$ tail bounds when ϵ is near 0. Here, it is imperative to point out that the “ $\log n$ ” appearing in the denominator is essentially from the $\sqrt{\rho \log t}$ optimistic bias term of UCB (see Algorithm 1), and therefore the convergence will, as such, hold also for other variations of the policy that have “less aggressive” $\omega(\log \log t)$ exploration functions vis-à-vis $\log t$. However, this will be achieved at the expense of the policy’s expected regret performance,

as noted in Remark 1. We also note that the extremely slow $\mathcal{O}\left(\sqrt{\frac{\log \log n}{\log n}}\right)$ convergence is not an artifact of our analysis, but in fact, supported by the numerical evidence in Figure 1(a), suggestive of a plausible non-convergence (to $1/2$) in the limit. We believe such observations in previous works likely led to incorrect folk conjectures ruling out the existence of a deterministic limit under UCB à la Theorem 1 (see, e.g., [10] and references therein). The proof for a general Δ in the “small” and “moderate” gap regimes is skeletally similar to that for $\Delta = 0$, albeit guessing a candidate limit for $N_{i^*}(n)/n$ is non-trivial; a closed-form expression for $\lambda_\rho^*(\theta)$ is provided in Appendix A. Full details of the proof of Theorem 1 are provided in Appendix D,E,F. \square

Remark 2 (Possible generalizations of Theorem 1) *A simple extension to the K -armed setting is provided in Appendix B as Theorem 5. The behavior of UCB policies is largely governed by their optimistic bias. While Theorem 1 only covers the generic UCB policy with $\sqrt{\rho \log t}$ bias, results of the form $N_{i^*}(n)/n \xrightarrow{P} c_i$ for some $c_i \in (0, 1)$ continue to hold also under smaller $\omega(\sqrt{\rho \log \log t})$ bias (driven by the Law of the Iterated Logarithm). We believe this observation will be useful when examining more complicated UCB-inspired policies such as KL-UCB [13], DMED [18], etc.*

What about Thompson Sampling? Results such as those discussed above for other popular adaptive algorithms like Thompson Sampling⁴ are only arable in “well-separated” instances where $N_{i^*}(n)/n \xrightarrow{P} 1$ as $n \rightarrow \infty$ follows as a trivial consequence of its $\mathcal{O}(\sqrt{n})$ minimax regret bound [2]. For smaller gaps, theoretical understanding of the distribution of arm-pulls under Thompson Sampling remains largely absent even for its most widely-studied variants. In this paper, we provide a first result in this direction: Theorem 2 formalizes a revealing observation for classical Thompson Sampling (Algorithm 2) in instances with zero gap, and elucidates its instability in view of the numerical evidence reported in Figure 1(b) and 1(c). This result also offers an explanation for the sharp contrast with the statistical behavior of canonical UCB (Algorithm 1) à la Theorem 1, also evident from Figure 1(a). In what follows, rewards are assumed to be Bernoulli, and S_i (respectively F_i) counts the number of successes/1’s (respectively failures/0’s) associated with arm $i \in \{1, 2\}$.

Algorithm 2 Thompson Sampling for the two-armed Bernoulli bandit.

- 1: **Initialize:** Number of successes (1’s) and failures (0’s) for each arm $i \in \{1, 2\}$, $(S_i, F_i) = (0, 0)$.
 - 2: **for** $t \in \{1, 2, \dots\}$ **do**
 - 3: Sample for each $i \in \{1, 2\}$, $\mathcal{T}_i \sim \text{Beta}(S_i + 1, F_i + 1)$.
 - 4: Play arm $\pi_t \in \arg \max_{i \in \{1, 2\}} \mathcal{T}_i$ and observe reward $r_t \in \{0, 1\}$.
 - 5: Update success-failure counts: $S_{\pi_t} \leftarrow S_{\pi_t} + r_t$, $F_{\pi_t} \leftarrow F_{\pi_t} + 1 - r_t$.
-

Theorem 2 (Incomplete learning under Thompson Sampling) *In a two-armed model where both arms yield rewards distributed as Bernoulli(q), the following holds under Algorithm 2 as $n \rightarrow \infty$:*

- (I) *If $q = 0$, then $N_1(n)/n \Rightarrow 1/2$.*
- (II) *If $q = 1$, then $N_1(n)/n \Rightarrow \text{Uniform distribution on } [0, 1]$.*

Proof sketch. The proof of Theorem 2 relies on a careful application of two subtle properties of the Beta distribution (Fact 2 and Fact 3), stated and proved in Appendix C,K. For part (I), we invoke symmetry to deduce $\mathbb{E}N_1(n) = n/2$, and use Fact 2 to show that the standard deviation of $N_1(n)$ is sub-linear in n , thus proving the stated assertion in (I). More elaborately, Fact 2 states for the reward configuration in (I) that the probability of playing arm 1 after it has already been played n_1 times, and arm 2 n_2 times, equals $(n_2 + 1)/(n_1 + n_2 + 2)$. This probability is smaller than $1/2$ if $n_1 > n_2$, which provides an intuitive explanation for the fast convergence of $N_1(n)/n$ to $1/2$ observed in Figure 1(c). In fact, we conjecture that the result in (I) holds also with probability 1 based on the aforementioned “self-balancing” property. The conclusion in part (II) hinges on an application of Fact 3 to show the *stronger* result: $N_1(n)$ is *uniformly distributed* over $\{0, 1, \dots, n\}$ for any $n \in \mathbb{N}$. Contrary to Fact 2, Fact 3 states that quite the opposite is true for the reward configuration in (II): the probability of playing arm 1 after it has already been played n_1 times, and arm 2 n_2 times, equals $(n_1 + 1)/(n_1 + n_2 + 2)$, which is greater than $1/2$ when $n_1 > n_2$. That is, the posterior distributions

⁴This is the version based on Gaussian priors and Gaussian likelihoods, not the classical version based on Beta priors and Bernoulli likelihoods which has a minimax regret of $\mathcal{O}(\sqrt{n \log n})$ [2].

evolve in such a way that the algorithm is “deceived” into incorrectly believing one of the arms (arm 2 in this case) to be inferior. This leads to large sojourn times between successive visitations of arm 2 on such a sample-path, thereby resulting in a perpetual “imbalance” in the sample-counts. This provides an intuitive explanation for the non-degeneracy observed in Figure 1(b) and 2(a), which additionally, also indicates that such behavior, in fact, persists also for general (non-deterministic) reward distributions, as well as under the Gaussian prior-based version of the algorithm. Full proof of Theorem 2 is provided in Appendix G. \square

More on “incomplete learning.” The zero-gap setting is a special case of the “small gap” regime where canonical UCB guarantees a $(1/2, 1/2)$ sample-split in probability (Theorem 1). On the other hand, Theorem 2 suggests that second order factors such as the mean signal strength (magnitude of the mean reward) could significantly affect the nature of the resulting sample-split under Thompson Sampling. Note that even though the result only presupposes deterministic 0/1 rewards, the aforementioned claim is, in fact, borne out by the numerical evidence in Figure 1(b) and 1(c). The sampling distribution seemingly flattens rapidly from the Dirac measure at $1/2$ to the Uniform distribution on $[0, 1]$ as the mean rewards move away from 0. This uncertainty in the limiting sampling behavior has non-trivial implications for a variety of application areas of such learning algorithms. For instance, a uniform distribution of arm-sampling rates on $[0, 1]$ indicates that the sample-split could be arbitrarily imbalanced along a sample-path, despite, as in the setting of Theorem 2, the two arms being statistically identical; this phenomenon is typically referred to as “incomplete learning” (see [26, 24] for the original context). Non-degeneracy in the limiting distribution is also observable numerically up to diffusion-scale gaps of $\mathcal{O}(1/\sqrt{n})$ under other versions of Thompson Sampling (see [32] for examples); our focus on the more extreme zero-gap setting simplifies the illustration of these effects.

A brief survey of Thompson Sampling. While extant literature does not provide any explicit result for Thompson Sampling characterizing its arm-sampling behavior in instances with “small” and “moderate” gaps, there has been recent work on its analysis in the $\Delta \asymp 1/\sqrt{n}$ regime under what is known as the *diffusion approximation* lens (see [32, 12]). Cited works study Thompson Sampling primarily under the assumption that the prior variance associated with the mean reward of any arm vanishes in the horizon of play at an “appropriate” rate.⁵ Such a scaling, however, is not ideal from a regret standpoint and indeed, the versions of Thompson Sampling optimized for regret performance use fixed (non-vanishing) prior variances, e.g., Algorithm 2 and its Gaussian prior-based counterpart (see [2]). On a high level, [32, 12] establish that as $n \rightarrow \infty$, the pre-limit $(N_i(nt)/n)_{t \in [0,1]}$ under Thompson Sampling converges weakly to a “diffusion-limit” stochastic process on $t \in [0, 1]$. Recall from earlier discussion that $\Delta \asymp 1/\sqrt{n}$ is covered under the “small gap” regime; consequently, it follows from Theorem 1 that the analogous limit for UCB is, in fact, the deterministic process $t/2$. In sharp contrast, the diffusion-limit process under Thompson Sampling may at best be characterizable only as a solution (possibly non-unique) to an appropriate stochastic differential equation or ordinary differential equation driven by a suitably (random) time-changed Brownian motion. Consequently, the diffusion limit under Thompson Sampling is more difficult to interpret vis-à-vis UCB (see Theorem 4), and it is much harder to obtain lucid insights as to the nature of the distribution of $N_i(n)/n$ as $n \rightarrow \infty$. The asymptotic distribution of $N_i(n)/n$ under Thompson Sampling is also investigated in [19], albeit in a significantly different setting. Cited paper considers the Bayesian setting where a prior distribution exists over problem instances, and the Thompson Sampling algorithm is “well-specified,” i.e., information about said prior is baked into the algorithm. Specifically, a sample path of the algorithm in their model involves a *random* problem instance from the instance-space. In contrast, the derivation of the asymptotic distribution of arm-pulls in our work is for specific (fixed) problem instances, viz., reward configurations (I) and (II) described in Theorem 2, and under the classical version of Thompson Sampling (Algorithm 2).

3.2 Beyond arm-sampling rates

This part of the paper is dedicated to a more fine-grained analysis of the “stochastic” regret of UCB (defined in (1) in §2). Results are largely facilitated by insights on the sampling behavior of UCB in instances with “small” gaps, attributable to Theorem 1; however, we believe they are of interest in their own right. We commence with an application of Theorem 1 which provides the first complete characterization of the worst-case (minimax) performance of canonical UCB. A full diffusion-limit characterization of the two-armed bandit problem under UCB is provided thereafter in Theorem 4.

⁵The only result applicable to the case of non-vanishing prior variances is Theorem 4.2 of [12].

Theorem 3 (Minimax regret complexity of UCB) *In the “moderate gap” regime referenced in Theorem 1 where $\Delta \sim \sqrt{\frac{\theta \log n}{n}}$, the regret of the policy π given by Algorithm 1 with $\rho > 1$ satisfies*

$$\frac{R_n^\pi}{\sqrt{n \log n}} \Rightarrow \sqrt{\theta} (1 - \lambda_\rho^*(\theta)) =: h_\rho(\theta) \text{ as } n \rightarrow \infty, \quad (3)$$

where $\lambda_\rho^*(\theta)$ is the (unique) solution to (2).

To the best of our knowledge, this is the first *algorithm-specific* result (sharp asymptotic) that is distinct from the general $\Omega(\sqrt{n})$ information-theoretic lower bound by a horizon-dependent factor.⁶

Discussion. A closed-form expression for $\lambda_\rho^*(\theta)$ and $h_\rho(\theta)$ is provided in Appendix A. The behavior of $h_\rho(\theta)$ is illustrated below in Figure 3. For a fixed ρ , the function $h_\rho(\theta)$ is numerically observed to be uni-modal in θ and admit a global maximum at a unique $\theta_\rho^* := \arg \sup_{\theta \geq 0} h_\rho(\theta)$, bounded away from 0. Theorem 3 establishes that the worst-case (instance-independent) regret admits the sharp asymptotic $R_n^\pi \sim h_\rho(\theta_\rho^*) \sqrt{n \log n}$. In standard bandit parlance, this substantiates that the

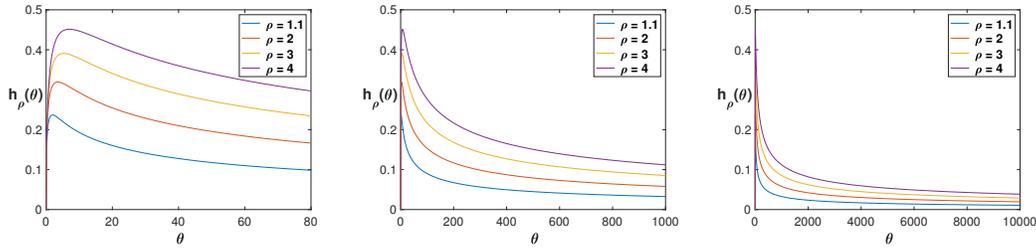


Figure 3: $h_\rho(\theta)$ vs. θ for different values of the exploration coefficient ρ in Algorithm 1. The graphs exhibit a unique global maximizer θ_ρ^* for each ρ . The $(\theta_\rho^*, h_\rho(\theta_\rho^*))$ pairs for $\rho \in \{1.1, 2, 3, 4\}$ in that order are: (1.9, 0.2367), (3.5, 0.3192), (5.3, 0.3909), (7, 0.4514).

$\mathcal{O}(\sqrt{n \log n})$ worst-case (minimax) performance guarantee of canonical UCB cannot be improved in terms of its horizon-dependence. In addition, the result also specifies the precise asymptotic constants achievable in the worst-case setting. This can alternately be viewed as a direct approach to proving the $\mathcal{O}(\sqrt{n \log n})$ performance bound for UCB vis-à-vis conventional minimax analyses such as [9].

Proof sketch. On a high level, note that when $\Delta = \sqrt{(\theta \log n)/n}$, it follows from Theorem 1 that $\mathbb{E}R_n^\pi = \sqrt{(\theta \log n)/n} \mathbb{E}[n - N_{i^*}(n)] \sim h_\rho(\theta) \sqrt{n \log n}$ (convergence in probability implies that in mean since $|N_{i^*}(n)/n| \leq 1$). That R_n^π also admits the same sharp asymptotic can be shown via a finer analysis. In other regimes of Δ , viz., “small” and “large” gaps, we already know that $R_n^\pi = o_p(\sqrt{n \log n})$. This is obvious for “small” Δ since $\mathbb{E}R_n^\pi \leq \Delta n = o(\sqrt{n \log n})$, while for “large” Δ , we use $\mathbb{E}R_n^\pi \leq C\rho((\log n)/\Delta + 1/(\rho - 1))$ for some absolute constant C (given $\rho > 1$, $\Delta \leq 1$) [6], followed by Markov’s inequality. Thus, the multiplicative constant $\sup_{\theta \geq 0} h_\rho(\theta)$ obtained in the “moderate” gap regime must correspond to the worst-case performance of the algorithm. \square

Towards diffusion asymptotics. Diffusion scaling is a classical stochastic analysis technique widely used in the mathematics and operations research literature, see, e.g., steady-state analyses of queuing systems in [15], and a recent application to certain sequential testing problems in [3]. Under this lens, time is accelerated linearly in n , space contracted by a factor of \sqrt{n} , and a sequence of systems indexed by n is considered. In our problem, the n^{th} such system refers to an instance of the two-armed bandit with: n as the horizon of play; a gap that vanishes in the horizon as $\Delta = c/\sqrt{n}$ for some fixed c ; and fixed reward variances given by σ_1^2, σ_2^2 . This is a natural scaling for MAB experiments in that it “preserves” the hardness of the learning problem as n sweeps over the sequence of systems. Recall also from previous discussion that the “hardest” information-theoretic instances have a $\Theta(1/\sqrt{n})$ gap; in short, the diffusion limit is an appropriate asymptotic lens for observing interesting process-level behavior in the MAB problem. However, despite the aforementioned reasons, the diffusion limit behavior of bandit algorithms remains poorly understood and largely unexplored. A recent foray

⁶Previous work establishes matching $\Theta(\sqrt{Kn \log K})$ upper and lower bounds for the minimax *expected* regret of the Gaussian prior-based Thompson Sampling algorithm in the K -armed problem, see [2].

was made in [32, 12], however, deterministic algorithms like UCB remain outside the ambit of such analysis, on account of discontinuities. Theorem 4 provides a complete characterization of this limit.

Theorem 4 (Diffusion asymptotics for canonical UCB) *Suppose that the mean reward of arm $i \in \{1, 2\}$ is given by $\mu_i = \mu + \theta_i/\sqrt{n}$, where n is the horizon of play and $\mu, \theta_1, \theta_2 \geq 0$ are fixed constants, and reward variances are σ_1^2, σ_2^2 . Define $\Delta_0 := |\theta_1 - \theta_2|$. Denote the cumulative reward earned from arm i until time m by $S_{i,m} := \sum_{j=1}^{N_i(m)} X_{i,j}$, and let $\tilde{S}_{i,m} := S_{i,m} - \mu N_i(m)$. Then, the following process-level convergences hold under the policy π given by Algorithm 1 with $\rho > 1$:*

$$(I) \left(\frac{\tilde{S}_{1,\lfloor nt \rfloor}}{\sqrt{n}}, \frac{\tilde{S}_{2,\lfloor nt \rfloor}}{\sqrt{n}} \right) \Rightarrow \left(\frac{\theta_1 t}{2} + \frac{\sigma_1}{\sqrt{2}} B_1(t), \frac{\theta_2 t}{2} + \frac{\sigma_2}{\sqrt{2}} B_2(t) \right),$$

$$(II) \frac{R_{\lfloor nt \rfloor}^\pi}{\sqrt{n}} \Rightarrow \frac{\Delta_0 t}{2} + \sqrt{\frac{\sigma_1^2 + \sigma_2^2}{2}} \tilde{B}(t),$$

where the process-level convergence is over $t \in [0, 1]$, and $B_1(t)$ and $B_2(t)$ are independent standard Brownian motions in \mathbb{R} , and $\tilde{B}(t) := -\sqrt{\frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2}} B_1(t) - \sqrt{\frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2}} B_2(t)$.

Proof sketch. Note that if the arms are played $\lfloor n/2 \rfloor$ times each *independently* over the horizon of play n (resulting in $N_i(n) = \lfloor n/2 \rfloor$, $i \in \{1, 2\}$), part (I) of the stated assertion would immediately follow from Donsker’s Theorem (see [8], Section 14). However, since the sequence of plays, and hence also the eventual allocation $(N_1(n), N_2(n))$, is determined *adaptively* by the policy, the aforementioned convergence may no longer be true. Here, the result hinges crucially on the observation from Theorem 1 that for any arm $i \in \{1, 2\}$ as $n \rightarrow \infty$, $N_i(n)/n \xrightarrow{p} 1/2$ under UCB when $\Delta \asymp 1/\sqrt{n}$ (diffusion-scaled gaps are covered under the “small gap” regime). This observation facilitates a standard “random time-change” argument $t \leftarrow N_i(\lfloor nt \rfloor)/n$, $i \in \{1, 2\}$, which followed upon by an application of Donsker’s Theorem, leads to the stated assertion in (I). This has the profound implication that for diffusion-scaled gaps, a two-armed bandit under UCB is, in fact, well-approximated by a classical system with *independent* samples (sample-interdependence due to the adaptive nature of the policy is washed away in the limit). The conclusion in (II) follows after a direct application of the Continuous Mapping Theorem (see [8], Theorem 2.7) to (I). \square

Discussion. An immediate observation following Theorem 4 is that the normalized regret R_n^π/\sqrt{n} is asymptotically Normal with mean $\Delta_0/2$ and variance $(\sigma_1^2 + \sigma_2^2)/2$ under UCB. Apart from aiding in obvious inferential tasks like construction of (asymptotically valid) confidence intervals (see, e.g., the binary hypothesis testing example referenced in Figure 2(c)), etc., such information provides new insights as to the problem’s minimax complexity as well. This is because $\Delta \asymp 1/\sqrt{n}$ is known to be the information-theoretic “worst-case” for the problem; the smallest achievable regret in this regime must asymptotically be dominated by that under UCB, i.e., $\Delta_0\sqrt{n}/2$. It is also noteworthy that while the diffusion limit in Theorem 4 does not itself depend on the exploration coefficient ρ , the rate at which the system converges to said limit indeed depends on ρ . Theorem 4 will continue to hold only as long as $\rho = \omega((\log \log n)/\log n)$; for smaller ρ , the convergence of $N_i(n)/n$ to $1/2$ may no longer be true (refer to the proof of Theorem 1 in the “small” gap regime in Appendix E).

4 Concluding remarks and open problems

Our results for the two-armed problem under canonical UCB (Algorithm 1) may be generalizable to the K -armed setting, leveraging the observations and insights from Theorem 1 (a simple extension is provided in Appendix B). The K -armed problem under UCB is of interest in its own right: we postulate a division of sampling effort within and across clusters of “similar” and “separated” arms, determined by their relative sub-optimality gaps in the spirit of Theorem 1. We expect that similar generalizations are possible also for Theorem 3 and Theorem 4. For Thompson Sampling, on the other hand, things are less obvious even in the two-armed setting. For example, in spite of compelling numerical evidence (refer, e.g., to Figure 1(b)) suggesting a plausibly non-degenerate distribution of arm-sampling rates for bounded rewards in $[0, 1]$ with means away from 0, the proof of Theorem 2 relies heavily on the rewards being deterministic 0/1, and cannot be extended to the general case. In addition, similar results are conjectured also for the more widely used Gaussian prior-based version of the algorithm. Such results may, in future, shed light on several important diffusion-limit performance metrics of Thompson Sampling, and elucidate its normalized minimax (stochastic) regret behavior.

Acknowledgments and Disclosure of Funding

The authors sincerely thank the anonymous reviewers for their insightful remarks and constructive feedback on the initial submission, and hope that this revision satisfactorily addresses all their points. The authors declare an absence of any competing interests, financial or otherwise.

References

- [1] AGRAWAL, S., AND GOYAL, N. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on learning theory* (2012), JMLR Workshop and Conference Proceedings, pp. 39–1.
- [2] AGRAWAL, S., AND GOYAL, N. Near-optimal regret bounds for thompson sampling. *Journal of the ACM (JACM)* 64, 5 (2017), 1–24.
- [3] ARAMAN, V. F., AND CALDENTEY, R. Diffusion approximations for a class of sequential testing problems. *arXiv preprint arXiv:2102.07030* (2021).
- [4] AUDIBERT, J.-Y., BUBECK, S., ET AL. Minimax policies for adversarial and stochastic bandits. In *COLT* (2009), vol. 7, pp. 1–122.
- [5] AUDIBERT, J.-Y., BUBECK, S., AND MUNOS, R. Best arm identification in multi-armed bandits. In *COLT* (2010), pp. 41–53.
- [6] AUDIBERT, J.-Y., MUNOS, R., AND SZEPESVÁRI, C. Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science* 410, 19 (2009), 1876–1902.
- [7] AUER, P., CESA-BIANCHI, N., AND FISCHER, P. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47, 2-3 (2002), 235–256.
- [8] BILLINGSLEY, P. *Convergence of probability measures*. John Wiley & Sons, 2013.
- [9] BUBECK, S., AND CESA-BIANCHI, N. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *arXiv preprint arXiv:1204.5721* (2012).
- [10] DESHPANDE, Y., MACKEY, L., SYRGKANIS, V., AND TADDY, M. Accurate inference for adaptive linear models. *arXiv preprint arXiv:1712.06695* (2017).
- [11] DURRETT, R. *Probability: theory and examples*, vol. 49. Cambridge university press, 2019.
- [12] FAN, L., AND GLYNN, P. W. Diffusion approximations for thompson sampling. *arXiv preprint arXiv:2105.09232* (2021).
- [13] GARIVIER, A., AND CAPPÉ, O. The kl-ucb algorithm for bounded stochastic bandits and beyond. In *Proceedings of the 24th annual conference on learning theory* (2011), pp. 359–376.
- [14] GARIVIER, A., LATTIMORE, T., AND KAUFMANN, E. On explore-then-commit strategies. In *Advances in Neural Information Processing Systems* (2016), pp. 784–792.
- [15] GLYNN, P. W. Diffusion approximations. *Handbooks in Operations research and management Science* 2 (1990), 145–198.
- [16] HADAD, V., HIRSHBERG, D. A., ZHAN, R., WAGER, S., AND ATHEY, S. Confidence intervals for policy evaluation in adaptive experiments. *arXiv preprint arXiv:1911.02768* (2019).
- [17] Hoeffding, W. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58, 301 (1963), 13–30.
- [18] HONDA, J., AND TAKEMURA, A. An asymptotically optimal bandit algorithm for bounded support models. In *COLT* (2010), Citeseer, pp. 67–79.
- [19] KALKANLI, C., AND OZGUR, A. Asymptotic convergence of thompson sampling. *arXiv preprint arXiv:2011.03917* (2020).
- [20] KALVIT, A., AND ZEEVI, A. From finite to countable-armed bandits. *Advances in Neural Information Processing Systems* 33 (2020), 8259–8269.
- [21] KAUFMANN, E., CAPPÉ, O., AND GARIVIER, A. On the complexity of best-arm identification in multi-armed bandit models. *The Journal of Machine Learning Research* 17, 1 (2016), 1–42.
- [22] LAI, T. L., AND ROBBINS, H. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics* 6, 1 (1985), 4–22.

- [23] LATTIMORE, T., AND SZEPESVÁRI, C. *Bandit algorithms*. Cambridge University Press, 2020.
- [24] MCLENNAN, A. Price dispersion and incomplete learning in the long run. *Journal of Economic dynamics and control* 7, 3 (1984), 331–347.
- [25] MEHRABI, N., MORSTATTER, F., SAXENA, N., LERMAN, K., AND GALSTYAN, A. A survey on bias and fairness in machine learning. *arXiv preprint arXiv:1908.09635* (2019).
- [26] ROTHSCHILD, M. A two-armed bandit theory of market pricing. *Journal of Economic Theory* 9, 2 (1974), 185–202.
- [27] SHIN, J., RAMDAS, A., AND RINALDO, A. Are sample means in multi-armed bandits positively or negatively biased? *arXiv preprint arXiv:1905.11397* (2019).
- [28] SHIN, J., RAMDAS, A., AND RINALDO, A. On the bias, risk and consistency of sample means in multi-armed bandits. *arXiv preprint arXiv:1902.00746* (2019).
- [29] SHIN, J., RAMDAS, A., AND RINALDO, A. On conditional versus marginal bias in multi-armed bandits. In *International Conference on Machine Learning* (2020), PMLR, pp. 8852–8861.
- [30] THOMPSON, W. R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25, 3/4 (1933), 285–294.
- [31] VILLAR, S. S., BOWDEN, J., AND WASON, J. Multi-armed bandit models for the optimal design of clinical trials: benefits and challenges. *Statistical science: a review journal of the Institute of Mathematical Statistics* 30, 2 (2015), 199.
- [32] WAGER, S., AND XU, K. Diffusion asymptotics for sequential experiments. *arXiv preprint arXiv:2101.09855* (2021).
- [33] ZHANG, K. W., JANSON, L., AND MURPHY, S. A. Inference for batched bandits. *arXiv preprint arXiv:2002.03217* (2020).

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#)
 - (c) Did you discuss any potential negative societal impacts of your work? [\[N/A\]](#)
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[Yes\]](#)
 - (b) Did you include complete proofs of all theoretical results? [\[Yes\]](#) Detailed proofs are provided in the supplementary material.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[N/A\]](#)
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[N/A\]](#)
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[N/A\]](#)
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[N/A\]](#)
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [\[N/A\]](#)
 - (b) Did you mention the license of the assets? [\[N/A\]](#)
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[N/A\]](#)

- (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Supplementary material

General organization:

1. Appendix A provides closed-form expressions for the $\lambda_\rho^*(\theta)$ and $h_\rho(\theta)$ functions that appear in Theorem 1 and Theorem 3.
2. Appendix B contains the statement of Theorem 5, which is a simple extension of Theorem 1 to the K -armed setting.
3. Appendix C states three ancillary results that will be used in other proofs.
4. Appendix D provides the proof of Theorem 1 in the “large gap” regime.
5. Appendix E provides the proof of Theorem 1 in the “small gap” regime.
6. Appendix F provides the proof of Theorem 1 in the “moderate gap” regime.
7. Appendix G provides the proof of Theorem 2.
8. Appendix H provides the proof of Theorem 3.
9. Appendix I provides the proof of Theorem 4.
10. Appendix J provides the proof of Theorem 5 from Appendix B.
11. Appendix K provides proofs for the ancillary results stated in Appendix C.

Additional notation. In the proofs that follow, $\lceil \cdot \rceil$ has been used to denote the “ceiling operator,” i.e., $\lceil x \rceil = \inf \{ \nu \in \mathbb{N} : \nu \geq x \}$ for any $x \in \mathbb{R}$. Similarly, $\lfloor \cdot \rfloor$ denotes the “floor operator,” i.e., $\lfloor x \rfloor = \sup \{ \nu \in \mathbb{N} : \nu \leq x \}$ for any $x \in \mathbb{R}$.

A Closed-form expressions for $\lambda_\rho^*(\theta)$ and $h_\rho(\theta)$

$\lambda_\rho^*(\theta)$ is given by:

$$\lambda_\rho^*(\theta) = \frac{1}{2} + \sqrt{\frac{1}{4} - \frac{1}{\left(1 + \sqrt{1 + \frac{\theta}{\rho}}\right)^2}}. \quad (4)$$

$h_\rho(\theta)$ is given by:

$$h_\rho(\theta) = \sqrt{\theta} \left(\frac{1}{2} - \sqrt{\frac{1}{4} - \frac{1}{\left(1 + \sqrt{1 + \frac{\theta}{\rho}}\right)^2}} \right). \quad (5)$$

B Arm-sampling rates under UCB in the K -armed bandit with multiple optimal arms

Algorithm 3 The canonical UCB policy for K -armed bandits.

- 1: **Input:** Exploration coefficient $\rho \in \mathbb{R}_+$.
 - 2: At $t = 1, \dots, K$, play each arm $i \in \{1, \dots, K\}$ once.
 - 3: **for** $t \in \{K + 1, K + 2, \dots\}$ **do**
 - 4: Play arm $\pi_t \in \arg \max_{i \in \{1, \dots, K\}} \left(\bar{X}_i(t-1) + \sqrt{\frac{\rho \log(t-1)}{N_i(t-1)}} \right)$.
-

Theorem 5 (Asymptotic sampling rate of optimal arms) Fix $K \in \mathbb{N}$, and consider a K -armed model with arms indexed by $[K] := \{1, \dots, K\}$. Let $\mathcal{I} \subseteq [K]$ be the set of optimal arms, i.e., arms with mean $\max_{i \in [K]} \mu_i$. If $\mathcal{I} \neq [K]$, define $\Delta_{\min} := \max_{i \in [K]} \mu_i - \max_{i \in [K] \setminus \mathcal{I}} \mu_i$. Then, there exists a finite $\rho_0 > 1$ that depends only on $|\mathcal{I}|$, such that the following results hold for any arm $i \in \mathcal{I}$ as $n \rightarrow \infty$ under Algorithm 3 initialized with $\rho \geq \rho_0$:

(I) If $\mathcal{I} = [K]$, then

$$\frac{N_i(n)}{n} \xrightarrow{p} \frac{1}{K}.$$

(II) If $\mathcal{I} \neq [K]$ and optimal arms are “well-separated,” i.e., $\Delta_{\min} = \omega\left(\sqrt{\frac{\log n}{n}}\right)$, then

$$\frac{N_i(n)}{n} \xrightarrow{p} \frac{1}{|\mathcal{I}|}.$$

Discussion. The main observation here is that if the set of optimal arms is “sufficiently separated” from the sub-optimal arms, then classical UCB policies eventually allocate the sampling effort over the set of optimal arms *uniformly*, in probability. This is a desirable property to have from a fairness standpoint, and also markedly different from the instability and imbalance results for Thompson Sampling discussed earlier in Theorem 2. We remark that the $\rho \geq \rho_0$ condition is only necessary for tractability of the proof, and conjecture the result to hold, in fact, for any $\rho > 1$, akin to the result for the two-armed setting (Theorem 1). We also conjecture analogous results for “small gap” and “moderate gap” regimes, in the spirit of Theorem 1; proofs, however, can be unwieldy in the general K -armed setting. The detailed proof of Theorem 5 is provided in Appendix J.

C Auxiliary results

We will frequently use the following version of the Chernoff-Hoeffding inequality [17] in our proofs:

Fact 1 (Chernoff-Hoeffding bound) *Suppose that $\{Y_{i,j} : i \in \{1, 2\}, j \in \mathbb{N}\}$ is a collection of independent, zero-mean random variables such that $\forall i \in \{1, 2\}, j \in \mathbb{N}, Y_{i,j} \in [c_i, 1 + c_i]$ almost surely, for some fixed $c_1, c_2 \leq 0$. Then, for any $m_1, m_2 \in \mathbb{N}$ and $\alpha > 0$,*

$$\mathbb{P}\left(\frac{\sum_{j=1}^{m_1} Y_{1,j}}{m_1} - \frac{\sum_{j=1}^{m_2} Y_{2,j}}{m_2} \geq \alpha\right) \leq \exp\left(\frac{-2\alpha^2 m_1 m_2}{m_1 + m_2}\right).$$

Proof. Let $[n] := \{1, \dots, n\}$ for $n \in \mathbb{N}$. The Chernoff-Hoeffding inequality in its standard form states that for independent, zero-mean, bounded random variables $\{Z_j : j \in [n]\}$ with $Z_j \in [a_j, b_j] \forall j \in [n]$, the following holds for any $\varepsilon > 0$,

$$\mathbb{P}\left(\sum_{j=1}^n Z_j \geq \varepsilon n\right) \leq \exp\left(\frac{-2\varepsilon^2 n^2}{\sum_{j=1}^n (b_j - a_j)^2}\right). \quad (6)$$

The desired form of the inequality can be obtained by making the following substitutions in (6): $n \leftarrow m_1 + m_2$; $Z_j \leftarrow \frac{Y_{1,j}}{m_1}$, $a_j \leftarrow \frac{c_1}{m_1}$, $b_j \leftarrow \frac{1+c_1}{m_1}$ for $j \in [m_1]$; $Z_j \leftarrow \frac{-Y_{2,j}-m_1}{m_2}$, $a_j \leftarrow \frac{-(1+c_2)}{m_2}$, $b_j \leftarrow \frac{-c_2}{m_2}$ for $j \in [m_1 + m_2] \setminus [m_1]$; and $\varepsilon \leftarrow \frac{\alpha}{m_1 + m_2}$, in that order. \square

In addition, we will use in the proof of Theorem 2 the following two properties of the Beta distribution:

Fact 2 *If $\theta_k, \tilde{\theta}_k$ are Beta($1, k + 1$)-distributed, and $\theta_k, \tilde{\theta}_l$ are independent $\forall k, l \in \mathbb{N} \cup \{0\}$, then*

$$\mathbb{P}\left(\theta_k > \tilde{\theta}_l\right) = \frac{l + 1}{k + l + 2} \quad \text{for any } k, l \in \mathbb{N} \cup \{0\}.$$

Fact 3 *If $\theta_k, \tilde{\theta}_k$ are Beta($k + 1, 1$)-distributed, and $\theta_k, \tilde{\theta}_l$ are independent $\forall k, l \in \mathbb{N} \cup \{0\}$, then*

$$\mathbb{P}\left(\theta_k > \tilde{\theta}_l\right) = \frac{k + 1}{k + l + 2} \quad \text{for any } k, l \in \mathbb{N} \cup \{0\}.$$

The proofs for Fact (2) and Fact (3) are elementary, and provided in Appendix K.

D Proof of Theorem 1 in the “large gap” regime

The proof is straightforward in this regime. We know that for $\rho > 1$, $\mathbb{E}R_n^\pi \leq C\rho \left(\frac{\log n}{\Delta} + \frac{\Delta}{\rho-1} \right)$ for some absolute constant C (see [6], Theorem 7). Since $\mathbb{E}R_n^\pi = \Delta \mathbb{E}[n - N_{i^*}(n)]$, it follows that $\mathbb{E} \left[\frac{n - N_{i^*}(n)}{n} \right] = o(1)$ in the “large gap” regime. Using Markov’s inequality, we then conclude that $\frac{n - N_{i^*}(n)}{n} = o_p(1)$, or equivalently, $\lim_{n \rightarrow \infty} \frac{N_{i^*}(n)}{n} = 1$. Results for “small” and “moderate” gaps are provided separately in Appendix E and Appendix F respectively. \square

E Proof of Theorem 1 in the “small gap” regime

Without loss of generality, suppose that arm 1 is optimal, i.e., $\mu_1 \geq \mu_2$. We will show that for any $\epsilon > 0$, it follows that $\lim_{n \rightarrow \infty} \mathbb{P} \left(\frac{N_1(n)}{n} \geq \frac{1}{2} + \epsilon \right) = 0$. Then, since arm 2 is inferior, an identical result would naturally hold for it as well. Combining the two would prove our assertion as desired. To this end, pick an arbitrary $\epsilon \in (0, 1/2)$, define $u(n) := \lceil (\frac{1}{2} + \epsilon)n \rceil$, and consider the following:

$$\begin{aligned}
N_1(n) &\leq u(n) + \sum_{t=u(n)+1}^n \mathbb{1} \{ \pi_t = 1, N_1(t-1) \geq u(n) \} && \text{(this is always true)} \\
&= u(n) + \sum_{t=u(n)}^{n-1} \mathbb{1} \{ \pi_{t+1} = 1, N_1(t) \geq u(n) \} \\
&\leq u(n) + \sum_{t=u(n)}^{n-1} \mathbb{1} \{ \pi_{t+1} = 1, N_1(t) \geq u(t) \} \\
&\leq u(n) + \sum_{t=u(n)}^{n-1} \mathbb{1} \left\{ \bar{X}_1(t) - \bar{X}_2(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_2(t)}} - \frac{1}{\sqrt{N_1(t)}} \right), N_1(t) \geq u(t) \right\} \\
&= u(n) + \underbrace{\sum_{t=u(n)}^{n-1} \mathbb{1} \left\{ \bar{Y}_1(t) - \bar{Y}_2(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_2(t)}} - \frac{1}{\sqrt{N_1(t)}} \right) - \Delta, N_1(t) \geq u(t) \right\}}_{=: Z(n)}, \tag{7}
\end{aligned}$$

where $\bar{Y}_i(t) := \frac{\sum_{j=1}^{N_i(t)} Y_{i,j}}{N_i(t)}$ with $Y_{i,j} := X_{i,j} - \mu_i$, $i \in \{1, 2\}, j \in \mathbb{N}$. Clearly, $Y_{i,j}$ ’s are independent, zero-mean, and $Y_{i,j} \in [-\mu_i, 1 - \mu_i] \forall i \in \{1, 2\}, j \in \mathbb{N}$.

E.1 An almost sure lower bound on the arm-sampling rates

As a meta-result, we will first show that $N_i(n)/n$, for both arms $i \in \{1, 2\}$, is bounded away from 0 by a positive constant, almost surely. To this end, consider n large enough such that for the ϵ selected earlier, we have $\Delta < \sqrt{\frac{\rho \log n}{n}} \left(\frac{1}{\sqrt{1/2-\epsilon}} - \frac{1}{\sqrt{1/2+\epsilon}} \right)$; this is possible since $\Delta = o \left(\sqrt{\frac{\log n}{n}} \right)$ in the “small gap” regime. Working with a large enough n will allow us to use the Chernoff-Hoeffding bound (Fact 1) in step (\star) in the forthcoming analysis. Observe from (7) that

$$\begin{aligned}
& \mathbb{E}Z(n) \\
&= \sum_{t=u(n)}^{n-1} \mathbb{P} \left(\bar{Y}_1(t) - \bar{Y}_2(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_2(t)}} - \frac{1}{\sqrt{N_1(t)}} \right) - \Delta, N_1(t) \geq u(t) \right) \\
&= \sum_{t=u(n)}^{n-1} \sum_{m=u(t)}^{t-1} \mathbb{P} \left(\frac{\sum_{j=1}^m Y_{1,j}}{m} - \frac{\sum_{j=1}^{t-m} Y_{2,j}}{t-m} \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{t-m}} - \frac{1}{\sqrt{m}} \right) - \Delta, N_1(t) = m \right) \\
&\stackrel{(\star)}{\leq} \sum_{t=u(n)}^{n-1} \sum_{m=u(t)}^{t-1} \mathbb{P} \left(\frac{\sum_{j=1}^m Y_{1,j}}{m} - \frac{\sum_{j=1}^{t-m} Y_{2,j}}{t-m} \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{t-m}} - \frac{1}{\sqrt{m}} \right) - \Delta \right) \\
&\stackrel{(\dagger)}{\leq} \sum_{t=u(n)}^{n-1} \sum_{m=u(t)}^{t-1} \exp \left[-2\rho \left(1 - 2\sqrt{\frac{m}{t} \left(1 - \frac{m}{t} \right)} \right) \log t \right] \exp \left[4\Delta \sqrt{\rho t \log t} \left(\sqrt{\frac{m}{t}} - \sqrt{1 - \frac{m}{t}} \right) \sqrt{\frac{m}{t} \left(1 - \frac{m}{t} \right)} \right] \\
&\stackrel{(\ddagger)}{\leq} \sum_{t=u(n)}^{n-1} \sum_{m=u(t)}^{t-1} \exp \left[-2\rho \left(1 - \sqrt{1 - 4\epsilon^2} \right) \log t \right] \exp \left[4\Delta \sqrt{\rho t \log t} \left(\sqrt{\frac{m}{t}} - \sqrt{1 - \frac{m}{t}} \right) \sqrt{\frac{m}{t} \left(1 - \frac{m}{t} \right)} \right] \\
&\leq \sum_{t=u(n)}^{n-1} \sum_{m=u(t)}^{t-1} \exp \left[-2\rho \left(1 - \sqrt{1 - 4\epsilon^2} \right) \log t \right] \exp \left[4\Delta \sqrt{\rho t \log t} \right] \\
&\leq \sum_{t=u(n)}^{n-1} \sum_{m=u(t)}^{t-1} \exp \left[-2\rho \left(1 - \sqrt{1 - 4\epsilon^2} \right) \log t \right] \exp \left[4\Delta \sqrt{\rho m \log n} \right] \\
&\leq \exp \left[4\Delta \sqrt{\rho m \log n} \right] \sum_{t=u(n)}^{n-1} t^{-(2\rho-1-2\rho\sqrt{1-4\epsilon^2})} \\
&= \exp \left[o(4\sqrt{\rho} \log n) \right] \sum_{t=u(n)}^{n-1} t^{-(2\rho-1-2\rho\sqrt{1-4\epsilon^2})} \quad \left(\because \Delta = o \left(\sqrt{\frac{\log n}{n}} \right) \right) \\
&\leq n^{\frac{1}{2}-\epsilon} \sum_{t=u(n)}^{n-1} t^{-(2\rho-1-2\rho\sqrt{1-4\epsilon^2})}, \tag{8}
\end{aligned}$$

where (\dagger) follows after an application of the Chernoff-Hoeffding bound (Fact 1), (\ddagger) since $\frac{m}{t} \left(1 - \frac{m}{t} \right) \leq \frac{1}{4} - \epsilon^2$ on the interval $\{m : u(t) \leq m \leq t-1\}$, and the last inequality in (8) holds for n large enough. Now consider an arbitrary $\delta > 0$. Then,

$$\begin{aligned}
\mathbb{P}(N_1(n) - u(n) \geq \delta n) &\leq \mathbb{P}(Z(n) \geq \delta n) && \text{(using (7))} \\
&\leq \frac{\mathbb{E}Z(n)}{\delta n} && \text{(Markov's inequality)} \\
&\leq \left(\frac{n^{-(\frac{1}{2}+\epsilon)}}{\delta} \right) \sum_{t=u(n)}^{n-1} t^{-(2\rho-1-2\rho\sqrt{1-4\epsilon^2})} && \text{(using (8))} \\
\implies \mathbb{P} \left(\frac{N_1(n)}{n} \geq \frac{1}{2} + \epsilon + \delta + \frac{1}{n} \right) &\leq \left(\frac{n^{-(\frac{1}{2}+\epsilon)}}{\delta} \right) \sum_{t=\lceil \frac{n}{2} \rceil}^{n-1} t^{-(2\rho-1-2\rho\sqrt{1-4\epsilon^2})}. && (9)
\end{aligned}$$

Define $g(\rho, \epsilon) := \frac{1}{2} + \epsilon + 2\rho - 1 - 2\rho\sqrt{1 - 4\epsilon^2}$. Since $\rho > 1$ is fixed, and $\epsilon \in (0, 1/2)$ is arbitrary, it is possible to push ϵ close to $1/2$ to ensure that $g(\rho, \epsilon) > 2$. Therefore, $\exists \epsilon_\rho \in (0, 1/2)$ s.t. $g(\rho, \epsilon) > 2$ for $\epsilon \geq \epsilon_\rho$. Plugging in $\epsilon = \epsilon_\rho$ in (9), we obtain

$$\mathbb{P} \left(\frac{N_1(n)}{n} \geq \frac{1}{2} + \epsilon_\rho + \delta + \frac{1}{n} \right) \leq \left(\frac{2^{2\rho-1}}{\delta} \right) n^{-(g(\rho, \epsilon_\rho)-1)}.$$

Note that since $\epsilon_\rho < 1/2$, $\exists \epsilon'_\rho < 1/2$ s.t. $\epsilon_\rho + 1/n < \epsilon'_\rho$ for n large enough, i.e., the following holds for all n large enough:

$$\mathbb{P}\left(\frac{N_1(n)}{n} \geq \frac{1}{2} + \epsilon'_\rho + \delta\right) \leq \left(\frac{2^{2\rho-1}}{\delta}\right) n^{-(g(\rho, \epsilon_\rho)-1)}.$$

Finally, since $\delta > 0$ is arbitrary, and $g(\rho, \epsilon_\rho) > 2$, it follows from the Borel-Cantelli Lemma that

$$\limsup_{n \rightarrow \infty} \frac{N_1(n)}{n} \leq \frac{1}{2} + \epsilon'_\rho < 1 \quad \text{w.p. 1.}$$

By assumption, arm 2 is inferior; the above result thus holds, in fact, for both the arms (An almost identical proof can be replicated for rigor). Therefore, we conclude

$$\liminf_{n \rightarrow \infty} \frac{N_i(n)}{n} \geq \frac{1}{2} - \epsilon'_\rho > 0 \quad \text{w.p. 1} \quad \forall i \in \{1, 2\}. \quad (10)$$

E.2 Closing the loop

In this part of the proof, we will leverage (10) to finally show that $N_i(n)/n = 1/2 + o_p(1)$ for $i \in \{1, 2\}$. To this end, recall from (7) that

$$\begin{aligned} & \mathbb{E}Z(n) \\ &= \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_1(t) - \bar{Y}_2(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_2(t)}} - \frac{1}{\sqrt{N_1(t)}}\right) - \Delta, N_1(t) \geq u(t)\right) \\ &\leq \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_1(t) - \bar{Y}_2(t) \geq \sqrt{\frac{\rho \log t}{t}} \left(\frac{1}{\sqrt{\frac{1}{2} - \epsilon}} - \frac{1}{\sqrt{\frac{1}{2} + \epsilon}}\right) - \Delta, N_1(t) \geq u(t)\right) \\ &\leq \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_1(t) - \bar{Y}_2(t) \geq \sqrt{\frac{\rho \log t}{t}} \left(\frac{1}{\sqrt{\frac{1}{2} - \epsilon}} - \frac{1}{\sqrt{\frac{1}{2} + \epsilon}}\right) - \Delta\right) \\ &= \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\underbrace{\sqrt{\frac{t}{\rho \log t}} (\bar{Y}_1(t) - \bar{Y}_2(t))}_{=: W_t} \geq \sqrt{\frac{2}{1-2\epsilon}} - \sqrt{\frac{2}{1+2\epsilon}} - \Delta \sqrt{\frac{t}{\rho \log t}}\right) \\ &\leq \sum_{t=u(n)}^{n-1} \mathbb{P}\left(W_t \geq \frac{1}{\sqrt{1-2\epsilon}} - \frac{1}{\sqrt{1+2\epsilon}}\right), \end{aligned} \quad (11)$$

for n large enough; the last inequality following since $\Delta = o\left(\sqrt{\frac{\log n}{n}}\right)$ and $u(n) > n/2$. Now,

$$\begin{aligned} & |W_t| \\ &\leq \sqrt{\frac{t}{\rho \log t}} \left(\left| \frac{\sum_{j=1}^{N_1(t)} Y_{1,j}}{N_1(t)} \right| + \left| \frac{\sum_{j=1}^{N_2(t)} Y_{2,j}}{N_2(t)} \right| \right) \\ &= \sqrt{\frac{2t}{\rho \log t}} \left(\sqrt{\frac{\log \log N_1(t)}{N_1(t)}} \left| \frac{\sum_{j=1}^{N_1(t)} Y_{1,j}}{\sqrt{2N_1(t) \log \log N_1(t)}} \right| + \sqrt{\frac{\log \log N_2(t)}{N_2(t)}} \left| \frac{\sum_{j=1}^{N_2(t)} Y_{2,j}}{\sqrt{2N_2(t) \log \log N_2(t)}} \right| \right) \\ &\leq \sqrt{\frac{2t}{\rho \log t}} \left(\sqrt{\frac{\log \log t}{N_1(t)}} \left| \frac{\sum_{j=1}^{N_1(t)} Y_{1,j}}{\sqrt{2N_1(t) \log \log N_1(t)}} \right| + \sqrt{\frac{\log \log t}{N_2(t)}} \left| \frac{\sum_{j=1}^{N_2(t)} Y_{2,j}}{\sqrt{2N_2(t) \log \log N_2(t)}} \right| \right) \\ &= \sqrt{\frac{2 \log \log t}{\rho \log t}} \left(\sqrt{\frac{t}{N_1(t)}} \left| \frac{\sum_{j=1}^{N_1(t)} Y_{1,j}}{\sqrt{2N_1(t) \log \log N_1(t)}} \right| + \sqrt{\frac{t}{N_2(t)}} \left| \frac{\sum_{j=1}^{N_2(t)} Y_{2,j}}{\sqrt{2N_2(t) \log \log N_2(t)}} \right| \right). \end{aligned} \quad (12)$$

We know that $N_i(t)$, for both arms $i \in \{1, 2\}$, can be lower bounded *path-wise* by a deterministic monotone increasing function of t , say $f(t)$, that grows to $+\infty$ as $t \rightarrow \infty$. This is a trivial consequence of the structure of canonical UCB (Algorithm 1), and the fact that the rewards are uniformly bounded. We therefore have for any $i \in \{1, 2\}$ that

$$\left| \frac{\sum_{j=1}^{N_i(t)} Y_{i,j}}{\sqrt{2N_i(t) \log \log N_i(t)}} \right| \leq \sup_{m \geq f(t)} \left| \frac{\sum_{j=1}^m Y_{i,j}}{\sqrt{2m \log \log m}} \right|.$$

For a fixed arm $i \in \{1, 2\}$, $\{Y_{i,j} : j \in \mathbb{N}\}$ is a collection of i.i.d. random variables with $\mathbb{E}Y_{i,1} = 0$ and $\text{Var}(Y_{i,1}) = \text{Var}(X_{i,1}) \leq 1$. Also, $f(t)$ is monotone increasing and coercive in t . Therefore, the *Law of the Iterated Logarithm* (see [11], Theorem 8.5.2) implies

$$\limsup_{t \rightarrow \infty} \left| \frac{\sum_{j=1}^{N_i(t)} Y_{i,j}}{\sqrt{2N_i(t) \log \log N_i(t)}} \right| \leq 1 \quad \text{w.p. 1 } \forall i \in \{1, 2\}. \quad (13)$$

Using (10), (12) and (13), we conclude that

$$\lim_{t \rightarrow \infty} W_t = 0 \quad \text{w.p. 1.} \quad (14)$$

Now consider an arbitrary $\delta > 0$. Then,

$$\begin{aligned} \mathbb{P}(N_1(n) - u(n) \geq \delta n) &\leq \mathbb{P}(Z(n) \geq \delta n) && \text{(using (7))} \\ &\leq \frac{\mathbb{E}Z(n)}{\delta n} && \text{(Markov's inequality)} \\ &\leq \frac{1}{\delta n} \sum_{t=u(n)}^{n-1} \mathbb{P}\left(W_t \geq \frac{1}{\sqrt{1-2\epsilon}} - \frac{1}{\sqrt{1+2\epsilon}}\right). && \text{(using (11))} \\ &\leq \frac{1}{\delta} \sup_{t > n/2} \mathbb{P}\left(W_t \geq \frac{1}{\sqrt{1-2\epsilon}} - \frac{1}{\sqrt{1+2\epsilon}}\right) \\ \implies \mathbb{P}\left(\frac{N_1(n)}{n} \geq \frac{1}{2} + \epsilon + \delta + \frac{1}{n}\right) &\leq \frac{1}{\delta} \sup_{t > n/2} \mathbb{P}\left(W_t \geq \frac{1}{\sqrt{1-2\epsilon}} - \frac{1}{\sqrt{1+2\epsilon}}\right). \end{aligned}$$

Since $\epsilon, \delta > 0$ are arbitrary, it follows that for n large enough,

$$\mathbb{P}\left(\frac{N_1(n)}{n} \geq \frac{1}{2} + 2(\epsilon + \delta)\right) \leq \frac{1}{\delta} \sup_{t > n/2} \mathbb{P}\left(W_t \geq \frac{1}{\sqrt{1-2\epsilon}} - \frac{1}{\sqrt{1+2\epsilon}}\right). \quad (15)$$

Using (14) and (15), we conclude that for any arbitrary $\epsilon, \delta > 0$,

$$\limsup_{n \rightarrow \infty} \mathbb{P}\left(\frac{N_1(n)}{n} \geq \frac{1}{2} + 2(\epsilon + \delta)\right) \leq \frac{1}{\delta} \limsup_{n \rightarrow \infty} \mathbb{P}\left(W_n \geq \frac{1}{\sqrt{1-2\epsilon}} - \frac{1}{\sqrt{1+2\epsilon}}\right) = 0.$$

It therefore follows that for any $\epsilon' > 0$, $\lim_{n \rightarrow \infty} \mathbb{P}\left(\frac{N_1(n)}{n} \geq \frac{1}{2} + \epsilon'\right) = 0$; equivalently,

$\lim_{n \rightarrow \infty} \mathbb{P}\left(\frac{N_2(n)}{n} \leq \frac{1}{2} - \epsilon'\right) = 0$. Since arm 2 is inferior by assumption, it naturally holds that $\lim_{n \rightarrow \infty} \mathbb{P}\left(\frac{N_2(n)}{n} \geq \frac{1}{2} + \epsilon'\right) = 0$ (The steps in E.2 can be replicated near-identically for rigor).

Thus, the stated assertion $\frac{N_i(n)}{n} \xrightarrow[n \rightarrow \infty]{P} \frac{1}{2} \forall i \in \{1, 2\}$, follows. \square

F Proof of Theorem 1 in the ‘‘moderate gap’’ regime

Firstly, note that the $\lambda_\rho^*(\theta)$ that solves (2), satisfies the following properties: (i) Continuous and monotone increasing in $\theta \geq 0$, (ii) $\lambda_\rho^*(\theta) \geq 1/2$ for all $\theta \geq 0$, (iii) $\lambda_\rho^*(0) = 1/2$ and $\lambda_\rho^*(\theta) \rightarrow 1$ as $\theta \rightarrow \infty$.

Secondly, because we are only interested in asymptotics, the $\Delta \sim \sqrt{\frac{\theta \log n}{n}}$ condition is as good as $\Delta = \sqrt{\frac{\theta \log n}{n}}$, since for any arbitrarily small $\epsilon' > 0$, $\Delta \in \left(\sqrt{\frac{(\theta - \epsilon') \log n}{n}}, \sqrt{\frac{(\theta + \epsilon') \log n}{n}} \right)$ for n large enough; the stated assertion would follow in the limit as ϵ' approaches 0. In what follows, we will therefore assume for readability of the proof, and without loss of generality, that $\Delta = \sqrt{\frac{\theta \log n}{n}}$.

Thirdly, without loss of generality, suppose that arm 1 is optimal, i.e., $\mu_1 \geq \mu_2$.

F.1 Focusing on arm 1

Consider an arbitrary $\epsilon \in (0, 1 - \lambda_\rho^*(\theta))$, and define $u(n) := \lceil (\lambda_\rho^*(\theta) + \epsilon)n \rceil$. We know that

$$\begin{aligned}
N_1(n) &\leq u(n) + \sum_{t=u(n)+1}^n \mathbb{1}\{\pi_t = 1, N_1(t-1) \geq u(n)\} && \text{(this is always true)} \\
&= u(n) + \sum_{t=u(n)}^{n-1} \mathbb{1}\{\pi_{t+1} = 1, N_1(t) \geq u(n)\} \\
&\leq u(n) + \sum_{t=u(n)}^{n-1} \mathbb{1}\{\pi_{t+1} = 1, N_1(t) \geq u(t)\} \\
&\leq u(n) + \sum_{t=u(n)}^{n-1} \mathbb{1}\left\{ \bar{X}_1(t) - \bar{X}_2(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_2(t)}} - \frac{1}{\sqrt{N_1(t)}} \right), N_1(t) \geq u(t) \right\} \\
&= u(n) + \underbrace{\sum_{t=u(n)}^{n-1} \mathbb{1}\left\{ \bar{Y}_1(t) - \bar{Y}_2(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_2(t)}} - \frac{1}{\sqrt{N_1(t)}} \right) - \Delta, N_1(t) \geq u(t) \right\}}_{=: Z(n)},
\end{aligned} \tag{16}$$

where $\bar{Y}_i(t) := \frac{\sum_{j=1}^{N_i(t)} Y_{i,j}}{N_i(t)}$ with $Y_{i,j} := X_{i,j} - \mu_i$, $i \in \{1, 2\}$, $j \in \mathbb{N}$. Clearly, $Y_{i,j}$'s are independent, zero-mean, and $Y_{i,j} \in [-\mu_i, 1 - \mu_i] \forall i \in \{1, 2\}, j \in \mathbb{N}$.

F.1.1 An almost sure lower bound on the arm-sampling rates

Consider n large enough such that $\sqrt{\frac{\log n}{n}}$ is monotone decreasing in n ($n \geq 3$ suffices). This will enable the inequality in step (\dagger) below. From (16), we have

$$\begin{aligned}
&\mathbb{E}Z(n) \\
&= \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_1(t) - \bar{Y}_2(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_2(t)}} - \frac{1}{\sqrt{N_1(t)}} \right) - \Delta, N_1(t) \geq u(t) \right) \\
&= \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_1(t) - \bar{Y}_2(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_2(t)}} - \frac{1}{\sqrt{N_1(t)}} \right) - \sqrt{\frac{\theta \log n}{n}}, N_1(t) \geq u(t) \right) \\
&\stackrel{(\dagger)}{\leq} \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_1(t) - \bar{Y}_2(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_2(t)}} - \frac{1}{\sqrt{N_1(t)}} \right) - \sqrt{\frac{\theta \log t}{t}}, N_1(t) \geq u(t) \right) \\
&= \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_1(t) - \bar{Y}_2(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_2(t)}} - \frac{1}{\sqrt{N_1(t)}} - \sqrt{\frac{\theta}{\rho t}} \right), N_1(t) \geq u(t) \right) \tag{17}
\end{aligned}$$

$$\leq \sum_{t=u(n)}^{n-1} \sum_{m=u(t)}^{t-1} \mathbb{P}\left(\frac{\sum_{j=1}^m Y_{1,j}}{m} - \frac{\sum_{j=1}^{t-m} Y_{2,j}}{t-m} \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{t-m}} - \frac{1}{\sqrt{m}} - \sqrt{\frac{\theta}{\rho t}} \right) \right). \tag{18}$$

Notice that in the interval $m \in [u(t), t-1]$,

$$\frac{1}{\sqrt{t-m}} - \frac{1}{\sqrt{m}} - \sqrt{\frac{\theta}{2t}} \geq \frac{1}{\sqrt{t-u(t)}} - \frac{1}{\sqrt{u(t)}} - \sqrt{\frac{\theta}{\rho t}} \geq \frac{1}{\sqrt{t}} \left(\frac{1}{\sqrt{1-\lambda_\rho^*(\theta)-\epsilon}} - \frac{1}{\sqrt{\lambda_\rho^*(\theta)+\epsilon}} - \sqrt{\frac{\theta}{\rho}} \right) > 0,$$

where the final inequality follows since $\lambda_\rho^*(\theta)$ is the solution to (2). We can therefore apply the Chernoff-Hoeffding bound (Fact 1) to (18) to conclude

$$\begin{aligned} & \mathbb{E}Z(n) \\ & \leq \sum_{t=u(n)}^{n-1} \sum_{m=u(t)}^{t-1} \exp \left[-2\rho \log t \left(\frac{1}{\sqrt{t-m}} - \frac{1}{\sqrt{m}} - \sqrt{\frac{\theta}{\rho t}} \right)^2 \frac{m(t-m)}{t} \right] \\ & = \sum_{t=u(n)}^{n-1} \sum_{m=u(t)}^{t-1} \exp \left[-2\rho \log t \left(\sqrt{\frac{m}{t}} - \sqrt{1-\frac{m}{t}} - \sqrt{\frac{\theta}{\rho}} \sqrt{\frac{m}{t}} \left(1-\frac{m}{t}\right) \right)^2 \right] \\ & = \sum_{t=u(n)}^{n-1} \sum_{m=u(t)}^{t-1} \exp \left[-2\rho \log t \left(f\left(\frac{m}{t}\right) \right)^2 \right], \end{aligned} \quad (19)$$

where the function $f(x) := \sqrt{x} - \sqrt{1-x} - \sqrt{\theta x(1-x)}/\rho$. Notice that $f(x)$ is monotone increasing over the interval $(1/2, 1)$ ($\cdot: \theta, \rho \geq 0$). Also, note that $1/2 < \lambda_\rho^*(\theta) + \epsilon < m/t < 1$ in (19). Thus, we have in (19) that $f\left(\frac{m}{t}\right) \geq \min_{x \in [\lambda_\rho^*(\theta) + \epsilon, 1]} f(x) = f(\lambda_\rho^*(\theta) + \epsilon)$. An expression for $f(\lambda_\rho^*(\theta) + \epsilon)$ is provided in (22) below. Observe that $f(\lambda_\rho^*(\theta) + \epsilon) > 0$; this follows since $\lambda_\rho^*(\theta)$ is the solution to (2). Using these facts in (19), we conclude

$$\begin{aligned} \mathbb{E}Z(n) & \leq \sum_{t=u(n)}^{n-1} \sum_{m=u(t)}^{t-1} \exp \left[-2\rho \log t \left(\min_{x \in [\lambda_\rho^*(\theta) + \epsilon, 1]} f(x) \right)^2 \right] \\ & = \sum_{t=u(n)}^{n-1} \sum_{m=u(t)}^{t-1} \exp \left[-2\rho \log t \left(f(\lambda_\rho^*(\theta) + \epsilon) \right)^2 \right] \\ & \leq \sum_{t=u(n)}^{n-1} t^{1-2\rho(f(\lambda_\rho^*(\theta) + \epsilon))^2} \\ & = \sum_{t=\lceil (\lambda_\rho^*(\theta) + \epsilon)n \rceil}^{n-1} t^{1-2\rho(f(\lambda_\rho^*(\theta) + \epsilon))^2}. \end{aligned} \quad (20)$$

Now consider an arbitrary $\delta > 0$. We then have

$$\begin{aligned} \mathbb{P}(N_1(n) - u(n) \geq \delta n) & \leq \mathbb{P}(Z(n) \geq \delta n) && \text{(using (16))} \\ & \leq \frac{\mathbb{E}Z(n)}{\delta n} && \text{(Markov's inequality)} \\ \implies \mathbb{P}(N_1(n) \geq \lceil (\lambda_\rho^*(\theta) + \epsilon)n \rceil + \delta n) & \leq \frac{1}{\delta n} \sum_{t=\lceil (\lambda_\rho^*(\theta) + \epsilon)n \rceil}^{n-1} t^{1-2\rho(f(\lambda_\rho^*(\theta) + \epsilon))^2}. && \text{(using (20))} \end{aligned} \quad (21)$$

Note that $f(\lambda_\rho^*(\theta) + \epsilon)$ is given by

$$f(\lambda_\rho^*(\theta) + \epsilon) = \sqrt{\lambda_\rho^*(\theta) + \epsilon} - \sqrt{1 - \lambda_\rho^*(\theta) - \epsilon} - \sqrt{\frac{\theta}{\rho}} \sqrt{(\lambda_\rho^*(\theta) + \epsilon)(1 - \lambda_\rho^*(\theta) - \epsilon)}. \quad (22)$$

Setting $\epsilon = 0$ in (22) yields $f(\lambda_\rho^*(\theta)) = 0$ (follows from (2)), whereas setting $\epsilon = 1 - \lambda_\rho^*(\theta)$ yields $f(1) = 1$. Since $\rho > 1$, and $f(\lambda_\rho^*(\theta) + \epsilon)$ is continuous and monotone increasing in ϵ ,

$\exists \epsilon_{\theta, \rho} \in (0, 1 - \lambda_\rho^*(\theta))$ s.t. $f(\lambda_\rho^*(\theta) + \epsilon) > 1/\sqrt{\rho}$ for $\epsilon \geq \epsilon_{\theta, \rho}$. Substituting $\epsilon = \epsilon_{\theta, \rho}$ in (21) and using the aforementioned fact, we obtain

$$\begin{aligned} \mathbb{P}(N_1(n) \geq \lceil (\lambda_\rho^*(\theta) + \epsilon_{\theta, \rho})n \rceil + \delta n) &\leq \frac{1}{\delta n} \sum_{t=\lceil (\lambda_\rho^*(\theta) + \epsilon_{\theta, \rho})n \rceil}^{n-1} t^{1-2\rho(f(\lambda_\rho^*(\theta) + \epsilon_{\theta, \rho}))^2} \\ &\leq \left(\frac{2^{2\rho-1}}{\delta}\right) n^{-(2\rho(f(\lambda_\rho^*(\theta) + \epsilon_{\theta, \rho}))^2-1)}, \end{aligned} \quad (23)$$

where the last inequality follows since $1/\sqrt{\rho} < f(\lambda_\rho^*(\theta) + \epsilon_{\theta, \rho}) < 1$, and $\lambda_\rho^*(\theta) + \epsilon_{\theta, \rho} > 1/2$. Finally since $\delta > 0$ is arbitrary, we conclude from (23) using the Borel-Cantelli Lemma that

$$\limsup_{n \rightarrow \infty} \frac{N_1(n)}{n} \leq \lambda_\rho^*(\theta) + \epsilon_{\theta, \rho} < 1 \quad \text{w.p. 1.}$$

The above result naturally holds for arm 2 as well, since it is inferior by assumption (we resort to the cop-out that a near-identical argument handles its case). Therefore, in conclusion,

$$\liminf_{n \rightarrow \infty} \frac{N_i(n)}{n} \geq 1 - \lambda_\rho^*(\theta) - \epsilon_{\theta, \rho} > 0 \quad \text{w.p. 1} \quad \forall i \in \{1, 2\}. \quad (24)$$

F.1.2 Closing the loop

From (17), we know that

$$\begin{aligned} &\mathbb{E}Z(n) \\ &\leq \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_1(t) - \bar{Y}_2(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_2(t)}} - \frac{1}{\sqrt{N_1(t)}} - \sqrt{\frac{\theta}{\rho t}}\right), N_1(t) \geq u(t)\right) \\ &\leq \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_1(t) - \bar{Y}_2(t) \geq \sqrt{\frac{\rho \log t}{t}} \left(\frac{1}{\sqrt{1 - \lambda_\rho^*(\theta) - \epsilon}} - \frac{1}{\sqrt{\lambda_\rho^*(\theta) + \epsilon}} - \sqrt{\frac{\theta}{\rho}}\right)\right) \\ &= \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\underbrace{\sqrt{\frac{t}{\rho \log t}} (\bar{Y}_1(t) - \bar{Y}_2(t))}_{=: W_t} \geq \frac{1}{\sqrt{1 - \lambda_\rho^*(\theta) - \epsilon}} - \frac{1}{\sqrt{\lambda_\rho^*(\theta) + \epsilon}} - \sqrt{\frac{\theta}{\rho}}\right), \end{aligned} \quad (25)$$

where we already know that $\frac{1}{\sqrt{1 - \lambda_\rho^*(\theta) - \epsilon}} - \frac{1}{\sqrt{\lambda_\rho^*(\theta) + \epsilon}} - \sqrt{\frac{\theta}{\rho}} > 0$ (since $\lambda_\rho^*(\theta)$ is the solution to (2)).

Now,

$$\begin{aligned} &|W_t| \\ &\leq \sqrt{\frac{t}{\rho \log t}} \left(\left| \frac{\sum_{j=1}^{N_1(t)} Y_{1,j}}{N_1(t)} \right| + \left| \frac{\sum_{j=1}^{N_2(t)} Y_{2,j}}{N_2(t)} \right| \right) \\ &= \sqrt{\frac{2t}{\rho \log t}} \left(\sqrt{\frac{\log \log N_1(t)}{N_1(t)}} \left| \frac{\sum_{j=1}^{N_1(t)} Y_{1,j}}{\sqrt{2N_1(t) \log \log N_1(t)}} \right| + \sqrt{\frac{\log \log N_2(t)}{N_2(t)}} \left| \frac{\sum_{j=1}^{N_2(t)} Y_{2,j}}{\sqrt{2N_2(t) \log \log N_2(t)}} \right| \right) \\ &\leq \sqrt{\frac{2t}{\rho \log t}} \left(\sqrt{\frac{\log \log t}{N_1(t)}} \left| \frac{\sum_{j=1}^{N_1(t)} Y_{1,j}}{\sqrt{2N_1(t) \log \log N_1(t)}} \right| + \sqrt{\frac{\log \log t}{N_2(t)}} \left| \frac{\sum_{j=1}^{N_2(t)} Y_{2,j}}{\sqrt{2N_2(t) \log \log N_2(t)}} \right| \right) \\ &= \sqrt{\frac{2 \log \log t}{\rho \log t}} \left(\sqrt{\frac{t}{N_1(t)}} \left| \frac{\sum_{j=1}^{N_1(t)} Y_{1,j}}{\sqrt{2N_1(t) \log \log N_1(t)}} \right| + \sqrt{\frac{t}{N_2(t)}} \left| \frac{\sum_{j=1}^{N_2(t)} Y_{2,j}}{\sqrt{2N_2(t) \log \log N_2(t)}} \right| \right). \end{aligned} \quad (26)$$

We know that $N_i(t)$, for both arms $i \in \{1, 2\}$, can be lower bounded *path-wise* by a deterministic monotone increasing function of t , say $g(t)$, that grows to $+\infty$ as $t \rightarrow \infty$. This is a trivial consequence

of the structure of the canonical UCB policy (Algorithm 1), and the fact that the rewards are uniformly bounded. Therefore, for any arm $i \in \{1, 2\}$, we have

$$\left| \frac{\sum_{j=1}^{N_i(t)} Y_{i,j}}{\sqrt{2N_i(t) \log \log N_i(t)}} \right| \leq \sup_{m \geq g(t)} \left| \frac{\sum_{j=1}^m Y_{i,j}}{\sqrt{2m \log \log m}} \right|.$$

For a fixed $i \in \{1, 2\}$, $\{Y_{i,j} : j \in \mathbb{N}\}$ is a collection of i.i.d. random variables with $\mathbb{E}Y_{i,1} = 0$ and $\text{Var}(Y_{i,1}) = \text{Var}(X_{i,1}) \leq 1$. Also, $g(t)$ is a monotone increasing and coercive function of t . Therefore, the *Law of the Iterated Logarithm* (see [11], Theorem 8.5.2) implies

$$\limsup_{t \rightarrow \infty} \left| \frac{\sum_{j=1}^{N_i(t)} Y_{i,j}}{\sqrt{2N_i(t) \log \log N_i(t)}} \right| \leq 1 \quad \text{w.p. 1 } \forall i \in \{1, 2\}. \quad (27)$$

Using (24), (26) and (27), we conclude that

$$\lim_{t \rightarrow \infty} W_t = 0 \quad \text{w.p. 1.} \quad (28)$$

Now consider an arbitrary $\delta > 0$. We have

$$\begin{aligned} \mathbb{P}(N_1(n) - u(n) \geq \delta n) &\leq \mathbb{P}(Z(n) \geq \delta n) && \text{(using (16))} \\ &\leq \frac{\mathbb{E}Z(n)}{\delta n} && \text{(Markov's inequality)} \\ &\leq \frac{1}{\delta n} \sum_{t=u(n)}^{n-1} \mathbb{P}\left(W_t \geq \frac{1}{\sqrt{1 - \lambda_\rho^*(\theta) - \epsilon}} - \frac{1}{\sqrt{\lambda_\rho^*(\theta) + \epsilon}} - \sqrt{\frac{\theta}{\rho}}\right) && \text{(using (25))} \\ &\leq \frac{1}{\delta} \sup_{t > n/2} \mathbb{P}\left(W_t \geq \frac{1}{\sqrt{1 - \lambda_\rho^*(\theta) - \epsilon}} - \frac{1}{\sqrt{\lambda_\rho^*(\theta) + \epsilon}} - \sqrt{\frac{\theta}{\rho}}\right). \end{aligned} \quad (29)$$

Using (28) and (29), it follows that

$$\limsup_{n \rightarrow \infty} \mathbb{P}(N_1(n) - u(n) \geq \delta n) \leq \frac{1}{\delta} \limsup_{n \rightarrow \infty} \mathbb{P}\left(W_n \geq \frac{1}{\sqrt{1 - \lambda_\rho^*(\theta) - \epsilon}} - \frac{1}{\sqrt{\lambda_\rho^*(\theta) + \epsilon}} - \sqrt{\frac{\theta}{\rho}}\right) = 0.$$

Since $u(n) = \lceil (\lambda_\rho^*(\theta) + \epsilon) n \rceil$ and $\epsilon, \delta > 0$ are arbitrary, we conclude that for any $\epsilon > 0$, it holds that $\lim_{n \rightarrow \infty} \mathbb{P}\left(\frac{N_1(n)}{n} \geq \lambda_\rho^*(\theta) + \epsilon\right) = 0$. Equivalently, $\lim_{n \rightarrow \infty} \mathbb{P}\left(\frac{N_2(n)}{n} \leq 1 - \lambda_\rho^*(\theta) - \epsilon\right) = 0$ holds for any $\epsilon > 0$. \square

F.2 Focusing on arm 2 and concluding

We will essentially replicate here the proof for arm 1 given in F.1, albeit with a few subtle modifications to account for the fact that arm 2 is inferior. Consistent with previous approach and notation, we consider an arbitrary $\epsilon \in (0, \lambda_\rho^*(\theta))$ and set $u(n) := \lceil (1 - \lambda_\rho^*(\theta) + \epsilon) n \rceil$, where $\lambda_\rho^*(\theta)$ is the solution to (2) (Note that the definition of $u(n)$ here is different from the one used in the proof for

arm 1.). We know that

$$\begin{aligned}
N_2(n) &\leq u(n) + \sum_{t=u(n)+1}^n \mathbb{1}\{\pi_t = 2, N_2(t-1) \geq u(n)\} && \text{(this is always true)} \\
&= u(n) + \sum_{t=u(n)}^{n-1} \mathbb{1}\{\pi_{t+1} = 2, N_2(t) \geq u(n)\} \\
&\leq u(n) + \sum_{t=u(n)}^{n-1} \mathbb{1}\left\{\bar{X}_2(t) - \bar{X}_1(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_1(t)}} - \frac{1}{\sqrt{N_2(t)}} \right), N_2(t) \geq u(n)\right\} \\
&= u(n) + \underbrace{\sum_{t=u(n)}^{n-1} \mathbb{1}\left\{\bar{Y}_2(t) - \bar{Y}_1(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_1(t)}} - \frac{1}{\sqrt{N_2(t)}} \right) + \Delta, N_2(t) \geq u(n)\right\}}_{=:Z(n)},
\end{aligned} \tag{30}$$

where $\bar{Y}_i(t) := \frac{\sum_{j=1}^{N_i(t)} Y_{i,j}}{N_i(t)}$ with $Y_{i,j} := X_{i,j} - \mu_i$, $i \in \{1, 2\}, j \in \mathbb{N}$ (Notice that these definitions of $\bar{Y}_i(t)$ and $Y_{i,j}$ are identical to their counterparts from the proof for arm 1.). From (30), it follows that

$$\begin{aligned}
&\mathbb{E}Z(n) \\
&= \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_2(t) - \bar{Y}_1(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_1(t)}} - \frac{1}{\sqrt{N_2(t)}} \right) + \Delta, N_2(t) \geq u(n)\right) \\
&= \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_2(t) - \bar{Y}_1(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_1(t)}} - \frac{1}{\sqrt{N_2(t)}} \right) + \sqrt{\frac{\theta \log n}{n}}, N_2(t) \geq u(n)\right) \\
&\leq \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_2(t) - \bar{Y}_1(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{t-u(n)}} - \frac{1}{\sqrt{u(n)}} \right) + \sqrt{\frac{\theta \log n}{n}}\right) \\
&\leq \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_2(t) - \bar{Y}_1(t) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{n-u(n)}} - \frac{1}{\sqrt{u(n)}} \right) + \sqrt{\frac{\theta \log t}{n}}\right) \\
&\leq \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_2(t) - \bar{Y}_1(t) \geq \sqrt{\frac{\rho \log t}{n}} \left(\frac{1}{\sqrt{\lambda_\rho^*(\theta) - \epsilon}} - \frac{1}{\sqrt{1 - \lambda_\rho^*(\theta) + \epsilon}} + \sqrt{\frac{\theta}{\rho}} \right)\right),
\end{aligned}$$

where $\frac{1}{\sqrt{\lambda_\rho^*(\theta) - \epsilon}} - \frac{1}{\sqrt{1 - \lambda_\rho^*(\theta) + \epsilon}} + \sqrt{\frac{\theta}{\rho}} > 0$ is guaranteed since $\lambda_\rho^*(\theta)$ is the solution to (2). Also, $t \geq u(n) = \lceil (1 - \lambda_\rho^*(\theta) + \epsilon)n \rceil \implies n \leq \frac{t}{1 - \lambda_\rho^*(\theta) + \epsilon}$. Therefore,

$$\begin{aligned}
&\mathbb{E}Z(n) \\
&\leq \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\bar{Y}_2(t) - \bar{Y}_1(t) \geq \sqrt{1 - \lambda_\rho^*(\theta) + \epsilon} \sqrt{\frac{\rho \log t}{t}} \left(\frac{1}{\sqrt{\lambda_\rho^*(\theta) - \epsilon}} - \frac{1}{\sqrt{1 - \lambda_\rho^*(\theta) + \epsilon}} + \sqrt{\frac{\theta}{\rho}} \right)\right) \\
&= \sum_{t=u(n)}^{n-1} \mathbb{P}\left(\underbrace{\sqrt{\frac{t}{\rho \log t}} (\bar{Y}_2(t) - \bar{Y}_1(t))}_{=:W_t} \geq \underbrace{\sqrt{1 - \lambda_\rho^*(\theta) + \epsilon} \left(\frac{1}{\sqrt{\lambda_\rho^*(\theta) - \epsilon}} - \frac{1}{\sqrt{1 - \lambda_\rho^*(\theta) + \epsilon}} + \sqrt{\frac{\theta}{\rho}} \right)}_{=: \varepsilon(\theta, \rho, \epsilon)}\right).
\end{aligned} \tag{31}$$

Recall that we have already handled W_t (albeit a negated version thereof) in the proof for arm 1 in (25) and shown that $W_t \rightarrow 0$ almost surely in (28). Now consider an arbitrary $\delta > 0$. We then have

$$\begin{aligned}
\mathbb{P}(N_2(n) - u(n) \geq \delta n) &\leq \mathbb{P}(Z(n) \geq \delta n) && \text{(using (30))} \\
&\leq \frac{\mathbb{E}Z(n)}{\delta n} && \text{(Markov's inequality)} \\
&\leq \frac{1}{\delta n} \sum_{t=u(n)}^{n-1} \mathbb{P}(W_t \geq \varepsilon(\theta, \rho, \epsilon)) && \text{(using (31))} \\
&\leq \frac{1}{\delta} \sup_{t > (1-\lambda_\rho^*(\theta))n} \mathbb{P}(W_t \geq \varepsilon(\theta, \rho, \epsilon)). && (32)
\end{aligned}$$

Taking limits on both sides of (32), we obtain

$$\limsup_{n \rightarrow \infty} \mathbb{P}(N_2(n) - u(n) \geq \delta n) \leq \frac{1}{\delta} \limsup_{n \rightarrow \infty} \mathbb{P}(W_n \geq \varepsilon(\theta, \rho, \epsilon)) = 0,$$

where the final conclusion follows since $W_n \rightarrow 0$ almost surely, and hence also in probability. Now since $u(n) = \lceil (1 - \lambda_\rho^*(\theta) + \epsilon)n \rceil$ and $\epsilon, \delta > 0$ are arbitrary, it follows that for any $\epsilon > 0$, we have $\lim_{n \rightarrow \infty} \mathbb{P}\left(\frac{N_2(n)}{n} \geq 1 - \lambda_\rho^*(\theta) + \epsilon\right) = 0$. From the proof for arm 1, we already know that $\lim_{n \rightarrow \infty} \mathbb{P}\left(\frac{N_2(n)}{n} \leq 1 - \lambda_\rho^*(\theta) - \epsilon\right) = 0$ holds for any $\epsilon > 0$. Therefore, it must be the case that $\frac{N_2(n)}{n} \xrightarrow{p} 1 - \lambda_\rho^*(\theta)$ and $\frac{N_1(n)}{n} \xrightarrow{p} \lambda_\rho^*(\theta)$, as desired. \square

G Proof of Theorem 2

G.1 Proof of part (I)

Let $\theta_k, \tilde{\theta}_k$ be Beta($1, k+1$)-distributed, with $\theta_k, \tilde{\theta}_l$ independent $\forall k, l \in \mathbb{N} \cup \{0\}$. In the two-armed bandit with deterministic 0 rewards, at any time $n+1$, the probability of playing arm 1 conditioned on the entire history up to that point, is given by $\mathbb{P}(\pi_{n+1} = 1 \mid \mathcal{F}_n) = \mathbb{P}(\theta_{N_1(n)} > \tilde{\theta}_{N_2(n)} \mid \mathcal{F}_n) = \frac{n - N_1(n) + 1}{n+2}$ (using Fact (2)). Since the arms are identical, and $N_1(n) + N_2(n) = n$, we must have $\mathbb{E}(N_1(n)/n) = 1/2 \forall n \in \mathbb{N}$ by symmetry. Define $Z_n := N_1(n)/n$. Then, Z_n evolves according to the following Markovian rule:

$$Z_{n+1} = \left(\frac{n}{n+1}\right) Z_n + \frac{Y(Z_n, n, \xi_n)}{n+1},$$

where $\{\xi_n\}$ is an independent noise process that is such that $Y(Z_n, n, \xi_n) \mid Z_n$ is distributed as Bernoulli $\left(\frac{n(1-Z_n)+1}{n+2}\right)$. Note that $Y(\cdot, \cdot, \cdot) \in \{0, 1\}$. Then,

$$\begin{aligned}
Z_{n+1}^2 &= \left(\frac{n}{n+1}\right)^2 Z_n^2 + \left(\frac{Y(Z_n, n, \xi_n)}{n+1}\right)^2 + \frac{2nZ_n Y(Z_n, n, \xi_n)}{(n+1)^2} \\
&= \left(\frac{n}{n+1}\right)^2 Z_n^2 + \frac{Y(Z_n, n, \xi_n)}{(n+1)^2} + \frac{2nZ_n Y(Z_n, n, \xi_n)}{(n+1)^2}.
\end{aligned}$$

Solving the recursion for Z_{n+1}^2 , we obtain

$$Z_{n+1}^2 = \frac{Z_1^2 + \sum_{t=1}^n [Y(Z_t, t, \xi_t) + 2tZ_t Y(Z_t, t, \xi_t)]}{(n+1)^2} = \frac{Z_1 + \sum_{t=1}^n [Y(Z_t, t, \xi_t) + 2tZ_t Y(Z_t, t, \xi_t)]}{(n+1)^2},$$

where the last equality follows since $Z_1 \in \{0, 1\}$. Taking expectations and using the fact that $\mathbb{E}Z_t = 1/2 \forall t \in \mathbb{N}$, yields

$$\mathbb{E}Z_{n+1}^2 = \frac{\frac{1}{2} + \sum_{t=1}^n \left(\frac{1}{2} + 2t\mathbb{E}\left[\frac{tZ_t(1-Z_t)+Z_t}{t+2}\right]\right)}{(n+1)^2}.$$

Using $Z_t(1 - Z_t) \leq 1/4$, we get the relation

$$\mathbb{E}Z_{n+1}^2 \leq \frac{1 + \sum_{t=1}^n \left(1 + t\mathbb{E}\left[\frac{t+4Z_t}{t+2}\right]\right)}{2(n+1)^2} = \frac{n+1 + \sum_{t=1}^n t}{2(n+1)^2} = \frac{n+2}{4(n+1)}.$$

Thus, $\text{Var}\left(\frac{N_1(n)}{n}\right) = \text{Var}(Z_n) \leq \frac{n+1}{4n} - \frac{1}{4} = \frac{1}{4n}$. Since $\mathbb{E}\left(\frac{N_1(n)}{n}\right) = \frac{1}{2}$, we conclude using Chebyshev's inequality that $\frac{N_1(n)}{n} \rightarrow \frac{1}{2}$ in probability as $n \rightarrow \infty$. \square

G.2 Proof of part (II)

Our proof of this part is essentially pivoted on showing the stronger result that $\mathbb{P}(N_1(n) = m) = \frac{1}{n+1}$ for any $m \in \{0, \dots, n\}$ and $n \in \mathbb{N}$. To this end, for an arbitrary m in said interval, let S_m be the set of sample-paths of length n such that $N_1(n) = m$ on each sample-path $s_m \in S_m$. Clearly, $|S_m| = \binom{n}{m}$. Let $i(s_m, t) \in \{1, 2\}$ denote the index of the arm pulled at time $t \in \{1, \dots, n\}$ on s_m , and let $\tilde{N}_j(t)$ denote the number of pulls of arm $j \in \{1, 2\}$ up to (and including) time t on s_m (with $\tilde{N}_1(0) = \tilde{N}_2(0) := 0$). Note that $i(s_m, t)$, $\tilde{N}_1(t)$ and $\tilde{N}_2(t)$ are deterministic for all $t \in \{1, \dots, n\}$, once s_m is fixed. Let $\theta_k, \tilde{\theta}_k$ be Beta($k+1, 1$)-distributed, with $\theta_k, \tilde{\theta}_k$ independent $\forall k, l \in \mathbb{N} \cup \{0\}$. It then follows that

$$\begin{aligned} \mathbb{P}(N_1(n) = m) &= \sum_{s_m \in S_m} \prod_{t=1}^n \mathbb{P}\left(\theta_{\tilde{N}_{i(s_m, t)}(t-1)} > \tilde{\theta}_{\tilde{N}_{\{1, 2\} \setminus i(s_m, t)}(t-1)}\right) \\ &= \sum_{s_m \in S_m} \prod_{t=1}^n \left(\frac{\tilde{N}_{i(s_m, t)}(t-1) + 1}{t+1}\right) && \text{(using Fact (3))} \\ &= \frac{1}{(n+1)!} \sum_{s_m \in S_m} \prod_{t=1}^n \left(\tilde{N}_{i(s_m, t)}(t-1) + 1\right) \\ &= \frac{1}{(n+1)!} \sum_{s_m \in S_m} m!(n-m)!, \end{aligned}$$

where the last equality follows since $\tilde{N}_1(n) = m$, $\tilde{N}_2(n) = n - m$, $\tilde{N}_1(0) = \tilde{N}_2(0) = 0$ on s_m . Therefore, we have for all $m \in \{0, \dots, n\}$ that

$$\mathbb{P}(N_1(n) = m) = \frac{\binom{n}{m} m!(n-m)!}{(n+1)!} = \frac{n!}{(n+1)!} = \frac{1}{n+1}. \quad (33)$$

This, in fact, proves a stronger result that $N_1(n)/n$ is uniformly distributed on $\{0, \frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n}, 1\}$ for any $n \in \mathbb{N}$. The desired result now follows as a corollary in the limit $n \rightarrow \infty$; for an arbitrary $x \in [0, 1]$, consider

$$\mathbb{P}\left(\frac{N_1(n)}{n} \leq x\right) = \sum_{m=0}^{\lfloor xn \rfloor} \mathbb{P}(N_1(n) = m) = \frac{\lfloor xn \rfloor + 1}{n+1},$$

where the last equality follows using (33). Thus, we have $\lim_{n \rightarrow \infty} \mathbb{P}\left(\frac{N_1(n)}{n} \leq x\right) = x$ for any $x \in [0, 1]$, i.e., $\frac{N_1(n)}{n}$ converges in law to the Uniform distribution on $[0, 1]$. \square

H Proof of Theorem 3

We essentially need to bound the growth rate of R_n^π under the policy π given by Algorithm 1 with $\rho > 1$, in three (exhaustive) regimes, viz., (i) $\Delta = o\left(\sqrt{\frac{\log n}{n}}\right)$ (“small gap”), (ii) $\Delta = \omega\left(\sqrt{\frac{\log n}{n}}\right)$ (“large gap”), and (iii) $\Delta = \Theta\left(\sqrt{\frac{\log n}{n}}\right)$ (“moderate gap”). We handle the three cases below separately.

H.1 The “small gap” regime

Here, we have

$$\frac{\mathbb{E}R_n^\pi}{\sqrt{n \log n}} \leq \frac{\Delta n}{\sqrt{n \log n}} = \sqrt{\frac{\Delta^2 n}{\log n}}.$$

Since $\Delta = o\left(\sqrt{\frac{\log n}{n}}\right)$, it follows that $\mathbb{E}R_n^\pi = o(\sqrt{n \log n})$. Therefore, we conclude using Markov’s inequality that $R_n^\pi = o_p(\sqrt{n \log n})$ whenever $\Delta = o\left(\sqrt{\frac{\log n}{n}}\right)$.

H.2 The “large gap” regime

In this regime, we have

$$\frac{\mathbb{E}R_n^\pi}{\sqrt{n \log n}} \leq \frac{C\rho\left(\frac{\log n}{\Delta} + \frac{\Delta}{\rho-1}\right)}{\sqrt{n \log n}},$$

where C is some absolute constant (follows from [6], Theorem 7). Since $\Delta = \omega\left(\sqrt{\frac{\log n}{n}}\right)$ and $\Delta \leq 1$ (rewards bounded in $[0, 1]$), it follows that $\mathbb{E}R_n^\pi = o(\sqrt{n \log n})$. Thus, we again conclude using Markov’s inequality that $R_n^\pi = o_p(\sqrt{n \log n})$ whenever $\Delta = \omega\left(\sqrt{\frac{\log n}{n}}\right)$.

H.3 The “moderate gap” regime

Since $\Delta = \Theta\left(\sqrt{\frac{\log n}{n}}\right)$, there exists some $\theta \in \mathbb{R}_+$ and a diverging sequence of natural numbers $\{n_k\}_{k \in \mathbb{N}}$ such that Δ scales with the horizon of play n_k along this sequence as $\Delta = \sqrt{\frac{\theta \log n_k}{n_k}}$. Without loss of generality, suppose that arm 1 is optimal, i.e., $\mu_1 \geq \mu_2$. We then have

$$\begin{aligned} \frac{R_{n_k}^\pi}{\sqrt{n_k \log n_k}} &= \frac{\mu_1 n_k - \sum_{j=1}^{N_1(n_k)} X_{1,j} - \sum_{j=1}^{N_2(n_k)} X_{2,j}}{\sqrt{n_k \log n_k}} && \text{(using (1))} \\ &= \frac{\mu_1 n_k - \mu_1 N_1(n_k) - \mu_2 N_2(n_k) - \sum_{j=1}^{N_1(n_k)} Y_{1,j} - \sum_{j=1}^{N_2(n_k)} Y_{2,j}}{\sqrt{n_k \log n_k}}, \end{aligned}$$

where $Y_{i,j} := X_{i,j} - \mu_i$, $i \in \{1, 2\}$, $j \in \mathbb{N}$. Therefore,

$$\begin{aligned} \frac{R_{n_k}^\pi}{\sqrt{n_k \log n_k}} &= \frac{\Delta N_2(n_k) - \sum_{j=1}^{N_1(n_k)} Y_{1,j} - \sum_{j=1}^{N_2(n_k)} Y_{2,j}}{\sqrt{n_k \log n_k}} \\ &= \Delta \sqrt{\frac{n_k}{\log n_k}} \left(\frac{N_2(n_k)}{n_k} \right) - \left(\frac{\sum_{j=1}^{N_1(n_k)} Y_{1,j} + \sum_{j=1}^{N_2(n_k)} Y_{2,j}}{\sqrt{n_k \log n_k}} \right) \\ &= \sqrt{\theta} \left(\frac{N_2(n_k)}{n_k} \right) - \left(\frac{\sum_{j=1}^{N_1(n_k)} Y_{1,j} + \sum_{j=1}^{N_2(n_k)} Y_{2,j}}{\sqrt{n_k \log n_k}} \right). \end{aligned} \quad (34)$$

Consider the summation terms above. We have

$$\begin{aligned} \left| \frac{\sum_{j=1}^{N_1(n_k)} Y_{1,j} + \sum_{j=1}^{N_2(n_k)} Y_{2,j}}{\sqrt{n_k \log n_k}} \right| &\leq \left| \frac{\sum_{j=1}^{N_1(n_k)} Y_{1,j}}{\sqrt{n_k \log n_k}} \right| + \left| \frac{\sum_{j=1}^{N_2(n_k)} Y_{2,j}}{\sqrt{n_k \log n_k}} \right| \\ &\leq \left| \frac{\sum_{j=1}^{N_1(n_k)} Y_{1,j}}{\sqrt{N_1(n_k) \log N_1(n_k)}} \right| + \left| \frac{\sum_{j=1}^{N_2(n_k)} Y_{2,j}}{\sqrt{N_2(n_k) \log N_2(n_k)}} \right|. \end{aligned} \quad (35)$$

Since $N_i(n_k)$, for each $i \in \{1, 2\}$, can be lower-bounded *path-wise* by a deterministic monotone increasing coercive function of k , say $f(k)$, we have

$$\left| \frac{\sum_{j=1}^{N_i(n_k)} Y_{i,j}}{\sqrt{N_i(n_k) \log N_i(n_k)}} \right| \leq \sup_{m \geq f(k)} \left| \frac{\sum_{j=1}^m Y_{i,j}}{\sqrt{m \log m}} \right| \quad \forall i \in \{1, 2\}. \quad (36)$$

Since $Y_{i,j}$'s are independent, zero-mean, bounded random variables, it follows from the Law of the Iterated Logarithm (see [11], Theorem 8.5.2) that

$$\sup_{m \geq f(k)} \left| \frac{\sum_{j=1}^m Y_{i,j}}{\sqrt{m \log m}} \right| \xrightarrow[k \rightarrow \infty]{\text{w.p. } 1} 0 \quad \forall i \in \{1, 2\}. \quad (37)$$

Combining (34), (35), (36) and (37), we conclude

$$\frac{R_{n_k}^\pi}{\sqrt{n_k \log n_k}} = \sqrt{\theta} \left(\frac{N_2(n_k)}{n_k} \right) + o_p(1).$$

From Theorem 1, we know that when $\Delta \sim \sqrt{\frac{\theta \log n_k}{n_k}}$, $\frac{N_2(n_k)}{n_k} \xrightarrow[k \rightarrow \infty]{p} 1 - \lambda_\rho^*(\theta)$. Thus, it follows that

$$\frac{R_{n_k}^\pi}{\sqrt{n_k \log n_k}} \xrightarrow[k \rightarrow \infty]{p} \sqrt{\theta} (1 - \lambda_\rho^*(\theta)) = h_\rho(\theta).$$

Since $\theta \in \mathbb{R}_+$ is arbitrary, the worst-case regret in the $\Delta = \Theta\left(\sqrt{\frac{\log n}{n}}\right)$ regime corresponds to the choice of θ given by $\theta_\rho^* = \arg \max_{\theta \geq 0} h_\rho(\theta)$. Since we already know that $R_n^\pi = o_p(\sqrt{n \log n})$ in the other two regimes (“small” and “large” gaps), it must be that the θ_ρ^* so obtained indeed corresponds to the global (in Δ) worst-case regret of Algorithm 1. \square

I Proof of Theorem 4

Notation. Let \mathcal{C} be the space of continuous functions $[0, 1] \mapsto \mathbb{R}^2$, endowed with the uniform metric. Let \mathcal{D} be the space of right-continuous functions with left limits, mapping $[0, 1] \mapsto \mathbb{R}^2$, and endowed with the Skorohod metric (see [8], Chapters 2 and 3, for an overview). Let \mathcal{D}_0 be the set of elements of \mathcal{D} of the form (ϕ_1, ϕ_2) , where ϕ_i is a non-decreasing real-valued function satisfying $0 \leq \phi_i(t) \leq 1$ for $i \in \{1, 2\}$ and $t \in [0, 1]$. For $t \in [0, 1]$, denote the identity map by $\epsilon(t) := t$.

For $i \in \{1, 2\}$ and $t \in [0, 1]$, define $\Psi_{i,n}(t) := \frac{\sum_{j=1}^{\lfloor nt \rfloor} X_{i,j} - \mu nt}{\sqrt{n}}$. Then, $(\Psi_{1,n}, \Psi_{2,n}) \in \mathcal{D}$. Also for $i \in \{1, 2\}$ and $t \in [0, 1]$, define $W_i(t) := \theta_i t + \sigma_i B'_i(t)$, where B'_1 and B'_2 are independent standard Brownian motions in \mathbb{R} . Note that $\mathbb{P}(W_i \in \mathcal{C}) = 1$ for $i \in \{1, 2\}$. Since $(X_{i,j})_{i \in \{1,2\}, j \in \mathbb{N}}$'s are independent random variables (i.i.d. within and independent across sequences), we know from Donsker's Theorem (see [8], Section 14, for details) that as $n \rightarrow \infty$,

$$(\Psi_{1,n}, \Psi_{2,n}) \Rightarrow (W_1, W_2) \text{ in } \mathcal{D}.$$

For $i \in \{1, 2\}$ and $t \in [0, 1]$, define $\phi_{i,n}(t) := \frac{N_i(\lfloor nt \rfloor)}{n}$. Thus, $(\phi_{1,n}, \phi_{2,n}) \in \mathcal{D}_0$, and it follows from the result for the “small gap” regime in Theorem 1 that as $n \rightarrow \infty$,

$$(\phi_{1,n}, \phi_{2,n}) \xrightarrow{p} \left(\frac{\epsilon}{2}, \frac{\epsilon}{2} \right) \text{ in } \mathcal{D}_0.$$

Thus, we have convergence in the product space (see [8], Theorem 3.9), i.e., as $n \rightarrow \infty$,

$$(\Psi_{1,n}, \Psi_{2,n}, \phi_{1,n}, \phi_{2,n}) \Rightarrow \left(W_1, W_2, \frac{\epsilon}{2}, \frac{\epsilon}{2} \right) \text{ in } \mathcal{D} \times \mathcal{D}_0.$$

For $i \in \{1, 2\}$ and $t \in [0, 1]$, define the composition $(\Psi_{i,n} \circ \phi_{i,n})(t) := \Psi_{i,n}(\phi_{i,n}(t))$, and $(W_i \circ \frac{\epsilon}{2})(t) := W_i\left(\frac{\epsilon(t)}{2}\right) = W_i\left(\frac{t}{2}\right)$. Since $W_1, W_2, \epsilon \in \mathcal{C}$ w.p. 1, it follows from the *random time-change lemma* (see [8], Section 14, for details) that as $n \rightarrow \infty$

$$(\Psi_{1,n} \circ \phi_{1,n}, \Psi_{2,n} \circ \phi_{2,n}) \Rightarrow \left(W_1 \circ \frac{\epsilon}{2}, W_2 \circ \frac{\epsilon}{2} \right) \text{ in } \mathcal{D}.$$

The stated assertion on cumulative rewards now follows by recognizing for $i \in \{1, 2\}$ and $t \in [0, 1]$ that $(\Psi_{i,n} \circ \phi_{i,n})(t) = \frac{\tilde{S}_{i,\lfloor nt \rfloor}}{\sqrt{n}}$, and defining $B_i(t) := \sqrt{2}B'_i\left(\frac{t}{2}\right)$. To prove the assertion on regret, assume without loss of generality that arm 1 is optimal, i.e., $\theta_1 \geq \theta_2$. Then, the result follows after a direct application of the Continuous Mapping Theorem (see [8], Theorem 2.7), to wit,

$$R_{\lfloor nt \rfloor}^\pi = \left(\mu + \frac{\theta_1}{\sqrt{n}} \right) \lfloor nt \rfloor - S_{1,\lfloor nt \rfloor} - S_{2,\lfloor nt \rfloor} = \frac{\theta_1 \lfloor nt \rfloor}{\sqrt{n}} - \left(\tilde{S}_{1,\lfloor nt \rfloor} + \tilde{S}_{2,\lfloor nt \rfloor} \right),$$

and therefore as $n \rightarrow \infty$,

$$\left(\frac{R_{\lfloor nt \rfloor}^\pi}{\sqrt{n}} \right)_{t \in [0,1]} \Rightarrow \left(\theta_1 t - \left(\left(\frac{\theta_1 + \theta_2}{2} \right) t + \frac{\sigma_1}{\sqrt{2}} B_1(t) + \frac{\sigma_2}{\sqrt{2}} B_2(t) \right) \right)_{t \in [0,1]} = \left(\frac{\Delta_0 t}{2} + \sqrt{\frac{\sigma_1^2 + \sigma_2^2}{2}} \tilde{B}(t) \right)_{t \in [0,1]},$$

$$\text{where } \tilde{B}(t) := -\sqrt{\frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2}} B_1(t) - \sqrt{\frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2}} B_2(t). \quad \square$$

J Proof of Theorem 5

We will prove this result in two parts; the preamble in J.1 below will prove a meta-result stating that $N_i(n)/n > 1/(2|\mathcal{I}|)$ with high probability (approaching 1 as $n \rightarrow \infty$) for any arm $i \in \mathcal{I}$. We will then leverage this meta-result to prove the assertions of the theorem in J.2.

J.1 Preamble

Let $L := |\mathcal{I}|$. If $L = 1$, the result follows trivially from the standard logarithmic bound for the expected regret (Theorem 7 in [6]), followed by Markov's inequality. Therefore, without loss of generality, suppose that $|\mathcal{I}| \geq 2$, and fix an arbitrary arm $i \in \mathcal{I}$. Then, we know that the following is true for any integer $u > 1$:

$$\begin{aligned} N_i(n) &\leq u + \sum_{t=u}^{n-1} \mathbb{1} \{ \pi_{t+1} = i, N_i(t) \geq u \} \\ &\leq u + \sum_{t=u}^{n-1} \mathbb{1} \left\{ \pi_{t+1} = i, N_i(t) \geq u, \sum_{j \in \mathcal{I} \setminus \{i\}} N_j(t) \leq t - u \right\}, \end{aligned}$$

where $\pi_{t+1} \in [K]$ indicates the arm played at time $t + 1$. In particular, the above holds also for $u = \lceil (\frac{1}{L} + \epsilon) n \rceil$, where $\epsilon \in (0, \frac{L-1}{L})$ is arbitrarily chosen. We will fix this u going forward, even though we may not always express its value explicitly for readability of the analysis that follows. We thus have

$$\begin{aligned} N_i(n) &\leq u + \sum_{t=u}^{n-1} \mathbb{1} \left\{ B_{i,N_i(t),t} \geq \max_{\hat{j} \in [K] \setminus \{i\}} B_{\hat{j},N_{\hat{j}}(t),t}, N_i(t) \geq u, \sum_{j \in \mathcal{I} \setminus \{i\}} N_j(t) \leq t - u \right\} \\ &\leq u + \sum_{t=u}^{n-1} \mathbb{1} \left\{ B_{i,N_i(t),t} \geq \max_{\hat{j} \in \mathcal{I} \setminus \{i\}} B_{\hat{j},N_{\hat{j}}(t),t}, N_i(t) \geq u, \sum_{j \in \mathcal{I} \setminus \{i\}} N_j(t) \leq t - u \right\}, \quad (38) \end{aligned}$$

where $B_{k,s,t} := \hat{X}_k(s) + \sqrt{(\rho \log t)/s}$ for $k \in [K]$, and $\hat{X}_k(s) := \sum_{l=1}^s X_{k,l}/s$ denotes the empirical mean reward from the "first s plays" of arm k (Note the distinction from $\bar{X}_k(s)$, which has been defined before as the empirical mean reward of arm k "at time s ," i.e., mean over its "first $N_k(s)$ plays"). Now observe that

$$\left\{ \sum_{j \in \mathcal{I} \setminus \{i\}} N_j(t) \leq t - u \right\} \subseteq \left\{ \exists j \in \mathcal{I} \setminus \{i\} : N_j(t) \leq \frac{t - u}{L - 1} \right\} \subseteq \left\{ \exists j \in \mathcal{I} \setminus \{i\} : N_j(t) \leq \left(\frac{1}{L} - \frac{\epsilon}{L - 1} \right) t \right\}, \quad (39)$$

where the last inclusion follows using $u = \lceil (\frac{1}{L} + \epsilon) n \rceil$ and $n \geq t$. Combining (38) and (39) using the Union bound, we obtain

$$\begin{aligned}
N_i(n) &\leq u + \sum_{t=u}^{n-1} \sum_{j \in \mathcal{I} \setminus \{i\}} \mathbb{1} \left\{ B_{i, N_i(t), t} \geq \max_{\hat{j} \in \mathcal{I} \setminus \{i\}} B_{\hat{j}, N_{\hat{j}}(t), t}, N_i(t) \geq u, N_j(t) \leq \left(\frac{1}{L} - \frac{\epsilon}{L-1} \right) t \right\} \\
&\leq u + \sum_{t=u}^{n-1} \sum_{j \in \mathcal{I} \setminus \{i\}} \mathbb{1} \left\{ B_{i, N_i(t), t} \geq B_{j, N_j(t), t}, N_i(t) \geq u, N_j(t) \leq \left(\frac{1}{L} - \frac{\epsilon}{L-1} \right) t \right\} \\
&\leq u + \underbrace{\sum_{t=u}^{n-1} \sum_{j \in \mathcal{I} \setminus \{i\}} \mathbb{1} \left\{ B_{i, N_i(t), t} \geq B_{j, N_j(t), t}, N_i(t) \geq \left(\frac{1}{L} + \epsilon \right) t, N_j(t) \leq \left(\frac{1}{L} - \frac{\epsilon}{L-1} \right) t \right\}}_{=: Z_n}, \tag{40}
\end{aligned}$$

where the last inequality again uses $u = \lceil (\frac{1}{L} + \epsilon) n \rceil$ and $n \geq t$. Define the events:

$E_i := \{N_i(t) \geq (\frac{1}{L} + \epsilon) t\}$, and $E_j := \{N_j(t) \leq (\frac{1}{L} - \frac{\epsilon}{L-1}) t\}$. Now,

$$\begin{aligned}
&\mathbb{E} Z_n \\
&= \sum_{t=u}^{n-1} \sum_{j \in \mathcal{I} \setminus \{i\}} \mathbb{P} (B_{i, N_i(t), t} \geq B_{j, N_j(t), t}, E_i, E_j) \\
&= \sum_{t=u}^{n-1} \sum_{j \in \mathcal{I} \setminus \{i\}} \mathbb{P} \left(\hat{Y}_i(N_i(t)) - \hat{Y}_j(N_j(t)) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_j(t)}} - \frac{1}{\sqrt{N_i(t)}} \right), E_i, E_j \right), \tag{41}
\end{aligned}$$

where $\hat{Y}_k(s) := \sum_{l=1}^s Y_{k,l}/s$ and $Y_{k,l} := X_{k,l} - \mathbb{E} X_{k,l}$ for $k \in [K]$, $s \in \mathbb{N}$, $l \in \mathbb{N}$. The last equality above follows since $i, j \in \mathcal{I}$ and the mean rewards of arms in \mathcal{I} are equal. Thus,

$$\mathbb{E} Z_n \leq \sum_{t=u}^{n-1} \sum_{j \in \mathcal{I} \setminus \{i\}} \sum_{m_i = \lceil (\frac{1}{L} + \epsilon) t \rceil}^t \sum_{m_j = 1}^{\lfloor (\frac{1}{L} - \frac{\epsilon}{L-1}) t \rfloor} \mathbb{P} \left(\hat{Y}_i(m_i) - \hat{Y}_j(m_j) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{m_j}} - \frac{1}{\sqrt{m_i}} \right) \right).$$

Since $\mathbb{E} [\hat{Y}_i(m_i) - \hat{Y}_j(m_j)] = 0$, and $m_j < m_i$ over the range of the summation above, we can use the Chernoff-Hoeffding bound (Fact 1) to obtain

$$\mathbb{E} Z_n \leq \sum_{t=u}^{n-1} \sum_{j \in \mathcal{I} \setminus \{i\}} \sum_{m_i = \lceil (\frac{1}{L} + \epsilon) t \rceil}^t \sum_{m_j = 1}^{\lfloor (\frac{1}{L} - \frac{\epsilon}{L-1}) t \rfloor} \exp \left[-2\rho \left(1 - 2\sqrt{\frac{m_i m_j}{(m_i + m_j)^2}} \right) \log t \right]. \tag{42}$$

Let $\gamma := m_i/(m_i + m_j)$. Then, $\gamma \geq \frac{\frac{1}{L} + \epsilon}{\frac{1}{L} + (\frac{L-2}{L-1})\epsilon} > 1/2$ over the range of the summation in (42).

Consequently, $\gamma(1 - \gamma)$ is maximized at $\gamma = \frac{\frac{1}{L} + \epsilon}{\frac{1}{L} + (\frac{L-2}{L-1})\epsilon}$, and therefore, $m_i m_j / (m_i + m_j)^2 = \gamma(1 - \gamma) \leq (f(\epsilon, L))^2 < 1/4$ in (42), where $f(\epsilon, L)$ as defined as:

$$f(\epsilon, L) := \sqrt{\frac{(L-1)(1 + \epsilon L)(L-1 - \epsilon L)}{(2(L-1) + L(L-2)\epsilon)^2}}. \tag{43}$$

Combining (42) and (43), we obtain

$$\mathbb{E} Z_n \leq (L-1) \sum_{t=u}^{n-1} t^{-2(\rho-1-2\rho f(\epsilon, L))}. \tag{44}$$

Now consider an arbitrary $\delta > 0$. From (40), we have

$$\mathbb{P} (N_i(n) \geq u + \delta n) \leq \mathbb{P} (Z_n \geq \delta n) \stackrel{(*)}{\leq} \frac{\mathbb{E} Z_n}{\delta n} \stackrel{(\dagger)}{\leq} \left(\frac{L-1}{\delta n} \right) \sum_{t=u}^{n-1} t^{-2(\rho(1-2f(\epsilon, L))-1)}, \tag{45}$$

where (\star) is due to Markov's inequality, and (\dagger) follows using (44). Observe from (43) that $f(\epsilon, L)$ is monotone decreasing in ϵ over the interval $\epsilon \in [0, \frac{L-1}{L}]$, with $f(0, L) = 1/2$ and $f(\frac{L-1}{L}, L) = 0$. Therefore, $1 - 2f(\epsilon, L) > 0$ in the interval $\epsilon \in (0, \frac{L-1}{L}]$. Thus, for ρ large enough, the exponent of t in (45) can be made arbitrarily small. That is, $\exists \rho_0 \in \mathbb{R}_+$ s.t. for all $\rho \geq \rho_0$, we have $2(\rho(1 - 2f(\epsilon, L)) - 1) > 0 \forall \epsilon \in [\frac{1}{2L(L-1)}, \frac{L-1}{L}]$. Now supposing $\rho \geq \rho_0$, plug in $\epsilon = \frac{1}{2L(L-1)}$ in (45) (this includes substituting $u = \lceil (\frac{1}{L} + \frac{1}{2L(L-1)})n \rceil$). Then since $\delta > 0$ and $i \in \mathcal{I}$ are arbitrary, it follows that for any $\delta > 0$ and $i \in \mathcal{I}$,

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(N_i(n) \geq \left(\frac{1}{L} + \frac{1}{2L(L-1)} + \delta \right) n \right) = \frac{L^{2\rho}}{\delta} \lim_{n \rightarrow \infty} n^{-2(\rho(1-2f(\frac{1}{2L(L-1)}, L)) - 1)} = 0. \quad (46)$$

Notice that for any $\delta > 0$ and $i \in \mathcal{I}$,

$$\begin{aligned} \mathbb{P} \left(N_i(n) + \sum_{j \in [K] \setminus \mathcal{I}} N_j(n) \leq \left(\frac{1}{2L} - (L-1)\delta \right) n \right) &= \mathbb{P} \left(\sum_{j \in \mathcal{I} \setminus \{i\}} N_j(n) \geq (L-1) \left(\frac{1}{L} + \frac{1}{2L(L-1)} + \delta \right) n \right) \\ &\leq \sum_{j \in \mathcal{I} \setminus \{i\}} \mathbb{P} \left(N_j(n) \geq \left(\frac{1}{L} + \frac{1}{2L(L-1)} + \delta \right) n \right), \end{aligned}$$

where the last inequality follows using the Union bound. Taking limits on both sides above, we conclude using (46) that for any $\delta > 0$ and $i \in \mathcal{I}$,

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(N_i(n) + \sum_{j \in [K] \setminus \mathcal{I}} N_j(n) \leq \left(\frac{1}{2L} - (L-1)\delta \right) n \right) = 0. \quad (47)$$

If $\mathcal{I} = [K]$, the conclusion that $N_i(n)/n > 1/(2L) = 1/(2K)$ with high probability (approaching 1 as $n \rightarrow \infty$) for all $i \in [K]$, is immediate from (47). If $\mathcal{I} \neq [K]$, then $\sum_{j \in [K] \setminus \mathcal{I}} \mathbb{E}(N_j(n)/n) \leq CK\rho \left[\left(\frac{1}{\Delta_{\min}^2} \right) \left(\frac{\log n}{n} \right) + \frac{1}{(\rho-1)n} \right]$ for some absolute constant $C > 0$ follows from [6], Theorem 7.

Consequently if $\Delta_{\min} = \omega \left(\sqrt{\frac{\log n}{n}} \right)$, Markov's inequality implies that $\sum_{j \in [K] \setminus \mathcal{I}} N_j(n)/n = o_p(1)$. Thus, it again follows using (47) that $N_i(n)/n > 1/(2L)$ with high probability (approaching 1 as $n \rightarrow \infty$) for all $i \in \mathcal{I}$. \square

J.2 Proof of part (I) and (II)

Note that the following holds for any integer $u > 1$ and any arm $i \in [K]$:

$$N_i(n) \leq u + \sum_{t=u+1}^n \mathbb{1} \{ \pi_t = i, N_i(t-1) \geq u \},$$

where $\pi_t \in [K]$ indicates the arm played at time t . In particular, the above is true also for $u = N_j(n) + \lceil \epsilon n \rceil$, where $j \in [K] \setminus \{i\}$ and $\epsilon > 0$ are arbitrarily chosen. Without loss of generality, suppose that $|\mathcal{I}| \geq 2$ (the result is trivial for $|\mathcal{I}| = 1$), and fix two arbitrary arms $i, j \in \mathcal{I}$. Then,

$$\begin{aligned} N_i(n) &\leq N_j(n) + \lceil \epsilon n \rceil + \sum_{t=N_j(n)+\lceil \epsilon n \rceil+1}^n \mathbb{1} \{ \pi_t = i, N_i(t-1) \geq N_j(n) + \lceil \epsilon n \rceil \} \\ &\leq N_j(n) + \lceil \epsilon n \rceil + \sum_{t=\lceil \epsilon n \rceil+1}^n \mathbb{1} \{ \pi_t = i, N_i(t-1) \geq N_j(n) + \epsilon n \} \\ &\leq N_j(n) + \lceil \epsilon n \rceil + \sum_{t=\lceil \epsilon n \rceil+1}^n \mathbb{1} \{ B_{i, N_i(t-1), t-1} \geq B_{j, N_j(t-1), t-1}, N_i(t-1) \geq N_j(n) + \epsilon n \}, \end{aligned}$$

where $B_{k,s,t} := \hat{X}_k(s) + \sqrt{(\rho \log t)/s}$ for $k \in [K]$, with $\hat{X}_k(s)$ denoting the empirical mean reward from the first s plays of arm k . Then,

$$\begin{aligned}
N_i(n) &\leq N_j(n) + \lceil \epsilon n \rceil + \sum_{t=\lceil \epsilon n \rceil}^{n-1} \mathbb{1} \{B_{i,N_i(t),t} \geq B_{j,N_j(t),t}, N_i(t) \geq N_j(n) + \epsilon n\} \\
&\leq N_j(n) + \lceil \epsilon n \rceil + \sum_{t=\lceil \epsilon n \rceil}^{n-1} \mathbb{1} \{B_{i,N_i(t),t} \geq B_{j,N_j(t),t}, N_i(t) \geq N_j(t) + \epsilon t\} \\
&\leq N_j(n) + \epsilon n + 1 + Z_n,
\end{aligned} \tag{48}$$

where $Z_n := \sum_{t=\lceil \epsilon n \rceil}^{n-1} \mathbb{1} \{B_{i,N_i(t),t} \geq B_{j,N_j(t),t}, N_i(t) \geq N_j(t) + \epsilon t\}$. Now,

$$\begin{aligned}
&\mathbb{E}Z_n \\
&= \sum_{t=\lceil \epsilon n \rceil}^{n-1} \mathbb{P} \{B_{i,N_i(t),t} \geq B_{j,N_j(t),t}, N_i(t) \geq N_j(t) + \epsilon t\} \\
&= \sum_{t=\lceil \epsilon n \rceil}^{n-1} \mathbb{P} \left\{ \hat{X}_i(N_i(t)) - \hat{X}_j(N_j(t)) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_j(t)}} - \frac{1}{\sqrt{N_i(t)}} \right), N_i(t) \geq N_j(t) + \epsilon t \right\} \\
&= \sum_{t=\lceil \epsilon n \rceil}^{n-1} \mathbb{P} \left\{ \hat{Y}_i(N_i(t)) - \hat{Y}_j(N_j(t)) \geq \sqrt{\rho \log t} \left(\frac{1}{\sqrt{N_j(t)}} - \frac{1}{\sqrt{N_i(t)}} \right), N_i(t) \geq N_j(t) + \epsilon t \right\},
\end{aligned}$$

where $\hat{Y}_k(s) := \sum_{l=1}^s Y_{k,l}/s$ for $k \in [K]$, $s \in \mathbb{N}$, with $Y_{k,l} := X_{k,l} - \mathbb{E}X_{k,l}$ for $l \in \mathbb{N}$. The last equality above follows since $i, j \in \mathcal{I}$ and the mean rewards of arms in \mathcal{I} are equal. Thus,

$$\begin{aligned}
&\mathbb{E}Z_n \\
&= \sum_{t=\lceil \epsilon n \rceil}^{n-1} \mathbb{P} \left(\hat{Y}_i(N_i(t)) - \hat{Y}_j(N_j(t)) \geq \sqrt{\frac{\rho \log t}{N_i(t)}} \left(\sqrt{\frac{N_i(t)}{N_j(t)}} - 1 \right), N_i(t) \geq N_j(t) + \epsilon t \right) \\
&\leq \sum_{t=\lceil \epsilon n \rceil}^{n-1} \mathbb{P} \left(\hat{Y}_i(N_i(t)) - \hat{Y}_j(N_j(t)) \geq \sqrt{\frac{\rho \log t}{t}} (\sqrt{1+\epsilon} - 1), N_i(t) \geq N_j(t) + \epsilon t \right) \\
&\leq \sum_{t=\lceil \epsilon n \rceil}^{n-1} \mathbb{P} (W_t \geq \sqrt{1+\epsilon} - 1),
\end{aligned} \tag{49}$$

where $W_t := \sqrt{\frac{t}{\rho \log t}} \left(\frac{\sum_{l=1}^{N_i(t)} Y_{i,l}}{N_i(t)} - \frac{\sum_{l=1}^{N_j(t)} Y_{j,l}}{N_j(t)} \right)$. Now,

$$\begin{aligned}
&|W_t| \\
&\leq \sqrt{\frac{t}{\rho \log t}} \left(\left| \frac{\sum_{l=1}^{N_i(t)} Y_{i,l}}{N_i(t)} \right| + \left| \frac{\sum_{l=1}^{N_j(t)} Y_{j,l}}{N_j(t)} \right| \right) \\
&= \sqrt{\frac{2t}{\rho \log t}} \left(\sqrt{\frac{\log \log N_i(t)}{N_i(t)}} \left| \frac{\sum_{l=1}^{N_i(t)} Y_{i,l}}{\sqrt{2N_i(t) \log \log N_i(t)}} \right| + \sqrt{\frac{\log \log N_j(t)}{N_j(t)}} \left| \frac{\sum_{l=1}^{N_j(t)} Y_{j,l}}{\sqrt{2N_j(t) \log \log N_j(t)}} \right| \right) \\
&\leq \sqrt{\frac{2t}{\rho \log t}} \left(\sqrt{\frac{\log \log t}{N_i(t)}} \left| \frac{\sum_{l=1}^{N_i(t)} Y_{i,l}}{\sqrt{2N_i(t) \log \log N_i(t)}} \right| + \sqrt{\frac{\log \log t}{N_j(t)}} \left| \frac{\sum_{l=1}^{N_j(t)} Y_{j,l}}{\sqrt{2N_j(t) \log \log N_j(t)}} \right| \right) \\
&= \sqrt{\frac{2 \log \log t}{\rho \log t}} \left(\sqrt{\frac{t}{N_i(t)}} \left| \frac{\sum_{l=1}^{N_i(t)} Y_{i,l}}{\sqrt{2N_i(t) \log \log N_i(t)}} \right| + \sqrt{\frac{t}{N_j(t)}} \left| \frac{\sum_{l=1}^{N_j(t)} Y_{j,l}}{\sqrt{2N_j(t) \log \log N_j(t)}} \right| \right).
\end{aligned} \tag{50}$$

We know that $N_k(t)$, for any arm $k \in [K]$, can be lower-bounded *path-wise* by a deterministic monotone increasing *coercive* function of t , say $h(t)$. This follows as a trivial consequence of the

structure of the policy, and the fact that the rewards are uniformly bounded. Therefore, we have for any arm $k \in \mathcal{I}$ that

$$\begin{aligned} & \left| \frac{\sum_{l=1}^{N_k(t)} Y_{k,l}}{\sqrt{2N_k(t) \log \log N_k(t)}} \right| \leq \sup_{m \geq h(t)} \left| \frac{\sum_{l=1}^m Y_{k,l}}{\sqrt{2m \log \log m}} \right| \\ \implies & \limsup_{t \rightarrow \infty} \left| \frac{\sum_{l=1}^{N_k(t)} Y_{k,l}}{\sqrt{2N_k(t) \log \log N_k(t)}} \right| \leq \limsup_{t \rightarrow \infty} \left| \frac{\sum_{l=1}^t Y_{k,l}}{\sqrt{2t \log \log t}} \right|. \end{aligned} \quad (51)$$

For any $k \in \mathcal{I}$, we know that $\{Y_{k,l} : l \in \mathbb{N}\}$ is a collection of i.i.d. random variables with $\mathbb{E}Y_{k,1} = 0$ and $\text{Var}(Y_{k,1}) = \text{Var}(X_{k,1}) \leq 1$. Therefore, we conclude using the Law of the Iterated Logarithm (see Theorem 8.5.2 in [11]) in (51) that

$$\limsup_{t \rightarrow \infty} \left| \frac{\sum_{l=1}^{N_k(t)} Y_{k,l}}{\sqrt{2N_k(t) \log \log N_k(t)}} \right| \leq 1 \quad \text{w.p. 1 } \forall k \in \mathcal{I}. \quad (52)$$

Using (50), (52), and the meta-result from the preamble in J.1 that $N_k(t)/t > 1/(2|\mathcal{I}|)$ with high probability (approaching 1 as $t \rightarrow \infty$) for any arm $k \in \mathcal{I}$, we conclude that

$$W_t \xrightarrow{p} 0 \quad \text{as } t \rightarrow \infty. \quad (53)$$

Now,

$$\mathbb{P} \left(\frac{N_i(n) - N_j(n)}{n} \geq 2\epsilon \right) \stackrel{(\dagger)}{\leq} \mathbb{P}(1 + Z_n \geq \epsilon n) \stackrel{(\ddagger)}{\leq} \frac{1 + \mathbb{E}Z_n}{\epsilon n} \stackrel{(\star)}{\leq} \frac{1}{\epsilon n} + \frac{1}{\epsilon n} \sum_{t=\lceil \epsilon n \rceil}^{n-1} \mathbb{P}(W_t > \sqrt{1+\epsilon} - 1),$$

where (\dagger) follows using (48), (\ddagger) using Markov's inequality, and (\star) from (49). Therefore,

$$\mathbb{P} \left(\frac{N_i(n) - N_j(n)}{n} \geq 2\epsilon \right) \leq \frac{1}{\epsilon n} + \left(\frac{1-\epsilon}{\epsilon} \right) \sup_{t \geq \epsilon n} \mathbb{P}(W_t > \sqrt{1+\epsilon} - 1). \quad (54)$$

Since $\epsilon > 0$ is arbitrary, we conclude using (53) and (54) that for any $\epsilon > 0$,

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\frac{N_i(n) - N_j(n)}{n} \geq 2\epsilon \right) = 0. \quad (55)$$

Our proof is symmetric w.r.t. the labels i, j , therefore, an identical result holds also with the labels interchanged in (55). Thus, we have $N_i(n)/n - N_j(n)/n \xrightarrow{p} 0$. Since i, j are arbitrary in \mathcal{I} , the aforementioned convergence holds for any pair of arms in \mathcal{I} . Now if $\mathcal{I} = [K]$, we are done. If $\mathcal{I} \neq [K]$, then $\sum_{i \in [K] \setminus \mathcal{I}} \mathbb{E}(N_i(n)/n) \leq CK\rho \left[\left(\frac{1}{\Delta_{\min}^2} \right) \left(\frac{\log n}{n} \right) + \frac{1}{(\rho-1)n} \right]$ for some absolute constant $C > 0$ follows from Theorem 7 in [6]. Consequently if $\Delta_{\min} = \omega \left(\sqrt{\frac{\log n}{n}} \right)$, it would follow from Markov's inequality that $\sum_{i \in [K] \setminus \mathcal{I}} N_i(n)/n = o_p(1)$. Thus, any arm $i \in \mathcal{I}$ must satisfy $N_i(n)/n \xrightarrow{p} 1/|\mathcal{I}|$. \square

K Proof of Fact 2 and Fact 3

K.1 Fact 2

Since θ_k is Beta($1, k+1$)-distributed, its PDF, say $f_k(\cdot)$, is given by

$$f_k(x) = (k+1)(1-x)^k; \quad x \in [0, 1]. \quad (56)$$

Thus,

$$\begin{aligned} \mathbb{P}(\theta_k > \tilde{\theta}_l) &= \int_0^1 \left(\int_y^1 f_k(x) dx \right) f_l(y) dy \\ &= \int_0^1 \left(\int_y^1 (k+1)(1-x)^k dx \right) (l+1)(1-y)^l dy \quad (\text{using (56)}) \\ &= \int_0^1 (l+1)(1-y)^{k+l+1} dy \\ &= \frac{l+1}{k+l+2}. \end{aligned}$$

□

K.2 Fact 3

Since θ_k is Beta($k + 1, 1$)-distributed, its PDF, say $f_k(\cdot)$, is given by

$$f_k(x) = (k + 1)x^k; \quad x \in [0, 1]. \quad (57)$$

Thus,

$$\begin{aligned} \mathbb{P}(\theta_k > \tilde{\theta}_l) &= \int_0^1 \left(\int_y^1 f_k(x) dx \right) f_l(y) dy \\ &= \int_0^1 \left(\int_y^1 (k + 1)x^k dx \right) (l + 1)y^l dy && \text{(using (57))} \\ &= \int_0^1 (l + 1) (1 - y^{k+1}) y^l dy \\ &= 1 - \frac{l + 1}{k + l + 2} \\ &= \frac{k + 1}{k + l + 2}. \end{aligned}$$

□