
Sequential Test for the Lowest Mean: From Thompson to Murphy Sampling

Emilie Kaufmann¹ Wouter M. Koolen² Aurélien Garivier³

¹ CNRS & U. Lille, CRIStAL / SequeL Inria Lille, emilie.kaufmann@univ-lille.fr

² Centrum Wiskunde & Informatica, Amsterdam, wmkoolen@cwi.nl

³ UMPA, École normale supérieure de Lyon, aurelien.garivier@ens-lyon.fr

Abstract

Learning the minimum/maximum mean among a finite set of distributions is a fundamental sub-task in planning, game tree search and reinforcement learning. We formalize this learning task as the problem of sequentially testing how the minimum mean among a finite set of distributions compares to a given threshold. We develop refined non-asymptotic lower bounds, which show that optimality mandates very different sampling behavior for a low vs high true minimum. We show that Thompson Sampling and the intuitive Lower Confidence Bounds policy each nail only one of these cases. We develop a novel approach that we call Murphy Sampling. Even though it entertains exclusively low true minima, we prove that MS is optimal for both possibilities. We then design advanced self-normalized deviation inequalities, fueling more aggressive stopping rules. We complement our theoretical guarantees by experiments showing that MS works best in practice.

1 Introduction

We consider a collection of core problems related to *minimums of means*. For a given finite collection of probability distributions parameterized by their means μ_1, \dots, μ_K , we are interested in learning about $\mu^* = \min_a \mu_a$ from adaptive samples $X_t \sim \mu_{A_t}$, where A_t indicates the distribution sampled at time t . We shall refer to these distributions as arms in reference to a multi-armed bandit model [28, 26]. Knowing about minima/maxima is crucial in reinforcement learning or game-playing, where the value of a state for an agent is the *maximum* over actions of the (expected) successor state value or the *minimum* over adversary moves of the next state value.

The problem of estimating $\mu^* = \min_a \mu_a$ was studied in [34] and subsequently [7, 31, 8]. It is known that no unbiased estimator exists for μ^* , and that estimators face an intricate bias-variance trade-off. Beyond estimation, the problem of constructing *confidence intervals* on minima/maxima naturally arises in (Monte Carlo) planning in Markov Decision Processes [15] and games [25]. Such confidence intervals are used hierarchically for Monte Carlo Tree Search (MCTS) in [32, 11, 17, 20]. The open problem of designing asymptotically optimal algorithms for MCTS led us to isolate one core difficulty that we study here, namely the construction of confidence intervals and associated sampling/stopping rules for learning minima (and, by symmetry, maxima).

Confidence intervals (that are uniform over time) can be naturally obtained from a (sequential) test of $\{\mu^* < \gamma\}$ versus $\{\mu^* > \gamma\}$, given a threshold γ . The main focus of the paper goes even further and investigates the minimum number of samples required for *adaptively* testing whether $\{\mu^* < \gamma\}$ or $\{\mu^* > \gamma\}$, that is sequentially sampling the arms in order to decide for one hypothesis as quickly as possible. Such a problem is interesting in its own right as it naturally arises in several statistical certification applications. As an example we may consider quality control testing in manufacturing, where we want to certify that in a batch of machines each has a guaranteed probability of successfully producing a widget. In e-learning, we may want to certify that a given student has

sufficient understanding of a range of subjects, asking as few questions as possible about the different subjects. Then in anomaly detection, we may want to flag the presence of an anomaly faster the more anomalies are present. Finally, in a crowdsourcing system, we may need to establish as quickly as possible whether a cohort of workers contains at least one unacceptably careless worker. Our own motivation for studying this problem is that it corresponds to an especially simple instance of the depth-two game tree search problem, as illustrated in Figure 1.

We thus study a particular example of sequential adaptive hypothesis testing problem, as introduced by Chernoff [5], in which multiple experiments (sampling from one arm) are available to the experimenter, each of which allows to gain different information about the hypotheses. The experimenter sequentially selects which experiment to perform, when to stop and then which hypothesis to recommend. Several recent works from the bandit literature fit into this framework, with the twist that they consider continuous, composite hypotheses and aim for δ -correct testing: the probability of guessing a wrong hypothesis has to be smaller than δ , while performing as few experiments as possible. The fixed-confidence *Best Arm Identification* problem (concerned with finding the arm with largest mean) is one such example [9, 23], of which several variants have been studied [19, 17, 12]. For example the Thresholding Bandit Problem [27] aims at finding the set of arms above a threshold, which is strictly harder than our testing problem. In the Ranking and Selection literature (see e.g., [14] for a survey) the related problem of finding systems whose expected performance is smaller than a known standard has been studied by [24], but if such system exist, the goal was to additionally identify the one with smallest expectation, which is strictly harder than our problem.




Figure 1: Game tree search problem of “depth $1\frac{1}{2}$ ”. We consider the scenario where it has been established that the right subtree (grey) of the root has value γ . Learning the optimal action at the root (orange) is equivalent to determining how the minimum (blue) of the leaf means (green) compares to γ .

A full characterization of the asymptotic complexity of the BAI problem was recently given in [11], highlighting the existence of an *optimal allocation of samples* across arms. The lower bound technique introduced therein can be generalized to virtually any testing problem in a bandit model (see, e.g. [20, 12]). Such an optimal allocation is also presented by [4] in the GENERAL-SAMP framework, which is quite generic and in particular encompasses testing on which side of γ the minimum falls. The proposed LPSAMPLE algorithm is thus a candidate to be applied to our testing problem. However, this algorithm is only proved to be *order-optimal*, that is to attain the minimal sample complexity up to a (large) multiplicative constant. Moreover, like other algorithms for special cases (e.g. TRACK-AND-STOP for BAI [11]), it relies on *forced exploration*, which may be harmful in practice and leads to unavoidably asymptotic analysis.

Our first contribution is a tight lower bound on the sample complexity that provides an oracle sample allocation, but also aims at reflecting the moderate-risk behavior of a δ -correct algorithm. Our second contribution is a new sampling rule for the minimum testing problem, under which the empirical fraction of selections converges to the optimal allocation without forced exploration. The algorithm is a variant of Thompson Sampling [33, 1] that is conditioning on the “worst” outcome $\mu^* < \gamma$, hence the name Murphy Sampling. This conditioning is inspired by the Top Two Thompson Sampling recently proposed by [29] for Best Arm Identification. As we shall see, the optimal allocation is very different whether $\mu^* < \gamma$ or $\mu^* > \gamma$ and yet Murphy Sampling automatically adopts the right behavior in each case. Our third contribution is a new stopping rule, that by aggregating samples from several arms that look small may lead to early stopping whenever $\mu^* < \gamma$. This stopping rule is based on a new self-normalized deviation inequality for exponential families (Theorem 7) of independent interest. It generalizes results obtained by [18, 23] in the Gaussian case and by [3] without the uniformity in time, and also handles subsets of arms.

The rest of the paper is structured as follows. In Section 2 we introduce our notation and formally define our objective. In Section 3, we present lower bounds on the sample complexity of sequential tests for minima. In particular, we compute the optimal allocations for this problem and discuss the limitation of naive benchmarks to attain them. In Section 4 we introduce Murphy sampling, and prove its optimality in conjunction with a simple stopping rule. Improved stopping rules (and

confidence intervals) are presented in Section 5. Finally, numerical experiments reported in Section 6 demonstrate the efficiency of Murphy Sampling paired with our new stopping rule.

2 Setup

We consider a family of K probability distributions that belong to a one-parameter canonical exponential family, that we shall call *arms* in reference to a multi-armed bandit model. Such exponential families include Gaussian with known variance, Bernoulli, Poisson, see [3] for details. For natural parameter ν , the density of the distribution w.r.t. carrier measure ρ on \mathbb{R} is given by $e^{x\nu - b(\nu)}\rho(dx)$, where the cumulant generating function $b(\nu) = \ln \mathbb{E}_\rho[e^{X\nu}]$ induces a bijection $\nu \mapsto \dot{b}(\nu)$ to the mean parameterization. We write $\text{KL}(\nu, \lambda)$ and $d(\mu, \theta)$ for the Kullback-Leibler divergence from natural parameters ν to λ and from mean parameters μ to θ . Specifically, with convex conjugate b_* ,

$$\text{KL}(\nu, \lambda) = b(\lambda) - b(\nu) + (\nu - \lambda)\dot{b}(\nu) \quad \text{and} \quad d(\mu, \theta) = b_*(\mu) - b_*(\theta) - (\mu - \theta)\dot{b}_*(\theta).$$

We denote by $\boldsymbol{\mu} = (\mu_1, \dots, \mu_K) \in \mathcal{I}^K$ the vector of arm means, which fully characterizes the model. In this paper, we are interested in the smallest mean (and the arm where it is attained)

$$\mu^* = \min_a \mu_a \quad \text{and} \quad a^* = a^*(\boldsymbol{\mu}) = \arg \min_a \mu_a.$$

Given a threshold $\gamma \in \mathcal{I}$, our goal is to decide whether $\mu^* < \gamma$ or $\mu^* > \gamma$. We introduce the hypotheses

$$\mathcal{H}_< = \{\boldsymbol{\mu} \in \mathcal{I}^K \mid \mu^* < \gamma\} \quad \text{and} \quad \mathcal{H}_> = \{\boldsymbol{\mu} \in \mathcal{I}^K \mid \mu^* > \gamma\}, \quad \text{and their union} \quad \mathcal{H} = \mathcal{H}_< \cup \mathcal{H}_>.$$

We want to propose a sequential and adaptive testing procedure, that consists in a *sampling rule* A_t , a *stopping rule* τ and a *decision rule* $\hat{m} \in \{<, >\}$. The algorithm samples $X_t \sim \mu_{A_t}$ while $t \leq \tau$, and then outputs a decision \hat{m} . We denote the information available after t rounds by $\mathcal{F}_t = \sigma(A_1, X_1, \dots, A_t, X_t)$. A_t is measurable with respect to \mathcal{F}_{t-1} and possibly some exogenous random variable, τ is a stopping time with respect to this filtration and \hat{m} is \mathcal{F}_τ -measurable.

Given a risk parameter $\delta \in (0, 1]$, we aim for a δ -correct algorithm, that satisfies $\mathbb{P}_\mu(\boldsymbol{\mu} \in \mathcal{H}_{\hat{m}}) \geq 1 - \delta$ for all $\boldsymbol{\mu} \in \mathcal{H}$. Our goal is to build δ -correct algorithms that use a small number of samples τ_δ in order to reach a decision. In particular, we want the *sample complexity* $\mathbb{E}_\mu[\tau]$ to be small.

Notation We let $N_a(t) = \sum_{s=1}^t \mathbb{1}_{(A_s=a)}$ be the number of selections of arm a up to round t , $S_a(t) = \sum_{s=1}^t X_s \mathbb{1}_{(A_s=a)}$ be the sum of the gathered observations from that arm and $\hat{\mu}_a(t) = S_a(t)/N_a(t)$ their empirical mean.

3 Lower Bounds

In this section we study information-theoretic sample complexity lower bounds, in particular to find out what the problem tells us about the behavior of oracle algorithms. [10] prove that for any δ -correct algorithm

$$\mathbb{E}_\mu[\tau] \geq T^*(\boldsymbol{\mu}) \text{kl}(\delta, 1 - \delta) \quad \text{where} \quad \frac{1}{T^*(\boldsymbol{\mu})} = \max_{\mathbf{w} \in \Delta} \min_{\boldsymbol{\lambda} \in \text{Alt}(\boldsymbol{\mu})} \sum_a w_a d(\mu_a, \lambda_a) \quad (1)$$

$\text{kl}(x, y) = x \ln \frac{x}{y} + (1 - x) \ln \frac{1-x}{1-y}$ and $\text{Alt}(\boldsymbol{\mu})$ is the set of bandit models where the correct recommendation differs from that on $\boldsymbol{\mu}$. The following result specialises the above to the case of testing $\mathcal{H}_<$ vs $\mathcal{H}_>$, and gives explicit expressions for the *characteristic time* $T^*(\boldsymbol{\mu})$ and *oracle weights* $\mathbf{w}^*(\boldsymbol{\mu})$.

Lemma 1. *Any δ -correct strategy satisfies (1) with*

$$T^*(\boldsymbol{\mu}) = \begin{cases} \frac{1}{d(\mu^*, \gamma)} & \mu^* < \gamma, \\ \sum_a \frac{1}{d(\mu_a, \gamma)} & \mu^* > \gamma, \end{cases} \quad \text{and} \quad \mathbf{w}_a^*(\boldsymbol{\mu}) = \begin{cases} \mathbf{1}_{a=a^*} & \mu^* < \gamma, \\ \frac{\frac{1}{d(\mu_a, \gamma)}}{\sum_j \frac{1}{d(\mu_j, \gamma)}} & \mu^* > \gamma. \end{cases}$$

Lemma 1 is proved in Appendix B. As explained by [10] the oracle weights correspond to the fraction of samples that should be allocated to each arm under a strategy matching the lower bound. The interesting feature here is that the lower bound indicates that an oracle algorithm should have very different behavior on $\mathcal{H}_<$ and $\mathcal{H}_>$. On $\mathcal{H}_<$ it should sample a^* (or all lowest means, if there are several) exclusively, while on $\mathcal{H}_>$ it should sample *all* arms with certain specific proportions.

3.1 Boosting the Lower Bounds

Following [13] (see also [30] and references therein), Lemma 1 can be improved under very mild assumptions on the strategies. We call a test *symmetric* if its sampling and stopping rules are invariant by conjugation under the action of the group of permutations on the arms. In that case, if all the arms are equal, then their expected numbers of draws are equal. For simplicity we assume $\mu_1 \leq \dots \leq \mu_K$.

Proposition 2. *Let $k = \max_a d(\mu_a, \gamma) = \max \{d(\mu_1, \gamma), d(\mu_K, \gamma)\}$. For any symmetric, δ -correct test, for all arms $a \in \{1, \dots, K\}$, the expected number of selections of arm a satisfies*

$$\mathbb{E}_\mu[N_a(\tau)] \geq \frac{2(1 - 2\delta K^3)}{27K^2k}.$$

Proposition 2 is proved in Appendix B. It is an open question to improve the dependency in K in this bound; moreover, one may expect a bound decreasing with δ , maybe in $\ln(\ln(1/\delta))$ (but certainly not in $\ln(1/\delta)$). This result already has two important consequences: first, it shows that even an optimal algorithm needs to draw all the arms a certain number of times, even on $\mathcal{H}_<$ where Lemma 1 may suggest otherwise. Second, this lower bound on the number of draws of each arm can be used to “boost” the lower bound on $\mathbb{E}_\mu[\tau]$: the following result is also proved in Appendix B.

Theorem 3. *When $\mu^* < \gamma$, for any symmetric, δ -correct strategy,*

$$\mathbb{E}_\mu[\tau] \geq \frac{\text{kl}(\delta, 1 - \delta)}{d(\mu_1, \gamma)} + \frac{2(1 - 2\delta K^3)}{27K^2k} \sum_a \left(1 - \frac{d(\mu_a, \gamma) \mathbb{1}_{(\mu_a \leq \gamma)}}{d(\mu_1, \gamma)}\right).$$

3.2 Lower Bound Inspired Matching Algorithms

In light of the lower bound in Lemma 1, we now investigate the design of optimal learning algorithms (sampling rule A_t and stopping rule τ). We start with the stopping rule. The first stopping rule that comes to mind consists in comparing *separately* each arm to the threshold and stopping when either one arm looks significantly below the threshold or all arms look significantly above. Introducing $d^+(u, v) = d(u, v) \mathbb{1}_{(u \leq v)}$ and $d^-(u, v) = d(u, v) \mathbb{1}_{(u \geq v)}$, we let

$$\tau_{\text{Box}} = \tau_< \wedge \tau_> \quad \text{where} \quad \begin{aligned} \tau_< &= \inf \{t \in \mathbb{N}^* : \exists a N_a(t) d^+(\hat{\mu}_a(t), \gamma) \geq C_<(\delta, N_a(t))\}, \\ \tau_> &= \inf \{t \in \mathbb{N}^* : \forall a N_a(t) d^-(\hat{\mu}_a(t), \gamma) \geq C_>(\delta, N_a(t))\}, \end{aligned} \quad (2)$$

and $C_<(\delta, r)$ and $C_>(\delta, r)$ are two *threshold functions* to be specified. Box refers to the fact that the decision to stop relies on individual “box” confidence intervals for each arm, whose endpoints are

$$\begin{aligned} U_a(t) &= \max\{q : N_a(t) d^+(\hat{\mu}_a(t), q) \geq C_<(\delta, N_a(t))\}, \\ L_a(t) &= \min\{q : N_a(t) d^-(\hat{\mu}_a(t), q) \geq C_>(\delta, N_a(t))\}. \end{aligned}$$

Indeed, $\tau_{\text{Box}} = \inf \{t \in \mathbb{N}^* : \min_a U_a(t) \leq \gamma \text{ or } \min_a L_a(t) \geq \gamma\}$. In particular, if $\forall a, \forall t \in \mathbb{N}^*, \mu_a \in [L_a(t), U_a(t)]$, any algorithm that stops using τ_{Box} is guaranteed to output a correct decision. In the Gaussian case, existing work [18, 23] permits to exhibit thresholds of the form $C_\leq(\delta, r) = \ln(1/\delta) + a \ln \ln(1/\delta) + b \ln(1 + \ln(r))$ for which this sufficient correctness condition is satisfied with probability larger than $1 - \delta$. Theorem 7 below generalizes this to exponential families.

Given that τ_{Box} can be proved to be δ -correct *whatever the sampling rule*, the next step is to propose sampling rules that, coupled with τ_{Box} , would attain the lower bound presented in Section 3. We now show that a simple algorithm, called LCB, can do that for all $\mu \in \mathcal{H}_>$. LCB selects at each round the arm with smallest Lower Confidence Bound:

$$\boxed{\text{LCB: Play } A_t = \arg \min_a L_a(t),} \quad (3)$$

which is intuitively designed to attain the stopping condition $\min_a L_a(t) \geq \gamma$ faster. In Appendix E we prove (Proposition 15) that LCB is optimal for $\mu \in \mathcal{H}_>$ however we show (Proposition 16) that on instances of $\mathcal{H}_<$ it draws all arms $a \neq a^*$ too much and cannot match our lower bound.

For $\mu \in \mathcal{H}_<$, the lower bound Lemma 1 can actually be a good guideline to design a matching algorithm: under such an algorithm, the empirical proportion of draws of the arm a^* with smallest mean should converge to 1. The literature on regret minimization in bandit models (see [2] for a survey) provides candidate algorithms that have this type of behavior, and we propose to use the

Thompson Sampling (TS) algorithm [1, 22]. Given independent prior distribution on the mean of each arm, this Bayesian algorithm selects an arm at random according to its posterior probability of being optimal (in our case, the arm with smallest mean). Letting π_a^t refer to the posterior distribution of μ_a after t samples, this can be implemented as

$$\text{TS: Sample } \forall a \in \{1, \dots, K\}, \theta_a(t) \sim \pi_a^{t-1}, \text{ then play } A_t = \arg \min_{a \in \{1, \dots, K\}} \theta_a(t).$$

It follows from Theorem 12 in Appendix 5 that if Thompson Sampling is run without stopping, $N_{a^*}(t)/t$ converges almost surely to 1, for every μ . As TS is an anytime sampling strategy (i.e. that does not depend on δ), Lemma 4 below permits to justify that on every instance of $\mathcal{H}_<$ with a unique optimal arm, under this algorithm $\tau_{\text{Box}} \simeq (1/d(\mu_1, \theta)) \ln(1/\delta)$. However, TS cannot be optimal for $\mu \in \mathcal{H}_>$, as the empirical proportions of draws cannot converge to $w^*(\mu) \neq \mathbf{1}_{a^*}$.

To summarize, we presented a simple stopping rule, τ_{Box} , that can be asymptotically optimal for every $\mu \in \mathcal{H}_<$ if it is used in combination with Thompson Sampling and for $\mu \in \mathcal{H}_>$ if it is used in combination with LCB. But neither of these two sampling rules are good for the other type of instances, which is a big limitation for a practical use of either of these. In the next section, we propose a new Thompson Sampling like algorithm that ensures the right exploration under both $\mathcal{H}_<$ and $\mathcal{H}_>$. In Section 5, we further present an improved stopping rule that may stop significantly earlier than τ_{Box} on instances of $\mathcal{H}_<$, by aggregating samples from multiple arms that look small.

We now argue that ensuring the sampling proportions converge to w^* is sufficient for reaching the optimal sample complexity, at least in an asymptotic sense. The proof can be found in Appendix C.

Lemma 4. Fix $\mu \in \mathcal{H}$. Fix an anytime sampling strategy (A_t) ensuring $\frac{N_t}{t} \rightarrow w^*(\mu)$. Let τ_δ be a stopping rule such that $\tau_\delta \leq \tau_\delta^{\text{Box}}$, for a Box stopping rule (2) whose threshold functions C_ξ satisfy the following: they are non-decreasing in r and there exists a function f such that,

$$\forall r \geq r_0, C_\xi(\delta, r) \leq f(\delta) + \ln r, \text{ where } f(\delta) = \ln(1/\delta) + o(\ln(1/\delta)).$$

Then $\limsup_{\delta \rightarrow 0} \frac{\tau_\delta}{\ln \frac{1}{\delta}} \leq T^*(\mu)$ almost surely.

4 Murphy Sampling

In this section we denote by $\Pi_n = \mathbb{P}(\cdot | \mathcal{F}_n)$ the posterior distribution of the mean parameters after n rounds. We introduce a new (randomised) sampling rule called *Murphy Sampling* after Murphy’s Law, as it performs some conditioning to the “worst event” ($\mu \in \mathcal{H}_<$):

$$\text{MS: Sample } \theta_t \sim \Pi_{t-1}(\cdot | \mathcal{H}_<), \text{ then play } A_t = a^*(\theta_t). \quad (4)$$

As we will argue below, the subtle difference of sampling from $\Pi_{n-1}(\cdot | \mathcal{H}_<)$ instead of Π_{n-1} (regular Thompson Sampling) ensures the required split personality behavior (see Lemma 1). Note that MS always conditions on $\mathcal{H}_<$ (and never on $\mathcal{H}_>$) regardless of the position of μ w.r.t. γ . This is different from the symmetric Top Two Thompson Sampling [29], which essentially conditions on $a^*(\theta) \neq a^*(\mu)$ a fixed fraction $1 - \beta$ of the time, where β is a parameter that needs to be tuned with knowledge of μ . MS on the other hand needs no parameters.

Also note that MS is an anytime sampling algorithm, being independent of the confidence level $1 - \delta$. The confidence will manifest only in the stopping rule.

MS is technically an instance of Thompson Sampling with a joint prior Π supported only on $\mathcal{H}_<$. This viewpoint is conceptually funky, as we will apply MS identically to $\mathcal{H}_<$ and $\mathcal{H}_>$. To implement MS, we use that independent conjugate per-arm priors induce likewise posteriors, admitting efficient (unconditioned) posterior sampling. Rejection sampling then achieves the required conditioning. Its computational cost is limited: the acceptance probability cannot be much smaller than the risk δ provided to the algorithm. Indeed, the fact that the stopping rule (see Section 5) has not yet fired, combined with the posterior concentration (Proposition 6) and the convergence of the sampling efforts to track the sampling proportions (Theorem 5) reveals that the MS rejection sampling step accepts with probability at least of order $\delta/(\ln t)^3$. So for reasonable values of δ , this can be small and require a few thousands of draws (not a big deal for today’s computers), but it cannot be *prohibitively* small.

The rest of this section is dedicated to the analysis of MS. First, we argue that the MS sampling proportions converge to the oracle weights of Lemma 1.

Assumption For purpose of analysis, we need to assume that the parameter space $\Theta \ni \boldsymbol{\mu}$ (or the support of the prior) is the interior of a bounded subset of \mathbb{R}^K . This ensures that $\sup_{\boldsymbol{\mu}, \theta \in \Theta} d(\boldsymbol{\mu}, \theta) < \infty$ and $\sup_{\boldsymbol{\mu}, \theta \in \Theta} \|\boldsymbol{\mu} - \theta\| < \infty$. This assumption is common [16, Section 7.1], [29, Assumption 1]. We also assume that the prior Π has a density π with bounded ratio $\sup_{\boldsymbol{\mu}, \theta \in \Theta} \frac{\pi(\theta)}{\pi(\boldsymbol{\mu})} < \infty$.

Theorem 5. *Under the above assumption, MS ensures $\frac{N_t}{t} \rightarrow \boldsymbol{w}^*(\boldsymbol{\mu})$ a.s. for any $\boldsymbol{\mu} \in \mathcal{H}$.*

We give a sketch of the proof below, the detailed argument can be found in Appendix D, Theorems 12 and 13. Given the convergence of the weights, the asymptotic optimality in terms of sample complexity follows by Lemma 4, if MS is used with an appropriate stopping rule (Box (2) or the improved Aggregate stopping rule discussed in Section 5).

Proof Sketch First, consider $\boldsymbol{\mu} \in \mathcal{H}_<$. In this case the conditioning in MS is asymptotically immaterial as $\Pi_n(\mathcal{H}_<) \rightarrow 1$, and the algorithm behaves like regular Thompson Sampling. As Thompson sampling has sublinear pseudo-regret [1], we must have $\mathbb{E}[N_1(t)]/t \rightarrow 1$. The crux of the proof in the appendix is to show the convergence occurs almost surely.

Next, consider $\boldsymbol{\mu} \in \mathcal{H}_>$. Following [29], we denote the sampling probabilities in round n by $\psi_a(n) = \Pi_{n-1}(a = \arg \min_j \theta_j | \mathcal{H}_<)$, and abbreviate $\Psi_a(n) = \sum_{t=1}^n \psi_a(t)$ and $\bar{\psi}_a(n) = \Psi_a(n)/n$. The main intuition is provided by

Proposition 6 ([29, Proposition 4]). *For any open subset $\tilde{\Theta} \subseteq \Theta$, the posterior concentrates at rate $\Pi_n(\tilde{\Theta}) \doteq \exp(-n \min_{\lambda \in \tilde{\Theta}} \sum_a \bar{\psi}_a(n) d(\mu_a, \lambda_a))$ a.s. where $a_n \doteq b_n$ means $\frac{1}{n} \ln \frac{a_n}{b_n} \rightarrow 0$.*

Let us use this to analyze $\psi_a(n)$. As we are on $\mathcal{H}_>$, the posterior $\Pi_n(\mathcal{H}_<) \rightarrow 0$ vanishes. Moreover, $\Pi_n(a = \arg \min_j \theta_j, \mathcal{H}_<) \sim \Pi_n(\theta_a < \gamma)$ as the probability that multiple arms fall below γ is negligible. Hence

$$\psi_a(n+1) \sim \frac{\Pi_n(\mu_a < \gamma)}{\sum_j \Pi_n(\mu_j < \gamma)} \doteq \frac{\exp(-n \bar{\psi}_a(n) d(\mu_a, \gamma))}{\sum_j \exp(-n \bar{\psi}_j(n) d(\mu_j, \gamma))}.$$

To get a good sense for what this means, let's analyze the version with equality. Using that $w_a^* d(\mu_a, \gamma)$ is constant (Lemma 1), we see

$$\psi_a(n+1) \leq e^{-n(\bar{\psi}_a(n) - w_a^*) d(\mu_a, \gamma)}.$$

Now this means that whenever $\bar{\psi}_a(n) \geq w_a^* + \epsilon$, we find that $\psi_a(n+1) \leq e^{-n\epsilon d_a} \approx 0$ is exponentially small, and hence $\bar{\psi}_a(n+1) \approx \frac{n}{n+1} \bar{\psi}_a(n)$ decays hyperbolically (i.e. without lower bound). Hence $\limsup_{n \rightarrow \infty} \bar{\psi}_a(n) \leq w_a^* + \epsilon$. As this holds for all arms a and $\epsilon > 0$, we must have $\lim_n \bar{\psi}_a(n) = w_a^*$.

5 Improved Stopping Rule and Confidence Intervals

Theorem 7 below provides a new self-normalized deviation inequality that given a *subset* of arms controls uniformly over time how the *aggregated mean* of the samples obtained from those arms can deviate from the smallest (resp. largest) mean in the subset. More formally for $\mathcal{S} \subseteq [K]$, we introduce

$$N_{\mathcal{S}}(t) = \sum_{a \in \mathcal{S}} N_a(t) \quad \text{and} \quad \hat{\mu}_{\mathcal{S}}(t) = \frac{\sum_{a \in \mathcal{S}} N_a(t) \hat{\mu}_a(t)}{N_{\mathcal{S}}(t)}$$

and recall $d^+(u, v) = d(u, v) \mathbb{1}_{(u \leq v)}$ and $d^-(u, v) = d(u, v) \mathbb{1}_{(u \geq v)}$. We prove the following for one-parameter exponential families.

Theorem 7. *Let $T : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ be the function defined by*

$$T(x) = 2h^{-1} \left(1 + \frac{h^{-1}(1+x) + \ln \zeta(2)}{2} \right) \quad (5)$$

where $h(u) = u - \ln(u)$ for $u \geq 1$ and $\zeta(s) = \sum_{n=1}^{\infty} n^{-s}$. For every subset \mathcal{S} of arms and $x \geq 0.04$,

$$\mathbb{P} \left(\exists t \in \mathbb{N} : N_{\mathcal{S}}(t) d^+ \left(\hat{\mu}_{\mathcal{S}}(t), \min_{a \in \mathcal{S}} \mu_a \right) \geq 3 \ln(1 + \ln(N_{\mathcal{S}}(t))) + T(x) \right) \leq e^{-x}, \quad (6)$$

$$\mathbb{P} \left(\exists t \in \mathbb{N} : N_{\mathcal{S}}(t) d^- \left(\hat{\mu}_{\mathcal{S}}(t), \max_{a \in \mathcal{S}} \mu_a \right) \geq 3 \ln(1 + \ln(N_{\mathcal{S}}(t))) + T(x) \right) \leq e^{-x}. \quad (7)$$

The proof of this theorem can be found in Section F and is sketched below. It generalizes in several directions the type of results obtained by [18, 23] for Gaussian distributions and $|\mathcal{S}| = 1$. Going beyond subsets of size 1 will be crucial here to obtain better confidence intervals on minimums, or stop earlier in tests. Note that the threshold function T introduced in (5) does not depend on the cardinality of the subset \mathcal{S} to which the deviation inequality is applied. Tight upper bounds on T can be given using Lemma 21 in Appendix F.3, which support the approximation $T(x) \simeq x + 3 \ln(x)$.

5.1 An Improved Stopping Rule

Fix a subset prior $\pi : \mathcal{P}(\{1, \dots, K\}) \rightarrow \mathbb{R}^+$ such that $\sum_{\mathcal{S} \in \{1, \dots, K\}} \pi(\mathcal{S}) = 1$ and let T be the threshold function defined in Theorem 7. We define the stopping rule $\tau^\pi := \tau_> \wedge \tau_<^\pi$, where

$$\begin{aligned} \tau_> &= \inf \{t \in \mathbb{N}^* : \forall a \in \{1, \dots, K\} N_a(t) d^- (\hat{\mu}_a(t), \gamma) \geq 3 \ln(1 + \ln(N_a(t))) + T(\ln(1/\delta))\}, \\ \tau_<^\pi &= \inf \{t \in \mathbb{N}^* : \exists \mathcal{S} : N_{\mathcal{S}}(t) d^+ (\hat{\mu}_{\mathcal{S}}(t), \gamma) \geq 3 \ln(1 + \ln(N_{\mathcal{S}}(t))) + T(\ln(1/(\delta\pi(\mathcal{S}))))\}. \end{aligned}$$

The associated recommendation rule selects $\mathcal{H}_>$ if $\tau^\pi = \tau_>$ and $\mathcal{H}_<$ if $\tau^\pi = \tau_<^\pi$. For the practical computation of $\tau_<^\pi$, the search over subsets can be reduced to nested subsets including arms sorted by increasing empirical mean and smaller than γ .

Lemma 8. *Any algorithm using the stopping rule τ^π and selecting $\hat{m} = >$ iff $\tau^\pi = \tau_>$, is δ -correct.*

From Lemma 8, proved in Appendix G, the prior π doesn't impact the correctness of the algorithm. However it may impact its sample complexity significantly. First it can be observed that picking π that is uniform over subset of size 1, i.e. $\pi(\mathcal{S}) = K^{-1} \mathbb{1}(|\mathcal{S}| = 1)$, one obtain a δ -correct τ_{Box} stopping rule with thresholds functions satisfying the assumptions of Lemma 4. However, in practice (especially more moderate δ), it may be more interesting to include in the support of π subsets of larger sizes, for which $N_{\mathcal{S}}(t) d^+ (\hat{\mu}_{\mathcal{S}}(t), \gamma)$ may be larger. We advocate the use of $\pi(\mathcal{S}) = K^{-1} \binom{K}{|\mathcal{S}|}^{-1}$, that puts the same weight on the set of subsets of each possible size.

Links with Generalized Likelihood Ratio Tests (GLRT). Assume we want to test \mathcal{H}_0 against \mathcal{H}_1 for composite hypotheses. A GLRT test based on t observations whose distribution depends on some parameter x rejects \mathcal{H}_0 if the test statistic $\max_{x \in \mathcal{H}_1} \ell(X_1, \dots, X_t; x) / \max_{x \in \mathcal{H}_0 \cup \mathcal{H}_1} \ell(X_1, \dots, X_t; x)$ has large values (where $\ell(\cdot; x)$ denotes the likelihood of the observations under the model parameterized by x). In our testing problem, the GLRT statistic for rejecting $\mathcal{H}_<$ is $\min_a N_a(t) d^- (\hat{\mu}_a(t), \gamma)$ hence $\tau_>$ is very close to a sequential GLRT test. However, the GLRT statistic for rejecting $\mathcal{H}_>$ is $\sum_{a=1}^K N_a(t) d^+ (\hat{\mu}_a(t), \gamma)$, which is quite different from the stopping statistic used by $\tau_<^\pi$. Rather than *aggregating samples* from arms, the GLRT statistic is *summing evidence* for exceeding the threshold. Using similar martingale techniques as for proving Theorem 7, one can show that replacing $\tau_<^\pi$ by

$$\tau_<^{\text{GLRT}} = \inf \left\{ t \in \mathbb{N}^* : \sum_{a: \hat{\mu}_a(t) \leq \gamma} [N_a(t) d^+ (\hat{\mu}_a(t), \gamma) - 3 \ln(1 + \ln(N_a(t)))]^+ \geq KT \left(\frac{\ln(1/\delta)}{K} \right) \right\}$$

also yields a δ -correct algorithm (see [21])¹. At first sight, $\tau_<^\pi$ and $\tau_<^{\text{GLRT}}$ are hard to compare: the stopping statistic used by the latter can be larger than that used by the former, but it is compared to a smaller threshold. In Section 6 we will provide empirical evidence in favor of aggregating samples.

5.2 A Confidence Intervals Interpretation

Inequality (6) (and a union bound over subsets) also permits building a tight upper confidence bound on the minimum μ^* . Indeed, defining

$$U_{\min}^\pi(t) := \max \left\{ q : \max_{\mathcal{S} \in \{1, \dots, K\}} [N_{\mathcal{S}}(t) d^+ (\hat{\mu}_{\mathcal{S}}(t), q) - 3 \ln(1 + \ln(1 + N_{\mathcal{S}}(t)))] \leq T \left(\ln \frac{1}{\delta\pi(\mathcal{S})} \right) \right\},$$

it is easy to show that $\mathbb{P}(\forall t \in \mathbb{N}, \mu^* \leq U_{\min}^\pi(t)) \geq 1 - \delta$. For general choices of π , this upper confidence bound may be much smaller than the naive bound $\min_a U_a(t)$ which corresponds to choosing π uniform over subset of size 1. We provide an illustration supporting this claim in Figure 2

¹In fact, we can slightly sharpen the bound by observing that we are controlling the deviation of a single composite arm, allowing us to replace (5) by $T(x) = 2h^{-1} \left(1 + \frac{x + \ln \zeta(2)}{2} \right)$, see [21, Appendix A.1]

below. The two type of upper confidence bounds (Aggregate corresponding to $\pi(\mathcal{S}) = K^{-1} \binom{K}{|\mathcal{S}|}^{-1}$ and Box corresponding to $\pi(\mathcal{S}) = K^{-1} \mathbb{1}_{(|\mathcal{S}|=1)}$) are compared under uniform sampling in a Bernoulli bandit model that has k arms with mean 0.1 plus 4 arms with means $[0.2 \ 0.3 \ 0.4 \ 0.5]$. The larger the number of arms close to minimum (here equal to it) is, the more UCB Aggregate beats UCB Box. Observe that using inequality (7) in Theorem 7 similarly allows to derive tighter lower confidence bounds on the maximum of several means.




Figure 2: Illustration of the Box versus Aggregate Upper Confidence Bounds as a function of time on Bernoulli instance for $k = 1$ (left), $k = 3$ (middle) and $k = 10$ (right) minimal arms.

5.3 Sketch of the Proof of Theorem 7

Fix $\eta \in [0, 1 + e[$. Introducing $X_\eta(t) = [N_{\mathcal{S}}(t)d^+(\hat{\mu}_{\mathcal{S}}(t), \min_{a \in \mathcal{S}} \mu_a) - 2(1 + \eta) \ln(1 + \ln N_{\mathcal{S}}(t))]$, the cornerstone of the proof (Lemma 17) consists in proving that for all $\lambda \in [0, (1 + \eta)^{-1}[$, there exists a martingale M_t^λ that “almost” upper bounds $e^{\lambda X_\eta(t)}$: there exists a function g_η such that

$$\mathbb{E}[M_0^\lambda] = 1 \quad \text{and} \quad \forall t \in \mathbb{N}^*, M_t^\lambda \geq e^{\lambda X_\eta(t) - g_\eta(\lambda)}. \quad (8)$$

From there, the proof easily follows from a combination of Chernoff method and Doob inequality:

$$\mathbb{P}(\exists t \in \mathbb{N}^* : X_\eta(t) > u) \leq \mathbb{P}(\exists t \in \mathbb{N}^* : M_t^\lambda > e^{\lambda u - g_\eta(\lambda)}) \leq \exp(-[\lambda u - g_\eta(\lambda)]).$$

Inequality (6) is then obtained by optimizing over λ , carefully picking η and inverting the bound.

The interesting part of the proof is to actually build a martingale satisfying (8). First, using the so-called method of mixtures [6] and some specific fact about exponential families already exploited by [3], we can prove that there exists a martingale \tilde{W}_t^x such that for some function f (see Equation (14))

$$\{X_\eta(t) - f(\eta) \geq x\} \subseteq \{\tilde{W}_t^x \geq e^{\frac{x}{1+\eta}}\}.$$

From there it follows that, for every λ and $z > 1$, $\{e^{\lambda(X_\eta(t) - f(\eta))} \geq z\} \subseteq \{e^{-\frac{\ln(z)}{\lambda(1+\eta)}} \tilde{W}_t^{\frac{1}{\lambda} \ln(z)} \geq 1\}$ and the trick is to introduce another mixture martingale,

$$\overline{M}_t^\lambda = 1 + \int_1^\infty e^{-\frac{\ln(z)}{\lambda(1+\eta)}} \tilde{W}_t^{\frac{1}{\lambda} \ln(z)} dz,$$

that is proved to satisfy $\overline{M}_t^\lambda \geq e^{\lambda[X_\eta(t) - f(\eta)]}$. We let $M_t^\lambda = \overline{M}_t^\lambda / \mathbb{E}[\overline{M}_t^\lambda]$.

6 Experiments

We discuss the results of numerical experiments performed on Gaussian bandits with variance 1, using the threshold $\gamma = 0$. Thompson and Murphy sampling are run using a flat (improper) prior on \mathbb{R} , which leads to a conjugate Gaussian posterior. The experiments demonstrate the flexibility of our MS sampling rule, which attains optimal performance on instances from both $\mathcal{H}_<$ and $\mathcal{H}_>$. Moreover, they show the advantage of using a stopping rule aggregating samples from subsets of arms when $\mu \in \mathcal{H}_<$. This aggregating stopping rule, that we refer to as τ^{Agg} is an instance of the τ^π stopping rule presented in Section 5 for $\pi(\mathcal{S}) = K^{-1} \binom{K}{|\mathcal{S}|}^{-1}$. We investigate the combined use of three sampling rules, MS, LCB and Thompson Sampling with three stopping rules, τ^{Agg} , τ^{Box} and τ^{GLRT} .

We first study an instance $\mu \in \mathcal{H}_<$ with $K = 10$ arms that are linearly spaced between -1 and 1 . We run the different algorithms (excluding the TS sampling rule, that essentially coincides with MS on $\mathcal{H}_<$) for different values of δ and report the estimated sample complexity in Figure 3 (left). For each sampling rule, it appears that $\mathbb{E}[\tau^{\text{Agg}}] \leq \mathbb{E}[\tau^{\text{Box}}] \leq \mathbb{E}[\tau^{\text{GLRT}}]$. Moreover, for each stopping rule MS is outperforming LCB, with a sample complexity of order $T^*(\mu) \ln(1/\delta) + C$. Then we study an instance $\mu \in \mathcal{H}_>$ with $K = 5$ arms that are linearly spaced between 0.5 and 1 , with τ^{Agg} as the sampling rule (which matters little as the algorithm mostly stops because of $\tau_>$ on $\mathcal{H}_>$). Results are reported in Figure 3 (right), in which we see that MS is performing very similarly to LCB (that is also proved optimal on $\mathcal{H}_>$), while vanilla TS fails dramatically. On those experiments, the empirical error was always zero, which shows that our theoretical thresholds are still quite conservative. More experimental results can be found in Appendix A: an illustration of the convergence properties of the MS sampling rule as well as a larger-scale comparison of stopping rules under $\mathcal{H}_<$.

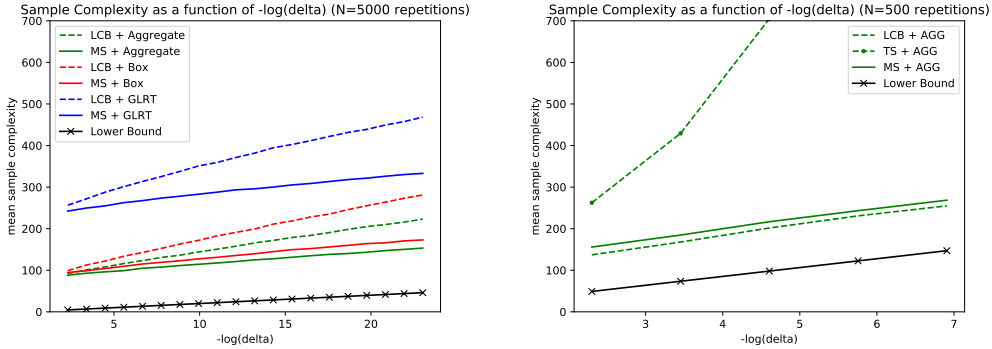


Figure 3: $\mathbb{E}[\tau_\delta]$ as a function of $\ln(1/\delta)$ for several algorithms on an instance $\mu \in \mathcal{H}_<$ (left) and $\mu \in \mathcal{H}_>$ (right), estimated using $N = 5000$ (resp. 500) repetitions.

7 Discussion

We propose new sampling and stopping rules for sequentially testing the minimum of means. As our guiding principle, we first prove sample complexity lower bounds, characterized by the emerging oracle sample allocation w^* , and develop the Murphy Sampling strategy to match it asymptotically. We observe in the experiments that the asymptotic regime does not necessarily kick in at moderate confidence δ (Figure 4, left) and that there is an important lower-order term in the practical sample complexity (Figure 3). It is an intriguing open problem of theoretical and practical importance to characterize and match optimal behavior at moderate confidence. We make first contributions in both directions: we prove tighter sample complexity lower bounds for symmetric algorithms (Proposition 2, Theorem 3) and we design aggregating confidence intervals which are tighter in practice (Figure 2).

The importance of this perspective arises, as highlighted in the introduction, from the *hierarchical* application of maxima/minima in learning applications. A better understanding of the moderate confidence regime for learning minima will very likely translate into new insights and methods for learning about hierarchical structures, where the benefits accumulate with depth.

References

- [1] S. Agrawal and N. Goyal. Analysis of Thompson Sampling for the multi-armed bandit problem. In *Proceedings of the 25th Conference On Learning Theory*, 2012.
- [2] S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- [3] O. Cappé, A. Garivier, O.-A. Maillard, R. Munos, and G. Stoltz. Kullback-Leibler upper confidence bounds for optimal sequential allocation. *Annals of Statistics*, 41(3):1516–1541, 2013.

- [4] L. Chen, A. Gupta, J. Li, M. Qiao, and R. Wang. Nearly optimal sampling algorithms for combinatorial pure exploration. In *Proceedings of the 30th Conference on Learning Theory (COLT)*, 2017.
- [5] H. Chernoff. Sequential design of Experiments. *The Annals of Mathematical Statistics*, 30(3): 755–770, 1959.
- [6] V. H. de la Peña, T. L. Lai, and S. Q. *Self-normalized processes. Limit Theory and Statistical applications*. Springer, 2009.
- [7] C. D’Eramo, M. Restelli, and A. Nuara. Estimating maximum expected value through Gaussian approximation. In *International Conference on Machine Learning (ICML)*, 2016.
- [8] C. D’Eramo, A. Nuara, M. Pirota, and M. Restelli. Estimating the maximum expected value in continuous reinforcement learning problems. In *AAAI*, 2017.
- [9] E. Even-Dar, S. Mannor, and Y. Mansour. Action Elimination and Stopping Conditions for the Multi-Armed Bandit and Reinforcement Learning Problems. *Journal of Machine Learning Research*, 7:1079–1105, 2006.
- [10] A. Garivier and E. Kaufmann. Optimal best arm identification with fixed confidence. In *Proceedings of the 29th Conference On Learning Theory (COLT)*, 2016.
- [11] A. Garivier, E. Kaufmann, and W. M. Koolen. Maximin action identification: A new bandit framework for games. In *Proceedings of the 29th Conference On Learning Theory (COLT)*, 2016.
- [12] A. Garivier, P. Ménard, and L. Rossi. Thresholding bandit for dose-ranging: The impact of monotonicity. *arXiv:1711.04454*, 2017.
- [13] A. Garivier, P. Ménard, and G. Stoltz. Explore first, exploit next: The true shape of regret in bandit problems. *Mathematics of Operations Research*, Jun. 2018.
- [14] D. Goldsman and B. L. Nelson. *Comparing Systems via Simulation, Chapter 9 in the Handbook of Simulation*. Wiley, 1998.
- [15] J.-B. Grill, M. Valko, and R. Munos. Blazing the trails before beating the path: Sample-efficient Monte-Carlo planning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 4680–4688, 2016.
- [16] P. D. Grünwald. *The minimum description length principle*. MIT press, 2007.
- [17] R. Huang, M. M. Ajallooeian, C. Szepesvári, and M. Müller. Structured best arm identification with fixed confidence. In *International Conference on Algorithmic Learning Theory (ALT)*, 2017.
- [18] K. Jamieson, M. Malloy, R. Nowak, and S. Bubeck. lil’UCB: an Optimal Exploration Algorithm for Multi-Armed Bandits. In *Proceedings of the 27th Conference on Learning Theory*, 2014.
- [19] S. Kalyan Krishnan, A. Tewari, P. Auer, and P. Stone. PAC subset selection in stochastic multi-armed bandits. In *International Conference on Machine Learning (ICML)*, 2012.
- [20] E. Kaufmann and W. M. Koolen. Monte-Carlo tree search by best arm identification. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [21] E. Kaufmann and W. M. Koolen. Mixture Martingales Revisited with Applications to Sequential Tests and Confidence Intervals. Preprint, 2018. URL <https://hal.archives-ouvertes.fr/hal-01886612v1/>.
- [22] E. Kaufmann, N. Korda, and R. Munos. Thompson Sampling : an Asymptotically Optimal Finite-Time Analysis. In *Proceedings of the 23rd conference on Algorithmic Learning Theory*, 2012.
- [23] E. Kaufmann, O. Cappé, and A. Garivier. On the Complexity of Best Arm Identification in Multi-Armed Bandit Models. *Journal of Machine Learning Research*, 17(1):1–42, 2016.

- [24] S.-H. Kim. Comparison with a standard via fully sequential procedures. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 15(2):155–174, 2005.
- [25] L. Kocsis and C. Szepesvári. Bandit based Monte-Carlo planning. In *Proceedings of the 17th European Conference on Machine Learning, ECML'06*, pages 282–293, Berlin, Heidelberg, 2006. Springer-Verlag. ISBN 3-540-45375-X, 978-3-540-45375-8.
- [26] T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- [27] A. Locatelli, M. Gutzeit, and A. Carpentier. An optimal algorithm for the thresholding bandit problem. In *International Conference on Machine Learning (ICML)*, 2016.
- [28] H. Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952.
- [29] D. Russo. Simple Bayesian algorithms for best arm identification. *CoRR*, abs/1602.08448, 2016. URL <http://arxiv.org/abs/1602.08448>.
- [30] M. Simchowitz, K. Jamieson, and B. Recht. The simulator: Understanding adaptive sampling in the moderate-confidence regime. In *Proceedings of the 30th Conference on Learning Theory (COLT)*, 2017.
- [31] I. Takahisa and T. Kaneko. Estimating the maximum expected value through upper confidence bound of likelihood. In *Conference on Technologies and Applications of Artificial Intelligence (TAAI)*, pages 202–207. IEEE, 2017.
- [32] K. Teraoka, K. Hatano, and E. Takimoto. Efficient sampling method for Monte Carlo tree search problem. *IEICE Transactions on Information and Systems*, pages 392–398, 2014.
- [33] W. R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25:285–294, 1933.
- [34] H. van Hasselt. Estimating the maximum expected value: An analysis of (nested) cross validation and the maximum sample average. *CoRR*, abs/1302.7175, 2013. URL <http://arxiv.org/abs/1302.7175>.

A Additional Experimental Results

We first report in Figure 4 further results regarding the convergence of the sampling proportions $N_a(\tau)/\tau$ under the two instances of $\mathcal{H}_<$ and $\mathcal{H}_>$ described in Section 6, for the smallest value of δ used in each experiment and under the stopping rule τ^{Agg} . Under $\mathcal{H}_<$ we see that MS has indeed spent a larger fraction of the time on the optimal arm, even if it does not yet reach the fraction 1 prescribed by the lower bound. One can also note that the empirical proportions of draws of the arms under LCB are very close to the sub-optimal weights obtained in Proposition 16 in Appendix E, which are added to the plot. Under $\mathcal{H}_>$, we see that the empirical fractions of draws of both MS and LCB converge to $\mathbf{w}^*(\boldsymbol{\mu})$ whereas the TS sampling rule departs significantly from those optimal weights, by drawing mostly arm 1.

Then we go deeper into investigating the impact of the proposed sampling rule under instances of $\mathcal{H}_<$. Indeed, we expect that grouping samples from several arms will help stop earlier as the number of arms under the threshold γ increases, which we illustrate with the following experiment. Consider $K = 100$ Gaussian arms with variance 1 and $\gamma = 0$. For several values of $k \in \{1, \dots, K\}$, we consider an instance in which there are k arms with mean -1 and $K - k$ arms with mean 0. Note that all those instances have the same (asymptotic) theoretical sample complexity, which is $T^*(\boldsymbol{\mu}) \ln(1/\delta)$, but in a regime with “large” δ (here we take $\delta = 0.1$), we expect this aggregating of samples to reduce significantly the sample complexity especially when there are a lot of arms below γ .

Figure 5 (left) reports the sample complexity of the Agg, Box and GLRT stopping rule, each used in combination with either the LCB or the MS sampling rule, for different values of k . One can note first that for a given stopping rule, MS is always outperforming LCB. Then, τ^{Agg} outperforms τ^{Box} for all the values of k , as well as τ^{GLRT} for values that are smaller than 70. GLRT is thus a better candidate only when the number of arms below the threshold is very large. This may be explained by the support plot displayed in Figure 5 (right): for each value of k , we report the number of arms in the subset \mathcal{S} that was used for stopping, that is which satisfies $N_{\mathcal{S}}(\tau) d^+(\hat{\mu}_{\mathcal{S}}(\tau), \gamma) \geq \ln(1 + \ln(N_{\mathcal{S}}(\tau))) + \ln(1/(\delta\pi(\mathcal{S})))$ in the case of Box and Agg. For the GLRT, the support is the number of arms for which $\hat{\mu}_a(\tau) \leq \gamma$, whose evidence for being below the threshold is included in the definition of the GLRT statistic. The support plot highlights that GLRT may sum evidence from more arms than the number of arms whose samples are aggregated by Agg, and in a regime in which the thresholds to which the two stopping statistics are compared are similar, this may favor GLRT. In Figure 6, we report similar experiments in instances in which $K = 100$ and for each k there are k linearly spaced arms below the threshold and $K - k$ arms with mean 0. In that case, even for large values of k , GLRT does not outperform the Aggregating stopping rule, which successfully combines samples from several arms below the threshold with different means.

B Proofs for the Sample Complexity Lower Bounds

We first need the following Lemma, that tells us that a δ -correct strategy stops with probability at most 2δ if all arms have mean exactly γ .

Lemma 9. *Let $\boldsymbol{\gamma} = (\gamma, \dots, \gamma)$. For any δ -correct test, $\mathbb{P}_{\boldsymbol{\gamma}}[\tau < \infty] \leq 2\delta$.*

Proof. Let $m > 0$, $\epsilon > 0$, $\boldsymbol{\mu} = (\gamma + \epsilon, \dots, \gamma + \epsilon)$ and $\boldsymbol{\mu}' = (\gamma - \epsilon, \dots, \gamma - \epsilon)$. Then the informational inequality of [23, Lemma 1] applied to the event $\{\tau \leq m, \hat{m} = >\}$, followed by $\text{kl}(p, q) \geq 2(p - q)^2$, implies that

$$\begin{aligned} md(\gamma + \epsilon, \gamma - \epsilon) &\geq \text{kl}(\mathbb{P}_{\boldsymbol{\mu}}(\tau \leq m, \hat{m} = >), \mathbb{P}_{\boldsymbol{\mu}'}(\tau \leq m, \hat{m} = >)) \\ &\geq 2(\mathbb{P}_{\boldsymbol{\mu}}(\tau \leq m, \hat{m} = >) - \mathbb{P}_{\boldsymbol{\mu}'}(\tau \leq m, \hat{m} = >))^2 \\ &= 2(\mathbb{P}_{\boldsymbol{\mu}}(\tau \leq m) - \mathbb{P}_{\boldsymbol{\mu}}(\tau \leq m, \hat{m} = <) - \mathbb{P}_{\boldsymbol{\mu}'}(\tau \leq m, \hat{m} = >))^2 \\ &\geq 2(\mathbb{P}_{\boldsymbol{\mu}}(\tau \leq m) - 2\delta)_+^2 \end{aligned}$$

and thus

$$\mathbb{P}_{\boldsymbol{\mu}}(\tau \leq m) \leq 2\delta + \sqrt{\frac{md(\gamma + \epsilon, \gamma - \epsilon)}{2}}.$$

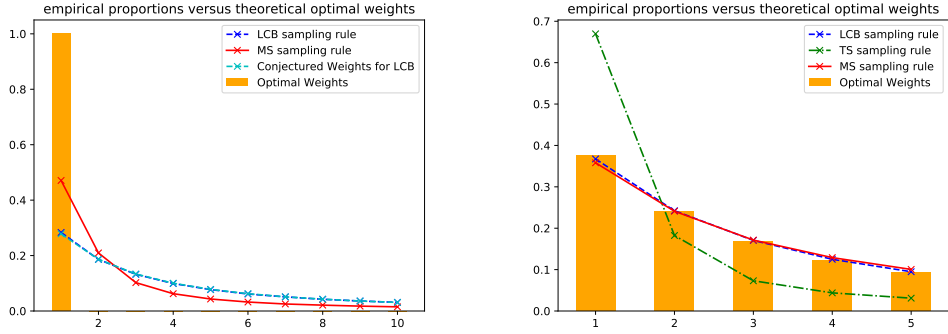


Figure 4: Empirical proportions of samples versus $w^*(\mu)$ for one instance in \mathcal{H}_L (left) and one instance in \mathcal{H}_S (right), in the same experimental setup as that of Figure 3.

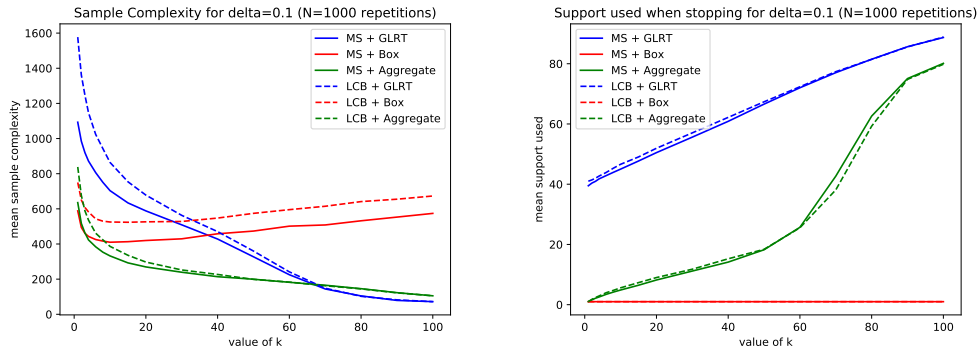


Figure 5: Sample complexity (left) and support when stopping (right) for different algorithms as a function of the number k of arms below the threshold $\gamma = 0$ on an instance for which $\mu_a \in \{-1, 0\}$.

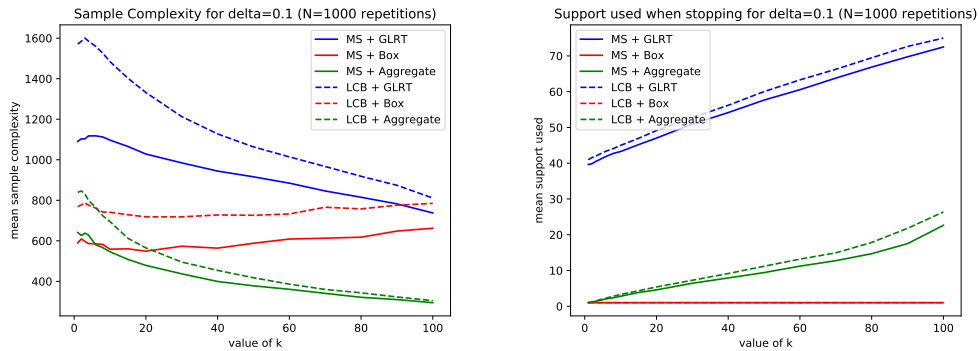


Figure 6: Sample complexity (left) and support when stopping (right) for different algorithms as a function of the number k of arms below the threshold $\gamma = 0$ on an instance for which the k arms below γ are linearly spaced between -1 and 0 .

Letting ϵ go to 0, one gets $\mathbb{P}_\gamma(\tau \leq m) \leq 2\delta$ and thus

$$\mathbb{P}_\gamma(\tau < \infty) = \mathbb{P}_\gamma\left(\bigcup_{m>0} (\tau < m)\right) = \lim_{m \rightarrow \infty} \mathbb{P}_\gamma(\tau < m) \leq 2\delta. \quad \square$$

B.1 Proof of Lemma 1

If $\mu^* < \gamma$, then we find

$$\frac{1}{T^*(\boldsymbol{\mu})} = \max_{\mathbf{w} \in \Delta} \sum_{a: \mu_a < \gamma} w_a d(\mu_a, \gamma) = \max_{a: \mu_a < \gamma} d(\mu_a, \gamma) = d(\mu^*, \gamma) \quad \text{where} \quad w_a^* = \mathbf{1}_{a=a^*}.$$

On the other hand, if $\mu^* > \gamma$, we find

$$\frac{1}{T^*(\boldsymbol{\mu})} = \max_{\mathbf{w} \in \Delta} \min_a w_a d(\mu_a, \gamma) = \frac{1}{\sum_a \frac{1}{d(\mu_a, \gamma)}} \quad \text{where} \quad w_a^* = \frac{\frac{1}{d(\mu_a, \gamma)}}{\sum_j \frac{1}{d(\mu_j, \gamma)}}.$$

B.2 Proof of Proposition 2

Let $\gamma = (\gamma, \dots, \gamma)$, and let $m > 0$. Fix $a \in \{1, \dots, K\}$. By Lemma 9,

$$\mathbb{E}_\gamma[N_a(\tau \wedge m)] \geq \mathbb{E}_\gamma[N_a(m)] - m \mathbb{P}_\gamma(\tau < m) \geq \mathbb{E}_\gamma[N_a(m)] - 2\delta m.$$

Then, by the informational lower bound (F-long) and by the generalized Pinsker inequality (Lemma 2 of [13]), one obtains

$$\begin{aligned} mk &\geq \sum_{j=1}^K \mathbb{E}_\gamma[N_j(\tau \wedge m)] d(\mu_a, \gamma) \\ &\geq \text{kl}\left(\frac{\mathbb{E}_\gamma[N_a(\tau \wedge m)]}{m}, \frac{\mathbb{E}_\mu[N_a(\tau \wedge m)]}{m}\right) \\ &\geq \text{kl}\left(\frac{1}{K} - 2\delta, \frac{\mathbb{E}_\mu[N_a(\tau \wedge m)]}{m} \wedge \left(\frac{1}{K} - 2\delta\right)\right) \\ &\geq \frac{K}{2} \left(\frac{1}{K} - 2\delta - \frac{\mathbb{E}_\mu[N_a(\tau \wedge m)]}{m}\right)_+^2. \end{aligned}$$

It follows that

$$\mathbb{E}_\mu[N_a(\tau \wedge m)] \geq \frac{m}{K} - 2\delta m - m \sqrt{\frac{2mk}{K}},$$

and the result follows from the choice $m = 2K(1/K - 2\delta)^2/(9k)$.

B.3 Proof of Theorem 3

By the informational lower bound (F-long) of [13],

$$\sum_a \mathbb{E}_\mu[N_a(\tau)] d_+(\mu_a, \gamma) = \sum_{a: \mu_a < \gamma} \mathbb{E}_\mu[N_a(\tau)] d(\mu_a, \gamma) \geq \text{kl}(\delta, 1 - \delta),$$

and by Proposition 2, for all $a \in \{1, \dots, K\}$,

$$\mathbb{E}_\mu[N_a(\tau)] \geq n := \frac{2(1 - 2\delta K^3)}{27K^2k}.$$

Hence,

$$\mathbb{E}_\mu[\tau] = \sum_a \mathbb{E}_\mu[N_a(\tau)] \geq \min \left\{ \sum_{a=1}^K n_a \quad \text{such that} \quad \sum_a n_a d_+(\mu_a, \gamma) \geq \text{kl}(\delta, 1 - \delta) \text{ and } \forall a, n_a \geq n \right\}.$$

The solution of this minimization problem is: $n_a^* = n$ for all $a > 1$, and

$$n_1^* = \frac{\text{kl}(\delta, 1 - \delta) - n \sum_{a>1} d_+(\mu_a, \gamma)}{d(\mu_1, \gamma)}.$$

Thus,

$$\mathbb{E}_\mu[\tau] \geq \sum_{a=1}^K n_a^* = \frac{\text{kl}(\delta, 1 - \delta)}{d(\mu_1, \gamma)} + n \sum_a \left(1 - \frac{d_+(\mu_a, \gamma)}{d(\mu_1, \gamma)}\right).$$

C Weight Convergence Implies Optimal Sample Complexity (Lemma 4)

Fix $\boldsymbol{\mu} \in \mathcal{H}_<$. Then there exists an event \mathcal{E} such that $N_1(t)/t \rightarrow w_1^*(\boldsymbol{\mu})$ and $\hat{\mu}_1(t) \rightarrow \mu_1$. On this event \mathcal{E} , for all $\epsilon > 0$, there exists t_0 such that for $t \geq t_0$, $N_1(t)d(\hat{\mu}_1(t), \gamma) \geq (1 - \epsilon)td(\mu_1, \gamma)$. We use (2) to write

$$\begin{aligned} \tau_\delta &\leq \tau_< \leq \inf\{t \in \mathbb{N}^*, N_1(t)d^-(\hat{\mu}_1(t), \gamma) \geq C_<(\delta, N_1(t))\} \\ &\leq \inf\{t \geq t_0 : t(1 - \epsilon)d(\mu_1, \gamma) \geq C_<(\delta, t)\} \\ &\leq \inf\{t \geq t_0 : t(1 - \epsilon)d(\mu_1, \gamma) \geq f(\delta) + \ln(t)\} \end{aligned}$$

hence

$$\tau_\delta \leq t_0 + \inf\left\{t \in \mathbb{N}^* : t \times [(1 - \epsilon)d(\mu_1, \gamma)] \geq \ln\left(\frac{t}{\delta}\right) + o(\ln(1/\delta))\right\}.$$

Simple algebra (e.g. Lemma 22 in [23]) yields

$$\tau_\delta \leq \frac{1}{(1 - \epsilon)d(\mu_1, \gamma)} \ln(1/\delta) + o(\ln(1/\delta))$$

hence $\limsup_{\delta \rightarrow 0} \tau_\delta / \ln(1/\delta) \leq T^*(\boldsymbol{\mu}) / (1 - \epsilon)$ for all ϵ , thus $\limsup_{\delta \rightarrow 0} \tau_\delta / \ln(1/\delta) \leq T^*(\boldsymbol{\mu})$.

Fix $\boldsymbol{\mu} \in \mathcal{H}_>$. As each a $w_a^*(\boldsymbol{\mu}) \neq 0$, all arms are drawn infinitely often, thus there exists an event \mathcal{E} of probability 1 such that $N_a(t)/t \rightarrow w_a^*(\boldsymbol{\mu})$ and $\hat{\mu}_a(t) \rightarrow \mu_a$. On \mathcal{E} , for all $\epsilon > 0$, there exists t_0 such that for all $t \geq t_0$, $\forall a, N_a(t)d^-(\hat{\mu}_a(t), \gamma) \geq (1 - \epsilon)tw_a^*(\boldsymbol{\mu})d(\mu_a, \gamma)$. This time

$$\begin{aligned} \tau_\delta &\leq \tau_> = \inf\{t \in \mathbb{N}^* : \forall a, N_a(t)d^-(\hat{\mu}_a(t), \gamma) \geq C_>(\delta, N_a(t))\} \\ &\leq \inf\{t \in \mathbb{N}^* : \forall a, N_a(t)d^-(\hat{\mu}_a(t), \gamma) \geq C_>(\delta, t)\} \\ &\leq \inf\{t \geq t_0 : \forall a, (1 - \epsilon)tw_a^*(\boldsymbol{\mu})d(\mu_a, \gamma) \geq f(\delta) + \ln(t)\} \end{aligned}$$

and

$$\tau_\delta \leq t_0 + \inf\left\{t \in \mathbb{N}^* : t \times \left[(1 - \epsilon) \min_a w_a^*(\boldsymbol{\mu})d(\mu_a, \gamma)\right] \geq \ln\left(\frac{t}{\delta}\right) + o(\ln(1/\delta))\right\}.$$

Similarly one obtains

$$\tau_\delta \leq \frac{T^*(\boldsymbol{\mu})}{(1 - \epsilon)} \ln(1/\delta) + o(\ln(1/\delta))$$

and $\limsup_{\delta \rightarrow 0} \tau_\delta / \ln(1/\delta) \leq T^*(\boldsymbol{\mu})$.

D Analysis of Murphy Sampling (Proof of Theorem 5)

In this section we analyse the Murphy Sampling (4) sampling rule. Throughout we will make the assumption stated in Section 4.

Let Π_n be the posterior on $\boldsymbol{\mu}$ after n rounds. Let $\psi_a(t)$ denote the probability of sampling arm a in round t , i.e.

$$\psi_a(t) = \mathbb{P}(A_t = a | \mathcal{F}_{t-1}) = \Pi_{t-1}\left(a = \arg \min_j \mu_j \mid \min_j \mu_j < \gamma\right).$$

let $\Psi_a(n) = \sum_{t=1}^n \psi_a(t)$ and $\bar{\psi}_a(n) = \Psi_a(n)/n$. We will make use of the following result

Proposition 10 ([29, Corollary 1]). *Let $\mathcal{S} \subseteq [K]$ be any subset of arms.*

$$\sum_{a \in \mathcal{S}} \Psi_a(t) \rightarrow \infty \implies \lim_{t \rightarrow \infty} \frac{\sum_{a \in \mathcal{S}} N_a(t)}{\sum_{a \in \mathcal{S}} \Psi_a(t)} = 1 \quad \text{a.s.}$$

Our main assumption is the following (see e.g. [29, Proposition 3]). Let $\Theta_a \subseteq \mathbb{R}$ be an open set. Then

$$\sup_t N_a(t) = \infty \implies \Pi_t(\theta_a \in \Theta_a) \rightarrow \mathbf{1}\{\mu_a \in \Theta_a\} \quad \text{a.s.} \quad (9)$$

$$\sup_t N_a(t) < \infty \implies \inf_t \Pi_t(\theta_a \in \Theta_a) > 0 \quad \text{a.s.} \quad (10)$$

We first show that every arm is drawn infinitely often

Proposition 11. Let $\mu \in \mathcal{H}$ with μ^* not on the boundary of Θ_{a^*} . Then the MS sampling rule ensures $N_a(t) \rightarrow \infty$ a.s. for all arm $a \in \{1, \dots, K\}$.

Proof. By Proposition 10, it suffices to show $\Psi_a(t) \rightarrow \infty$. Toward contradiction assume that $\mathcal{A} := \{a \mid \sup_t \Psi_a(t) < \infty\} \neq \emptyset$. Let $B = \{\theta \mid \theta < \mu^* - \epsilon\}$. Now for every arm $a \notin \mathcal{A}$, we have $\Pi_t(\theta_a \notin B) \rightarrow 1$ by (9). Let $C = \max_{a \in \mathcal{A}} \lim_t \Pi_t(\theta_a \in B)$. We have $C > 0$ by (10). But then

$$\begin{aligned} \sum_{a \in \mathcal{A}} \psi_a(t) &\geq \Pi_t \left(\arg \min_a \theta_a \in \mathcal{A}, \min_a \theta_a < \gamma \right) \\ &\geq \max_{a \in \mathcal{A}} \Pi_t(\theta_a \in B) \prod_{a \notin \mathcal{A}} \Pi_t(\theta_a \notin B) \rightarrow C > 0. \end{aligned}$$

But this means that $\sum_{a \in \mathcal{A}} \Psi_a(t) \rightarrow \infty$, a contradiction. \square

The analysis now splits in 2 cases, depending on the location of $\min_a \mu_a$ w.r.t. γ . First we consider the case $\mu \in \mathcal{H}_<$.

Theorem 12. Consider $\mu \in \mathcal{H}_<$ with minimal arms $\mathcal{A} = \{a \mid \mu_a = \mu_{*}\}$. Note that although Lemma 1 may not uniquely identify $w^*(\mu)$, all candidate $w^*(\mu)$ must satisfy $\sum_{a \in \mathcal{A}} w_a^*(\mu) = 1$. The MS sampling rule ensures that the sampling frequencies converge to $\frac{\sum_{a \in \mathcal{A}} N_a(t)}{t} \rightarrow \sum_{a \in \mathcal{A}} w_a^*(\mu)$ a.s.

Proof. Let $\zeta \in (\mu^*, \gamma \wedge \min_{a \notin \mathcal{A}} \mu_a)$. We have that

$$\sum_{a \in \mathcal{A}} \psi_a(t) \geq \Pi_t \left(\arg \min_j \mu_j \in \mathcal{A}, \min_j \mu_j \leq \gamma \right) \geq \underbrace{\max_{a \in \mathcal{A}} \Pi_t(\mu_a \leq \zeta)}_{\rightarrow 1 \text{ by (9)}} \prod_{a \notin \mathcal{A}} \underbrace{\Pi_t(\mu_a \geq \zeta)}_{\rightarrow 1 \text{ by (9)}} \rightarrow 1.$$

It follows that $\frac{\sum_{a \in \mathcal{A}} \Psi_a(t)}{t} \rightarrow 1$, and the result follows from Proposition 10. \square

Next we analyze the behavior of the MS sampling rule on $\mu \in \mathcal{H}_>$. We follow the proof strategy of [29, Section G.1].

Theorem 13. Let $\mu \in \mathcal{H}_>$. Then $\frac{N(t)}{t} \rightarrow w^*(\mu)$ a.s.

Proof. Let us abbreviate $w^* \equiv w^*(\mu)$. By Proposition 10, it suffices to show $\bar{\psi}(t) \rightarrow w^*$. We will show this by applying Proposition 14 below. First, recall from Lemma 1 that

$$T^*(\mu)^{-1} = \max_w \min_{\lambda: \min_a \lambda_a < \gamma} \sum_a w_a d(\mu_a, \lambda_a) = \max_w \min_a w_a d(\mu_a, \gamma) = w_a^* d(\mu_a, \gamma) \quad \forall a. \quad (11)$$

Furthermore, by Proposition 6, for any $a \in [K]$

$$\Pi_n(\theta_a < \gamma) \doteq \exp \left(-n \min_{\lambda: \lambda_a < \gamma} \sum_b \bar{\psi}_b(n) d(\mu_b, \lambda_b) \right) = \exp \left(-n \bar{\psi}_a(n) d(\mu_a, \gamma) \right).$$

In particular, there is a sequence ϵ_n decreasing to zero such that

$$\forall n: \quad \Pi_n(\theta_a < \gamma) \in \exp \left(-n \left(\bar{\psi}_a(n) d(\mu_a, \gamma) \pm \epsilon_n \right) \right).$$

To establish the precondition of Proposition 14 below, fix $a \in [K]$ and $c > 0$ and consider any round n where $\bar{\psi}_a(n) \geq w_a^* + c$. Then

$$\begin{aligned} \psi_a(n) &= \frac{\Pi_{n-1}(a = \arg \min_j \theta_j, \min_j \theta_j < \gamma)}{\Pi_{n-1}(\min_j \theta_j < \gamma)} \leq \frac{\Pi_{n-1}(\theta_a < \gamma)}{\max_a \Pi_{n-1}(\theta_a < \gamma)} \\ &\leq \frac{e^{-n(\bar{\psi}_a(n) d(\mu_a, \gamma) - \epsilon_n)}}{\max_a e^{-n(\bar{\psi}_a(n) d(\mu_a, \gamma) + \epsilon_n)}} = e^{-n(\bar{\psi}_a(n) d(\mu_a, \gamma) - \min_a \bar{\psi}_a(n) d(\mu_a, \gamma) - 2\epsilon_n)} \end{aligned}$$

By (11) $\min_a \bar{\psi}_a(n) d(\mu_a, \gamma) \leq \max_w \min_a w_a d(\mu_a, \gamma) = w_a^* d(\mu_a, \gamma)$. Also $\bar{\psi}_a(n) \geq w_a^* + c$ so

$$\psi_a(n) \leq e^{-n((w_a^* + c)d(\mu_a, \gamma) - w_a^* d(\mu_a, \gamma) - 2\epsilon_n)} = e^{-n(cd(\mu_a, \gamma) - 2\epsilon_n)}.$$

Now as $\epsilon_n \rightarrow 0$, this establishes eventual exponential decay, hence ensuring that

$$\sum_n \psi_a(n) \mathbb{1} \{ \bar{\psi}_a(n) \geq w_a^* + c \} < \infty$$

as required. The conclusion follows from Proposition 14. \square

Proposition 14 ([29, Simplified version of Lemma 11]). *Let $\mathbf{w}^* \equiv \mathbf{w}^*(\boldsymbol{\mu})$. Consider any sampling rule $(A_t)_t$. If for any arm $a \in [K]$ and all $c > 0$*

$$\sum_n \psi_a(n) \mathbb{1} \{ \bar{\psi}_a(n) \geq w_a^* + c \} < \infty$$

then $\bar{\psi}(n) \rightarrow \mathbf{w}^$.*

E Analysis of LCB

The LCB algorithm (3) constructs confidence intervals $[L_a(t), U_a(t)]$. With the Box stopping rule (2) it stops and recommends $<$ when there exists a such that $U_a(t) < \gamma$. It stops and recommends $>$ when for all a , $L_a(t) > \gamma$. When it has not stopped yet, it plays $A_{t+1} = \arg \min_a L_a(t)$ the arm of smallest lower confidence bound.

In this section we show that LCB works fine on $\mathcal{H}_>$, but has the wrong behavior on $\mathcal{H}_<$. For simplicity we only consider the Gaussian case, in which confidence intervals have the stylized form

$$[L_a(t), U_a(t)] = \left[\hat{\mu}_a(t) \mp \sqrt{\frac{2 \ln \frac{1}{\delta}}{N_a(t)}} \right].$$

Note that LCB is not anytime, as its sampling rule is also a function of the confidence level $1 - \delta$. We let τ_δ denote the stopping rule associated to the algorithm that combines the LCB sampling rule with the Box stopping rule, both tuned for the confidence level $1 - \delta$.

Let $\mathcal{E}_\delta = \left\{ \forall t \forall a : |\mu_a - \hat{\mu}_a(t)| \leq \sqrt{\frac{2 \ln \frac{1}{\delta}}{N_a(t)}} \right\}$. By design of the confidence intervals $\mathbb{P}(\mathcal{E}_\delta^c) \leq \delta^2$. Moreover, on \mathcal{E}_δ the algorithm stops and outputs the correct recommendation.

We first show that LCB/Box is sample efficient on $\mathcal{H}_>$.

Proposition 15. *There is a function $\epsilon_\delta \rightarrow 0$ decreasing as $\delta \rightarrow 0$ such that for every $\boldsymbol{\mu} \in \mathcal{H}_>$*

$$\lim_{\delta \rightarrow 0} \mathbb{P}_\boldsymbol{\mu} \left(\frac{\tau_\delta}{\ln \frac{1}{\delta}} \leq (1 + \epsilon_\delta) T^*(\boldsymbol{\mu}) \right) = 1.$$

Proof. Let $\kappa \in (0, 1)$. On the event $\mathcal{E}_{\delta\kappa} \subseteq \mathcal{E}_\delta$ the algorithm (for confidence δ) stops and outputs the correct recommendation, yielding

$$\forall a : L_a(\tau) > \gamma.$$

Moreover, by the sampling rule we have $\forall a : L_a(\tau) \approx \gamma$, and we will ignore the difference. We find

$$\gamma \approx L_a(\tau) = \hat{\mu}_a(\tau) - \sqrt{\frac{2 \ln \frac{1}{\delta}}{N_a(\tau)}} \geq \mu_a - \sqrt{\kappa \frac{2 \ln \frac{1}{\delta}}{N_a(\tau)}} - \sqrt{\frac{2 \ln \frac{1}{\delta}}{N_a(\tau)}} = \mu_a - (1 + \sqrt{\kappa}) \sqrt{\frac{2 \ln \frac{1}{\delta}}{N_a(\tau)}}.$$

We conclude $N_a(\tau) \leq \frac{2(1+\sqrt{\kappa})^2 \ln \frac{1}{\delta}}{(\mu_a - \gamma)^2}$ and hence

$$\tau = \sum_a N_a(\tau) \leq (1 + \sqrt{\kappa})^2 \ln \frac{1}{\delta} \frac{2}{(\mu_a - \gamma)^2} = (1 + \sqrt{\kappa})^2 \ln \frac{1}{\delta} T^*(\boldsymbol{\mu}).$$

The result follows by picking $\kappa = \frac{1}{-\ln \delta}$, achieving $\kappa \rightarrow 0$ as $\delta \rightarrow 0$ yet $\mathbb{P}(\mathcal{E}_{\delta\kappa}^c) \leq \delta^\kappa \rightarrow 0$. \square

Next we consider the behavior on $\mathcal{H}_<$. We characterize the inefficiency of LCB/Box on $\mathcal{H}_<$.

Proposition 16. *There is a function $\epsilon_\delta \rightarrow 0$ decreasing as $\delta \rightarrow 0$ such that for every bandit model $\boldsymbol{\mu} \in \mathcal{H}_<$ with $\mu_1 < \gamma < \mu_2 \leq \dots$ i.e. on which there is only a single arm below the threshold,*

$$\lim_{\delta \rightarrow 0} \mathbb{P}_\boldsymbol{\mu} \left(\forall a \neq a^*(\boldsymbol{\mu}) : \frac{N_a(\tau)}{\ln \frac{1}{\delta}} \geq (1 - \epsilon_\delta) \frac{2}{(\mu_a + \gamma - 2\mu_1)^2} \right) = 1.$$

²As it is written this inequality is not actually correct: in the definition of the confidence interval for arm a , $\ln(1/\delta)$ should be replaced by $\ln(1/\delta) + c \ln \ln(1/\delta) + d \ln(1 + \ln(N_a(t)))$ for some constants c and d (see the discussion in Section 3.2). However, the reasoning that we present with the stylized confidence intervals can be adapted to handle those correct threshold functions, at the price of extra technicalities (e.g., Lemma 22 in [23]).

Proof. Let $\kappa \in (0, 1)$. We analyse the algorithm on the event $\mathcal{E}_{\delta^\kappa} \subseteq \mathcal{E}_\delta$, on which it stops and recommends the correct output. At that time τ , we know

$$U_1(\tau) \leq \gamma \quad \text{and also} \quad \forall a : L_1(\tau) \leq L_a(\tau).$$

Since we are on the event $\mathcal{E}_{\delta^\kappa}$, we know

$$\gamma \geq U_1(\tau) = \hat{\mu}_1(\tau) + \sqrt{\frac{2 \ln \frac{1}{\delta}}{N_1(\tau)}} \geq \mu_1 + \sqrt{\kappa \frac{2 \ln \frac{1}{\delta}}{N_1(\tau)}} + \sqrt{\frac{2 \ln \frac{1}{\delta}}{N_1(\tau)}} = \mu_1 + (1 + \sqrt{\kappa}) \sqrt{\frac{2 \ln \frac{1}{\delta}}{N_1(\tau)}}.$$

On the other hand, we know for each other arm $a \neq 1$ that

$$\sqrt{\frac{2 \ln \frac{1}{\delta}}{N_a(\tau)}} = \hat{\mu}_a(\tau) - L_a(\tau) \leq \hat{\mu}_a(\tau) - L_1(\tau) \leq \mu_a + \sqrt{\kappa \frac{2 \ln \frac{1}{\delta}}{N_a(\tau)}} - L_1(\tau).$$

Finally, since $L_1(\tau) = U_1(\tau) - 2\sqrt{\frac{2 \ln \frac{1}{\delta}}{N_1(\tau)}}$ and $U_1(\tau) \approx \gamma$ (we will ignore the difference), we find

$$(1 - \sqrt{\kappa}) \sqrt{\frac{2 \ln \frac{1}{\delta}}{N_a(\tau)}} \leq \mu_a - \gamma + 2\sqrt{\frac{2 \ln \frac{1}{\delta}}{N_1(\tau)}} \leq \mu_a - \gamma + 2\frac{\gamma - \mu_1}{1 + \sqrt{\kappa}}$$

All in all, this shows

$$N_a(\tau) \geq \frac{2(1 - \sqrt{\kappa})^2 \ln \frac{1}{\delta}}{\left(\mu_a - \gamma + 2\frac{\gamma - \mu_1}{1 + \sqrt{\kappa}}\right)^2} \rightarrow \frac{2}{(\mu_a + \gamma - 2\mu_1)^2}.$$

The result follows by considering the sequence κ exhibited in the proof of Proposition 15. \square

Now this demonstrates a problem, since Lemma 1 shows that optimal algorithms necessarily have $\frac{N_a(\tau)}{\ln \frac{1}{\delta}} \rightarrow 0$, but instead for LCB it tends to a specific positive constant. In other words, a non-vanishing hence significant portion of the samples are wasted “exploring” suboptimal arms.

F Proof of the Deviation Inequality (Theorem 7)

To ease the notation, we introduce $\mu_S^{\min} = \min_{a \in S} \mu_a$ and $\mu_S^{\max} = \max_{a \in S} \mu_a$. Fix $\eta > 0$ and $c > 0$ and define

$$\begin{aligned} X_{\eta,c}(t)^+ &= [N_S(t)d^+(\hat{\mu}_S(t), \mu_S^{\min}) - c(1 + \eta) \ln(1 + \ln N_S(t))] \\ X_{\eta,c}(t)^- &= [N_S(t)d^-(\hat{\mu}_S(t), \mu_S^{\max}) - c(1 + \eta) \ln(1 + \ln N_S(t))] \end{aligned}$$

Throughout the proof we use the notation $X_{\eta,c}(t)$ to refer to either $X_{\eta,c}(t)^+$ or $X_{\eta,c}(t)^-$. The cornerstone of the proof is the following Lemma 17, that tells us that $e^{\lambda X_{\eta,c}(t)}$ can be “almost” upper-bounded by some martingale.

Lemma 17. *Assume $1 + \eta \leq e$. Fix $X_{\eta,c}(t) = X_{\eta,c}(t)^+$ or $X_{\eta,c}(t)^-$. For every $\lambda \in [0, (1 + \eta)^{-1}[$ there exists a martingale M_t^λ such that $\mathbb{E}[M_t^\lambda] = 1$ and*

$$\forall t \in \mathbb{N}^*, M_t^\lambda \geq e^{\lambda X_{\eta,c}(t) - g_{\eta,c}(\lambda)}, \quad (12)$$

with $g_{\eta,c}(\lambda) = \lambda(1 + \eta) \ln\left(\frac{\zeta(c)}{\ln(1 + \eta)^c}\right) - \ln(1 - \lambda(1 + \eta))$.

The deviation inequality follows by combining Chernoff’s method with Doob’s inequality, and then carefully picking η and c . For any $\lambda \in [0, (1 + \eta)^{-1}[$, by Lemma 17,

$$\begin{aligned} \mathbb{P}(\exists t \in \mathbb{N}^* : X_{\eta,c}(t) > u) &\leq \mathbb{P}(\exists t \in \mathbb{N}^* : e^{\lambda X_{\eta,c}(t)} > e^{\lambda u}) \\ &\leq \mathbb{P}(\exists t \in \mathbb{N}^* : M_t^\lambda > e^{\lambda u - g_{\eta,c}(\lambda)}) \\ &\leq \exp(-[\lambda u - g_{\eta,c}(\lambda)]). \end{aligned}$$

Then we want to apply this inequality to the best possible λ . Defining

$$g_{\eta,c}^*(u) = \max_{\lambda \in (0, \frac{1}{1+\eta})} [\lambda u - g_{\eta,c}(\lambda)],$$

a direct computation of this Fenchel conjugate (see Lemma 20 in Appendix F.3) yields

$$g_{\eta,c}^*(u) = h\left(\frac{u}{1+\eta} - \ln\left(\frac{\zeta(c)}{(\ln(1+\eta))^c}\right)\right) - 1,$$

for $\frac{u}{1+\eta} - \ln\left(\frac{\zeta(c)}{(\ln(1+\eta))^c}\right) > 1$, where we recall that $h(x) = x - \ln(x)$.

Using the inequality³ $\ln(1+\eta)^{-1} \leq 1 + \eta^{-1}$ implies that, for $\frac{u}{1+\eta} - \ln(\zeta(c)) - c \ln\left(1 + \frac{1}{\eta}\right) \geq 1$,

$$\mathbb{P}(\exists t \in \mathbb{N}^* : X_{\eta,c}(t) > u) \leq \exp\left(-\left[h\left(\frac{u}{1+\eta} - \ln(\zeta(c)) - c \ln\left(1 + \frac{1}{\eta}\right)\right) - 1\right]\right).$$

For the sake of clarity, we now pick $X_{\eta,c}(t) = X_{\eta,c}^+(t)$. Picking $\eta^* = \frac{c}{u-c}$ (that minimizes the right hand side) it holds that

$$\begin{aligned} & \mathbb{P}\left(\exists t \in \mathbb{N}^* : \left[N_{\mathcal{S}}(t)d^+(\hat{\mu}_{\mathcal{S}}(t), \mu_{\mathcal{S}}^{\min}) - \frac{cu}{u-c} \ln(1 + \ln N_{\mathcal{S}}(t))\right]^+ \geq u\right) \\ & \leq \exp\left(-\left[h\left(u - c - \ln(\zeta(c)) - c \ln\left(\frac{u}{c}\right)\right) - 1\right]\right) \\ & \leq \exp\left(-\left[h\left(ch\left(\frac{u}{c}\right) - c - \ln(\zeta(c))\right) - 1\right]\right) \end{aligned}$$

whenever u is such that $h\left(\frac{u}{c}\right) \geq 1 + \frac{1+\ln\zeta(c)}{c}$ and $1 + \eta^* = \frac{u}{u-c} \leq e$.

Picking $c = 2$, for all $u \geq 6$ the three conditions $\frac{cu}{u-c} \leq 3$, $h\left(\frac{u}{c}\right) \geq 1 + \frac{1+\ln\zeta(c)}{c}$ and $\frac{u}{u-c} \leq e$ are satisfied and one has

$$\mathbb{P}\left(\exists t \in \mathbb{N}^* : \left[N_{\mathcal{S}}(t)d^+(\hat{\mu}_{\mathcal{S}}(t), \mu_{\mathcal{S}}^{\min}) - 3 \ln(1 + \ln N_{\mathcal{S}}(t))\right]^+ \geq u\right) \leq e^{-[h(2h(\frac{u}{2}) - 2 - \ln(\zeta(2))) - 1]}$$

Picking u (large enough) such that

$$h\left(2h\left(\frac{u}{2}\right) - 2 - \ln(\zeta(2))\right) - 1 = x \Leftrightarrow u = T(x)$$

yields inequality (6) in Theorem 7, whenever $T(x) \geq 6$. It can be checked numerically this holds for $x \geq 0.04$. Inequality (7) can be obtained following the same lines by choosing $X_{\eta,c}(t) = X_{\eta,c}^-(t)$. \square

Proofs of intermediate results are now given in separate sections.

F.1 Building the martingale: proof of Lemma 17

Our goal is to propose a martingale M_t^λ that satisfies the assumptions of Lemma 17. Let

$$\phi_\mu(\lambda) = \ln \mathbb{E}_\mu[e^{\lambda X}] = b(\dot{b}^{-1}(\mu) + \lambda) - b(\dot{b}^{-1}(\mu)) \quad (13)$$

denote the cumulant generating function of the distribution that has mean μ . First, it can be checked that for all λ for which $\phi_\mu(\lambda)$ is defined, and for all arm a ,

$$\exp(\lambda S_a(t) - N_a(t)\phi_{\mu_a}(\lambda)) \quad \text{where} \quad S_a(t) = \sum_{s=1}^t \mathbf{1}\{A_s = a\} X_s$$

is a martingale. Due to the fact that only one arm is drawn at each round, the product of these martingales for all the arms in the subset \mathcal{S} is still a martingale, that can be rewritten

$$W_t^\lambda = \exp\left(\sum_{a \in \mathcal{S}} [S_a(t)\lambda - \phi_{\mu_a}(\lambda)N_a(t)]\right).$$

Moreover, $\mathbb{E}[W_t^\lambda] = 1$. We first prove the following result, that relates $X_{\eta,c}(t)$ exceeding a threshold to some W_t^λ martingale exceeding some other threshold, for a well-chosen λ .

³In [21, Theorem 16] we omit this inequality and sharpen the threshold to $T(x) = 2\tilde{h}\left(\frac{h^{-1}(x) + \ln\zeta(2)}{2}\right)$ for $x \geq 3.81$, where $\tilde{h}(x) = h^{-1}(x)e^{1/h^{-1}(x)} \leq h^{-1}(1+x)$.

Lemma 18. Let $i \in \mathbb{N}^*$ and $x > 0$. There exists $\lambda_i^+ = \lambda_i^+(x) < 0$ such that if $N_S(t) \in [(1+\eta)^{i-1}, (1+\eta)^i]$ then

$$\{N_S(t)d^+(\hat{\mu}_S(t), \mu_S^{\min}) \geq x\} \subseteq \{W_t^{\lambda_i^+} \geq e^{\frac{x}{1+\eta}}\}.$$

Moreover, there exists $\lambda_i^- = \lambda_i^-(x) > 0$ such that if $N_S(t) \in [(1+\eta)^{i-1}, (1+\eta)^i]$ then

$$\{N_S(t)d^-(\hat{\mu}_S(t), \mu_S^{\max}) \geq x\} \subseteq \{W_t^{\lambda_i^-} \geq e^{\frac{x}{1+\eta}}\}.$$

Lemma 18 shows that the event of interest is related to a martingale exceeding a threshold for t that belongs to some slice $N_S(t) \in [(1+\eta)^{i-1}, (1+\eta)^i]$. We now prove that for all $x > 0$ and $1+\eta \leq e$, there exists a martingale \tilde{W}_t^x such that $\mathbb{E}[\tilde{W}_t^x] = 1$ and

$$\left\{X_{\eta,c}(t) - (1+\eta) \ln \left(\frac{\zeta(c)}{(\ln(1+\eta))^c} \right) \geq x \right\} \subseteq \{\tilde{W}_t^x \geq e^{\frac{x}{1+\eta}}\}. \quad (14)$$

This martingale is one of the following *mixture martingales*:

$$\tilde{W}_t^{+,x} = \sum_{i=1}^{\infty} \gamma_i W_t^{\lambda_i^+(x+(1+\eta)\ln(1/\gamma_i))} \quad \text{and} \quad \tilde{W}_t^{-,x} = \sum_{i=1}^{\infty} \gamma_i W_t^{\lambda_i^-(x+(1+\eta)\ln(1/\gamma_i))},$$

where $\gamma_i = \frac{1}{\zeta(c)i^c}$ and $\lambda_i^\pm(x)$ are defined in Lemma 18. As $\sum_{i=1}^{\infty} \gamma_i = 1$, $\tilde{W}_t^{\pm,x}$ are martingales that satisfy $\mathbb{E}[\tilde{W}_t^{\pm,x}] = 1$. We first prove that

$$\begin{aligned} & \left\{ N_S(t)d^+(\hat{\mu}_S(t), \mu_S^{\min}) - c(1+\eta) \ln(1 + \ln N_S(t)) - (1+\eta) \ln \left(\frac{\zeta(c)}{(\ln(1+\eta))^c} \right) \geq x \right\} \\ & \subseteq \{\tilde{W}_t^{+,x} \geq e^{\frac{x}{1+\eta}}\}. \end{aligned}$$

If $N_S(t) \in [(1+\eta)^{i-1}, (1+\eta)^i]$, one can observe that $\frac{\ln N_S(t)}{\ln(1+\eta)} \geq i-1$, thus, for $1+\eta \leq e$,

$$\begin{aligned} & c(1+\eta) \ln(1 + \ln N_S(t)) + (1+\eta) \ln \left(\frac{\zeta(c)}{(\ln(1+\eta))^c} \right) = (1+\eta) \ln \left(\frac{\zeta(c)(1 + N_S(t))^c}{(\ln(1+\eta))^c} \right) \\ & \geq (1+\eta) \ln \left(\frac{\zeta(c)(\ln(1+\eta) + N_S(t))^c}{(\ln(1+\eta))^c} \right) = (1+\eta) \ln \left(\zeta(c) \left[1 + \frac{\ln N_S(t)}{\ln(1+\eta)} \right]^c \right) \\ & \geq (1+\eta) \ln(\zeta(c)i^c) = (1+\eta) \ln \frac{1}{\gamma_i} \end{aligned}$$

Thus for $N_S(t) \in [(1+\eta)^{i-1}, (1+\eta)^i]$, it holds using Lemma 18 that

$$\begin{aligned} & \left\{ N_S(t)d^+(\hat{\mu}_S(t), \mu_S^{\min}) - c(1+\eta) \ln(1 + \ln N_S(t)) - (1+\eta) \ln \left(\frac{\zeta(c)}{(\ln(1+\eta))^c} \right) \geq x \right\} \\ & \subseteq \left\{ N_S(t)d^+(\hat{\mu}_S(t), \mu_S^{\min}) \geq x + (1+\eta) \ln \frac{1}{\gamma_i} \right\} \\ & \subseteq \left\{ \tilde{W}_t^{\lambda_i^+(x+(1+\eta)\ln \frac{1}{\gamma_i})} \geq e^{\frac{x}{1+\eta} + \ln(1/\gamma_i)} \right\} = \left\{ \gamma_i \tilde{W}_t^{\lambda_i^+(x+(1+\eta)\ln \frac{1}{\gamma_i})} \geq e^{\frac{x}{1+\eta}} \right\} \\ & \subseteq \{\tilde{W}_t^{+,x} \geq e^{\frac{x}{1+\eta}}\}. \end{aligned}$$

Similarly, one can prove that

$$\begin{aligned} & \left\{ N_S(t)d^-(\hat{\mu}_S(t), \mu_S^{\max}) - c(1+\eta) \ln(1 + \ln N_S(t)) - (1+\eta) \ln \left(\frac{\zeta(c)}{(\ln(1+\eta))^c} \right) \geq x \right\} \\ & \subseteq \{\tilde{W}_t^{-,x} \geq e^{\frac{x}{1+\eta}}\}. \end{aligned}$$

Let $C(\eta) = \frac{\zeta(c)}{(\ln(1+\eta))^c}$. For all $\lambda > 0$ and $z > 1$ it follows from inequality (14) that

$$\left\{ e^{\lambda(X_{\eta,c}(t) - (1+\eta) \ln C(\eta))} \geq z \right\} \subseteq \left\{ \tilde{W}_t^{\frac{1}{\lambda} \ln(z)} \geq e^{\frac{\ln(z)}{\lambda(1+\eta)}} \right\} = \left\{ e^{-\frac{\ln(z)}{\lambda(1+\eta)}} \tilde{W}_t^{\frac{1}{\lambda} \ln(z)} \geq 1 \right\}$$

Letting $\overline{W}_t^{\lambda,z} = e^{-\frac{\ln(z)}{\lambda(1+\eta)}} \tilde{W}_t^{\frac{1}{\lambda} \ln(z)}$, $\overline{W}_t^{\lambda,z}$ is a martingale that satisfies $\mathbb{E}[\overline{W}_t^{\lambda,z}] = e^{-\frac{\ln(z)}{\lambda(1+\eta)}}$. For all $\lambda \in [0, 1/(1+\eta)[$, we now define

$$\overline{M}_t^\lambda = 1 + \int_1^\infty \overline{W}_t^{\lambda,z} dz$$

Using that $\overline{W}_t^{\lambda,z} \geq \mathbb{1}_{(e^{\lambda(X_{\eta,c}(t)-(1+\eta)\ln C(\eta)}) \geq z)}$ and the expression of $\mathbb{E}[\overline{W}_t^{\lambda,z}]$ yields

$$\overline{M}_t^\lambda \geq e^{\lambda(X_{\eta,c}(t)-(1+\eta)\ln C(\eta))} \quad \text{and} \quad \mathbb{E}[\overline{M}_t^\lambda] = \frac{1}{1-\lambda(1+\eta)}.$$

From there, we obtain that $M_t^\lambda = (1-\lambda(1+\eta))\overline{M}_t^\lambda$ satisfies $\mathbb{E}[M_t^\lambda] = 1$ and

$$M_t^\lambda \geq e^{\lambda X_{\eta,c}(t) - \lambda(1+\eta)\ln C(\eta) + \ln(1-\lambda(1+\eta))},$$

which concludes the proof.

F.2 Proof of Lemma 18

Let ν_{\min} be the natural parameter such that $\mu_S^{\min} = \dot{b}(\nu_{\min})$ and ν_{\max} be the natural parameter such that $\mu_S^{\max} = \dot{b}(\nu_{\max})$. Define $\lambda_i^- > 0$ and $\lambda_i^+ < 0$ such that

$$\text{KL}(\nu_{\min} + \lambda_i^+, \nu_{\min}) = \frac{x}{(1+\eta)^i} \quad \text{and} \quad \text{KL}(\nu_{\max} + \lambda_i^-, \nu_{\max}) = \frac{x}{(1+\eta)^i},$$

where $\text{KL}(\nu, \nu')$ is the Kullback-Leibler divergence between the distributions of natural parameter ν and ν' . Defining $\mu_i^+ := \dot{b}^{-1}(\nu + \lambda_i^+) < \mu_S^{\min}$ and $\mu_i^- := \dot{b}^{-1}(\nu + \lambda_i^-) > \mu_S^{\max}$ and using some properties of the KL-divergence for exponential families, one can write

$$\begin{aligned} d(\mu_i^+, \mu_S^{\min}) &= \text{KL}(\nu_{\min} + \lambda_i^+, \nu_{\min}) = \lambda_i^+ \mu_i^+ - \phi_{\mu_S^{\min}}(\lambda_i^+) \\ d(\mu_i^-, \mu_S^{\max}) &= \text{KL}(\nu_{\max} + \lambda_i^-, \nu_{\max}) = \lambda_i^- \mu_i^- - \phi_{\mu_S^{\max}}(\lambda_i^-). \end{aligned}$$

For $N_S(t) \in [(1+\eta)^{i-1}, (1+\eta)^i]$, one can write (using notably that λ_i^+ is negative)

$$\begin{aligned} \{N_S(t)d^+(\hat{\mu}_S(t), \mu_S^{\min}) \geq x\} &\subseteq \left\{d^+(\hat{\mu}_S(t), \mu_S^{\min}) \geq \frac{x}{(1+\eta)^i}\right\} \\ &\subseteq \{\hat{\mu}_S(t) \leq \mu_i^+\} \\ &\subseteq \left\{\lambda_i^+ \hat{\mu}_S(t) - \phi_{\mu_S^{\min}}(\lambda_i^+) \geq \lambda_i^+ \mu_i^+ - \phi_{\mu_S^{\min}}(\lambda_i^+) = \text{KL}(\nu_{\min} + \lambda_i^+, \nu_{\min})\right\} \\ &\subseteq \left\{\lambda_i^+ \hat{\mu}_S(t) - \phi_{\mu_S^{\min}}(\lambda_i^+) \geq \frac{x}{(1+\eta)^i}\right\} \\ &\subseteq \left\{\lambda_i^+ N_S(t) \hat{\mu}_S(t) - N_S(t) \phi_{\mu_S^{\min}}(\lambda_i^+) \geq \frac{x}{1+\eta}\right\} \\ &\subseteq \left\{\lambda_i^+ \sum_{a \in \mathcal{S}} S_a(t) - \left(\sum_{a \in \mathcal{S}} N_a(t)\right) \phi_{\mu_S^{\min}}(\lambda_i^+) \geq \frac{x}{1+\eta}\right\} \end{aligned}$$

Now using Lemma 19 below, that can easily be checked by differentiating the equality (13), one can use that as $\lambda_i^+ < 0$,

$$\forall a \in \mathcal{S}, \phi_{\mu_S^{\min}}(\lambda_i^+) \geq \phi_{\mu_a}(\lambda_i^+).$$

Therefore, it follows that

$$\{N_S(t)d^+(\hat{\mu}_S(t), \mu_S^{\min}) \geq x\} \subseteq \left\{\lambda_i^+ \sum_{a \in \mathcal{S}} S_a(t) - \sum_{a \in \mathcal{S}} N_a(t) \phi_{\mu_a}(\lambda_i^+) \geq \frac{x}{1+\eta}\right\} = \left\{W_t^{\lambda_i^+} \geq \frac{x}{1+\eta}\right\}$$

This proves the first inclusion in Lemma 18. The proof of the second inclusion follows exactly the same lines, using this time that $\lambda_i^- > 0$:

$$\begin{aligned}
\{N_{\mathcal{S}}(t)d^-(\hat{\mu}_{\mathcal{S}}(t), \mu_{\mathcal{S}}^{\max}) \geq x\} &\subseteq \{\hat{\mu}_{\mathcal{S}}(t) \geq \mu_i^-\} \\
&\subseteq \left\{ \lambda_i^- \hat{\mu}_{\mathcal{S}}(t) - \phi_{\mu_{\mathcal{S}}^{\max}}(\lambda_i^-) \geq \frac{x}{(1+\eta)^i} \right\} \\
&\subseteq \left\{ \lambda_i^- \sum_{a \in \mathcal{S}} S_a(t) - \left(\sum_{a \in \mathcal{S}} N_a(t) \right) \phi_{\mu_{\mathcal{S}}^{\max}}(\lambda_i^-) \geq \frac{x}{1+\eta} \right\} \\
&\subseteq \left\{ \lambda_i^- \sum_{a \in \mathcal{S}} S_a(t) - \sum_{a \in \mathcal{S}} N_a(t) \phi_{\mu_a}(\lambda_i^-) \geq \frac{x}{1+\eta} \right\},
\end{aligned}$$

where the last inequality uses that by Lemma 19 $\forall a \in \mathcal{S}, \phi_{\mu_{\mathcal{S}}^{\max}}(\lambda_i^-) \geq \phi_{\mu_a}(\lambda_i^-)$.

Lemma 19. *The mapping $\mu \mapsto \phi_{\mu}(\lambda)$ is non-increasing if $\lambda < 0$ and non-decreasing if $\lambda > 0$.*

F.3 Technical Results

Lemma 20. *Define $g(\lambda) = A\lambda - \ln(1 - \lambda B)$ for $\lambda \in [0, B^{-1}[$. Then if $\frac{u-A}{B} \geq 1$,*

$$g^*(u) = \max_{\lambda \in [0, B^{-1}[} [\lambda u - g(\lambda)] = h\left(\frac{u-A}{B}\right) - 1,$$

where $h(u) = u - \ln u$.

Proof. Differentiating shows that $\lambda \mapsto \lambda u - g(\lambda)$ attains its maximum on in $\lambda^* = \frac{1}{B} - \frac{1}{x-A}$, which also satisfies $1 - \lambda^* B = \frac{B}{x-A}$. If $\frac{u-A}{B} \geq 1$, $\lambda^* \in [0, B^{-1}[$ and one obtains

$$\begin{aligned}
g^*(u) &= \lambda^* u - g(\lambda^*) = \lambda^*(u - A) + \ln(1 - \lambda^* B) \\
&= \frac{x-A}{B} - \ln \frac{x-A}{B} - 1 \\
&= h\left(\frac{x-A}{B}\right) - 1.
\end{aligned}$$

□

The next result permits to derive a tight upper bound on the threshold function T featured in Theorem 7. Recall this function is defined in terms of the inverse mapping of $h : [1, +\infty[\rightarrow \mathbb{R}^*$ defined by $h(u) = u - \ln(u)$.

Lemma 21. *h is increasing on $[1, +\infty[$ and its inverse function, defined on $[1, +\infty[$ can be expressed in terms of negative branch of the Lambert function: $h^{-1}(x) = -W_{-1}(-e^{-x})$. The following inequality holds:*

$$\forall x \geq 1, \quad h^{-1}(x) \leq x + \ln(x + \sqrt{2(x-1)}).$$

Proof. We may write

$$h^{-1}(x) = \inf_{z \geq 1} z \left(x - 1 + \ln \frac{z}{z-1} \right)$$

Plugging in the sub-optimal feasible choice $z = 1 + \frac{1}{(x-1) + \sqrt{2(x-1)}}$ reveals

$$\begin{aligned}
h^{-1}(x) &\leq \left(1 + \frac{1}{(x-1) + \sqrt{2(x-1)}} \right) \left(x - 1 + \ln \left(x + \sqrt{2(x-1)} \right) \right) \\
&\leq 1 + (x-1) + \ln \left(x + \sqrt{2(x-1)} \right).
\end{aligned}$$

Where the last inequality uses $\ln \left(x + \sqrt{2(x-1)} \right) \leq \sqrt{2(x-1)}$ which holds with equality at $x = 1$ and whose gap is increasing (as can be checked by differentiation). □

G Aggregate Stopping Rule is δ -correct (Lemma 8)

First assume $\mu \in \mathcal{H}_>$. Then the probability of error is upper bounded by

$$\begin{aligned}
& \mathbb{P}(\exists t \in \mathbb{N}, \exists \mathcal{S} : N_{\mathcal{S}}(t) d^+(\hat{\mu}_{\mathcal{S}}(t), \theta) \geq 3 \ln(1 + \ln(N_{\mathcal{S}}(t))) + T(\ln(1/(\delta\pi(\mathcal{S})))) \\
& \leq \sum_{\mathcal{S}} \mathbb{P}(\exists t \in \mathbb{N} : N_{\mathcal{S}}(t) d^+(\hat{\mu}_{\mathcal{S}}(t), \theta) \geq 3 \ln(1 + \ln(N_{\mathcal{S}}(t))) + T(\ln(1/(\delta\pi(\mathcal{S})))) \\
& \leq \sum_{\mathcal{S}} \mathbb{P}\left(\exists t \in \mathbb{N} : N_{\mathcal{S}}(t) d^+\left(\hat{\mu}_{\mathcal{S}}(t), \min_{a \in \mathcal{S}} \mu_a\right) \geq 3 \ln(1 + \ln(N_{\mathcal{S}}(t))) + T(\ln(1/(\delta\pi(\mathcal{S}))))\right) \\
& \leq \sum_{\mathcal{S}} \delta\pi(\mathcal{S}) = \delta.
\end{aligned}$$

The second inequality uses that on $\mathcal{H}_<$, all μ_a are larger than γ and $x \mapsto d^+(\hat{\mu}_{\mathcal{S}}(t), x)$ is non-decreasing. The last inequality follows from the first inequality in Theorem 7.

Now assume $\mu \in \mathcal{H}_<$: there exists a such that $\mu_a < \gamma$. The probability of error is upper bounded by

$$\begin{aligned}
& \mathbb{P}(\exists t \in \mathbb{N}, \forall a, N_a(t) d^-(\hat{\mu}_a(t), \gamma) \geq 3 \ln(1 + \ln(N_a(t))) + T(\ln(1/\delta))) \\
& \leq \mathbb{P}(\exists t \in \mathbb{N} : N_a(t) d^-(\hat{\mu}_a(t), \gamma) \geq 3 \ln(1 + \ln(N_a(t))) + T(\ln(1/\delta))) \\
& \leq \mathbb{P}(\exists t \in \mathbb{N} : N_a(t) d^-(\hat{\mu}_a(t), \mu_a) \geq 3 \ln(1 + \ln(N_a(t))) + T(\ln(1/\delta))) \leq \delta.
\end{aligned}$$

The second inequality holds as $\mu_a < \gamma$ and $x \mapsto d^-(\hat{\mu}_a(t), \gamma)$ is non-increasing. The last inequality is an application of the second inequality of Theorem 7, for singleton $\mathcal{S} = \{a\}$.