



HAL
open science

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

Emilie Kaufmann, Wouter Koolen, Aurélien Garivier

► **To cite this version:**

Emilie Kaufmann, Wouter Koolen, Aurélien Garivier. Sequential Test for the Lowest Mean: From Thompson to Murphy Sampling. *Advances in Neural Information Processing Systems (NIPS)*, Dec 2018, Montréal, Canada. hal-01804581

HAL Id: hal-01804581

<https://hal.science/hal-01804581v1>

Submitted on 31 May 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

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

Emilie Kaufmann¹, Wouter M. Koolen² and Aurélien Garivier³

emilie.kaufmann@univ-lille1.fr, wmkoolen@cwi.nl, aurelien.garivier@math.univ-toulouse.fr

¹ CNRS & Univ. Lille, UMR 9189 CRIStAL, Inria SequeL, Lille, France

² Centrum Wiskunde & Informatica, Science Park 123, 1098 XG Amsterdam, The Netherlands

³ Institut de Mathématiques de Toulouse; CNRS UMR5219, Université de Toulouse, France

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 [29, 27]. 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 [35] and subsequently [10, 32, 9]. 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 [17] and games [26]. Such confidence intervals are used hierarchically for Monte Carlo Tree Search (MCTS) in [33, 15, 19, 24]. 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 interval (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.

We thus study a particular example of sequential adaptive hypothesis testing problem, as introduced by Chernoff [7], 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 [11, 22], of which several variants have been studied [21, 19, 14]. For example the Thresholding Bandit Problem [28] aims at finding the set of arms above a threshold, which is strictly harder than our testing problem.

A full characterization of the asymptotic complexity of the BAI problem was recently given in [13], 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. [25, 14]). Such an optimal allocation is also presented by [6] 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 [13]), 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 [34, 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 [30] 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 [20, 22] in the Gaussian case and by [5] 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 associated confidence intervals are presented in Section 5. Finally, numerical experiments reported in Section 6 demonstrate the efficiency of Murphy Sampling paired with our new Aggregate 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 [5] 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} an possibly some exogenous random variable, τ is a stopping time with respect to this filtration and \hat{m} is \mathcal{F}_τ -measurable.

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. [12] 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*

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

Lemma 1 is proved in Appendix B. As explained by [12] 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 [16] (see also [31] 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 $a \in \{1, \dots, K\}$,*

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

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,*

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

When $d(\mu_1, \gamma) \geq d(\mu_K, \gamma)$, this bound can be rewritten as:

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

The lower bound for the case $\mu^* > \gamma$ can also be boosted similarly, with a less explicit result.

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

$$\tau_{\text{Box}} = \tau_< \wedge \tau_> \quad \text{where} \quad \begin{aligned} \tau_< &= \inf \{t \in \mathbb{N}^* : \exists a N_a(t) d^+(\hat{\mu}_a(t), \gamma) \geq C_<(\delta, N_a(t))\}, \\ \tau_> &= \inf \{t \in \mathbb{N}^* : \forall a N_a(t) d^-(\hat{\mu}_a(t), \gamma) \geq C_>(\delta, N_a(t))\}, \end{aligned} \quad (3)$$

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 [20, 22] 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 $\boldsymbol{\mu} \in \mathcal{H}_>$. LCB selects at each round the arm with smallest Lower Confidence Bound:

$$\boxed{\text{LCB: Play } A_t = \operatorname{argmin}_a L_a(t),} \quad (4)$$

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 $\boldsymbol{\mu} \in \mathcal{H}_>$ 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 $\boldsymbol{\mu} \in \mathcal{H}_<$, 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 [4] for a survey) provides candidate algorithms that have this type of behavior, and we propose to use the Thompson Sampling (TS) algorithm [1, 23]. 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

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

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 $\boldsymbol{\mu}$. 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 $\tau_{\text{Box}} \simeq (1/d(\boldsymbol{\mu}_1, \boldsymbol{\theta})) \ln(1/\delta)$. However, TS cannot be optimal for $\boldsymbol{\mu} \in \mathcal{H}_>$, as the empirical proportions of draws cannot converge to $\boldsymbol{w}^*(\boldsymbol{\mu}) \neq \mathbf{1}_{a^*}$.

To summarize, we presented a simple stopping rule, τ_{Box} , that can be asymptotically optimal for every $\boldsymbol{\mu} \in \mathcal{H}_<$ if it is used in combination with Thompson Sampling and for $\boldsymbol{\mu} \in \mathcal{H}_>$ 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 \boldsymbol{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 $\boldsymbol{\mu} \in \mathcal{H}$. Fix an anytime sampling strategy (A_t) ensuring $\frac{N_t}{t} \rightarrow \boldsymbol{w}^*(\boldsymbol{\mu})$. Let τ_δ be a stopping rule such that $\tau_\delta \leq \tau_\delta^{\text{Box}}$, for a Box stopping rule (3) whose threshold functions C_\leq satisfy the following: they are non-decreasing in r and there exists a function f such that,

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

Then $\limsup_{\delta \rightarrow 0} \frac{\tau_\delta}{\ln \frac{1}{\delta}} \leq T^*(\boldsymbol{\mu})$ 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" ($\boldsymbol{\mu} \in \mathcal{H}_<$):

$$\boxed{\text{MS: Sample } \boldsymbol{\theta}_t \sim \Pi_{t-1}(\cdot | \mathcal{H}_<), \text{ then play } A_t = a^*(\boldsymbol{\theta}_t).} \quad (5)$$

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 $\boldsymbol{\mu}$ w.r.t. γ . This is different from the symmetric Top Two Thompson

Sampling [30], 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. In our experiments on $\mathcal{H}_>$ (with moderate δ), stopping rules kick in before the rejection probability becomes impractically high.

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 $\Theta \ni \mu$ (or the support of the prior) is the interior of a bounded subset of \mathbb{R}^K . This ensures that $\sup_{\mu, \theta \in \Theta} d(\mu, \theta) < \infty$ and $\sup_{\mu, \theta \in \Theta} \|\mu - \theta\| < \infty$. This assumption is common [18, Section 7.1], [30, Assumption 1]. We also assume that the prior Π has a density π with bounded ratio $\sup_{\mu, \theta \in \Theta} \frac{\pi(\theta)}{\pi(\mu)} < \infty$.

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

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 (3) or the improved Aggregate stopping rule discussed in Section 5).

Proof Sketch First, consider $\mu \in \mathcal{H}_<$. 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 $\mu \in \mathcal{H}_>$. Following [30], 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 ([30, 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

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

This is an asymptotic recurrence relation in $\psi_a(n)$. To get a good sense for what is happening we may solve the exact analogue. Abbreviating $d_a = d(\mu_a, \gamma)$, we find $\Psi_a(n) = \left(n - \sum_j \frac{\ln d_j}{d_j}\right) \frac{\frac{1}{d_a}}{\sum_j \frac{1}{d_j}} + \frac{\ln d_a}{d_a}$ and hence $\psi_a(t) = \Psi_a(t) - \Psi_a(t-1) = \frac{1/d_a}{\sum_j 1/d_j} = w_a^*(\mu)$. Proposition 10 then establishes that $N_a(t)/t \sim \Psi_a(t)/t \rightarrow w_a^*(\mu)$ as well.

In our proof in Appendix D we technically bypass solving the above approximate recurrence, and proceed to pin down the answer by composing the appropriate one-sided bounds. Yet as we were guided by the above picture of $w^*(\mu)$ as the eventually stable direction of the dynamical system governing the sampling proportions, we believe it is more revealing.

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*

$$T(x) = 2h^{-1} \left(1 + \frac{h^{-1}(1+x) + \ln \zeta(2)}{2} \right) \quad (6)$$

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 $x \geq 0.04$,

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

$$\mathbb{P} \left(\exists t \in \mathbb{N} : N_{\mathcal{S}}(t) d^- \left(\hat{\mu}_{\mathcal{S}}(t), \max_{a \in \mathcal{S}} \mu_a \right) \geq 3 \ln(1 + \ln(N_{\mathcal{S}}(t))) + T(x) \right) \leq e^{-x}. \quad (8)$$

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 [20, 22] 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 (6) 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 $\sum_{\mathcal{S} \subseteq \{1, \dots, K\}} \pi(\mathcal{S}) = 1$ 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 [2]). 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 (7) (and a union bound over subsets) also permits to build a tight upper confidence bound on the minimum μ^* . Indeed, defining

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

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 5 in Appendix A. Observe that using inequality (8) in Theorem 7 similarly allows to derive tighter lower confidence bounds on the maximum of several means.

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

$$\mathbb{E}[M_0^\lambda] = 1 \quad \text{and} \quad \forall t \in \mathbb{N}^*, M_t^\lambda \geq e^{\lambda X_\eta(t) - g_\eta(\lambda)}. \quad (9)$$

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 (7) 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 (9). First, using the so-called method of mixtures [8] and some specific fact about exponential families already exploited by [5], we can prove that there exists a martingale \tilde{W}_t^x such that for some function f (see Equation (15))

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

From there it follows that, for every λ and $z > 1$, $\{e^{\lambda(X_\eta(t) - f(\eta))} \geq z\} \subseteq \{e^{-\frac{\ln(z)}{\lambda(1+\eta)}} \tilde{W}_t^{\frac{1}{\lambda}} \geq 1\}$ 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 $M_t^\lambda = \overline{M}_t^\lambda / \mathbb{E}[\overline{M}_t^\lambda]$.