



HAL
open science

Dealing with Unknown Variances in Best-Arm Identification

Marc Jourdan, Rémy Degenne, Emilie Kaufmann

► **To cite this version:**

Marc Jourdan, Rémy Degenne, Emilie Kaufmann. Dealing with Unknown Variances in Best-Arm Identification. *Algorithmic Learning Theory (ALT)*, Feb 2023, Singapore (SG), Singapore. hal-04306221

HAL Id: hal-04306221

<https://hal.science/hal-04306221>

Submitted on 24 Nov 2023

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.



Distributed under a Creative Commons Attribution 4.0 International License

Dealing with Unknown Variances in Best-Arm Identification

Marc Jourdan

MARC.JOURDAN@INRIA.FR

Rémy Degenne

REMY.DEGENNE@INRIA.FR

Emilie Kaufmann

EMILIE.KAUFMANN@UNIV-LILLE.FR

Univ. Lille, CNRS, Inria, Centrale Lille, UMR 9198-CRISAL, F-59000 Lille, France

Editors: Shipra Agrawal and Francesco Orabona

Abstract

The problem of identifying the best arm among a collection of items having Gaussian rewards distribution is well understood when the variances are known. Despite its practical relevance for many applications, few works studied it for unknown variances. In this paper we introduce and analyze two approaches to deal with unknown variances, either by plugging in the empirical variance or by adapting the transportation costs. In order to calibrate our two stopping rules, we derive new time-uniform concentration inequalities, which are of independent interest. Then, we illustrate the theoretical and empirical performances of our two sampling rule wrappers on Track-and-Stop and on a Top Two algorithm. Moreover, by quantifying the impact on the sample complexity of not knowing the variances, we reveal that it is rather small.

Keywords: Gaussian Bandits, Unknown Variances, Best-arm Identification.

1. Introduction

In a bandit model, an agent sequentially collects samples from unknown probability distributions, called arms. These samples may be viewed as rewards that the agent seeks to maximize, or equivalently minimize its regret (Bubeck and Cesa-Bianchi, 2012). In this paper our focus is instead on a Best Arm Identification (BAI) problem in which the agent should identify the arm that has the largest expected reward using as few samples as possible, without incentive on maximizing rewards.

We are interested in a Gaussian bandit model in which the variances of the arms are *unknown*. Quite surprisingly, and despite its practical relevance, this problem has received little attention in the bandit literature. Gaussian distributions could indeed be used to model the revenue generated by different versions of a website in the context of A/B testing, or some biological indicator of the efficiency of a treatment in the context of an adaptive clinical trial comparing several treatments. In both case, assuming known variances is a limitation. Formally, we consider a bandit model with K arms ν_1, \dots, ν_K in which ν_a is a Gaussian distribution with mean μ_a and variance σ_a . The best arm (assumed unique) is defined as the arm with largest mean $a^*(\mu) \stackrel{\text{def}}{=} \arg \max_{a \in [K]} \mu_a$. We consider the fixed confidence setting, in which the parameter $\delta \in (0, 1)$ is an upper bound on the probability that the algorithm makes an error.

A fixed confidence BAI algorithm is made of a sampling rule and a stopping and recommendation rule. In each time $t \in \mathbb{N}$, an arm $a_t \in [K]$ is chosen by the *sampling rule*, then an observation $X_{t,a_t} \sim \nu_{a_t}$ is received. The choice of a_t may depend on a random variable U_{t-1} , independent of everything else, which models internal randomization. The σ -algebra generated by $(U_0, X_{1,a_1}, \dots, U_{t-1}, X_{t,a_t}, U_t)$ is denoted by \mathcal{F}_t . a_t is then an \mathcal{F}_{t-1} -measurable random variable, and X_t is independent of \mathcal{F}_{t-1} conditionally on a_t . The *stopping rule* is a stopping time with respect

to the filtration $(\mathcal{F}_t)_{t \in \mathbb{N}}$, denoted by τ_δ . When the algorithm stops, it recommends an arm $\hat{a}_{\tau_\delta} \in [K]$, which is measurable with respect to τ_δ . τ_δ is called the sample complexity of the algorithm.

The goal of a fixed confidence best arm identification algorithm is to return the best arm with high probability while having low sample complexity. The main requirement we impose on a fixed confidence identification method is δ -correctness.

Definition 1 (δ -correct) *Let \mathcal{D} be a set of distributions on \mathbb{R} . Given $\delta \in (0, 1)$, we say that an identification strategy is δ -correct on the problem class \mathcal{D}^K if for all $\nu = (\nu_a)_{a \in [K]} \in \mathcal{D}^K$, $\mathbb{P}_\nu(\tau_\delta < +\infty, \hat{a}_{\tau_\delta} \neq a^*(\mu)) \leq \delta$.*

We follow the approach pioneered by [Garivier and Kaufmann \(2016\)](#) and initially introduced for one-dimensional parametric models (e.g. Gaussian with known variance). They derived lower bounds on the expected sample complexity of δ -correct algorithms and introduced algorithms inspired by the maximization of those lower bounds. Extending their lower bound to our two-parameters setting allows us to quantify the impact on the expected sample complexity of not knowing the variances, and reveals that this impact is rather small. To leverage the stopping and sampling rules of existing algorithms, we propose two approaches to deal with unknown variances: plugging in the empirical variance or considering the transportation costs for unknown variance.

As it is common in previous work for the stopping rule, we will compare a Generalized Likelihood Ratio (GLR) to a well chosen threshold ([Kaufmann and Koolen, 2021](#)). Our two approaches yield the Empirical Variance GLR (EV-GLR) stopping rule, which plugs in the empirical variance in a GLR assuming known variance, and the GLR stopping rule, which corresponds to a GLR assuming unknown variance. Our main technical contribution lies in the derivation of (near) optimal stopping thresholds which ensure the δ -correctness of both the GLR and the EV-GLR stopping rules, regardless of the sampling rule. These thresholds are based on new time-uniform concentration inequalities for Gaussian with unknown variances, which are of independent interest ([Corollary 26](#) and [Theorem 43](#)).

When considering the sampling rule, each approach yields a wrapper which is a simple procedure that can be applied to any BAI algorithm for known variances. We illustrate each wrapper with Track-and-Stop ([Garivier and Kaufmann, 2016](#)) and the Top Two algorithm β -EB-TCI ([Jourdan et al., 2022](#)). By deriving upper bound on the expected sample complexity, we show that algorithms obtained by adapting the transportation costs enjoy stronger theoretical guarantees than the ones plugging in the empirical variance. In particular, we propose the first asymptotically optimal algorithms for Gaussian bandits with unknown variances. Our experiments reveal that both wrappers have comparable performance when applied to several BAI algorithms including the ones above, DKM ([Degenne et al., 2019](#)) and FWS ([Wang et al., 2021](#)). This reinforces our finding that not knowing the variances has a small impact on the sample complexity.

Related work Algorithms based on GLR stopping rules and aimed at matching a sample complexity lower bound were either studied for one-parameter exponential families ([Degenne et al., 2019](#)) or under generic heavy tails assumption ([Agrawal et al., 2020](#)). Other algorithms are either based on eliminations or on confidence intervals and have been mostly analyzed for sub-Gaussian distributions with a known variance proxy¹ ([Even-Dar et al., 2006](#); [Kalyanakrishnan et al., 2012](#); [Jamieson et al., 2014](#)). For the special case of bounded distributions, confidence intervals based

1. A random variable X with mean μ is σ^2 sub-Gaussian if $\mathbb{E}[\exp(\lambda(X - \mu))] \leq \frac{\lambda^2 \sigma^2}{2}$ for all $\lambda \in \mathbb{R}$.

on the empirical variance have been used (Gabillon et al., 2012; Lu et al., 2021) but the resulting algorithms cannot be applied to unbounded distributions as they rely on the empirical Bernstein inequality (Maurer and Pontil, 2009). In the fixed budget setting, in which the size of the exploration phase is fixed in advance, it is possible to upper bound the error probability of the Successive Reject algorithm of Audibert et al. (2010) when the variances are unknown, as we only need to upper bound the probability that one empirical mean is smaller than another, see also Faella et al. (2020). However, in the fixed-confidence setting elimination thresholds, confidence intervals or GLR tests need to be calibrated in a data-dependent way, which calls for the development of new time-uniform concentration inequalities, that we provide in this work.

In the related literature on ranking and selection (Hong et al., 2021), the problem of finding the Gaussian distribution with largest mean has been studied for unknown variances. This literature mostly seek to design algorithm that are δ -correct whenever the gap between the best and second best arm is larger than some specified indifference zone (Kim and Nelson, 2001). However the work of Fan et al. (2016) does not consider an indifference zone and their algorithm is therefore comparable to ours. They propose an elimination strategy which features the empirical variances and whose calibration is done based on simulation arguments (resorting to continuous-time approximations) and justified in an asymptotic regime only (when δ goes to zero). Our algorithms have better empirical performance and stronger theoretical guarantees.

2. Lower Bounds and GLR-based Stopping Rules

First, we introduce the lower bounds characterizing the complexity of the setting in Section 2.1. Then, we present the generalized log-likelihood ratios (GLR) stopping rules in Section 2.3.

2.1. Lower Bounds

In the following, all the distributions are Gaussian denoted by $\nu_{x,\sigma^2} = \mathcal{N}(x, \sigma^2)$. The class of Gaussian distributions with known variance σ^2 is denoted by $\mathcal{D}_{\sigma^2} = \{\nu_{x,\sigma^2} \mid \exists x \in \mathbb{R}\}$, and the class of Gaussian distributions with unknown variance by $\mathcal{D} = \bigcup_{\sigma^2 > 0} \mathcal{D}_{\sigma^2}$. We denote the Kullback-Leibler (KL) divergence between ν_{x_1, σ_1^2} and ν_{x_2, σ_2^2} by $\text{KL}((x_1, \sigma_1^2), (x_2, \sigma_2^2))$.

Let $(\mu, \sigma^2) \in \mathcal{M} = \mathbb{R}^K \times (\mathbb{R}_+^K)$ such that $|a^*(\mu)| = 1$. The alternative sets $\Lambda(\mu, \sigma^2) = \{(\lambda, \kappa^2) \in \mathcal{M} \mid a^*(\mu) \notin \arg \max_a \lambda_a\}$ and $\Lambda_{\sigma^2}(\mu) = \{\lambda \mid (\lambda, \sigma^2) \in \mathcal{M}, a^*(\mu) \notin \arg \max_a \lambda_a\}$ are the sets of parameter for which $a^*(\mu)$ is not the best arm. The $(K-1)$ -dimensional probability simplex is denoted by $\Delta_K = \{w \in \mathbb{R}_+^K \mid \sum_{a \in [K]} w_a = 1\}$.

For Gaussian with unknown (resp. known) variances, Lemma 2 shows that $T^*(\mu, \sigma^2)$ (resp. $T_{\sigma^2}^*(\mu)$) is the asymptotic complexity of the BAI problem on the instance $\nu \stackrel{\text{def}}{=} (\nu_{\mu_a, \sigma_a^2})_a$, where

$$T^*(\mu, \sigma^2)^{-1} = \sup_{w \in \Delta_K} \inf_{(\lambda, \kappa^2) \in \Lambda(\mu, \sigma^2)} \sum_{a \in [K]} w_a \text{KL}((\mu_a, \sigma_a^2), (\lambda_a, \kappa_a^2)),$$

$$T_{\sigma^2}^*(\mu)^{-1} = \sup_{w \in \Delta_K} \inf_{\lambda \in \Lambda_{\sigma^2}(\mu)} \sum_{a \in [K]} w_a \text{KL}((\mu_a, \sigma_a^2), (\lambda_a, \sigma_a^2)).$$

The maximizer over the simplex Δ_K in these complexities is denoted by $w^*(\mu, \sigma^2)$ and $w_{\sigma^2}^*(\mu)$. The rationale for the difference between the $T^*(\mu, \sigma^2)$ and $T_{\sigma^2}^*(\mu)$ is that when the variances are unknown, there exist instances of the form (λ, κ^2) for $\kappa \neq \sigma$ that are harder to differentiate from (μ, σ^2) than instances of the form (λ, σ^2) with respect to an information criterion.

Lemma 2 (Garivier and Kaufmann (2016)) *An algorithm which is δ -correct on all problems in $\mathcal{D}_{\sigma^2}^K$ satisfies that for all $\mu \in \mathbb{R}^K$, $\mathbb{E}_\nu[\tau_\delta] \geq T_{\sigma^2}^*(\mu) \log(1/(2.4\delta))$.*

An algorithm which is δ -correct on all problems in \mathcal{D}^K satisfies that for all $(\mu, \sigma^2) \in \mathcal{M}$, $\mathbb{E}_\nu[\tau_\delta] \geq T^(\mu, \sigma^2) \log(1/(2.4\delta))$.*

We say that an algorithm is asymptotically optimal on \mathcal{D}^K if it is δ -correct and its sample complexity matches that lower bound, i.e. $\liminf_{\delta \rightarrow 0} \mathbb{E}_\mu[\tau_\delta]/\log(1/\delta) \leq T^*(\mu, \sigma^2)$. A weaker notion of optimality is β -optimality (Qin et al., 2017; Shang et al., 2020). An algorithm is called asymptotically β -optimal on \mathcal{D}^K if it satisfies $\liminf_{\delta \rightarrow 0} \mathbb{E}_\mu[\tau_\delta]/\log(1/\delta) \leq T_\beta^*(\mu, \sigma^2)$ and is δ -correct, for $T_\beta^*(\mu, \sigma^2)$ defined as follows. For $\beta \in (0, 1)$, the definition of $T_\beta^*(\mu, \sigma^2)$ is the same as $T^*(\mu, \sigma^2)$ with the additional constraint on the outer maximization that $w_{a^*} = \beta$, hence $T^*(\mu, \sigma^2) = \min_{\beta \in (0, 1)} T_\beta^*(\mu, \sigma^2)$.

An asymptotically β -optimal algorithm is asymptotically minimizing the sample complexity among algorithms which allocate a β fraction of samples to the best arm. Russo (2016) shows that an asymptotically β -optimal algorithm with $\beta = 1/2$ also has an expected sample complexity which is asymptotically optimal, up to a multiplicative factor 2, i.e. $T_{1/2}^*(\mu, \sigma^2) \leq 2T^*(\mu, \sigma^2)$. The β -optimality on $\mathcal{D}_{\sigma^2}^K$ involves $T_{\sigma^2, \beta}^*(\mu)$, which is similarly related to $T_{\sigma^2}^*(\mu)$. While there is a rich literature on asymptotically (β -)optimal algorithms for Gaussian with known variance, we are the first to derive algorithms with those guarantees when the variances are unknown.

2.2. Comparing the Complexities

To compare $T_{\sigma^2}^*(\mu)$ and $T^*(\mu, \sigma^2)$, we first propose a more explicit expression of the infimum over the alternative set featured in their expression, in terms of appropriate *transportation costs*.

Lemma 3 *For every μ such that $a^*(\mu) = \{a^*\}$ and $w \in \mathbb{R}_+^K$,*

$$\inf_{(\lambda, \kappa^2) \in \Lambda(\mu, \sigma^2)} \sum_{a \in [K]} w_a \text{KL}((\mu_a, \sigma_a^2), (\lambda_a, \kappa_a^2)) = \min_{a \neq a^*} C(a^*, a; w),$$

where the transportation cost from a to b given an allocation w is defined by

$$\begin{aligned} C(a, b; w) &= \mathbf{1}\{\mu_a > \mu_b\} \inf_{\substack{\lambda_b \geq \lambda_a \\ \kappa_a \geq 0, \kappa_b \geq 0}} \sum_{c \in \{a, b\}} w_c \text{KL}((\mu_c, \sigma_c^2), (\lambda_c, \kappa_c^2)) \\ &= \mathbf{1}\{\mu_a > \mu_b\} \inf_{\lambda \in (\mu_b, \mu_a)} \sum_{c \in \{a, b\}} \frac{w_c}{2} \log \left(1 + \frac{(\mu_c - \lambda)^2}{\sigma_c^2} \right). \end{aligned} \quad (1)$$

From the proof (Appendix C.1) we note that the minimizer in κ is $\kappa_a = \sigma_a^2 + (\mu_a - \lambda)^2$, thus even if we want to identify the arm with largest mean, the closest alternatives have an increased variance. When the variances are known, computing the infimum over the alternative $\lambda \in \Lambda_{\sigma^2}(\mu)$ yields the same expression but with a different transportation cost, which has a convenient closed form:

$$C_{\sigma^2}(a, b; w) = \mathbf{1}\{\mu_a > \mu_b\} \inf_{\lambda \in (\mu_b, \mu_a)} \sum_{c \in \{a, b\}} w_c \frac{(\mu_c - \lambda)^2}{2\sigma_c^2} = \mathbf{1}\{\mu_a > \mu_b\} \frac{1}{2} \frac{(\mu_a - \mu_b)^2}{\sigma_a^2/w_a + \sigma_b^2/w_b}.$$

On the contrary, the infimum in the mean parameter λ in the transportation cost for unknown variance has no simple analytic form (see Appendix C.2 for details on its computation). Still, comparing

the two types of transportation costs (and using properties of the mapping $x \mapsto \log(1+x)/x$) permits to establish a link between $T_{\sigma^2}^*(\mu)$ and $T^*(\mu, \sigma^2)$ (resp. $T_{\sigma^2, \beta}^*(\mu)$ and $T_{\beta}^*(\mu, \sigma^2)$), hence to quantify the impact of not knowing the variances.

Lemma 4 *Let $d(\mu, \sigma^2) = \max_{a \neq a^*(\mu)} \frac{(\mu_{a^*(\mu)} - \mu_a)^2}{\min\{\sigma_a^2, \sigma_{a^*(\mu)}^2\}}$. Then,*

$$1 < \frac{T^*(\mu, \sigma^2)}{T_{\sigma^2}^*(\mu)} \leq \frac{d(\mu, \sigma^2)}{\log(1 + d(\mu, \sigma^2))} \quad \text{and} \quad 1 < \frac{T_{\beta}^*(\mu, \sigma^2)}{T_{\sigma^2, \beta}^*(\mu)} \leq \frac{d(\mu, \sigma^2)}{\log(1 + d(\mu, \sigma^2))}. \quad (2)$$

When $d(\mu, \sigma^2)$ is small, say $d(\mu, \sigma^2) \leq 1$, the two complexities are close since we then have $T_{\sigma^2}^*(\mu)/T^*(\mu, \sigma^2) \in [\log 2, 1)$. Observe that a small $d(\mu, \sigma^2)$ also implies that the BAI problem is hard: if $d(\mu, \sigma^2) \leq c \in \mathbb{R}_+$ then for all $a \in [K]$, $\frac{\min\{\sigma_a^2, \sigma_{a^*(\mu)}^2\}}{(\mu_{a^*(\mu)} - \mu_a)^2} \geq c^{-1}$. Since that ratio is roughly the number of samples needed to distinguish the two arms, the problem is hard when it is large. Still, there exist instances with an arbitrarily large complexity ratio $T^*(\mu, \sigma^2)/T_{\sigma^2}^*(\mu)$ (Lemma 12). We conjecture that they always correspond to easy problems, for which both $T_{\sigma^2}^*(\mu)$ and $T^*(\mu, \sigma^2)$ are small. Lemma 4 is not sufficient to prove this conjecture as there exists hard instances with a large value of $d(\mu, \sigma^2)$ and instances for which the upper bound in (2) is not tight (Appendix J.2.1).

2.3. GLR Stopping Rules

Given any sampling rule, constructing a stopping and recommendation rule for the BAI problem may be viewed as a sequential testing problem with multiple hypotheses $\{\mu_a = \max_{b \in [K]} \mu_b\}$. In one of the first papers on active hypothesis testing (in which the data collection process is further optimized), Chernoff (1959) proposed to rely on Generalized Likelihood Ratio Tests (GLRT) for stopping. This idea was later popularized by Garivier and Kaufmann (2016) for the BAI problem.

For all $a \in [K]$, let $N_{t,a} = \sum_{s \in [t]} \mathbf{1}\{a_s = a\}$, $\mu_{t,a}$ and $\sigma_{t,a}^2$ be the empirical count, mean and variance of arm a after time t , where

$$\mu_{t,a} \stackrel{\text{def}}{=} \frac{1}{N_{t,a}} \sum_{s \in [t]} \mathbf{1}\{a_s = a\} X_{s,a} \quad \text{and} \quad \sigma_{t,a}^2 \stackrel{\text{def}}{=} \frac{1}{N_{t,a}} \sum_{s \in [t]} \mathbf{1}\{a_s = a\} (X_{s,a} - \mu_{t,a})^2.$$

For Gaussian with unknown variances, the GLR to reject $(\mu, \sigma^2) \in \Lambda$, with $\Lambda \subseteq \mathcal{D}$, is written as

$$\text{GLR}_t^{\mathcal{D}}(\Lambda) = \inf_{(\lambda, \kappa^2) \in \Lambda} \sum_{a \in [K]} N_{t,a} \text{KL}((\mu_{t,a}, \sigma_{t,a}^2), (\lambda_a, \kappa_a^2)), \quad (3)$$

which is reminiscent to the expression in the lower bound. We let $\hat{a}_t \stackrel{\text{def}}{=} a^*(\mu_t)$ denote the empirical best arm (EB). Similar calculations as in the proof of Lemma 3 yield that $\text{GLR}_t^{\mathcal{D}}(\Lambda(\mu_t, \sigma_t^2)) = \min_{a \neq \hat{a}_t} Z_a(t)$ where the GLR statistic of arm $a \neq \hat{a}_t$ is defined as

$$Z_a(t) = \text{GLR}_t^{\mathcal{D}}(\{(\lambda, \kappa^2) \mid \lambda_a \geq \lambda_{\hat{a}_t}\}) = \inf_{\lambda \in [\mu_{t,a}, \mu_{t, \hat{a}_t}]} \sum_{b \in \{a, \hat{a}_t\}} \frac{N_{t,b}}{2} \log \left(1 + \frac{(\mu_{t,b} - \lambda)^2}{\sigma_{t,b}^2} \right),$$

which we refer to as the empirical transportation cost between arm \hat{a}_t and arm a .

GLR Stopping Rule In its general form, the GLR stopping rule triggers when $\text{GLR}_t^{\mathcal{D}}(\Lambda(\mu_t, \sigma_t^2))$ exceeds a threshold $c(t, \delta)$. Here we propose to further exploit the structure of the problem and use a family of thresholds $c_{a,b} : \mathbb{N}^K \times (0, 1] \rightarrow \mathbb{R}_+$ for all $(a, b) \in [K]^2$, leading to the stopping rule

$$\tau_\delta \stackrel{\text{def}}{=} \inf \{t \in \mathbb{N} \mid \forall a \neq \hat{a}_t, Z_a(t) > c_{\hat{a}_t, a}(N_t, \delta)\}. \quad (4)$$

EV-GLR Stopping Rule For known variances, we should consider $\text{GLR}^{\mathcal{D}_{\sigma^2}}(\{\lambda : \lambda_a \geq \lambda_{\hat{a}_t}\})$, which can be computed in closed-form and depends on the variance σ^2 . Replacing the variance vector σ^2 by its empirical estimate σ_t^2 yields the Empirical Variance GLR (EV-GLR) statistic

$$Z_a^{\text{EV}}(t) = \inf_{u \in [\mu_{t,a}, \mu_{t, \hat{a}_t}]} \sum_{b \in \{a, \hat{a}_t\}} N_{t,b} \frac{(\mu_{t,b} - u)^2}{2\sigma_{t,b}^2} = \frac{1}{2} \frac{(\mu_{t,a} - \mu_{t, \hat{a}_t})^2}{\sigma_{t,a}^2/N_{t,a} + \sigma_{t, \hat{a}_t}^2/N_{t, \hat{a}_t}}.$$

The EV-GLR stopping rule given a family of thresholds $(c_{a,b})_{(a,b) \in [K]^2}$ is defined as

$$\tau_\delta^{\text{EV}} \stackrel{\text{def}}{=} \inf \{t \in \mathbb{N} \mid \forall a \neq \hat{a}_t, Z_a^{\text{EV}}(t) > c_{\hat{a}_t, a}(N_t, \delta)\}. \quad (5)$$

Given their proximity with the lower bound –see (3)–, GLR stopping rules are good candidates to match T^* . Indeed, it is easy to prove that sampling arms from w^* and using the threshold $c_{a,b}(N, \delta) = \log(1/\delta)$, the lower bound would be matched. However, such a threshold is too good to be δ -correct (Section 3). Moreover, w^* needs to be estimated since it is unknown (Section 4).

3. Calibration of the Stopping Thresholds

We present ways of calibrating the thresholds used by the GLR stopping rule, by leveraging concentration arguments. Under any sampling rule, to obtain a δ -correct GLR stopping rule it suffices to show that the family of thresholds is such that the following time-uniform concentration inequality holds for all $\nu \in \mathcal{D}^K$: with probability $1 - \delta$, for all $t \in \mathbb{N}$ and for all $a \neq a^*(\mu)$,

$$\sum_{b \in \{a, a^*(\mu)\}} \frac{N_{t,b}}{2} \log \left(1 + \frac{(\mu_{t,b} - \mu_b)^2}{\sigma_{t,b}^2} \right) \leq c_{a, a^*(\mu)}(N_t, \delta). \quad (6)$$

Aiming at matching the lower bound, we want to derive a family of thresholds satisfying $c_{a,b}(N, \delta) \sim_{\delta \rightarrow 0} \log(1/\delta)$. As regards the time dependency, generalizations of the law of the iterated logarithm suggest we could achieve $\mathcal{O}(\log \log t)$. Both dependencies are achieved for known variances (Kaufmann and Koolen, 2021), and we are the first to show it for unknown variances (Theorem 6). While simple ideas yield δ -correct thresholds (Section 3.1), obtaining the ideal dependency in δ requires sophisticated concentration arguments (Section 3.2).

Similar arguments can be used to calibrate the thresholds used by the EV-GLR stopping rule (Appendix G). Moreover, δ -correct thresholds for the EV-GLR stopping rule can be obtained by using the ones calibrated for GLR stopping rule, and vice-versa (Lemma 47).

3.1. Simple Ideas

As per-arm concentration results are easier to obtain, we first control each term of the sum in (6).

Student thresholds Since $(\mu_{t,a} - \mu_a)/\sigma_{t,a}$ is an observation of the Student distribution $\mathcal{T}_{N_{t,a}-1}$, a first simple approach involves the quantiles of Student distributions with n degrees of freedom. A direct union bound over time and arms yield a δ -correct family of thresholds (Lemma 49).

Box thresholds As illustrated in Figure 1, the Student threshold suffers from a probably sub-optimal dependence in both $\log(1/\delta)$ and t . This is why we propose an alternative method where the union bound is replaced by time-uniform concentration (which has proved useful to improve both dependencies in different contexts) and the Student concentration by concentration on the mean and the variance separately. The resulting time-uniform upper and lower tail concentration inequalities for the empirical variance (Corollary 26) are of independent interest. Thanks to these “box” confidence regions on (μ_t, σ_t^2) , Lemma 5 yields a δ -correct family of thresholds.

Lemma 5 *Let $\eta_0 > 0$, $s > 1$, ζ be the Riemann ζ function and, for $i \in \{0, -1\}$, $\overline{W}_i(x) = -W_i(-e^{-x})$ for $x \geq 1$ where $(W_i)_{i \in \{0, -1\}}$ are the branches of the Lambert W function. Define*

$$\begin{aligned} \varepsilon_\mu(t, \delta) &= \frac{1}{t} \overline{W}_{-1} \left(1 + 2 \log \left(\frac{4(K-1)\zeta(s)}{\delta} \right) + 2s + 2s \log \left(1 + \frac{\log t}{2s} \right) \right), \\ 1 - \varepsilon_{-, \sigma}(t, \delta) &= \overline{W}_0 \left(1 + \frac{2(1+\eta_0)}{t} \left(\log \left(\frac{4(K-1)\zeta(s)}{\delta} \right) + s \log \left(1 + \log_{1+\eta_0}(t) \right) \right) \right) - \frac{1}{t}. \end{aligned}$$

The family of thresholds $c_{a,b}^{Box}(N_t, \delta)$ with value $+\infty$ if $t < \max_{c \in \{a,b\}} t_c^{Box}(\delta)$ and otherwise

$$c_{a,b}^{Box}(N_t, \delta) = \sum_{c \in \{a,b\}} \frac{N_{t,c}}{2} \log \left(1 + \frac{\varepsilon_\mu(N_{t,c}, \delta)}{1 - \varepsilon_{-, \sigma}(N_{t,c} - 1, \delta)} \right) \quad (7)$$

yields a δ -correct family of thresholds for the GLR stopping rule. The stochastic initial times are

$$t_a^{Box}(\delta) = \inf \left\{ t \mid N_{t,a} > 1 + e^{1+W_0 \left(\frac{2(1+\eta_0)}{e} \left(\log \left(\frac{4(K-1)\zeta(s)}{\delta} \right) + s \log \left(1 + \frac{\log(N_{t,a}-1)}{\log(1+\eta_0)} \right) \right) - e^{-1} \right)} \right\}. \quad (8)$$

To derive the Box threshold, we leverage a lower bound on the empirical variance which is ensured to be strictly positive (hence informative) thanks to the initial time condition (8). As $W_0(x) \in [-1, +\infty)$, it also yields that $N_{t,a} > 2$. Using that $W_0(x) \approx \log(x) - \log \log(x)$ (Appendix I), it is asymptotically equivalent to $\frac{2(1+\eta_0) \log(1/\delta)}{\log \log(1/\delta)}$. Since the lower bound in Lemma 2 suggests that the stopping time is asymptotically equivalent to $T^*(\mu, \sigma^2) \log(1/\delta)$, the condition (8) has a vanishing influence compared to the stopping time. For the parameters used in our simulations (see Section 3.3), (8) is empirically satisfied after sampling each arm 16 times for $\delta = 0.1$ and 20 times for $\delta = 0.001$. Recall that $\overline{W}_{-1}(x) \approx x + \log x$ and $\overline{W}_0(x) \approx e^{-x+e^{-x}}$ (see Appendix I).

3.2. Beyond Box

While being simpler to derive by controlling each arm independently, the above thresholds have a worse δ dependency than more sophisticated approach controlling directly the joint term (6). Since it is challenging to deal with (6), we consider as a proxy the KL divergences for which is easier to construct martingales, which can improve on the δ dependency. To do so, we consider the formulation (1), which removes the minimization step over variances, and apply the arguments used to obtain (6). Under any sampling rule, to obtain a δ -correct GLR stopping rule it suffices to show

that the family of thresholds is such that the following time-uniform concentration inequality holds for all $\nu \in \mathcal{D}^K$: with probability $1 - \delta$, for all $t \in \mathbb{N}$ and for all $a \neq a^*(\mu)$,

$$\sum_{b \in \{a, a^*(\mu)\}} N_{t,b} \text{KL}((\mu_{t,b}, \sigma_{t,b}^2), (\mu_b, \sigma_b^2)) \leq c_{a, a^*(\mu)}(N_t, \delta). \quad (9)$$

KL thresholds First, we derive time-uniform concentration results on the summation of KL divergences (Theorems 42 and 43), which are of independent interest. Then, applied to our setting, it yields a δ -correct family of thresholds (Theorem 6).

Theorem 6 Let $\eta_1 > 0$, $\gamma, s > 1$. Let $\varepsilon_\mu, \varepsilon_{-, \sigma}$ as in Lemma 5 with $\tilde{\delta} = \frac{\delta}{3}$ and $(t_a^{\text{Box}})_a$ as in (8),

$$1 + \varepsilon_{+, \sigma}(t, \delta) = \overline{W}_{-1} \left(1 + \frac{2(1 + \eta_1)}{t} \left(\log \left(\frac{12(K-1)\zeta(s)}{\delta} \right) + s \log(1 + \log_{1+\eta_1}(t)) \right) \right) - \frac{1}{t}.$$

For all t , define $i_{t,a} = \lfloor \log_\gamma N_{t,a} \rfloor$, $n_{t,a} = \gamma^{i_{t,a}}$, $\bar{t}_a = \inf \{t \mid N_{t,a} = n_{t,a}\}$,

$$\begin{aligned} \mu_{\pm, t, a}^2 &= \max_{\pm} \left(\mu_{\bar{t}_a, a} \pm 2\sigma_{\bar{t}_a, a} \sqrt{\frac{\varepsilon_\mu(n_{t,a}, \delta)}{1 - \varepsilon_{-, \sigma}(n_{t,a} - 1, \delta)}} \right)^2, \\ \sigma_{\pm, t, a}^2 &= \sigma_{\bar{t}_a, a}^2 \frac{1 \pm \varepsilon_{\pm, \sigma}(n_{t,a} - 1, \delta)}{1 \mp \varepsilon_{\mp, \sigma}(n_{t,a} - 1, \delta)} \quad \text{and} \quad R_{t,a}(\delta) = \frac{\sigma_{+, t, a}^3 f_+(g(\sigma_{+, t, a}^2, \mu_{+, t, a}^2))}{\sigma_{-, t, a}^3 f_-(g(\sigma_{-, t, a}^2, \mu_{-, t, a}^2))}, \end{aligned}$$

where $f_\pm(x) = \frac{1 \pm \sqrt{1-x}}{\sqrt{x}}$ and $g(x, y) = \frac{2x}{(x+2y+\frac{1}{2})^2}$. The family of thresholds $c_{a,b}^{\text{KL}}(N_t, \delta)$ with value $+\infty$ if $t < \max_{c \in \{a,b\}} \max\{t_c^{\text{Box}}(\delta/3), t_c^m(\delta)\}$ and otherwise

$$c_{a,b}^{\text{KL}}(N_t, \delta) = 4\overline{W}_{-1} \left(1 + \frac{\log \frac{2\zeta(s)^2}{\delta}}{4} + \frac{s}{4} \sum_{c \in \{a,b\}} \log(1 + \log_\gamma N_{t,c}) + \frac{1}{2} \sum_{c \in \{a,b\}} \log(\gamma R_{t,c}(\delta)) \right) \quad (10)$$

yields a δ -correct family of thresholds for the GLR stopping rule. The stochastic initial times are

$$t_a^m(\delta) = \inf \left\{ t \mid N_{t,a} > 1 + \max \left\{ \frac{e^{s/\log\left(\frac{12(K-1)\zeta(s)}{\delta}\right)}}{1 + \eta_0}, \frac{e^{s/\left(\log\left(\frac{12(K-1)\zeta(s)}{\delta}\right) - \frac{1}{2(1+\eta_1)}\right)}}{1 + \eta_1} \right\} \right\}. \quad (11)$$

As $\overline{W}_{-1}(x) \approx x + \log(x)$, Theorem 6 proves that we can obtain δ -correct threshold with the dependencies $c(t, \delta) \sim_{\delta \rightarrow 0} \log(1/\delta)$ and $c(t, \delta) \sim_{t \rightarrow +\infty} C \log \log(t)$, which are widely used in practice for BAI problems. While this dependency was already motivated when the variances are known (Kaufmann and Koolen, 2021), Theorem 6 legitimates its use for unknown variances.

To control the KL divergence between the true parameter and the MLE for Gaussian with unknown variances, our threshold combines two concentration results and is obtained by covering \mathbb{N} with slices of times with geometrically increasing size to cover (referred to as the ‘‘peeling’’ method). First, we use a crude per-arm concentration step to restrict the estimated parameters to a region around the true mean and variance. Then, a second result uses the knowledge of the restriction to get a finer concentration on the weighted sum of KL. It is proved for generic exponential

families by approximating the KL divergence by a quadratic function on this crude confidence region. In (10), $R_{t,a}(\delta)$ represents the cost of this approximation, while $\log(1 + \log_\gamma N_{t,c})$ is the cost of time-uniform. The initial time condition (11) ensures the monotonicity of the preliminary concentration, and it is of the form $N_{t,a} > 1 + c_0(\delta)$ where $c_0(\delta) > 0$. In our simulations (see Section 3.3), (11) is empirically satisfied after sampling each arm twice for all considered δ .

Degegne (2019) derives concentration on the KL divergence of sub-Gaussian d -dimensional exponential families defined on the natural parameter space $\Theta_D = \mathbb{R}^d$. This doesn't include Gaussian with unknown variance, but our proof builds on his method. The main challenge was to tackle $\Theta_D \neq \mathbb{R}^d$, and we solved it by truncation on the sequence of crude confidence regions. In generalized linear bandits, truncated Gaussians were also used to derive tail-inequalities for martingales "re-normalized" by their quadratic variation (Faury, 2021). For general d -dimensional exponential families, Chowdhury et al. (2022) derives concentrations on the KL divergence between the true parameter and a linear combination of the MLE and the true parameter. As we are interested in the KL divergence between the true parameter and the MLE, we cannot leverage their result.

BoB thresholds While the KL thresholds reach the desired dependency in (t, δ) , using (9) instead of (6) yields larger thresholds due to additive constants. To overcome this hurdle, we maximize (6) under the per-arm box constraints (Lemma 5) and the pairwise non-linear constraint (Theorem 6). The resulting family of thresholds is denoted by BoB (Best of Both) thresholds. While the BoB thresholds have no closed-form solution, they can be approximated with non-linear solvers, e.g. Ipopt (Wächter and Biegler, 2006).

Corollary 7 *Let $f(x, y) = (1 + y)x - 1 - \log(x)$ for all $(x, y) \in (\mathbb{R}_+^*)^2$. Let $(t_a^{\text{Box}})_a$ and $(t_a^m)_a$ as in (8,11). Let $\varepsilon_\mu, \varepsilon_{-, \sigma}$ as in Lemma 5 and $(c_{b,a}^{\text{KL}})_{b,a \in [K]}$ as in (10). The family of thresholds $c_{a,b}^{\text{BoB}}(N_t, \delta)$ with value $+\infty$ if $t < \max_{c \in \{a,b\}} \max\{t_c^{\text{Box}}(\delta/6), t_c^m(\delta/2)\}$ and otherwise solution of the optimization problem*

$$\begin{aligned} & \text{maximize } \frac{1}{2} \sum_{c \in \{a,b\}} N_{t,c} \log(1 + y_c) \\ & \text{such that } \forall c \in \{a,b\}, \quad y_c \geq 0, \quad x_c y_c \leq \varepsilon_\mu(N_{t,c}, \delta/2), \quad x_c \geq 1 - \varepsilon_{-, \sigma}(N_{t,c} - 1, \delta/2), \\ & \quad \text{and } \frac{1}{2} \sum_{c \in \{a,b\}} N_{t,c} f(x_c, y_c) \leq c_{b,a}^{\text{KL}}(N_t, \delta/2), \end{aligned}$$

yields a δ -correct family of thresholds for the GLR stopping rule.

Since (6) is smaller than (9), the KL constraint is an upper bound on the BoB threshold. Compared to the box threshold, the maximization underlying the BoB threshold has an additional constraint. Therefore, we have $c_{a,b}^{\text{BoB}}(N_t, \delta) \leq \min\{c_{b,a}^{\text{Box}}(N_t, \delta/2), c_{b,a}^{\text{KL}}(N_t, \delta/2)\}$. In particular, the BoB threshold combines the best of both thresholds in terms of (t, δ) dependencies.

3.3. Simulations

We perform numerical simulations to compare the family of thresholds introduced above for the GLR stopping rule (see Appendix J.2.2 for the EV-GLR stopping rule). Taking $K = 2$, we consider the instance $\mu = (0, -0.2)$ and $\sigma^2 = (1, 0.5)$. Since we are not interested in observing the influence




Figure 1: Thresholds for (4) as a function of (a) $\log(1/\delta)$ for $t = 5000$ and (b) t for $\delta = 0.01$.

of the sampling rule, the stream of data is uniform between both arms. For the thresholds, we set the parameters to $s = 2$, $\gamma = 1.2$ and $\eta_0 = \eta_1 = \log(1/\delta)^{-1}$.

Figure 1 plots the dependency of the thresholds in $\log(1/\delta)$ and t . In Figure 1(a), we are only interested by the slopes, and smaller slopes are equivalent to better dependency in $\log(1/\delta)$. As expected, Student thresholds have poor performance for both variables. While box thresholds improve in t , they suffer from a worse dependency in $\log(1/\delta)$. KL thresholds circumvent this issue with the best dependency in $\log(1/\delta)$ so far. However, they incur a large constant cost making it worse than the box threshold in practice. As hoped, BoB thresholds combine the good performance in t of the box threshold and the asymptotic dependency in $\log(1/\delta)$ of the KL threshold.

The improved theoretical dependency of the BoB threshold comes at the price of a higher computational cost: on average 400, 600 and 800 times larger than the ones for the KL threshold, the Box threshold and the Student threshold respectively. When the computational cost is a major concern, the Box threshold should be used since it has low computational cost and good empirical performance. Alternatively, we could use the BoB threshold and evaluate the stopping rule only on a predefined geometric grid of times. This “lazy” stopping rule is still δ -correct.

4. Sampling Rule Wrappers

After calibrating the stopping threshold to ensure δ -correctness, we need to design a sampling rule which requires few samples before stopping. Given any BAI algorithm for Gaussian with known variances, we propose two wrappers that can adapt the algorithm to tackle unknown variances: plugging in the empirical variance or adapting the transportation cost.

When the variances are unknown, a natural idea is to plug in the empirical variances instead of using the true variances which are now unknown. We can apply this wrapper to any BAI algorithm.

Section 2 discusses the differences and links between the transportation costs for known and unknown variances. Leveraging this interplay, we can adapt a BAI algorithm to use the transportation costs for unknown variances instead of the ones for known variances. We can apply this wrapper to any BAI algorithm relying on transportation costs.

We illustrate how to instantiate each wrapper (Section 4.1), derive guarantees on their asymptotic expected sample complexity (Section 4.2), and assess their empirical performance (Section 4.3).

4.1. Instantiating the Wrappers

As initialization, we start by pulling each arm $n_0 \geq 2$ times, and let $t_0 = n_0 K$.

Track-and-Stop The Track-and-Stop algorithm (Garivier and Kaufmann, 2016) computes at each time $t > t_0$ the optimal allocation for the considered transportation costs, i.e. $w_t = w_{\sigma_t^2}^*(\mu_t)$ for Gaussian with known variances. Given the vector w_t in the simplex, it uses a so-called tracking procedure to obtain an arm a_{t+1} to sample. We describe and use the one called C-tracking by Garivier and Kaufmann (2016). On top of this tracking a forced exploration is used to enforce convergence towards the optimal allocation for the true unknown parameters. Let $\varepsilon \in (0, 1/K]$ and $\Delta_K^\varepsilon = \{w \in [\varepsilon, 1]^K \mid \sum_{a \in [K]} w_a = 1\}$. Defining w_t^ε the L^∞ projection of w_t on Δ_K^ε , C-Tracking pulls $a_{t+1} \in \arg \max_{a \in [K]} \sum_{s=t_0}^t w_{s,a}^\varepsilon - N_{t,a}$.

Plugging in the empirical variance yields the EV-TaS (Empirical Variance Track-and-Stop) algorithm which computes $w_t = w_{\sigma_t^2}^*(\mu_t)$. Adapting the transportation cost yields the TaS algorithm which uses $w_t = w^*(\mu_t, \sigma_t^2)$. Computing $w_{\sigma_t^2}^*(\mu_t)$ and $w^*(\mu_t, \sigma_t^2)$ can be done by solving an equivalent optimization problem with one bounded variable (Theorem 13 in Appendix C.2), which can itself be numerically approximated with binary search.

Top Two algorithm At each time $t > t_0$, the Top Two algorithm β -EB-TCI (Jourdan et al., 2022) pulls the EB leader $B_{t+1}^{\text{EB}} = \hat{a}_t$ with probability β . If B_{t+1}^{EB} is not sampled, then it pulls the TCI challenger $A_{t+1}^{\text{TCI}} \in \arg \min_{a \neq B_{t+1}^{\text{EB}}} C_t(B_{t+1}^{\text{EB}}, a) + \log N_{t,a}$ for the considered transportation costs $C_t(a, b)$, i.e. $C_t(a, b) = \mathbf{1}\{\mu_{t,a} > \mu_{t,b}\} \frac{1}{2} \frac{(\mu_{t,a} - \mu_{t,b})^2}{\sigma_{t,a}^2/N_{t,a} + \sigma_{t,b}^2/N_{t,b}}$ for Gaussian with known variance.

Plugging in the empirical variance yields the β -EB-EVTCI algorithm which uses

$$C_t^{\text{EV}}(a, b) \stackrel{\text{def}}{=} \mathbf{1}\{\mu_{t,a} > \mu_{t,b}\} \frac{1}{2} \frac{(\mu_{t,a} - \mu_{t,b})^2}{\sigma_{t,a}^2/N_{t,a} + \sigma_{t,b}^2/N_{t,b}}.$$

Adapting the transportation cost yields the β -EB-TCI algorithm which computes

$$C_t(a, b) \stackrel{\text{def}}{=} \mathbf{1}\{\mu_{t,a} > \mu_{t,b}\} \inf_{\lambda \in \mathbb{R}} \sum_{c \in \{a, b\}} \frac{N_{t,c}}{2} \log \left(1 + \frac{(\mu_{t,c} - \lambda)^2}{\sigma_{t,c}^2} \right).$$

Since $C_t(\hat{a}_t, a) = Z_a(t)$, we can re-use computations of the GLR stopping rule.

4.2. Sample Complexity Upper Bound

Definition 8 introduces the notion of asymptotically tight family threshold (Jourdan et al., 2022), which corresponds informally to $c_{a,b}(N, \delta) \sim_{\delta \rightarrow 0} \log(1/\delta)$. As hinted in Figure 1(a), the KL and the BoB thresholds are asymptotically tight (Appendix G.5), but not the Student and Box thresholds.

Definition 8 A family of thresholds $(c_{a,b})_{(a,b) \in [K]^2}$ is said to be asymptotically tight if there exists $\alpha \in [0, 1]$, $\delta_0 \in (0, 1]$, functions $f, \bar{T} : (0, 1] \rightarrow \mathbb{R}_+$ and C independent of δ satisfying: (1) for all $(a, b) \in [K]^2$, $\delta \in (0, \delta_0]$ and $N \in \mathbb{N}^K$ such that $\|N\|_1 \geq \bar{T}(\delta)$, then $c_{a,b}(N, \delta) \leq f(\delta) + C\|N\|_1^\alpha$, (2) $\limsup_{\delta \rightarrow 0} f(\delta)/\log(1/\delta) \leq 1$ and $\limsup_{\delta \rightarrow 0} \bar{T}(\delta)/\log(1/\delta) = 0$.

When combined with the GLR stopping rule using the KL or the BoB thresholds, Theorem 9 shows that TaS (resp. β -EB-TCI) is a δ -correct and asymptotically (resp. β -)optimal algorithm.

Theorem 9 *Using the GLR stopping rule with an asymptotically tight family of thresholds, TaS (resp. β -EB-TCI with $n_0 \geq 4$) satisfies that, for all ν with $|a^*(\mu)| = 1$ (resp. $\min_{a \neq b} |\mu_a - \mu_b| > 0$),*

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\nu [\tau_\delta]}{\log(1/\delta)} \leq T^*(\mu, \sigma^2) \quad (\text{resp. } T_\beta^*(\mu, \sigma^2)).$$

In Appendix H, we derive a similar result involving $T_{\sigma^2}^*(\mu)$ (resp. $T_{\sigma^2, \beta}^*(\mu)$) for EV-TaS (resp. β -EB-EVTCI with $n_0 \geq 6$) combined with the EV-GLR stopping rule using an asymptotically tight threshold. However, since $T_{\sigma^2}^*(\mu) < T^*(\mu, \sigma^2)$ and $T_{\sigma^2, \beta}^*(\mu) < T_\beta^*(\mu, \sigma^2)$, neither of these algorithms can be δ -correct. Otherwise it would yield a contradiction with the lower bound in Lemma 2. Moreover, as there exist instances for which the ratios $T^*(\mu, \sigma^2)/T_{\sigma^2}^*(\mu)$ and $T_\beta^*(\mu, \sigma^2)/T_{\sigma^2, \beta}^*(\mu)$ are arbitrarily large (Lemma 12), multiplying the thresholds by a problem independent constant is not sufficient either to obtain δ -correctness, as expressed in Theorem 10.

Theorem 10 *There exists a sampling rule such that: for all asymptotically tight family of thresholds $(c_{a,b})_{(a,b) \in [K]^2}$ and problem independent constant $\alpha_0 > 0$, combining this sampling rule with the EV-GLR stopping rule using $(\alpha_0 c_{a,b})_{(a,b) \in [K]^2}$ yields an algorithm which is not δ -correct.*

Inspired by Section 3, we propose families of thresholds (EV-Student, EV-Box and EV-BoB) which are δ -correct for the EV-GLR stopping rule (see Appendix G) but are not asymptotically tight (Theorem 10). Still, in our experiments the empirical proportion of error is lower than δ even when using a heuristic, asymptotically tight threshold.

Based on Theorems 9 and 10, algorithms obtained by adapting the transportation costs enjoy stronger theoretical guarantees than the ones plugging in the empirical variance.

4.3. Experiments

We compare the empirical performance of the two wrappers for different BAI algorithms in the moderate regime ($\delta = 0.01$). As benchmarks, we consider FHN2 (procedure 2 in Fan et al. (2016), see Algorithm 1 in Appendix J), uniform sampling and “fixed” sampling which is an oracle playing with proportions $w^*(\mu, \sigma^2)$. FHN2 is an elimination strategy which repeatedly samples all arms until only one arm is left. Its elimination mechanism is calibrated by resorting to continuous-time approximations. Therefore, FHN2 is only asymptotically δ -correct and has no guaranties on the sample complexity. Based on Degenne et al. (2019) and Wang et al. (2021), plugging in the empirical variance yields EV-DKM and EV-FWS, while DKM and FWS refers to the algorithms using the transportation costs for unknown variances. Even though those instances are not analyzed, we believe that similar guarantees on the sample complexity can be shown.

Algorithms obtained by plugging in the empirical variance uses the EV-GLR stopping rule, while the GLR stopping rule is used by the ones with adapted transportation cost and the uniform sampling. We consider the stylized stopping threshold $c(t, \delta) = \log((1 + \log t)/\delta)$, which was proposed in Garivier and Kaufmann (2016). While it doesn’t ensure δ -correctness of the stopping threshold, it is asymptotically tight and yields an empirical error which is several order of magnitude lower than δ . Top Two algorithms use $\beta = 0.5$.

We assess the performance on 1000 random instances with $K = 10$ such that $(\mu_1, \sigma_1^2) = (0, 1)$. For $a \neq 1$, we set $(\mu_a, \sigma_a^2) = (-\Delta_a, r_a)$ where $\Delta_a \sim \mathcal{U}([0.2, 1.0])$ and $r_a \sim \mathcal{U}([0.1, 10])$. To illustrate the two regimes for $T^*(\mu, \sigma^2)/T_{\sigma^2}^*(\mu)$, we consider a *standard* instance ($T^*(\mu, \sigma^2)/T_{\sigma^2}^*(\mu) \approx 1.015$) and an *easy* instance ($T^*(\mu, \sigma^2)/T_{\sigma^2}^*(\mu) \approx 1.384$). We average over 5000 runs.

In Figure 2, we observe that algorithms obtained by plugging in the empirical variance yield similar result as the ones using the adapted transportation cost, and slightly better performance on the easy instance. Moreover, those wrapped BAI algorithms outperforms uniform sampling and are on par with “fixed” sampling. On random instances FHN2 has similar performance to the wrapped BAI algorithms, but it wastes precious samples on easy instances.




Figure 2: Empirical stopping time on Gaussian (top) random instances with $K = 10$, (left) standard instance $(\mu, \sigma^2) = ((1.0, 0.85, 0.8, 0.7, 0.65), (1.0, 0.6, 0.5, 0.4, 0.35))$ and (right) easy instance $(\mu, \sigma^2) = ((1.0, 0.2, 0.15, 0.1, 0.05), (1.0, 0.05, 0.05, 0.05, 0.05))$. Lower bound is $T^*(\mu) \log(1/\delta)$.

5. Conclusion

In this paper we provided two approaches to deal with unknown variances, either by plugging in the empirical variance or by adapting the transportation costs. New time-uniform concentration results were derived to calibrate our two stopping rules. Then, we showed theoretical guarantees and competitive empirical performance of our two sampling rule wrappers on two existing algorithms.

While the literature abounds with designs of sampling rule, the optimal calibration of stopping rules is a most pressing issue as it leads to lower empirical stopping time. While calibrated thresholds have been derived with (near) optimal dependency in δ , those thresholds are known to be too conservative in the moderate confidence regime where their empirical error rate is orders of magnitude lower than δ .

Finally, in the fixed-budget setting, characterizing the impact of not knowing the variances on the probability of misidentifying the best-arm is still an open problem. While similar approaches

might be used to deal with the unknown variances, the resulting algorithms might not enjoy similar theoretical guarantees and empirical performance.

Acknowledgments

Experiments presented in this paper were carried out using the Grid’5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr>). This work has been partially supported by the THIA ANR program “AI-PhD@Lille”. The authors acknowledge the funding of the French National Research Agency under the projects BOLD (ANR-19-CE23-0026-04) and FATE (ANR-22-CE23-0016-01).

References

- Shubhada Agrawal, Sandeep Juneja, and Peter W. Glynn. Optimal δ -correct best-arm selection for heavy-tailed distributions. In *Algorithmic Learning Theory (ALT)*, 2020.
- Jean-Yves Audibert, Sébastien Bubeck, and Rémi Munos. Best arm identification in multi-armed bandits. In *Conference on Learning Theory (COLT)*, 2010.
- Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5:1–122, 2012.
- Hock Peng Chan. The multi-armed bandit problem: An efficient nonparametric solution. *The Annals of Statistics*, 48(1):346–373, 2020.
- Herman Chernoff. Sequential Design of Experiments. *The Annals of Mathematical Statistics*, 30: 755–770, 1959.
- Sayak Ray Chowdhury, Patrick Saux, Odalric-Ambrym Maillard, and Aditya Gopalan. Bregman deviations of generic exponential families. *arXiv preprint arXiv:2201.07306*, 2022.
- Rémy Degenne. *Impact of structure on the design and analysis of bandit algorithms*. PhD thesis, Université de Paris, 2019.
- Rémy Degenne and Wouter M. Koolen. Pure exploration with multiple correct answers. In *Advances in Neural Information Processing Systems*, 2019.
- Rémy Degenne, Wouter M. Koolen, and Pierre Ménard. Non-Asymptotic Pure Exploration by Solving Games. In *Advances in Neural Information Processing Systems*, 2019.
- Amir Dembo and Ofer Zeitouni. *Large Deviations Techniques and Applications*. Applications of mathematics. Springer, 1998.
- Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action Elimination and Stopping Conditions for the Multi-Armed Bandit and Reinforcement Learning Problems. *Journal of Machine Learning Research*, 7:1079–1105, 2006.
- Marco Faella, Alberto Finzi, and Luigi Sauro. Rapidly finding the best arm using variance. In *European Conference on Artificial Intelligence*, 2020.

- Weiwei Fan, L. Jeff Hong, and Barry L. Nelson. Indifference-zone-free selection of the best. *Operations Research*, 64:1499–1514, 2016.
- Louis Faury. *Variance-sensitive confidence intervals for parametric and offline bandits*. PhD thesis, Institut Polytechnique de Paris, 2021.
- Victor Gabillon, Mohammad Ghavamzadeh, and Alessandro Lazaric. Best Arm Identification: A Unified Approach to Fixed Budget and Fixed Confidence. In *Advances in Neural Information Processing Systems*, 2012.
- Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *Conference on Learning Theory*, 2016.
- Junya Honda and Akimichi Takemura. Optimality of thompson sampling for gaussian bandits depends on priors. In *Artificial Intelligence and Statistics*, 2014.
- L. Jeff Hong, Weiwei Fan, and Jun Luo. Review on ranking and selection: A new perspective. *Frontiers of Engineering Management*, 8:321–343, 2021.
- Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 17:257–317, 2020.
- Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49:1055 – 1080, 2021.
- Kevin Jamieson, Matthew Malloy, Robert Nowak, and Sébastien Bubeck. lil’UCB: an Optimal Exploration Algorithm for Multi-Armed Bandits. In *Conference on Learning Theory*, 2014.
- Marc Jourdan, Rémy Degenne, Dorian Baudry, Rianne De Heide, and Emilie Kaufmann. Top two algorithms revisited. *Advances in Neural Information Processing Systems*, 2022.
- Shivaram Kalyanakrishnan, Ambuj Tewari, Peter Auer, and Peter Stone. PAC subset selection in stochastic multi-armed bandits. In *International Conference on Machine Learning*, 2012.
- Emilie Kaufmann and Wouter M. Koolen. Mixture martingales revisited with applications to sequential tests and confidence intervals. *Journal of Machine Learning Research*, 22, 2021.
- Seong Hee Kim and Barry L. Nelson. A fully sequential procedure for indifference-zone selection in simulation. *ACM Transactions on Modeling and Computer Simulation*, 11:251–273, 2001.
- Pinyan Lu, Chao Tao, and Xiaojin Zhang. Variance-dependent best arm identification. In *Uncertainty in Artificial Intelligence*, 2021.
- Andreas Maurer and Massimiliano Pontil. Empirical bernstein bounds and sample-variance penalization. In *Conference on Learning Theory*, 2009.
- Chao Qin, Diego Klabjan, and Daniel Russo. Improving the expected improvement algorithm. In *Advances in Neural Information Processing Systems*, 2017.
- Daniel Russo. Simple Bayesian algorithms for best arm identification. In *Proceedings of the 29th Conference on Learning Theory (COLT)*, 2016.

Xuedong Shang, Rianne de Heide, Emilie Kaufmann, Pierre Ménard, and Michal Valko. Fixed-confidence guarantees for bayesian best-arm identification. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.

Andreas Wächter and Lorenz T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106:25–57, 2006.

Po-An Wang, Ruo-Chun Tzeng, and Alexandre Proutiere. Fast pure exploration via frank-wolfe. *Advances in Neural Information Processing Systems*, 2021.

Appendix A. Outline

- Notation are summarized in Appendix B.
- In Appendix C, we study the characteristic times $T_{\sigma^2}^*(\mu)$, $T^*(\mu, \sigma^2)$ and how to compute the optimal allocation oracles $w_{\sigma^2}^*(\mu)$, $w^*(\mu, \sigma^2)$.
- Properties on the GLR and EV-GLR statistics are detailed in Appendix D.
- In Appendix E, we prove time-uniform and fixed time upper and lower tail concentrations for sub-exponential process and the empirical mean and variance of Gaussian observations.
- For d -dimensional exponential families, a time-uniform concentration inequality on the sum of KL is shown in Appendix F.
- In Appendix G, we show how to calibrate the GLR and EV-GLR stopping rules and derive several family of thresholds.
- Asymptotic guarantees on the expected sample complexity (Theorem 9) and the impossibility results (Theorem 10) are proved in Appendix H.
- The functions $(\overline{W}_i)_{i \in \{-1, 0\}}$ based on Lambert’s branches and their properties are presented in Appendix I.
- Implementation details and additional experiments are available in Appendix J.

Appendix B. Notation

We recall some commonly used notation: the set of integers $[K] \stackrel{\text{def}}{=} \{1, \dots, K\}$, the complement X^c of a set X , the $(K - 1)$ -dimensional probability simplex $\triangle_K \stackrel{\text{def}}{=} \left\{ w \in \mathbb{R}_+^K \mid \sum_{a \in [K]} w_a = 1 \right\}$, the Gaussian distribution $\mathcal{N}(x, v)$ with mean x and variance v , the Kullback-Leibler (KL) divergence $\text{KL}((x_1, \sigma_1^2), (x_2, \sigma_2^2))$ between two distributions $\mathcal{N}(x_1, \sigma_1^2)$ and $\mathcal{N}(x_2, \sigma_2^2)$, Landau’s notation o and \mathcal{O} and the two main branches W_{-1} (negative) and W_0 (positive) of the Lambert W function which is implicitly defined by the equation $W(x)e^{W(x)} = x$. Problem-specific notation are grouped in Table 1. Table 2 gathers notation for the algorithms.

Table 1: Notation for the setting.

Notation	Type	Description
K	\mathbb{N}	Number of arms
\mathcal{D}		Set of Gaussian distributions
\mathcal{D}_v		Set of Gaussian distributions with variance v
ν	\mathcal{D}^K	Vector of distributions, $\nu = (\nu_a)_{a \in [K]}$
μ	\mathbb{R}^K	Vector of means, $\mu = (\mu_a)_{a \in [K]}$
σ^2	$(\mathbb{R}_+^*)^K$	Vector of variances, $\sigma^2 = (\sigma_a^2)_{a \in [K]}$
ν_{μ, σ^2}	\mathcal{D}^K	Problem in which each arm has distribution $\mathcal{N}(\mu_a, \sigma_a^2)$
θ_a	$\mathbb{R} \times \mathbb{R}_-^*$	Natural parameter of the distribution of arm $a \in [K]$
θ	$(\mathbb{R} \times \mathbb{R}_-^*)^K$	Vector of natural parameters, $\theta = (\theta_a)_{a \in [K]}$
F	$\mathbb{R} \rightarrow \mathbb{R} \times \mathbb{R}_+$	Sufficient statistic
ϕ	$(\mathbb{R} \times \mathbb{R}_-^*) \rightarrow \mathbb{R}$	Log-partition function
$T_{\sigma^2}^*(\mu), w_{\sigma^2}^*(\mu)$		Characteristic time and optimal allocation, known σ^2
$T_{\sigma^2, \beta}^*(\mu), w_{\sigma^2, \beta}^*(\mu)$		Characteristic time and β -optimal allocation, known σ^2
$T^*(\mu, \sigma^2), w^*(\mu, \sigma^2)$		Characteristic time and optimal allocation
$T_\beta^*(\mu, \sigma^2), w_\beta^*(\mu, \sigma^2)$		Characteristic time and β -optimal allocation

Table 2: Notation for the algorithms.

Notation	Type	Description
δ	$(0, 1)$	Confidence parameter
a_t	$[K]$	Arm sampled at time t
X_{t, a_t}	\mathbb{R}	Observation at time t , $X_{t, a_t} \sim \nu_{a_t}$
\mathcal{F}_t		History up to time t , $\sigma(a_1, X_{1, a_1}, \dots, a_t, X_{t, a_t})$
$N_{t, a}$	\mathbb{N}	Empirical count, $N_{t, a} = \sum_{s=1}^t \mathbf{1}\{a_s = a\}$
$\mu_{t, a}$	\mathbb{R}	Empirical mean, $\mu_{t, a} = \frac{1}{N_{t, a}} \sum_{s=1}^t X_{s, a_s} \mathbf{1}\{a_s = a\}$
$\sigma_{t, a}^2$	\mathbb{R}_+^*	Empirical variance, $\sigma_{t, a}^2 = \frac{1}{N_{t, a}} \sum_{s=1}^t (X_{s, a_s} - \hat{\mu}_{t, a})^2 \mathbf{1}\{a_s = a\}$
$\tau_\delta, \tau_\delta^{\text{EV}}$	\mathbb{N}	Stopping times (sample complexity)
\hat{a}_t	$[K]$	Candidate arm at time t , $\hat{a}_t \in \arg \max_{a \in [K]} \mu_{t, a}$
$Z_a(t), Z_a^{\text{EV}}(t)$	\mathbb{R}_+^*	GLR and EV-GLR statistic of arm a at time t
$c_{a, b}(N_t, \delta)$	\mathbb{R}	Stopping threshold at time t for the arm pair (a, b)
$F_{t, a}$	$\mathbb{R} \times \mathbb{R}_+$	Averaged statistic, $\frac{1}{N_{t, a}} \sum_{s=1}^t F(X_{s, a_s}) \mathbf{1}\{a_s = a\}$
$\theta_{t, a}$	$\mathbb{R} \times \mathbb{R}_-^*$	Maximum likelihood estimator of θ_a , $\nabla \phi^{-1}(F_{t, a})$ for $N_{t, a} \geq 2$

Appendix C. Characteristic Times

In Appendix C.1, we show properties on the characteristic times $T_{\sigma^2}^*(\mu)$ and $T^*(\mu, \sigma^2)$, i.e. explicit formulas and relationships. In Appendix C.2, we show those characteristic times can be obtained as solution of a simpler optimization problem, allowing to compute the associated optimal allocation.

Theorem 1 in [Garivier and Kaufmann \(2016\)](#) yields the first part of Lemma 2 (Gaussian with known variances). Its second part (Gaussian with unknown variances) is a direct consequence of the arguments in [Garivier and Kaufmann \(2016\)](#), hence we omit the proof.

C.1. Explicit Formulas and Inequalities

We can obtain slightly more explicit formulas for the characteristic times $T_{\sigma^2}^*(\mu)$ and $T^*(\mu, \sigma^2)$ (Lemma 11). It is direct to see that similar explicit formulas can be shown for $T_{\sigma^2, \beta}^*(\mu)$ and $T_{\beta}^*(\mu, \sigma^2)$.

Lemma 11 *Let $\mathcal{M} = \mathbb{R}^K \times (\mathbb{R}_+^*)^K$ and $a^*(\mu) = a^*$. Then,*

$$T_{\sigma^2}^*(\mu)^{-1} = \max_{w \in \Delta_K} \min_{a \neq a^*} \inf_{y \in [\mu_a, \mu_{a^*}]} \sum_{b \in \{a, a^*\}} w_b \frac{(y - \mu_b)^2}{2\sigma_b^2} = \max_{w \in \Delta_K} \min_{a \neq a^*} \frac{(\mu_a - \mu_{a^*})^2}{2 \left(\frac{\sigma_a^2}{w_a} + \frac{\sigma_{a^*}^2}{w_{a^*}} \right)},$$

$$T^*(\mu, \sigma^2)^{-1} = \max_{w \in \Delta_K} \min_{a \neq a^*} \inf_{y \in [\mu_a, \mu_{a^*}]} \sum_{b \in \{a, a^*\}} \frac{w_b}{2} \log \left(1 + \frac{(y - \mu_b)^2}{\sigma_b^2} \right).$$

Proof While an explicit formula for $T_{\sigma^2}^*(\mu)$ was already proven in [Garivier and Kaufmann \(2016\)](#), we derive the equivalent result for $T^*(\mu, \sigma^2)$. For Gaussian with different variances, the KL has the following expression

$$\text{KL}((\mu, \sigma^2), (\lambda, \kappa^2)) = \frac{1}{2} \left(\frac{(\mu - \lambda)^2}{\kappa^2} + \frac{\sigma^2}{\kappa^2} - 1 - \log \left(\frac{\sigma^2}{\kappa^2} \right) \right).$$

Direct computations yield that

$$\inf_{\kappa^2 > 0} \text{KL}((\mu, \sigma^2), (\lambda, \kappa^2)) = \text{KL}((\mu, \sigma^2), (\lambda, \sigma^2 + (\mu - \lambda)^2)) = \frac{1}{2} \log \left(1 + \frac{(\mu - \lambda)^2}{\sigma^2} \right)$$

The set $\Lambda(\mu, \sigma^2) = \{(\lambda, \kappa^2) \in \mathbb{R}^K \times (\mathbb{R}_+^*)^K \mid a^* \notin \arg \max_{a \in [K]} \lambda^a\}$ can be rewritten as $\bigcup_{a \neq a^*} \{\lambda \in \mathbb{R}^K \mid \lambda^a > \lambda^{a^*}\} \times (\mathbb{R}_+^*)^K$. Therefore, taking $\lambda_b = \mu_b$ and $\kappa_b^2 = \sigma_b^2$ for $b \notin \{a, a^*\}$, we obtain

$$\begin{aligned} T^*(\mu, \sigma^2)^{-1} &= \max_{w \in \Delta_K} \min_{a \neq a^*} \inf_{(\lambda, \kappa^2) \in \mathbb{R}^2 \times (\mathbb{R}_+^*)^2: \lambda^a > \lambda^{a^*}} \sum_{b \in \{a, a^*\}} w_b \text{KL}((\mu_b, \sigma_b^2), (\lambda_b, \kappa_b^2)) \\ &= \max_{w \in \Delta_K} \min_{a \neq a^*} \inf_{\lambda \in \mathbb{R}^2: \lambda^a > \lambda^{a^*}} \sum_{b \in \{a, a^*\}} \frac{w_b}{2} \log \left(1 + \frac{(\mu_b - \lambda_b)^2}{\sigma_b^2} \right) \\ &= \max_{w \in \Delta_K} \min_{a \neq a^*} \inf_{\lambda \in [\mu_a, \mu_{a^*}]} \sum_{b \in \{a, a^*\}} \frac{w_b}{2} \log \left(1 + \frac{(\mu_b - \lambda)^2}{\sigma_b^2} \right) \end{aligned}$$

where the last equality is obtained since at the infimum there is equality, i.e. $\lambda^a = \lambda^{a^*}$, and Lemma 14. ■

Using Lemma 11 and algebraic manipulation, we derive inequalities between $T_{\sigma^2}^*(\mu)$ and $T^*(\mu, \sigma^2)$ (Lemma 4). It is direct to see that the same proof would hold for $T_{\sigma^2, \beta}^*(\mu)$ and $T_{\beta}^*(\mu, \sigma^2)$.

Proof For all $a \neq a^*$, $\lambda \in [\mu_a, \mu_{a^*}]$ and $\{a, a^*\}$, we have

$$\frac{(\mu_b - \lambda)^2}{\sigma_b^2} \leq d(\mu, \sigma^2) \stackrel{\text{def}}{=} \max_{a \neq a^*} \frac{(\mu_{a^*} - \mu_a)^2}{\min\{\sigma_{a^*}^2, \sigma_a^2\}}$$

Using that $x \mapsto \frac{\log(1+x)}{x}$ is decreasing for $x \geq 0$, we obtain, for all $\lambda \in [\mu_a, \mu_{a^*}]$ and $b \in \{a, a^*\}$,

$$\log\left(1 + \frac{(\mu_b - \lambda)^2}{\sigma_b^2}\right) \geq \frac{(\mu_b - \lambda)^2 \log(1 + d(\mu, \sigma^2))}{\sigma_b^2 d(\mu, \sigma^2)}.$$

Since $T_{\sigma^2}^*(\mu)^{-1} = \max_{w \in \Delta_K} \min_{a \neq a^*} \inf_{\lambda \in [\mu_a, \mu_{a^*}]} \sum_{b \in \{a, a^*\}} w_b \frac{(\mu_b - \lambda)^2}{2\sigma_b^2}$, this yields

$$T^*(\mu, \sigma^2)^{-1} \geq \frac{\log(1 + d(\mu, \sigma^2))}{d(\mu, \sigma^2)} T_{\sigma^2}^*(\mu)^{-1}$$

The same arguments as above also yield the formulation

$$2T^*(\mu, \sigma^2)^{-1} = \max_{w \in \Delta_K} \inf_{\lambda \in \Lambda_{\sigma^2}(\mu)} \sum_{b \in [K]} w_b \log\left(1 + \frac{(\mu_b - \lambda_b)^2}{\sigma_b^2}\right),$$

and a similar one for $T_{\sigma^2}^*(\mu)$. Since $x \mapsto \log(1+x)$ is concave, we obtain that

$$2T^*(\mu, \sigma^2)^{-1} \leq \log(1 + 2T_{\sigma^2}^*(\mu)^{-1}) < 2T_{\sigma^2}^*(\mu)^{-1}$$

where the strict inequality uses that $T_{\sigma^2}^*(\mu)^{-1} > 0$. ■

Using Lemma 11, Lemma 12 shows that $T^*(\mu, \sigma^2)/T_{\sigma^2}^*(\mu)$ can become arbitrarily large by taking instances with large gaps.

Lemma 12 *For all $K \geq 2$, there exists a sequence of instances $(\nu_n)_{n \in \mathbb{N}}$ with $|a^*(\nu_n)| = 1$ such that $\lim_{n \rightarrow +\infty} T^*(\mu_n, \sigma_n^2)/T_{\sigma_n^2}^*(\mu_n) = +\infty$.*

Proof Let $K \geq 2$. We construct explicitly the sequence instances. Let $(\Delta_n)_{n \in \mathbb{N}}$ such that $\Delta_n > 0$ for all $n \in \mathbb{N}$ and $\lim_{n \rightarrow +\infty} \Delta_n = +\infty$. For all $n \in \mathbb{N}$, we define $\nu_n = \nu_{\mu_n, \sigma_n^2}$ where $(\mu_{n,1}, \sigma_{n,1}^2) = (0, 1)$ and $(\mu_{n,a}, \sigma_{n,a}^2) = (-\Delta_n, 1)$ for all $a \neq 1$.

Using Lemma 11, we obtain

$$\frac{T_{\sigma_n^2}^*(\mu_n) \Delta_n^2}{2} = \min_{\beta \in (0,1)} \min_{w \in \Delta_{K-1}} \max_{a \neq 1} \frac{1}{(1-\beta)w_a} + \frac{1}{\beta} = \min_{\beta \in (0,1)} \frac{K-1}{1-\beta} + \frac{1}{\beta} = (1 + \sqrt{K-1})^2,$$

and, taking $y = -\Delta_n/2$, we obtain

$$\begin{aligned} 2T^*(\mu_n, \sigma_n^2)^{-1} &= \max_{\beta \in (0,1)} \max_{w \in \Delta_{K-1}} \min_{a \neq 1} \inf_{y \in [-\Delta_n, 0]} \{\beta \log(1+y^2) + (1-\beta)w_a \log(1+(\Delta_n+y)^2)\} \\ &\leq \log(1 + \Delta_n^2/4) \max_{\beta \in (0,1)} \left\{ \beta + (1-\beta) \max_{w \in \Delta_{K-1}} \min_{a \neq 1} w_a \right\} \\ &= \log(1 + \Delta_n^2/4) \max_{\beta \in (0,1)} \left\{ \beta + \frac{1-\beta}{K-1} \right\} \leq \log(1 + \Delta_n^2/4). \end{aligned}$$

Therefore, we have

$$\frac{T^*(\mu_n, \sigma_n^2)}{T_{\sigma_n^2}^*(\mu_n)} \geq \frac{\Delta_n^2}{(1 + \sqrt{K-1})^2 \log(1 + \Delta_n^2/4)} \xrightarrow{n \rightarrow +\infty} +\infty,$$

since $\lim_{n \rightarrow +\infty} \Delta_n = +\infty$ and $\log(1+x) =_{x \rightarrow +\infty} o(x)$. \blacksquare

It is slightly more technical to construct a sequence of instances $(\nu_n)_{n \in \mathbb{N}}$ with $\min_{a \neq b} |\mu_{a,n} - \mu_{b,n}| > 0$ such that $\lim_{n \rightarrow +\infty} T_{\beta}^*(\mu_n, \sigma_n^2)/T_{\sigma_n^2, \beta}^*(\mu_n) = +\infty$. Since the means have to be distinct, it is not possible to use directly closed-form formulas. However, the same type of construction with increasing gaps will allow to conclude the proof. The main argument behind both proofs is that $\log(1+x) =_{x \rightarrow +\infty} o(x)$.

C.2. Optimal Allocation Oracles

In [Garivier and Kaufmann \(2016\)](#), they show that $w_{\sigma^2}^*(\mu)$ can be computed as the solution of an optimization problem with one bounded variable. By using similar arguments, [Theorem 13](#) gives a similar equivalent optimization problem for $w^*(\mu, \sigma^2)$. Without loss of generality we consider $a^*(\mu) = 1$ to ease the notations in the following arguments.

Theorem 13 *For all $a \in [K]$ and $\lambda \in \mathbb{R}$, let $d_a(\lambda) = \log\left(1 + \frac{(\lambda - \mu_a)^2}{\sigma_a^2}\right)$. For all $a \neq 1$ and $x \in [0, +\infty)$, let $g_a(x) = \min_{\lambda \in \mathbb{R}} (d_1(\lambda) + x d_a(\lambda))$, with $\lambda_a(x)$ being the minimizer realizing $g_a(x)$. The functions g_a are strictly increasing function with values on $[0, d_1(\mu_a))$ with inverse function $x_a(y) = g^{-1}(y)$. Then, we have $w(\mu, \sigma^2)_a = x_a(y^*) w(\mu, \sigma^2)_1$ for all $a \neq 1$,*

$$2T^*(\mu, \sigma^2)^{-1} = \frac{y^*}{1 + \sum_{a \neq 1} x_a(y^*)} \quad \text{and} \quad w(\mu, \sigma^2)_1 = \frac{1}{1 + \sum_{a \neq 1} x_a(y^*)}$$

Moreover, $y^* \in [0, \min_{a \neq 1} d_1(\mu_a))$ is a solution of the equation $F(y) = 1$ where

$$F(y) = \sum_{a=2}^K \frac{d_1(\lambda_a(x_a(y)))}{d_a(\lambda_a(x_a(y)))}$$

is an increasing function such that $F(0) = 0$ and $\lim_{+\infty} F(y) = +\infty$.

Proof When $w_1 = 0$, we have $\min_{a \neq 1} (w_1 d_1(\lambda) + w_a d_a(\lambda)) = 0$, hence this is not the maximum of a positive quantity, i.e. for all $w \in w^*(\mu, \sigma^2)$, $w_1 > 0$. Therefore, by dividing by w_1 and using that $g_a(x) = \inf_{\lambda \in \mathbb{R}} (d_1(\lambda) + x d_a(\lambda))$, we obtain directly

$$2T(\mu, \sigma^2)^{-1} = \max_{w \in \Delta_K} \min_{a \neq 1} \inf_{\lambda \in \mathbb{R}} (w_1 d_1(\lambda) + w_a d_a(\lambda)) = \max_{w \in \Delta_K} w_1 \min_{a \neq 1} g_a\left(\frac{w_a}{w_1}\right).$$

Let $w^* \in w^*(\mu, \sigma^2)$. Then, using the above result, we obtain

$$w^* \in \arg \max_{w \in \Delta_K} w_1 \min_{a \neq 1} g_a\left(\frac{w_a}{w_1}\right)$$

Introducing $x_a^* = \frac{w_a^*}{w_1^*}$ for all $a \neq 1$, using that $\sum_{a \in [K]} w_a^* = 1$, one has

$$w_1^* = \frac{1}{1 + \sum_{a=2}^K x_a^*} \quad \text{and, for } a \geq 2, w_a^* = \frac{x_a^*}{1 + \sum_{a=2}^K x_a^*}.$$

Moreover, $\{x_a^*\}_{a=2}^K \in \mathbb{R}^{K-1}$ belongs to

$$\arg \max_{\{x_a\}_{a=2}^K \in \mathbb{R}^{K-1}} \frac{\min_{a \neq 1} g_a(x_a)}{1 + \sum_{a=2}^K x_a} \quad (12)$$

Let $\mathcal{O} = \{b \in [K] \setminus \{1\} \mid g_b(x_b^*) = \min_{a \neq 1} g_a(x_a^*)\}$ and $\mathcal{A} = [K] \setminus (\{1\} \cup \mathcal{O})$. Let's show that all the $g_a(x_a^*)$ have to be equal. Assume that $\mathcal{A} \neq \emptyset$. For all $a \in \mathcal{A}$ and $b \in \mathcal{O}$, one has $g_a(x_a^*) > g_b(x_b^*)$. Using the continuity of the g_a functions and the fact that they are strictly increasing (Lemma 14), there exists $\varepsilon > 0$ such that

$$\forall a \in \mathcal{A}, b \in \mathcal{O}, \quad g_a(x_a^* - \varepsilon/|\mathcal{A}|) > g_b(x_b^* + \varepsilon/|\mathcal{O}|) > g_b(x_b^*).$$

We introduce $\bar{x}_a = x_a^* - \varepsilon/|\mathcal{A}|$ for all $a \in \mathcal{A}$ and $\bar{x}_b = x_b^* + \varepsilon/|\mathcal{O}|$ for all $b \in \mathcal{O}$, hence $\sum_{a=2}^K \bar{x}_a = \sum_{a=2}^K x_a^*$. There exists $b \in \mathcal{O}$ such that $\min_{a \neq 1} g_a(\bar{x}_a) = g_b(\bar{x}_b)$, hence

$$\frac{\min_{a \neq 1} g_a(\bar{x}_a)}{1 + \bar{x}_2 + \dots + \bar{x}_K} = \frac{g_b(\bar{x}_b)}{1 + x_2^* + \dots + x_K^*} > \frac{g_b(x_b^*)}{1 + x_2^* + \dots + x_K^*} = \frac{\min_{a \neq 1} g_a(x_a^*)}{1 + x_2^* + \dots + x_K^*},$$

This is a contradiction with the fact that x^* belongs to (12). Hence $\mathcal{A} = \emptyset$ and there exists $y^* \in [0, \min_{a \neq 1} d_1(\mu_a))$ such that, for all $a \in [K] \setminus \{1\}$,

$$g_a(x_a^*) = y^* \iff x_a^* = x_a(y^*)$$

with the function x_a introduced in Lemma 14. From (12), y^* belongs to

$$\arg \max_{y \in [0, \min_{a \neq 1} d_1(\mu_a))} \frac{y}{1 + \sum_{a=2}^K x_a(y)}$$

Using Lemma 14, we have that y^* is a solution of $F(y) = 1$ with

$$F(y) = \sum_{a=2}^K \frac{d_1(\lambda_a(x_a(y)))}{d_a(\lambda_a(x_a(y)))}.$$

■

A key theoretical and computational difference between the oracle for $w_{\sigma^2}^*(\mu)$ and the one for $w^*(\mu, \sigma^2)$ is that $\lambda_a(x)$ is defined implicitly as one of the real solution of a third order polynomial equation (Lemma 14). Fortunately, there exists closed form solutions for the roots of a third order polynomial equation, namely Cardano's formula. Therefore, we only need to compute three roots, among which at least one is real (two might be complex), and find the real one minimizing our original functions.

Lemma 14 gathers technical results used to prove Theorem 13.

Lemma 14 *Let d_a , g_a and λ_a as in Theorem 13.*

1. *For all $x \in (0, +\infty)$, $\lambda_a(x) \in (\mu_a, \mu_1)$, $\lambda_a(0) = \mu_1$ and $\lim_{x \rightarrow +\infty} \lambda_a(x) = \mu_a$. Moreover, the functions $\lambda_a(x)$ are among the (at least one) real solutions of $P_a(\lambda, x) = 0$ where*

$$\begin{aligned} P_a(\lambda, x) &= \lambda^3 - \alpha_{a,2}(x)\lambda^2 + \alpha_{a,1}(x)\lambda - \alpha_{a,0}(x), \\ \alpha_{a,2}(x) &= \mu_a + \mu_1 + \frac{\mu_a + x\mu_1}{1+x}, \\ \alpha_{a,1}(x) &= \frac{\sigma_a^2 + x\sigma_1^2}{1+x} + \mu_1\mu_a + (\mu_1 + \mu_a)\frac{\mu_a + x\mu_1}{1+x}, \\ \alpha_{a,0}(x) &= \frac{\mu_1(\mu_a^2 + \sigma_a^2) + \mu_a(\mu_1^2 + \sigma_1^2)x}{1+x}. \end{aligned}$$

2. *The function g_a is a concave and strictly increasing one-to-one mapping from $[0, +\infty)$ to $[0, d_1(\mu_a))$, such that $g'_a(x) = d_a(\lambda_a(x))$ and $g''_a(x) = \lambda'_a(x)d'_a(\lambda_a(x))$. In particular, $x \mapsto \lambda_a(x)$ is decreasing. Moreover, the function x_a is strictly increasing and $x'_a(y) = \frac{1}{d_a(\lambda_a(x(y)))}$.*
3. *Defining $G(y) = \frac{y}{1 + \sum_{a=2}^K x_a(y)}$ and $F(y) = \sum_{a=2}^K \frac{d_1(\lambda_a(x_a(y)))}{d_a(\lambda_a(x_a(y)))}$ for $y \in [0, \min_{a \neq 1} d_1(\mu_a))$. Then $G'(y) = 0$ if and only if $F(y) = 1$. The function F is increasing such that $F(0) = 0$ and $\lim_{y \rightarrow \min_{a \neq 1} d_1(\mu_a)} F(y) = +\infty$.*

Proof For all $a \in [K]$ and all $\lambda \in \mathbb{R}$, let $d_a^{\text{KV}}(\lambda) = \frac{(\lambda - \mu_a)^2}{\sigma_a^2}$. Then, we have

$$d'_a(\lambda) = \frac{(d_a^{\text{KV}})'(\lambda)}{1 + d_a^{\text{KV}}(\lambda)} = \frac{2(\lambda - \mu_a)}{\sigma_a^2 + (\lambda - \mu_a)^2}.$$

(1) Let $h_a(x, \lambda) = d_1(\lambda) + xd_a(\lambda)$. Since $\frac{\partial h_a}{\partial \lambda}(x, \lambda) = d'_1(\lambda) + xd'_a(\lambda)$, we have $\frac{\partial h_a}{\partial \lambda}(x, \lambda) > 0$ for $\lambda > \mu_1$ and $\frac{\partial h_a}{\partial \lambda}(x, \lambda) < 0$ for $\lambda < \mu_a$. Therefore, the solution of $\frac{\partial h_a}{\partial \lambda}(x, \lambda) = 0$ is in $[\mu_a, \mu_1]$ if it exists. As $h_a(x, \lambda)$ is continuous and bounded on $[\mu_a, \mu_1]$, $\lambda_a(x)$ exists.

The functions d_a have a unique minimizer μ_a such that $d_a(\mu_a) = 0$. Using that $h_a(0, \lambda) = d_1(\lambda)$, we obtain $\lambda_a(0) = \mu_1$. Introducing $w = w_1x$ such that $w_1 + w \leq 1$, we have $\lambda_a(x) = \tilde{\lambda}_a(w) = \arg \min_{\lambda \in [\mu_a, \mu_1]} w_1d_1(\lambda) + wd_a(\lambda)$. The same argument as above shows that $\tilde{\lambda}_a(1) = \mu_a$, as $w = 1$ implies w_0 , we have $\lim_{x \rightarrow +\infty} \lambda_a(x) = \tilde{\lambda}_a(1) = \mu_a$.

Since d'_a is continuous, we have $d'_a(\mu_a) = 0$. For all $x > 0$ and $a \in [K] \setminus \{1\}$, we have $\frac{\partial h_a}{\partial \lambda}(x, \mu_a) = d'_1(\mu_a) + xd'_a(\mu_a) = d'_1(\mu_a) < 0$ and $\frac{\partial h_a}{\partial \lambda}(x, \mu_1) = d'_1(\mu_1) + xd'_a(\mu_1) = xd'_a(\mu_1) > 0$. Therefore, we have $\lambda_a(x) \in (\mu_a, \mu_1)$ for all $x \in (0, +\infty)$.

Writing the condition that $\lambda_a(x)$ is a minimum, rewrites as $d'_1(\lambda_a(x)) + xd'_a(\lambda_a(x)) = 0$ and $d''_1(\lambda_a(x)) + xd''_a(\lambda_a(x)) > 0$. Direct computations yield

$$\begin{aligned} d'_1(\lambda_a(x)) + xd'_a(\lambda_a(x)) &= 0 \\ \iff (\lambda - \mu_1)(\lambda - \mu_a) \left(\lambda - \frac{\mu_a + x\mu_1}{1+x} \right) + \frac{\sigma_a^2 + x\sigma_1^2}{1+x} (\lambda - \lambda_a^{\text{KV}}(x)) &= 0 \iff P_a(\lambda, x) = 0 \end{aligned}$$

where $P_a(\lambda, x)$ is defined in the statement of Lemma 14. When $x \rightarrow +\infty$, we already know that $\lambda_a(x) = \mu_a + o(1)$. To prove $\lim_{x \rightarrow +\infty} xd_a(\lambda_a(x)) = 0$, we need a finer dependency in x . Writing

the change of variable $y = \lambda - \mu_a$, we obtain $P_a(\lambda, x) = 0$ if and only if $Q_a(y, x) = 0$ where

$$Q_a(y, x) = y(y + \mu_a - \mu_1) \left(y - \frac{x(\mu_1 - \mu_a)}{1+x} \right) + \frac{\sigma_a^2 + x\sigma_1^2}{1+x} \left(y - \frac{\sigma_a^2(\mu_1 - \mu_a)}{\sigma_a^2 + x\sigma_1^2} \right)$$

When $x \rightarrow +\infty$, we have $y(x) = o(1)$, $\frac{x(\mu_1 - \mu_a)}{1+x} = (\mu_1 - \mu_a)(1 + \mathcal{O}(\frac{1}{x}))$, $\frac{\sigma_a^2(\mu_1 - \mu_a)}{\sigma_a^2 + x\sigma_1^2} = \mathcal{O}(\frac{1}{x})$ and $\frac{\sigma_a^2 + x\sigma_1^2}{1+x} = \sigma_1^2 + \mathcal{O}(\frac{1}{x})$, we obtain the following

$$Q_a(y(x), x) = 0 \implies y(x) (\sigma_1^2 + (\mu_1 - \mu_a)^2 + o(1)) + \mathcal{O}(\frac{1}{x}) = 0$$

This yields that $y(x) = \mathcal{O}(\frac{1}{x})$. Therefore, $d_a^{\text{KV}}(\lambda_a(x)) = \frac{y(x)^2}{\sigma_a^2} = \mathcal{O}(\frac{1}{x^2})$ and

$$x d_a(\lambda_a(x)) = x \log(1 + d_a^{\text{KV}}(\lambda_a(x))) \sim_\infty x d_a^{\text{KV}}(\lambda_a(x)) \sim_\infty x \frac{y(x)^2}{\sigma_a^2} = 0$$

Therefore, we have shown $\lim_{x \rightarrow +\infty} x d_a(\lambda_a(x)) = 0$.

(2) Since $g_a(x) = \min_{\lambda \in [\mu_a, \mu_1]} d_1(\lambda) + x d_a(\lambda)$, g_a is a concave function as infimum of an infinite number of linear functions. Since $g_a(x) = h_a(x, \lambda_a(x))$, $\frac{\partial h_a}{\partial \lambda}(x, \lambda_a(x)) = 0$ (optimality condition) and $\frac{\partial h_a}{\partial x}(x, \lambda) = d_a(\lambda)$, we have

$$g'_a(x) = \frac{\partial h_a}{\partial x}(x, \lambda_a(x)) + \lambda'_a(x) \frac{\partial h_a}{\partial \lambda}(x, \lambda_a(x)) = \frac{\partial h_a}{\partial x}(x, \lambda_a(x)) = d_a(\lambda_a(x)) > 0$$

where the last inequality is strict since $\lambda_a(x) \in (\mu_a, \mu_1]$ for $x \in [0, +\infty)$. Therefore, g_a is a strictly increasing one-to-one mapping from $[0, +\infty)$ to $[0, \lim_{x \rightarrow +\infty} g_a(x))$. Since $\lim_{x \rightarrow +\infty} \lambda_a(x) = \mu_a$ and $\lim_{x \rightarrow +\infty} x d_a(\lambda_a(x)) = 0$, we obtain $\lim_{x \rightarrow +\infty} g_a(x) = d_1(\mu_a)$.

Directly computing the derivative of $d_a(\lambda_a(x))$, we obtain $g''_a(x) = \lambda'_a(x) d'_a(\lambda_a(x))$. Using that $\lambda_a(x) > \mu_a$ for all $x \in [0, +\infty)$, we have $d'_a(\lambda_a(x)) > 0$. Since g_a is concave, we know that $g''_a(x) \geq 0$, hence $\lambda'_a(x) \leq 0$, i.e. $x \mapsto \lambda_a(x)$ is decreasing.

Using the above result, $x_a(y)$ is well defined. Using the derivative of the inverse function and $g'_a(x) = d_a(\lambda_a(x))$, we obtain $x'_a(y) = \frac{1}{g'_a(x_a(y))} = \frac{1}{d_a(\lambda_a(x_a(y)))}$. Since $d_a(\lambda_a(x)) > 0$ for all $x \in [0, +\infty)$, we have $x'_a(y) > 0$ for all $y \in [0, d_1(\mu_a))$.

(3) Using that $x'_a(y) = \frac{1}{d_a(\lambda_a(x_a(y)))}$ and $d_1(\lambda_a(x_a(y))) + x_a(y) d_a(\lambda_a(x_a(y))) = y$. Direct computations yield that $G'(y) = 0$ if and only if

$$\begin{aligned} \frac{1}{1 + \sum_{a=2}^K x_a(y)} &= \frac{y}{(1 + \sum_{a=2}^K x_a(y))^2} \sum_{a=2}^K x'_a(y) \\ \iff \sum_{a=2}^K \frac{y}{d_a(\lambda_a(x_a(y)))} &= 1 + \sum_{a=2}^K x_a(y) \iff \sum_{a=2}^K \frac{d_1(\lambda_a(x_a(y)))}{d_a(\lambda_a(x_a(y)))} = 1 \iff F(y) = 1 \end{aligned}$$

Using that $\lambda_a(0) = \mu_1$ and $d_1(\mu_1) = 0$, we obtain $F(0) = 0$. Using that $\lim_{x \rightarrow +\infty} \lambda(x) = \mu_a$, $d_1(\mu_a) > 0$, $d_a(\mu_a) = 0$ and $\lim_{y \rightarrow d_1(\mu_a)} x_a(y) = +\infty$, we obtain that $\lim_{y \rightarrow \min_{a \neq 1} d_1(\mu_a)} F(y) = +\infty$. Let $H(y) = \sum_{a=2}^K \frac{1}{d_a(\lambda_a(x_a(y)))}$ for $y \in [0, \min_{a \neq 1} d_1(\mu_a))$. Since $d'_1(\lambda_a(x)) + x d'_a(\lambda_a(x)) =$

0, $d_1(\lambda_a(x_a(y))) + x_a(y)d_a(\lambda_a(x_a(y))) = y$, direct computations yield

$$\begin{aligned} H'(y) &= - \sum_{a=2}^K x'_a(y) \lambda'_a(x_a(y)) \frac{d'_a(\lambda_a(x_a(y)))}{d_a(\lambda_a(x_a(y)))^2} \\ F'(y) &= \sum_{a=2}^K x'_a(y) \lambda'_a(x_a(y)) \frac{d'_1(\lambda_a(x_a(y))) d_a(\lambda_a(x_a(y))) - d_1(\lambda_a(x_a(y))) d'_a(\lambda_a(x_a(y)))}{d_a(\lambda_a(x_a(y)))^2} \\ &= - \sum_{a=2}^K x'_a(y) \lambda'_a(x_a(y)) d'_a(\lambda_a(x_a(y))) \frac{x_a(y) d_a(\lambda_a(x_a(y))) + d_1(\lambda_a(x_a(y)))}{d_a(\lambda_a(x_a(y)))^2} \\ &= -y \sum_{a=2}^K x'_a(y) \lambda'_a(x_a(y)) \frac{d'_a(\lambda_a(x_a(y)))}{d_a(\lambda_a(x_a(y)))^2} = yH'(y) \end{aligned}$$

Note that for all $a \in [K] \setminus \{1\}$, $y \mapsto \frac{1}{d_a(\lambda_a(x_a(y)))}$ is an increasing function of y since $d_a(\lambda)$ is strictly increasing on $(\mu_a, \mu_1]$, $\lambda_a(x)$ is decreasing on $[0, +\infty)$ and $x_a(y)$ is strictly increasing on $y \in [0, \min_{a \neq 1} \lim_{x \rightarrow \infty} g_a(x))$. As a summation of increasing functions, H is an increasing function, i.e. $H'(y) \geq 0$. Since $F'(y) = yH'(y) \geq 0$, this yield that F is an increasing function. \blacksquare

Appendix D. Generalized Log-Likelihood Ratios

The GLR to reject the hypothesis $\mathcal{H}_0 = \{(\mu, \sigma^2) \in \Lambda\}$ for a subset $\Lambda \subseteq \mathcal{M}$ of the whole parameter space \mathcal{M} is given by

$$\text{GLR}_t^{\mathcal{M}}(\Lambda) \stackrel{\text{def}}{=} \log \frac{\sup_{(\mu, \sigma^2) \in \mathcal{M}} d\nu_{\mu, \sigma^2}(X_{1, a_1}, \dots, X_{t, a_t})}{\sup_{(\lambda, \kappa^2) \in \Lambda} d\nu_{\lambda, \kappa^2}(X_{1, a_1}, \dots, X_{t, a_t})}$$

where $d\nu_{\mu, \sigma^2}(X_{1, a_1}, \dots, X_{t, a_t})$ denotes the likelihood of the observations $X_{1, a_1}, \dots, X_{t, a_t}$ for a Gaussian bandit with parameters (μ, σ^2) . The empirical mean and variance are defined as

$$\mu_{t, a} = \frac{1}{N_{t, a}} \sum_{s \in [t]} \mathbf{1}\{a_s = a\} X_{s, a} \quad \text{and} \quad \sigma_{t, a}^2 = \frac{1}{N_{t, a}} \sum_{s \in [t]} \mathbf{1}\{a_s = a\} (X_{s, a} - \mu_{t, a})^2.$$

Using the formula of the KL for Gaussian bandits with different variances, the GLR rewrites as

$$\begin{aligned} \text{GLR}_t^{\mathcal{M}}(\Lambda) &= \inf_{(\lambda, \kappa^2) \in \Lambda} \sum_{a \in [K]} \frac{N_{t, a}}{2} \left(\log \left(\frac{\kappa_a^2}{\sigma_{t, a}^2} \right) + \frac{\sigma_{t, a}^2 + (\mu_{t, a} - \lambda_a)^2}{\kappa_a^2} \right) \\ &= \inf_{(\lambda, \kappa^2) \in \Lambda} \sum_{a \in [K]} N_{t, a} \text{KL}((\mu_{t, a}, \sigma_{t, a}^2), (\lambda_a, \kappa_a^2)). \end{aligned}$$

When $\Lambda = \mathcal{M} = \mathbb{R}^K \times (\mathbb{R}_+^*)^K$, the minimizers of the log-likelihood are (μ_t, σ_t^2) , i.e. the empirical estimators (μ_t, σ_t^2) are exactly the MLE.

Explicit formulas The GLR has a more explicit formula when considering the alternative to the subset of interest $\Lambda(\mu_t, \sigma_t^2)$ (Lemma 15).

Lemma 15 Let $\mathcal{M} = \mathbb{R}^K \times (\mathbb{R}_+^*)^K$. Then, $\text{GLR}_t^{\mathcal{D}}(\Lambda(\mu_t, \sigma_t^2)) = \min_{a \neq \hat{a}_t} Z_a(t)$ where

$$Z_a(t) = \text{GLR}_t^{\mathcal{D}}(\{(\lambda, \kappa^2) \mid \lambda_a \geq \lambda_{\hat{a}_t}\}) = \inf_{\lambda: \lambda_a \geq \lambda_{\hat{a}_t}} \sum_{b \in \{a, \hat{a}_t\}} \frac{N_{t,b}}{2} \log \left(1 + \frac{(\mu_{t,b} - \lambda_b)^2}{\sigma_{t,b}^2} \right).$$

Proof Since $a^*(\mu_t) = \hat{a}_t$ and $\Lambda(\mu_t, \sigma_t^2) = \bigcup_{a \neq \hat{a}_t} \{(\lambda, \kappa^2) \mid \lambda_a \geq \lambda_{\hat{a}_t}\}$, we obtain

$$\text{GLR}_t^{\mathcal{M}}(\Lambda(\mu_t, \sigma_t^2)) = \min_{a \neq \hat{a}_t} Z_a(t) \quad \text{with} \quad Z_a(t) = \text{GLR}_t^{\mathcal{D}}(\{(\lambda, \kappa^2) \mid \lambda_a \geq \lambda_{\hat{a}_t}\}).$$

For $b \notin \{a, \hat{a}_t\}$, we can take $(\lambda_b, \kappa_b^2) = (\mu_{t,b}, \sigma_{t,b}^2)$, hence we obtain

$$\begin{aligned} Z_a(t) &= \inf_{(\lambda, \kappa^2): \lambda_a \geq \lambda_{\hat{a}_t}} \sum_{b \in \{a, \hat{a}_t\}} N_{t,b} \text{KL}((\mu_{t,b}, \sigma_{t,b}^2), (\lambda_b, \kappa_b^2)) \\ &= \inf_{\lambda: \lambda_a \geq \lambda_{\hat{a}_t}} \sum_{b \in \{a, \hat{a}_t\}} \frac{N_{t,b}}{2} \log \left(1 + \frac{(\mu_{t,b} - \lambda_b)^2}{\sigma_{t,b}^2} \right) \end{aligned}$$

The second equality uses that the objectives and the constraints are separate and that

$$\arg \min_{\kappa_b^2} \text{KL}((\mu_{t,b}, \sigma_{t,b}^2), (\lambda_b, \kappa_b^2)) = (\mu_{t,b} - \lambda_b)^2 + \sigma_{t,b}^2,$$

since $f_a(x) = \frac{a}{x} + \log(x)$ has $f'_a(x) = \frac{1}{x} - \frac{a}{x^2}$ and $\arg \min_{x>0} f_a(x) = a$, for $a = (\mu_{t,b} - \lambda_b)^2 + \sigma_{t,b}^2$. ■

Inequalities Lemma 16 gives inequalities between $Z_a(t)$ and $Z_a^{\text{EV}}(t)$.

Lemma 16 For $a \neq \hat{a}_t$, let $C_a(\mu_t, \sigma_t^2) = \frac{(\mu_{t,\hat{a}_t} - \mu_{t,a})^2}{\min\{\sigma_{t,\hat{a}_t}^2, \sigma_{t,a}^2\}}$. The statistics $Z_a^{\text{EV}}(t)$ and $Z_a(t)$ satisfy

$$Z_a^{\text{EV}}(t) \geq Z_a(t) \geq \frac{\log \left(1 + C_a(\mu_t, \sigma_t^2) \right)}{C_a(\mu_t, \sigma_t^2)} Z_a^{\text{EV}}(t). \quad (13)$$

The proof is omitted since it was obtained with similar manipulations on $Z_a(t)$ and $Z_a^{\text{EV}}(t)$ as done on $T^*(\mu, \sigma^2)$ and $T_{\sigma^2}^*(\mu)$ in Appendix C.1.

Appendix E. Box Concentration

In Appendix E.1, we derive time-uniform upper and lower tail concentrations for sub-exponential processes, as well as their fixed-time counterpart. In Appendix E.2, for Gaussian observations, we use those results to obtain similar concentrations for the empirical variance and derive time-uniform (and fixed-time) upper and lower tail concentrations for the empirical mean.

E.1. Sub-Exponential Processes

We prove time-uniform and fixed-time concentration results for 1-sub- $\psi_{E,c}$ process with variance process $V_t = ct$ (Appendix E.1.1) and 1-sub- $\psi_{E,-c}$ process with variance process $V_t = ct$ (Appendix E.1.2). The concept of sub- ψ process (Definition 17) was introduced in Howard et al. (2020). This concept is particularly useful to derive time-uniform concentration results.

Definition 17 Let $(S_t)_{t \in \mathcal{T} \cup \{0\}}$ and $(V_t)_{t \in \mathcal{T} \cup \{0\}}$ be two real-valued processes adapted to an underlying filtration $(\mathcal{F}_t)_{t \in \mathcal{T} \cup \{0\}}$ with $S_0 = 0$ and $V_0 = 0$ a.s. and $V_t \geq 0$ a.s. for all $t \in \mathcal{T}$. For a function $\psi : [0, \lambda_{\max}) \mapsto \mathbb{R}$ and a scalar $l_0 \in [1, +\infty)$, we say that (S_t) is l_0 -sub- ψ with variance process (V_t) if, for each $\lambda \in [0, \lambda_{\max})$, there exists a supermartingale $(L_t(\lambda))_{t \in \mathcal{T} \cup \{0\}}$ with respect to (\mathcal{F}_t) such that $L_0(\lambda) \leq l_0$ a.s. and

$$\exp \{ \lambda S_t - \psi(\lambda) V_t \} \leq L_t(\lambda) \quad \text{a.s. for all } t \in \mathcal{T} .$$

Lemma 18 (Ville's inequality) Let $\mathbb{P}_0[\cdot] = \mathbb{P}_0[\cdot \mid \mathcal{F}_0]$. Let $\mathcal{T} \subseteq \mathbb{N}^*$, such that $|\mathcal{T}| = \infty$. If $(L_t)_{t \in \mathcal{T} \cup \{0\}}$ is a non-negative supermartingale with respect to the filtration $(\mathcal{F}_t)_{t \in \mathcal{T} \cup \{0\}}$, then

$$\forall a > 0, \quad \mathbb{P}_0(\exists t \in \mathcal{T} : L_t \geq a) \leq L_0/a .$$

Since we aim at deriving one-sided bounds on scalar martingales, we have $l_0 = 1$. Using Ville's inequality (Lemma 18) on a sub- ψ process yields time-uniform concentration results. Let (S_t) be a 1-sub- ψ with variance process (V_t) , then for all $\lambda \in [0, \lambda_{\max})$, with probability greater than $1 - \delta$,

$$\forall t \in \mathcal{T}, \quad \lambda S_t - \psi(\lambda) V_t < \log(1/\delta) .$$

Let $\lambda \in [0, \lambda_{\max})$. Direct manipulations show the above result,

$$\mathbb{P}(\exists t \in \mathcal{T} : \lambda S_t - \psi(\lambda) V_t \geq \log(1/\delta)) \leq \mathbb{P}(\exists t \in \mathcal{T} : L_t(\lambda) \geq 1/\delta) \leq \delta .$$

In the following, we are interested by 1-sub- $\psi_{E,c}$ processes for $c \in \mathbb{R}$, where $\psi_{E,c}$ is defined as

$$\forall \lambda \in [0, 1/(c \vee 0)), \quad \psi_{E,c}(\lambda) = \frac{-\log(1 - c\lambda) - c\lambda}{c^2} . \quad (14)$$

The derived upper and lower tails concentrations involve the positive ($i = 0$) and negative ($i = -1$) Lambert's branches W_i solutions of $W(x)e^{W(x)} = x$. We refer the reader to Appendix I for more details and corresponding technical results.

E.1.1. UPPER TAIL CONCENTRATION

We derive time-uniform and fixed-time upper tail concentration for 1-sub- $\psi_{E,c}$ process with variance process $V_t = ct$. While the time-uniform result require using the peeling method, the proof of the fixed-time concentration is simpler. To use the peeling method, we need to control the deviation of the process on slices of time (Lemma 19).

Lemma 19 Let $c > 0$ and S_t a 1-sub- $\psi_{E,c}$ process with variance process $V_t = ct$. Let $N > 0$. For all $x > 1$, there exists $\lambda = \lambda(x)$ such that for all $t \geq N$,

$$\{S_t + t \geq tx\} \subseteq \left\{ \lambda S_t - ct\psi_{E,c}(\lambda) \geq \frac{N}{c} (h(x) - 1) \right\}$$

where $\lambda(x) = \arg \max_{\lambda \in [0, 1/c)} \left(x\lambda + \frac{\log(1 - c\lambda)}{c} \right)$ and $h(x) = x - \log(x)$ for $x > 1$.

Proof Defining $\psi_U(\lambda) = \lambda + c\psi_{E,c}(\lambda) = -\frac{\log(1-c\lambda)}{c}$ and $\lambda(x) = \arg \max_{\lambda \in [0, 1/c]} x\lambda - \psi_U(\lambda)$, we have $x\lambda(x) - \psi_U(\lambda(x)) = \psi_U^*(x)$ where ψ_U^* is the convex conjugate of ψ_U . Note that $\psi_U^*(x) \geq 0$ (see below), hence $t\psi_U^*(x) \geq N\psi_U^*(x)$ for $t \geq N$. Direct computations yield

$$\begin{aligned} S_t + t \geq tx &\iff \lambda S_t - ct\psi_{E,c}(\lambda) \geq tx\lambda - t(\lambda + c\psi_{E,c}(\lambda)) \\ &\implies \lambda S_t - ct\psi_{E,c}(\lambda) \geq t(x\lambda - \psi_U(\lambda)) = t\psi_U^*(x) \\ &\implies \lambda S_t - ct\psi_{E,c}(\lambda) \geq N\psi_U^*(x) = \frac{N}{c}(h(x) - 1) \end{aligned}$$

Note that for $f(\lambda) = \lambda x + \frac{\log(1-c\lambda)}{c}$, we have $f'(\lambda) = x - \frac{1}{1-c\lambda} = 0 \iff \lambda = \frac{1}{c}(1 - \frac{1}{x})$ and $\frac{1}{c}(1 - \frac{1}{x}) \in [0, \frac{1}{c}] \iff x > 1$. Since $f''(\lambda) = -\frac{c}{(1-c\lambda)^2} \leq 0$, the function is concave hence this is a maximum. This yields that for all $x > 1$, $\psi_U^*(x) = f(\frac{1}{c}(1 - \frac{1}{x})) = \frac{1}{c}(x - 1 - \log(x)) = \frac{1}{c}(h(x) - 1) \geq 0$ where $h(x) = x - \log(x)$. \blacksquare

Let $\eta > 0$. Applying Lemma 19 on slices of time with geometric growth rate $(N_i)_{i \in \mathbb{N}^*}$ with $N_i = (1 + \eta)^{i-1}$, we obtain Lemma 20.

Lemma 20 Let $\bar{W}_{-1}(x) = -W_{-1}(-e^{-x})$ for $x \geq 1$, $\delta \in (0, 1)$, $\eta > 0$, $s > 1$, $c > 0$, and ζ be the Riemann ζ function. Let S_t a 1-sub- $\psi_{E,c}$ process with variance process $V_t = ct$. Then, with probability greater than $1 - \delta$, for all $t \in \mathbb{N}^*$,

$$S_t + t \leq t\bar{W}_{-1} \left(1 + \frac{c(1+\eta)}{t} \left(\log \left(\frac{\zeta(s)}{\delta} \right) + s \log \left(1 + \frac{\log(t)}{\log(1+\eta)} \right) \right) \right).$$

Proof Let $g(t, \delta)$ such that $g(t, \delta) \geq x_i(\delta)$ for $t \in [N_i, N_{i+1})$ and $x_i(\delta) > 1$. Using Lemma 19 with $x_i(\delta) > 1$ and $g(t, \delta) \geq x_i(\delta)$ on $[N_i, N_{i+1})$, we obtain

$$\begin{aligned} \mathbb{P}(\exists t \in \mathbb{N}^* : S_t + t \geq tg(t, \delta)) &\leq \sum_{i \in \mathbb{N}^*} \mathbb{P}(\exists t \in [N_i, N_{i+1}) : S_t + t \geq tx_i(\delta)) \\ &\leq \sum_{i \in \mathbb{N}^*} \mathbb{P} \left(\exists t \in [N_i, N_{i+1}) : \lambda S_t - ct\psi_{E,c}(\lambda) \geq \frac{N_i}{c}(h(x_i(\delta)) - 1) \right) \\ &\leq \sum_{i \in \mathbb{N}^*} e^{-\frac{N_i}{c}(h(x_i(\delta)) - 1)}, \end{aligned}$$

where the last inequality uses that S_t a 1-sub- $\psi_{E,c}$ process with variance process $V_t = ct$. Taking

$$g(t, \delta) = \bar{W}_{-1} \left(1 + \frac{c(1+\eta)}{t} \left(\log \left(\frac{\zeta(s)}{\delta} \right) + s \log \left(1 + \frac{\log(t)}{\log(1+\eta)} \right) \right) \right)$$

and $x_i(\delta) = \bar{W}_{-1} \left(1 + \frac{c}{N_i} \log \left(\frac{i^s \zeta(s)}{\delta} \right) \right)$ satisfies the required properties. First, we have $x_i(\delta) > 1$ (Lemma 75). Second, since \bar{W}_{-1} is increasing on $(1, +\infty)$ (Lemma 75), $t \in [N_i, N_{i+1})$ and $i = 1 + \frac{\log(N_i)}{\log(1+\eta)}$, we obtain

$$g(t, \delta) \geq \bar{W}_{-1} \left(1 + \frac{c \log \left(\frac{\zeta(s)}{\delta} \right) + cs \log \left(1 + \frac{\log(t)}{\log(1+\eta)} \right)}{N_i} \right) \geq \bar{W}_{-1} \left(1 + \frac{c}{N_i} \log \left(\frac{i^s \zeta(s)}{\delta} \right) \right)$$

Using Lemma 75 for each $i \in \mathbb{N}^*$ yields

$$\mathbb{P}(\exists t \in \mathbb{N}^* : S_t + t \geq tg(t, \delta)) \leq \sum_{i \in \mathbb{N}^*} e^{-\frac{N_i}{c}(h(x_i(\delta)) - 1)} \leq \frac{\delta}{\zeta(s)} \sum_{i \in \mathbb{N}^*} \frac{1}{i^s} = \delta$$

■

Fixed-time concentration When the time is fixed and not random, there is no need to consider slices of time and we can directly control the deviation of the process (Lemma 21).

Lemma 21 *Let $h(x) = x - \log(x)$ for $x > 1$. Let $c > 0$ and S_t a 1-sub- $\psi_{E,c}$ process with variance process $V_t = ct$. Then,*

$$\forall t \in \mathbb{N}^*, \forall x > 1, \quad \mathbb{P}(S_t + t \geq tx) \leq \exp\left(-\frac{t}{c}(h(x) - 1)\right).$$

Proof With similar computations as in the proof of Lemma 19, the fact that S_t a 1-sub- $\psi_{E,c}$ process with variance process $V_t = ct$ and the Chernoff inequality yield the first result. ■

E.1.2. LOWER TAIL CONCENTRATION

We derive time-uniform and fixed-time lower tail concentration for 1-sub- $\psi_{E,-c}$ process with variance process $V_t = ct$. Likewise, we use the peeling method and control the deviation of the process on slices of time (Lemma 22).

Lemma 22 *Let $c > 0$ and $-S_t$ a 1-sub- $\psi_{E,-c}$ process with variance process $V_t = ct$. Let $N > 0$. For all $x \in (0, 1)$, there exists $\lambda = \lambda(x)$ such that for all $t \geq N$,*

$$\{-S_t - t \geq -tx\} \subseteq \left\{ \lambda(-S_t) - ct\psi_{E,-c}(\lambda) \geq \frac{N}{c}(h(x) - 1) \right\}$$

where $\lambda(x) = \arg \max_{\lambda \in [0, +\infty)} \left(-x\lambda + \frac{\log(1+c\lambda)}{c}\right)$ and $h(x) = x - \log(x)$ for $x \in (0, 1)$.

Proof Defining $\psi_L(\lambda) = -\lambda + c\psi_{E,-c}(\lambda) = -\frac{\log(1+c\lambda)}{c}$ and $\lambda(x) = \arg \max_{\lambda \in [0, +\infty)} -x\lambda - \psi_L(\lambda)$, we have $-x\lambda(x) - \psi_L(\lambda(x)) = \psi_L^*(-x) \geq 0$ (see below), hence $t\psi_L^*(-x) \geq N\psi_L^*(-x)$ for $t \geq N$. Direct computations yield

$$\begin{aligned} -S_t - t \geq -tx &\iff \lambda(-S_t) - ct\psi_{E,-c}(\lambda) \geq -tx\lambda - t(-\lambda + c\psi_{E,-c}(\lambda)) \\ &\iff \lambda(-S_t) - ct\psi_{E,-c}(\lambda) \geq t(-x\lambda - \psi_L(\lambda)) = t\psi_L^*(-x) \\ &\iff \lambda(-S_t) - ct\psi_{E,-c}(\lambda) \geq N\psi_L^*(-x) = \frac{N}{c}(h(x) - 1) \end{aligned}$$

Note that for $f(\lambda) = -x\lambda + \frac{\log(1+c\lambda)}{c}$, we have $f'(\lambda) = -x + \frac{1}{1+c\lambda} = 0 \iff \lambda = \frac{1}{c}\left(\frac{1}{x} - 1\right)$ and $\frac{1}{c}\left(\frac{1}{x} - 1\right) \in [0, +\infty) \iff x \in (0, 1)$. Since $f''(\lambda) = -\frac{c}{(1+c\lambda)^2} \leq 0$, the function is concave hence this is a maximum. This yields that for all $x \in (0, 1)$, $\psi_L^*(-x) = f\left(\frac{1}{c}\left(\frac{1}{x} - 1\right)\right) = \frac{1}{c}(x - 1 - \log(x)) = \frac{1}{c}(h(x) - 1) \geq 0$ where $h(x) = x - \log(x)$ for $x \in (0, 1)$. ■

Let $\eta > 0$. Applying Lemma 22 on slices of time with geometric growth rate $(N_i)_{i \in \mathbb{N}^*}$ with $N_i = (1 + \eta)^{i-1}$, we obtain Lemma 23.

Lemma 23 Let $\bar{W}_0(x) = -W_0(-e^{-x})$ for $x \geq 1$, $\delta \in (0, 1)$, $\eta > 0$, $s > 1$, $c > 0$, and ζ be the Riemann ζ function. Let S_t a 1-sub- $\psi_{E,-c}$ process with variance process $V_t = ct$. Then, with probability greater than $1 - \delta$, for all $t \in \mathbb{N}^*$,

$$S_t + t \geq t\bar{W}_0 \left(1 + \frac{c(1+\eta)}{t} \left(\log \left(\frac{\zeta(s)}{\delta} \right) + s \log \left(1 + \frac{\log(t)}{\log(1+\eta)} \right) \right) \right).$$

Proof Let $g(t, \delta)$ positive such that $g(t, \delta) \leq x_i(\delta)$ for $t \in [N_i, N_{i+1})$ and $x_i(\delta) \in (0, 1)$. Using Lemma 22 with $x_i(\delta) < 1$ and $g(t, \delta) \leq x_i(\delta)$ for $t \in [N_i, N_{i+1})$, we obtain

$$\begin{aligned} \mathbb{P}(\exists t \in \mathbb{N}^* : S_t + t \leq tg(t, \delta)) &= \mathbb{P}(\exists t \in \mathbb{N} : -S_t - t \geq -tg(t, \delta)) \\ &\leq \sum_{i \in \mathbb{N}^*} \mathbb{P}(\exists t \in T_i : -S_t - t \geq -tx_i(\delta)) \\ &\leq \sum_{i \in \mathbb{N}^*} \mathbb{P} \left(\exists t \in T_i : \lambda(-S_t) - ct\psi_{E,-c}(\lambda) \geq \frac{N_i}{c} (h(x_i(\delta)) - 1) \right) \\ &\leq \sum_{i \in \mathbb{N}^*} e^{-\frac{N_i}{c} (h(x_i(\delta)) - 1)} \end{aligned}$$

where the last inequality uses that $-S_t$ a 1-sub- $\psi_{E,-c}$ process with variance process $V_t = ct$. Taking

$$g(t, \delta) = \bar{W}_0 \left(1 + \frac{c(1+\eta)}{t} \left(\log \left(\frac{\zeta(s)}{\delta} \right) + s \log \left(1 + \frac{\log(t)}{\log(1+\eta)} \right) \right) \right)$$

and $x_i(\delta) = \bar{W}_0 \left(1 + \frac{c}{N_i} \log \left(\frac{i^s \zeta(s)}{\delta} \right) \right)$ satisfies the required properties. First, we have $x_i(\delta) \in (0, 1)$ and $g(t, \delta) > 0$ (Lemma 75). Second, since \bar{W}_0 is decreasing on $(1, +\infty)$ (Lemma 75), $t \in [N_i, N_{i+1})$ and $i = 1 + \frac{\log(N_i)}{\log(1+\eta)}$, we obtain

$$g(t, \delta) \leq \bar{W}_0 \left(1 + \frac{c \log \left(\frac{\zeta(s)}{\delta} \right) + cs \log \left(1 + \frac{\log(t)}{\log(1+\eta)} \right)}{N_i} \right) \leq \bar{W}_0 \left(1 + \frac{c}{N_i} \log \left(\frac{i^s \zeta(s)}{\delta} \right) \right)$$

Using Lemma 75 for each $i \in \mathbb{N}^*$ yields

$$\mathbb{P}(\exists t \in \mathbb{N}^* : S_t + t \leq tg(t, \delta)) \leq \sum_{i \in \mathbb{N}^*} e^{-\frac{N_i}{c} (h(x_i(\delta)) - 1)} \leq \frac{\delta}{\zeta(s)} \sum_{i \in \mathbb{N}^*} \frac{1}{i^s} = \delta. \quad \blacksquare$$

Fixed-time concentration When the time is fixed and not random, there is no need to consider slices of time and we can directly control the deviation of the process (Lemma 24).

Lemma 24 Let $h(x) = x - \log(x)$ for $x \in (0, 1)$. Let $c > 0$ and $-S_t$ a 1-sub- $\psi_{E,-c}$ process with variance process $V_t = ct$. Then,

$$\forall t \in \mathbb{N}^*, \forall x \in (0, 1), \quad \mathbb{P}(S_t + t \leq tx) \leq \exp \left(-\frac{t}{c} (h(x) - 1) \right).$$

Proof With similar computations as in the proof of Lemma 22, the fact that $-S_t$ a 1-sub- $\psi_{E,-c}$ process with variance process $V_t = ct$ and the Chernoff inequality yield the first result. \blacksquare

E.2. Univariate Gaussian

We prove time-uniform and fixed-time upper and lower concentration results for the empirical variance (Appendix E.2.1) and empirical mean (Appendix E.2.2) of Gaussian observations.

E.2.1. EMPIRICAL VARIANCE

We first prove Lemma 25 which shows that the empirical variance is closely linked with a sub-exponential process. This is obtained with manipulations derived in the Appendix H of Howard et al. (2021), in which they consider martingales with χ^2 increments.

Lemma 25 *Let σ_t^2 be the empirical variance of t i.i.d. samples from a Gaussian distribution with variance σ^2 . Then, $\frac{\sigma_t^2}{\sigma^2} = \frac{S_{t-1}-1}{t} + 1$ with $S_{t-1} + t - 1 = \sum_{i=1}^{t-1} Y_i^2$ where (Y_i) are i.i.d. with distributions $\mathcal{N}(0, 1)$. In particular, S_t is a 1-sub- $\psi_{E,2}$ process and $-S_t$ is a 1-sub- $\psi_{E,-2}$ process, both with variance process $V_t = 2t$.*

Proof Let $(X_i)_{i \in [n]}$ the n samples from a Gaussian distribution with parameters (μ, σ^2) . Let $\hat{\mu}_n$ and $\hat{\sigma}_n^2$ be the empirical mean and variance. Let $Z_i = \frac{X_i - \mu}{\sigma}$ for all $i \in [n]$, $\hat{Z}_n = \frac{1}{n} \sum_{i \in [n]} Z_i$ and $S_{n-1} = \sum_{i=1}^n (Z_i - \hat{Z}_n)^2 - (n-1)$. Then, $S_0 = 0$ and for all $n \geq 2$

$$S_{n-1} = \frac{1}{\sigma^2} \sum_{i=1}^n (X_i - \hat{\mu}_n)^2 - (n-1) = n \frac{\sigma_n^2}{\sigma^2} - (n-1).$$

Rewriting the increment of S_n , we obtain for all $n \geq 2$

$$\begin{aligned} S_{n-1} - S_{n-2} &= (Z_n - \hat{Z}_n)^2 + \sum_{i=1}^{n-1} ((Z_i - \hat{Z}_n)^2 - (Z_i - \hat{Z}_{n-1})^2) - 1 \\ &= Z_n^2 + \hat{Z}_n^2 - 2Z_n \hat{Z}_n + \sum_{i=1}^{n-1} (-2Z_i \hat{Z}_n + 2Z_i \hat{Z}_{n-1}) + (n-1)(\hat{Z}_n^2 - \hat{Z}_{n-1}^2) - 1 \\ &= Z_n^2 - n\hat{Z}_n^2 + (n-1)\hat{Z}_{n-1}^2 - 1 = \frac{n-1}{n}(Z_n - \hat{Z}_{n-1})^2 - 1. \end{aligned}$$

Since $S_{n-1} = \sum_{s=1}^{n-1} (S_s - S_{s-1})$ and $S_0 = 0$ a.s., we obtain $S_{n-1} = \sum_{i=1}^{n-1} (Y_i^2 - 1)$ where $Y_{n-1} = \sqrt{\frac{n-1}{n}}(Z_n - \hat{Z}_{n-1})$. The (Y_i) are iid with distribution $\mathcal{N}(0, 1)$ and the CGF of $(Y_i^2 - 1)$ is

$$\log \mathbb{E} e^{\lambda(Y_i^2 - 1)} = -\frac{\log(1 - 2\lambda)}{2} - \lambda = 2\psi_{E,2}(\lambda) \quad \text{for } \lambda \in (-\infty, 1/2).$$

By Definition 17, we have that S_n is 1-sub- $\psi_{E,2}$ and that $-S_n$ is 1-sub- $\psi_{E,-2}$, both with variance process $V_n = 2(n-1)$. \blacksquare

Thanks to Lemmas 20-23-25, Corollary 26 gives time-uniform upper and lower tails concentrations on the empirical variance of Gaussian observation.

Corollary 26 For $i \in \{0, -1\}$, let $\overline{W}_i(x) = -W_i(-e^{-x})$ for $x \geq 1$, $\delta \in (0, 1)$, $\eta_0, \eta_1 > 0$, $s > 1$ and ζ be the Riemann ζ function. Let σ_{t+1}^2 be the empirical variance of $t + 1$ i.i.d. samples from a Gaussian distribution with variance σ^2 . Then, with probability greater than $1 - \delta$, for all $t \in \mathbb{N}^*$,

$$\sigma_{t+1}^2 \leq \sigma^2 \left(\overline{W}_{-1} \left(1 + \frac{2(1 + \eta_1)}{t} \left(\log \left(\frac{\zeta(s)}{\delta} \right) + s \log \left(1 + \frac{\log(t)}{\log(1 + \eta_1)} \right) \right) \right) - \frac{1}{t} \right).$$

Moreover, with probability $1 - \delta$, for all $t \geq t_0(\delta)$,

$$\sigma_{t+1}^2 \geq \sigma^2 \left(\overline{W}_0 \left(1 + \frac{2(1 + \eta_0)}{t} \left(\log \left(\frac{\zeta(s)}{\delta} \right) + s \log \left(1 + \frac{\log(t)}{\log(1 + \eta_0)} \right) \right) \right) - \frac{1}{t} \right),$$

where the initial time condition, which ensures the lower bound is positive, is

$$t_0(\delta) = \inf \left\{ t \mid t > e^{1 + W_0 \left(\frac{2(1 + \eta_0)}{e} \left(\log \left(\frac{\zeta(s)}{\delta} \right) + s \log \left(1 + \frac{\log(t)}{\log(1 + \eta_0)} \right) \right) - e^{-1}} \right) \right\}.$$

Proof Combining Lemmas 20-23-25 yields the desired result. Using Lemma 75, we know that the upper bound is always positive. The initial time condition, after which the lower bound is positive, is obtained by Lemma 75

$$\overline{W}_0 \left(1 + \frac{2(1 + \eta_0)}{t} \left(\log \left(\frac{\zeta(s)}{\delta} \right) + s \log \left(1 + \frac{\log(t)}{\log(1 + \eta_0)} \right) \right) \right) > \frac{1}{t} \iff t \geq t_0(\delta). \quad \blacksquare$$

Fixed-time concentration To our knowledge, the first fixed-time upper tail concentration result for the empirical variance dates back to Lemma 3 of Honda and Takemura (2014) and to Lemma 7 of ? for fixed-time lower tail concentration results. Corollary 27 is obtained as a direct consequence of Lemmas 21-24-25, hence the proof is omitted.

Corollary 27 Let $h(x) = x - \log(x)$ for $x > 0$. Let σ_{t+1}^2 be the empirical variance of $t + 1$ i.i.d. samples from a Gaussian distribution with variance σ^2 . Then,

$$\begin{aligned} \forall t \geq 1, \forall x > 1, \quad \mathbb{P}(\sigma_{t+1}^2 \geq \sigma^2 x) &\leq \exp \left(-\frac{t}{2} \left(h \left(x + \frac{1}{t} \right) - 1 \right) \right) \\ \forall t \geq 1, \forall x \in \left(0, 1 - \frac{1}{t} \right), \quad \mathbb{P}(\sigma_{t+1}^2 \leq \sigma^2 x) &\leq \exp \left(-\frac{t}{2} \left(h \left(x + \frac{1}{t} \right) - 1 \right) \right). \end{aligned}$$

On sub-Gaussian distributions with unknown variances While it is well-known that (time-uniform) concentration for Gaussian distributions with known σ^2 apply to σ^2 -sub-Gaussian distributions, thus the same bandit algorithms can be used in both settings, we believe that there is no counterpart of this phenomenon when σ^2 is unknown. For regret minimization, some papers have provided examples of sub-Gaussian arms with unknown σ^2 under which the regret can be linear.

Extending our algorithms to sub-Gaussian distributions would require time-uniform concentration results for the empirical variance of sub-Gaussian distributions with unknown variance. Our concentration on the empirical variance relies on Lemma 25, which leverages the fact that the empirical mean and empirical variances of Gaussian distributions are independent, which does not extend to the sub-Gaussian case.

Moreover, if the focus is on asymptotically optimal algorithms, we note that it is difficult to express the characteristic time T^* for the non-parametric class of sub-Gaussian distributions with unknown variances. Optimal BAI has however been studied under other interesting non-parametric assumptions (see, e.g., [Agrawal et al. \(2020\)](#)).

E.2.2. EMPIRICAL MEAN

While time-uniform concentration results for the empirical mean of Gaussian observations already exist in the literature (e.g. [Kaufmann and Koolen \(2021\)](#)), Lemma 28 is proved for completeness and to present unified concentration results as it also involves \overline{W}_{-1} . The empirical mean of a Gaussian after a given number of observations is a sub-Gaussian random variable. It is in fact exactly Gaussian, but the sub-Gaussian hypothesis will be easier to handle for a random number of samples.

Lemma 28 *Let $\overline{W}_{-1}(x) = -W_{-1}(-e^{-x})$ for $x \geq 1$, $\delta \in (0, 1)$, $s > 1$ and ζ be the Riemann ζ function. Let μ_t be the empirical mean of t i.i.d. samples from a Gaussian distribution with parameter (μ, σ^2) . Then, with probability greater than $1 - \delta$, for all $t \in \mathbb{N}^*$,*

$$|\mu_t - \mu| \leq \sqrt{\frac{\sigma^2}{t} \overline{W}_{-1} \left(1 + 2 \log \left(\frac{1}{\delta} \right) + 2g(s) + 2s \log(2s + \log t) \right)},$$

where $g(s) = \log(\zeta(s)) + s(1 - \log(2s))$.

Proof Let $(X_s)_{s \in [t]}$ the observations from a standard normal distributions and denote $S_t = \sum_{s \in [t]} X_s$. We will derive a concentration result for $S_t = t \frac{\hat{\mu}_t - \mu}{\sigma}$, which implies a concentration result for $\hat{\mu}_t$.

Let $\eta > 0$ and $s > 1$. For all $i \in \mathbb{N}^*$, let $\gamma_i > 0$ and $N_i = (1 + \eta)^{i-1}$. For all $i \in \mathbb{N}^*$, we define the family of priors $f_{N_i, \gamma_i}(x) = \sqrt{\frac{\gamma_i N_i}{2\pi}} \exp\left(-\frac{x^2 \gamma_i N_i}{2}\right)$ with weights $w_i = \frac{1}{i^s \zeta(s)}$ and process

$$\overline{M}(t) = \sum_{i \in \mathbb{N}^*} w_i \int f_{N_i, \gamma_i}(x) \exp\left(xS_t - \frac{1}{2}x^2t\right) dx,$$

which satisfies $\overline{M}(0) = 1$ since $\sum_{i \in \mathbb{N}^*} w_i$ and $\int f_{N_i, \gamma_i}(x) dx = 1$. A test martingale is defined as a non-negative martingale of unit initial value. It is direct to see that $M(t) = \exp\left(xS_t - \frac{1}{2}x^2t\right)$ is a test martingale, as $M(0) = 1$ and $\mathbb{E}[M(t) | \mathcal{F}_{t-1}] = M(t-1) \mathbb{E}_{Y \sim \mathcal{N}(0,1)}[e^{xY - \frac{1}{2}x^2}] = M(t-1)$ since $\mathbb{E}_{Y \sim \mathcal{N}(0,1)}[e^{xY}] = e^{\frac{1}{2}x^2}$. By Tonelli's theorem and using that $M(t)$ is a martingale

$$\mathbb{E}[\overline{M}(t) | \mathcal{F}_{t-1}] = \sum_{i \in \mathbb{N}^*} w_i \int f_{N_i, \gamma_i}(x) M(t-1) dx = \overline{M}(t-1).$$

Therefore, $\overline{M}(t)$ is also a test martingale. Let $i \in \mathbb{N}^*$ and consider $t \in [N_i, N_{i+1})$. For all x ,

$$f_{N_i, \gamma_i}(x) \geq \sqrt{\frac{N_i}{t}} f_{t, \gamma_i}(x) \geq \frac{1}{\sqrt{1 + \eta}} f_{t, \gamma_i}(x)$$

Direct computations shows that

$$\int f_{t, \gamma_i}(x) \exp\left(xS_t - \frac{1}{2}x^2t\right) dx = \frac{1}{\sqrt{1 + \gamma_i^{-1}}} \exp\left(\frac{S_t^2}{2(1 + \gamma_i)t}\right).$$

Combining those results with the fact that $\overline{M}(t) \geq w_i \int f_{N_i, \gamma_i}(x) \exp(xS_t - \frac{1}{2}x^2t) dx$, we obtain

$$\overline{M}(t) \geq \frac{1}{i^s \zeta(s)} \frac{1}{\sqrt{(1 + \gamma_i^{-1})(1 + \eta)}} \exp\left(\frac{S_t^2}{2(1 + \gamma_i)t}\right),$$

Using Ville's maximal inequality, we have that with probability greater than $1 - \delta$, $\log \overline{M}(t) \leq \log\left(\frac{1}{\delta}\right)$. Therefore, with probability greater than $1 - \delta$, for all $i \in \mathbb{N}^*$ and $t \in [N_i, N_{i+1})$,

$$\frac{|S_t|}{\sqrt{t}} \leq \sqrt{(1 + \gamma_i) \left(2 \log\left(\frac{1}{\delta}\right) + 2 \log(i^s \zeta(s)) + \log(1 + \gamma_i^{-1}) + \log(1 + \eta)\right)}.$$

Since this upper bound is independent of t , we can optimize it and choose γ_i as in Lemma 76 for all $i \in \mathbb{N}^*$. Therefore, with probability greater than $1 - \delta$, for all $i \in \mathbb{N}^*$ and $t \in [N_i, N_{i+1})$,

$$\begin{aligned} \frac{S_t^2}{t} &\leq \overline{W}_{-1} \left(1 + 2 \log\left(\frac{\zeta(s)}{\delta}\right) + 2s \log(i) + \log(1 + \eta)\right) \\ &\leq \overline{W}_{-1} \left(1 + 2 \log\left(\frac{\zeta(s)}{\delta}\right) + 2s \log(\log(1 + \eta) + \log t) - 2s \log \log(1 + \eta) + \log(1 + \eta)\right) \\ &= \overline{W}_{-1} \left(1 + 2 \log\left(\frac{1}{\delta}\right) + 2s \log(2s + \log t) + 2g(s)\right) \end{aligned}$$

where $g(s) = \log(\zeta(s)) + s(1 - \log(2s))$. The second inequality is obtained since $i \leq 1 + \frac{\log t}{\log(1 + \eta)}$ for $t \in [N_i, N_{i+1})$. The last equality is obtained for the choice $\eta^* = e^{2s} - 1$, which minimizes $\eta \mapsto \log(1 + \eta) - 2s \log(\log(1 + \eta))$. Since $\mathbb{N}^* \subseteq \bigcup_{i \in \mathbb{N}^*} [N_i, N_{i+1})$ and $\frac{\hat{\mu}_t - \mu}{\sigma} = \frac{S_t}{t}$ this yields the result. \blacksquare

Fixed-time concentration Lemma 29 is a known result for the deviation of the empirical mean of Gaussian observations, hence we omit the proof (see Ex. 2.2.23 in Dembo and Zeitouni (1998)).

Lemma 29 *Let μ_t be the empirical mean of t i.i.d. samples from a Gaussian distribution with parameter (μ, σ^2) . Then, for all $t \in \mathbb{N}^*$ and all $x > 0$*

$$\mathbb{P}(\mu_t \geq \mu + x) \leq \exp\left(-\frac{tx^2}{2\sigma^2}\right) \quad \text{and} \quad \forall x > 0, \quad \mathbb{P}(\mu_t \leq \mu - x) \leq \exp\left(-\frac{tx^2}{2\sigma^2}\right).$$

Appendix F. Kullback-Leibler Concentration

We prove our time-uniform upper tail concentration for d -dimensional exponential families (Theorem 42). The two key novelties compared to previous work Degenne (2019) are that we consider: (1) a sum over $\mathcal{S} \subseteq [K]$ arms and (2) $\Theta_D \subseteq \mathbb{R}^d$. To go from one arm to \mathcal{S} arms, it is enough to consider the product of the the mixture of martingales used for each arm. Dealing with a support different from \mathbb{R}^d is the real challenge of the proof. This is the reason why we start by showing the result for one arm in Appendix F.1, and then generalize it to \mathcal{S} arms in Appendix F.2. For Gaussian with unknown variance, which are a 2-dimensional exponential family with support $\Theta_D = \mathbb{R} \times \mathbb{R}_+^*$, we obtain Theorem 43 (Appendix F.3).

In the following, each arm $a \in [K]$ has a parameter θ_a and distribution ν_{θ_a} belonging to an exponential family with parameter space $\Theta_a \subseteq \mathbb{R}^d$ (we call d the dimension of the family), sufficient statistic $F_a : \mathbb{R} \rightarrow \mathbb{R}^d$ and log-partition function $\phi_a : \Theta_a \rightarrow \mathbb{R}$. That is, there exists a distribution ν_0 such that ν_{θ_a} is defined by $\frac{d\nu_{\theta_a}}{d\nu_0}(X) = \exp(\theta_a^\top F_a(X) - \phi_a(\theta_a))$.

We define the average statistic $F_{t,a} := \frac{1}{N_{t,a}} \sum_{s=1}^t \mathbf{1}\{a_s = a\} F_a(X_s)$. We denote the maximum likelihood estimator (MLE) of θ_a by $\theta_{t,a}$. It is defined as $\theta_{t,a} := \arg \max_{\lambda \in \Theta_a} \lambda^\top F_{t,a} - \phi_a(\lambda)$ (and may not exist). When $F_{t,a} \in \nabla \phi_a(\Theta_a)$, we have $\theta_{t,a} = (\nabla \phi_a)^{-1}(F_{t,a})$.

Let $d_{\phi_a}(\theta_a, \lambda_a)$ denote the Bregman divergence of ϕ_a between parameters θ_a and λ_a . It is equal to the Kullback-Leibler divergence between the distributions with parameters λ_a and θ_a (note the reversed order of the parameters). For a subset of arms $\mathcal{S} \subseteq [K]$, we seek high probability bounds on $\sum_{a \in \mathcal{S}} N_{t,a} d_{\phi_a}(\theta_a, \theta_{t,a})$.

F.1. One Arm

In this section we consider one arm $a \in [K]$. The first step of the proof consists in linking the KL divergence to a mixture of martingales in order to obtain a time-uniform upper tail concentration (Lemma 30).

Lemma 30 *Let $\rho_{0,a}$ be a distribution supported on $\Theta_{D,a}$ and τ be an almost surely bounded stopping time. With probability $1 - \delta$, either $F_{N_{\tau,a},a} \notin \nabla \phi_a(\Theta_{D,a})$ or*

$$N_{\tau,a} d_{\phi_a}(\theta_a, \theta_{N_{\tau,a},a}) \leq -\log \mathbb{E}_{y \sim \rho_{0,a}} \exp(-N_{\tau,a} d_{\phi_a}(y, \theta_{N_{\tau,a},a})) + \log \frac{1}{\delta}.$$

Proof We first remark that for all $y \in \Theta_{D,a}$, the log-likelihood ratio $\prod_{s \leq t, a_s = a} \frac{d\nu_y}{d\nu_{\theta_a}}(X_s)$ is a martingale with expectation 1 under ν_{θ_a} . This is also true for $\mathbb{E}_{y \sim \rho_{0,a}} \left[\prod_{s \leq t, a_s = a} \frac{d\nu_y}{d\nu_{\theta_a}}(X_s) \right]$. By the optional stopping theorem, its stopped version at τ also has expectation 1. Then by Markov's inequality, with probability $1 - \delta$,

$$\mathbb{E}_{y \sim \rho_{0,a}} \left[\prod_{s \leq \tau, a_s = a} \frac{d\nu_y}{d\nu_{\theta_a}}(X_s) \right] \leq \frac{1}{\delta}. \quad (15)$$

We now rewrite the product as $\exp\left(\sum_{s=1}^{\tau} \mathbb{I}\{a_s = a\} \log \frac{d\nu_y}{d\nu_{\theta_a}}(X_s)\right)$ and detail its value:

$$\begin{aligned} \sum_{s=1}^t \mathbb{I}\{a_s = a\} \log \frac{d\nu_y}{d\nu_{\theta_a}}(X_s) &= \sum_{s=1}^t \mathbb{I}\{a_s = a\} (\phi_a(\theta_a) - \phi_a(y) - (\theta_a - y)^\top F_a(X_s)) \\ &= N_{t,a} (\phi_a(\theta_a) - \phi_a(y) - (\theta_a - y)^\top F_{N_{t,a},a}). \end{aligned}$$

Now if $F_{N_{\tau,a},a} \in \nabla \phi_a(\Theta_{D,a})$, we have $F_{N_{t,a},a} = \nabla \phi_a(\theta_{N_{t,a},a})$ and

$$\begin{aligned} \sum_{s=1}^t \mathbb{I}\{a_s = a\} \log \frac{d\nu_y}{d\nu_{\theta_a}}(X_s) &= N_{t,a} (\phi_a(\theta_a) - \phi_a(y) - (\theta_a - y)^\top \nabla \phi_a(\theta_{N_{t,a},a})) \\ &= N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) - N_{t,a} d_{\phi_a}(y, \theta_{N_{t,a},a}). \end{aligned}$$

We can write (15) again with that expression to get

$$\mathbb{E}_{y \sim \rho_{0,a}} \exp(N_{\tau,a} d_{\phi_a}(\theta_a, \theta_{N_{\tau,a},a}) - N_{\tau,a} d_{\phi_a}(y, \theta_{N_{\tau,a},a})) \leq \frac{1}{\delta}.$$

The first term in the subtraction does not depend on y and can be brought outside of the expectation. Taking the logarithm then proves the lemma. \blacksquare

Corollary 31 generalizes Lemma 30 to all times.

Corollary 31 *Let $\rho_{0,a}$ be a distribution supported on $\Theta_{D,a}$. With probability $1 - \delta$, for all $t \in \mathbb{N}$, either $F_{N_{t,a},a} \notin \nabla \phi_a(\Theta_{D,a})$ or*

$$N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) \leq -\log \mathbb{E}_{y \sim \rho_{0,a}} \exp(-N_{t,a} d_{\phi_a}(y, \theta_{N_{t,a},a})) + \log \frac{1}{\delta}.$$

Proof The extension of the result of Lemma 30 from a stopping time to all times is standard. \blacksquare

We will now use Corollary 31 with a truncated Gaussian prior (Lemma 32). The reason for the truncation is that the domain $\Theta_{D,a}$ may not be \mathbb{R}^d but only a subset. In that case, we have to restrict the prior. We also truncate to ensure a control the hessian $\nabla^2 \phi_a$ on the support with positive semi-definite upper and lower bound.

For M a positive-definite matrix and a set S , let $V_M(S)$ be the volume of set S according to measure $\mathcal{N}(0, M^{-1})$. Note that $V_M(S) = V_{I_d}(M^{1/2}S)$.

Lemma 32 *Let A be a convex set containing 0, H and G be upper bounds and lower bounds of $\nabla^2 \phi$ in the positive semi-definite (PSD) sense on $\theta_a + A$. Let M_0 be a positive definite matrix. For $n \in \mathbb{N}$, let $\xi_{n,a} = ((nH)^{-1} + M_0^{-1})^{-1}(nH\theta_n + M_0\theta_a)$. With probability $1 - \delta$, for all $t \in \mathbb{N}$, if $F_{N_{t,a},a} \in \nabla \phi_a(\Theta_{D,a})$ and $\theta_{N_{t,a},a} \in \theta_a + A$ then*

$$\begin{aligned} N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) &\leq \frac{1}{2} \|\theta_a - \theta_{N_{t,a},a}\|_{((N_{t,a}H)^{-1} + M_0^{-1})^{-1}}^2 + \frac{1}{2} \log \det(1 + N_{t,a}HM_0^{-1}) \\ &\quad + \log \frac{1}{\delta} + \log \frac{V_{M_0}(A)}{V_{N_{t,a}H + M_0}(\theta_a - \xi_{N_{t,a},a} + A)}. \end{aligned}$$

Proof We use Corollary 31 with prior $\rho_{0,a}$ equal to $\mathcal{N}(\theta_a, M_0^{-1})$ truncated to A and rescaled to have mass 1.

Suppose that $F_{N_{t,a},a} \in \nabla \phi_a(\Theta_{D,a})$ and $\theta_{N_{t,a},a} \in \theta_a + A$. For all $y \in \theta_a + A$, since $\theta_{N_{t,a},a} \in \theta_a + A$ as well and A is convex, we can conclude from the definition of H that $d_{\phi_a}(y, \theta_{N_{t,a},a}) \leq \frac{1}{2} \|y - \theta_{N_{t,a},a}\|_H^2$. We use this to compute an upper bound to the expectation in Corollary 31.

$$\begin{aligned} \mathbb{E}_{y \sim \rho_{0,a}} [\exp(-nd_{\phi_a}(y, \theta_n))] &\geq \mathbb{E}_{y \sim \rho_{0,a}} \left[\exp\left(-\frac{1}{2} \|y - \theta_n\|_{nH}^2\right) \right] \\ &= \frac{\sqrt{\det(M_0)}}{(2\pi)^{d/2} V_{M_0}(A)} \int_{y \in \mathbb{R}^d} \mathbf{1}_{\theta+A} \exp\left(-\frac{1}{2} \|y - \theta\|_{M_0}^2 - \frac{1}{2} \|y - \theta_n\|_{nH}^2\right) dy \end{aligned}$$

Let $\xi_n = (nH + M_0)^{-1}(nH\theta_n + M_0\theta_a)$ and $C_n = \frac{1}{2} \|\theta_a - \theta_n\|_{((nH)^{-1} + M_0^{-1})^{-1}}^2$. Remark that

$$-\frac{1}{2} \|y - \theta\|_{M_0}^2 - \frac{1}{2} \|y - \theta_n\|_{nH}^2 = -\frac{1}{2} \|y - \xi_n\|_{nH + M_0}^2 - C_n.$$

We use this equality:

$$\begin{aligned} \mathbb{E}_{y \sim \rho_{0,a}} [\exp(-nd_\phi(y, \theta_n))] &\geq \frac{\sqrt{\det(M_0)}}{(2\pi)^{d/2} V_{M_0}(A)} e^{-C_n} \int_{y \in \mathbb{R}^d} \mathbf{1}_{\theta_a + A} \exp\left(-\frac{1}{2} \|y - \xi_n\|_{nH+M_0}^2\right) dy \\ &= \frac{\sqrt{\det(M_0)}}{\sqrt{\det(nH + M_0)}} e^{-C_n} \frac{V_{nH+M_0}(\theta_a - \xi_n + A)}{V_{M_0}(A)}. \end{aligned}$$

The result of the lemma is obtained by taking the logarithm of this expression. \blacksquare

Lemma 33 starts by choosing for H and G multiples of the identity, hence we will compare the divergence d_{ϕ_a} with the euclidean distance $(x, y) \mapsto \|x - y\|^2/2$. This choice is possible by using the maximal and minimal eigenvalues of $\nabla^2 \phi_a(\lambda)$ on A , which we denote by

$$\lambda_{+,a} = \max \left\{ \lambda_+ \left(\nabla^2 \phi_a(\tilde{\theta}_a) \right) \mid \tilde{\theta}_a \in \theta_a + A \right\}, \quad \lambda_{-,a} = \min \left\{ \lambda_- \left(\nabla^2 \phi_a(\tilde{\theta}_a) \right) \mid \tilde{\theta}_a \in \theta_a + A \right\}.$$

For $v > 0$ and a set S , let $V_v(S)$ be the volume of set S according to measure $\mathcal{N}(0, v^{-1}I_d)$.

Lemma 33 *Let A be a convex set containing 0 , $\lambda_{+,a}$ and $\lambda_{-,a}$ be the maximum and minimum of the eigenvalues of $\nabla^2 \phi_a(\lambda)$ on $\theta_a + A$ and let $v_0 > 0$. For $n \in \mathbb{N}$, let $\kappa_n = \frac{v_0}{n\lambda_{+,a} + v_0}$. With probability $1 - \delta$, for all $t \in \mathbb{N}$, if $F_{N_{t,a},a} \in \nabla \phi_a(\Theta_{D,a})$ and $\theta_{N_{t,a},a} \in \theta_a + A$ then*

$$\begin{aligned} N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) &\leq \frac{1}{2} \|\theta_a - \theta_{N_{t,a},a}\|^2 \frac{1}{(N_{t,a} \lambda_{+,a})^{-1} + v_0^{-1}} + \frac{d}{2} \log(1 + N_{t,a} \lambda_{+,a} v_0^{-1}) \\ &\quad + \log \frac{1}{\delta} + \log \frac{V_{v_0}(A)}{V_{N_{t,a} \lambda_{+,a} + v_0}(\kappa_{N_{t,a}} A)}. \end{aligned}$$

Note that since $\rho_{0,a}$ is a prior and has to ensure the martingale property used in the proof of Lemma 30, it cannot depend on the observations and in particular it cannot depend on the random variable $N_{t,a}$.

Proof We use Corollary 31 with prior $\rho_{0,a}$ equal to $\mathcal{N}(\theta_a, v_0^{-1}I_d)$ truncated to A and rescaled to have mass 1.

Suppose that $F_{N_{t,a},a} \in \nabla \phi_a(\Theta_{D,a})$ and $\theta_{N_{t,a},a} \in \theta_a + A$. As in Lemma 32, we obtain

$$\mathbb{E}_{y \sim \rho_{0,a}} \left[e^{-nd_\phi(y, \theta_n)} \right] \geq \frac{v_0^{d/2}}{(2\pi)^{d/2} V_{v_0}(A)} e^{-C_n} \int_{y \in \mathbb{R}^d} \mathbf{1}_{\theta_a + A} \exp\left(-\frac{n\lambda_{+,a} + v_0}{2} \|y - \xi_n\|^2\right) dy.$$

We define $\kappa_n = \frac{v_0}{n\lambda_{+,a} + v_0}$ and show that $\xi_n + \kappa_n A \subseteq \theta_a + A$. We can rewrite $\xi_n = \kappa_n \theta_a + (1 - \kappa_n) \theta_n$. Let then $y \in \xi_n + \kappa_n A$. We need to prove that $y - \theta_a \in A$, but $y - \theta_a = y - \xi_n + (1 - \kappa_n)(\theta_n - \theta_a)$. By hypothesis $y - \xi_n \in \kappa_n A$ and $(1 - \kappa_n)(\theta_n - \theta_a) \in (1 - \kappa_n)A$. By convexity of A , we conclude that $\kappa_n A + (1 - \kappa_n)A \subseteq A$ and $y \in \theta_a + A$.

We use that set inclusion to lower bound the integral over $\theta_a + A$ by the same integral over the subset $\xi_n + \kappa_n A$.

$$\begin{aligned} \mathbb{E}_{y \sim \rho_{0,a}} \left[e^{-nd_\phi(y, \theta_n)} \right] &\geq \frac{v_0^{d/2}}{(2\pi)^{d/2} V_{v_0}(A)} e^{-C_n} \int_{y \in \mathbb{R}^d} \mathbf{1}_{\xi_n + \kappa_n A} \exp \left(-\frac{n\lambda_{+,a} + v_0}{2} \|y - \xi_n\|^2 \right) dy \\ &= \frac{v_0^{d/2} e^{-C_n}}{(n\lambda_{+,a} + v_0)^{d/2} V_{v_0}(A)} \frac{(n\lambda_{+,a} + v_0)^{d/2}}{(2\pi)^{d/2}} \int_{y \in \mathbb{R}^d} \mathbf{1}_{\xi_n + \kappa_n A} \exp \left(-\frac{n\lambda_{+,a} + v_0}{2} \|y - \xi_n\|^2 \right) dy \\ &= e^{-C_n} \frac{v_0^{d/2}}{(n\lambda_{+,a} + v_0)^{d/2}} \frac{V_{n\lambda_{+,a} + v_0}(\kappa_n A)}{V_{v_0}(A)}. \end{aligned}$$

The result of the lemma is obtained by taking the logarithm of this expression. ■

Lemma 34 gives a bound on the ratio of volumes.

Lemma 34 *Let $\lambda, n, v > 0$ and $\kappa_n = \frac{v}{n\lambda + v}$. Then*

$$\frac{V_v(A)}{V_{n\lambda + v}(\kappa_n A)} \leq \kappa_n^{-d/2}.$$

Proof We first prove that for all $\alpha \geq 1$, $\frac{V_1(\alpha A)}{V_1(A)} \leq \alpha^d$. For a convex set $S \subseteq \mathbb{R}^d$ containing 0, let $\gamma_S(x) = \arg \max\{r \geq 0 \mid rx \in S\}$. We have that for all $y \in \mathbb{R}^d$ and $\beta \in \mathbb{R}^*$, $\gamma_S(\beta y)\beta y = \gamma_S(y)y$. Note also that $\gamma_S(x) \geq 1 \iff x \in S$.

$$\begin{aligned} \frac{V_1(\alpha A)}{V_1(A)} &= 1 + \frac{V_1(\alpha A \setminus A)}{V_1(A)} = 1 + \frac{\int_{x \in \mathbb{R}^d} \mathbf{1}_{\alpha A \setminus A}(x) \exp \left(-\frac{1}{2} \|x\|_2^2 \right) dx}{\int_{x \in \mathbb{R}^d} \mathbf{1}_A(x) \exp \left(-\frac{1}{2} \|x\|_2^2 \right) dx} \\ &\leq 1 + \frac{\int_{x \in \mathbb{R}^d} \mathbf{1}_{\alpha A \setminus A}(x) \exp \left(-\frac{1}{2} \|\gamma_A(x)x\|_2^2 \right) dx}{\int_{x \in \mathbb{R}^d} \mathbf{1}_A(x) \exp \left(-\frac{1}{2} \|\gamma_A(x)x\|_2^2 \right) dx} \\ &= \frac{\int_{x \in \mathbb{R}^d} \mathbf{1}_{\alpha A}(x) \exp \left(-\frac{1}{2} \|\gamma_A(x)x\|_2^2 \right) dx}{\int_{x \in \mathbb{R}^d} \mathbf{1}_A(x) \exp \left(-\frac{1}{2} \|\gamma_A(x)x\|_2^2 \right) dx} \\ &= \frac{\alpha^d \int_{y \in \mathbb{R}^d} \mathbf{1}_A(y) \exp \left(-\frac{1}{2} \|\gamma_A(\alpha y)\alpha y\|_2^2 \right) dy}{\int_{x \in \mathbb{R}^d} \mathbf{1}_A(x) \exp \left(-\frac{1}{2} \|\gamma_A(x)x\|_2^2 \right) dx} \\ &= \frac{\alpha^d \int_{y \in \mathbb{R}^d} \mathbf{1}_A(y) \exp \left(-\frac{1}{2} \|\gamma_A(y)y\|_2^2 \right) dy}{\int_{x \in \mathbb{R}^d} \mathbf{1}_A(x) \exp \left(-\frac{1}{2} \|\gamma_A(x)x\|_2^2 \right) dx} = \alpha^d. \end{aligned}$$

Using that $\frac{\sqrt{v}}{\sqrt{n\lambda + v\kappa_n}} = \kappa_n^{-1/2} \geq 1$, we obtain the result

$$\frac{V_v(A)}{V_{n\lambda + v}(\kappa_n A)} = \frac{V_1(\sqrt{v}A)}{V_1(\sqrt{n\lambda + v\kappa_n}A)} \leq \left(\frac{\sqrt{v}}{\sqrt{n\lambda + v\kappa_n}} \right)^d = \kappa_n^{-d/2}.$$

Combining Lemma 33 and Lemma 34 yield Corollary 35. ■

Corollary 35 *Let A be a convex set containing θ , $\lambda_{+,a}$ and $\lambda_{-,a}$ be the maximum and minimum of the eigenvalues of $\nabla^2 \phi_a(\lambda)$ on $\theta_a + A$ and let $v_0 > 0$. For $n \in \mathbb{N}$, let $\kappa_n = \frac{v_0}{n\lambda_{+,a} + v_0}$. With probability $1 - \delta$, for all $t \in \mathbb{N}$, if $F_{N_{t,a},a} \in \nabla \phi_a(\Theta_{D,a})$ and $\theta_{N_{t,a},a} \in \theta_a + A$ then*

$$\begin{aligned} N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) &\leq \frac{1}{2} \|\theta_a - \theta_{N_{t,a},a}\|^2 \frac{1}{(N_{t,a}\lambda_{+,a})^{-1} + v_0^{-1}} + \frac{d}{2} \log(1 + N_{t,a}\lambda_{+,a}v_0^{-1}) \\ &\quad + \log \frac{1}{\delta} + \frac{d}{2} \log \kappa_{N_{t,a}}^{-1}. \end{aligned}$$

Proof This is simply Lemma 33 combined with the bound on the ratio of volumes of Lemma 34. ■

We will now choose $\gamma > 1$ and define values $n_i = \gamma^i$ for $i \in \mathbb{N}$. For each interval $[n_i, n_{i+1})$, we will choose a good value of the variance v_0^{-1} to get a bound for $N_{t,a}$ in this interval. The reason we restrict to an interval is that the best choice for that variance scales with $N_{t,a}$, but we are not allowed to do that since the prior cannot depend on $N_{t,a}$. Thanks to this smart choice, we can control the KL divergence on this interval (Lemma 36).

The ratio of the eigenvalues represents the approximation error in considering that d_{ϕ_a} is locally quadratic. It yields that $\lambda_{-,a}I_d \leq \nabla^2 \phi_a(\lambda) \leq \lambda_{+,a}I_d$ on the set, which for $x, y \in \Theta_a$ translates to

$$\lambda_{-,a} \leq d_{\phi_a}(x, y) / (\|x - y\|^2 / 2) \leq \lambda_{+,a} \quad .$$

Lemma 36 *Let A be a convex set containing θ , $\lambda_{+,a}$ and $\lambda_{-,a}$ be the maximum and minimum of the eigenvalues of $\nabla^2 \phi_a(\lambda)$ on $\theta_a + A$ and let $\eta > 0$, $i \in \mathbb{N}$. With probability $1 - \delta$, for all $t \in \mathbb{N}$, if $F_{N_{t,a},a} \in \nabla \phi_a(\Theta_{D,a})$, $N_{t,a} \in [n_i, n_{i+1})$ and $\theta_{N_{t,a},a} \in \theta_a + A$ then*

$$N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) \leq (1 + \eta) \left(d \log(1 + \frac{1}{\eta}) + d \log(\gamma \frac{\lambda_{+,a}}{\lambda_{-,a}}) + \log \frac{1}{\delta} \right).$$

Proof We assume in this proof that the three conditions of the lemma are true.

Choose $v_{0,i}^{-1} = \frac{1}{\lambda_{-,a}n_i} ((1 + \frac{1}{\eta}) - \frac{\lambda_{-,a}}{\lambda_{+,a}})$. Then for $N_{t,a} \in [n_i, n_{i+1})$,

$$\begin{aligned} \frac{1}{(N_{t,a}\lambda_{+,a})^{-1} + v_{0,i}^{-1}} &\leq \frac{1}{(N_{t,a}\lambda_{+,a})^{-1} + \frac{1}{\lambda_{-,a}N_{t,a}} ((1 + \frac{1}{\eta}) - \frac{\lambda_{-,a}}{\lambda_{+,a}})} = N_{t,a}\lambda_{-,a} (1 - \frac{1}{1 + \eta}), \\ \log(1 + N_{t,a}\lambda_{+,a}v_{0,i}^{-1}) &\leq \log(1 + \gamma n_i \lambda_{+,a} v_{0,i}^{-1}) \\ &= \log(1 - \gamma + \gamma \frac{\lambda_{+,a}}{\lambda_{-,a}} (1 + \frac{1}{\eta})) \leq \log(1 + \frac{1}{\eta}) + \log(\gamma \frac{\lambda_{+,a}}{\lambda_{-,a}}). \end{aligned}$$

With that value of $v_{0,i}$, we also have

$$\begin{aligned} \kappa_{N_{t,a}}^{-1} &= 1 + N_{t,a}\lambda_{+,a}v_{0,i}^{-1} \leq 1 + \gamma n_i \lambda_{+,a} v_{0,i}^{-1} = 1 - \gamma + \gamma \frac{\lambda_{+,a}}{\lambda_{-,a}} (1 + \frac{1}{\eta}) \\ &\leq \gamma \frac{\lambda_{+,a}}{\lambda_{-,a}} (1 + \frac{1}{\eta}). \end{aligned}$$

We use those inequalities in Corollary 35 to get

$$\begin{aligned} N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) &\leq \frac{1}{2} \|\theta_a - \theta_{N_{t,a},a}\|^2 N_{t,a}\lambda_{-,a} (1 - \frac{1}{1 + \eta}) + d \log(1 + \frac{1}{\eta}) + d \log(\gamma \frac{\lambda_{+,a}}{\lambda_{-,a}}) + \log \frac{1}{\delta}. \end{aligned}$$

Thanks to the fact that $\lambda_{-,a}$ is a lower bound on the eigenvalues of $\nabla^2 \phi_a$ on $\theta_a + A$, we have $\frac{1}{2} \|\theta_a - \theta_{N_{t,a},a}\|^2 \lambda_{-,a} \leq d_{\phi_a}(\theta_a, \theta_{N_{t,a},a})$. We group the divergences on the left of the inequality to finally obtain

$$N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) \leq (1 + \eta) \left(d \log\left(1 + \frac{1}{\eta}\right) + d \log\left(\gamma \frac{\lambda_{+,a}}{\lambda_{-,a}}\right) + \log \frac{1}{\delta} \right). \quad \blacksquare$$

A good choice of η yields Lemma 37.

Lemma 37 *Let A be a convex set containing 0 , $\lambda_{+,a}$ and $\lambda_{-,a}$ be the maximum and minimum of the eigenvalues of $\nabla^2 \phi_a$ on $\theta_a + A$ and let $i \in \mathbb{N}$. With probability $1 - \delta$, for all $t \in \mathbb{N}$, if $F_{N_{t,a},a} \in \nabla \phi_a(\Theta_{D,a})$, $N_{t,a} \in [n_i, n_{i+1}]$ and $\theta_{N_{t,a},a} \in \theta_a + A$ then*

$$N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) \leq d \overline{W}_{-1} \left(1 + \frac{1}{d} \log \frac{1}{\delta} + \log\left(\gamma \frac{\lambda_{+,a}}{\lambda_{-,a}}\right) \right).$$

Proof Optimize η in Lemma 36 by using Lemma 76. \(\blacksquare\)

Considering a union bound over the intervals and applying Lemma 37, we obtain Lemma 38. Therefore, we have obtained a time-uniform upper tail concentration of the KL divergence of a d -dimensional exponential family with support $\Theta_D \subseteq \mathbb{R}^d$.

Lemma 38 *Let $(A_a(n, \delta))_{n \in \mathbb{N}}$ be a sequence of non-increasing convex sets containing 0 , $\lambda_{+,a}(n, \delta)$ and $\lambda_{-,a}(n, \delta)$ be the maximum and minimum of the eigenvalues of $\nabla^2 \phi_a(\lambda)$ on $\theta_a + A_a(n, \delta)$. Let $s > 1$ and let ζ be the Riemann ζ function. For all t , let $i(t) = \lfloor \log_\gamma N_{t,a} \rfloor$. With probability $1 - \delta$, for all $t \in \mathbb{N}$, if $F_{N_{t,a},a} \in \nabla \phi_a(\Theta_{D,a})$ and $\theta_{N_{t,a},a} \in \theta_a + A_a(N_{t,a}, \delta)$, then*

$$N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) \leq d \overline{W}_{-1} \left(1 + \log \frac{\zeta(s)}{\delta} + \frac{s}{d} \log(1 + \log_\gamma N_{t,a}) + \log \left(\gamma \frac{\lambda_{+,a}(n_{i(t)}, \delta)}{\lambda_{-,a}(n_{i(t)}, \delta)} \right) \right).$$

Proof Suppose that $F_{N_{t,a},a} \in \nabla \phi_a(\Theta_{D,a})$.

For all $i \in \mathbb{N}$, Lemma 37 gives that with probability $1 - \frac{\delta}{\zeta(s)(i+1)^s}$, if $N_{t,a} \in [n_i, n_{i+1}]$ and $\theta_{N_{t,a},a} \in \theta_a + A_a(n_i, \delta)$, then

$$N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) \leq d \overline{W}_{-1} \left(1 + \frac{1}{d} \log \frac{\zeta(s)(i+1)^s}{\delta} + \log\left(\gamma \frac{\lambda_{+,a}(n_i, \delta)}{\lambda_{-,a}(n_i, \delta)}\right) \right).$$

With probability $1 - \delta$, this is true for all $i \in \mathbb{N}$. In particular we have the inequality for $i(t)$: if $\theta_{N_{t,a},a} \in \theta_a + A_a(n_{i(t)}, \delta)$, then

$$N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) \leq d \overline{W}_{-1} \left(1 + \frac{1}{d} \log \frac{\zeta(s)(i(t)+1)^s}{\delta} + \log\left(\gamma \frac{\lambda_{+,a}(n_{i(t)}, \delta)}{\lambda_{-,a}(n_{i(t)}, \delta)}\right) \right).$$

Since the sets $A_a(n, \delta)$ are non-increasing, a sufficient condition for $\theta_{N_{t,a},a} \in \theta_a + A_a(n_{i(t)}, \delta)$ is $\theta_{N_{t,a},a} \in \theta_a + A_a(N_{t,a}, \delta)$. By definition of $i(t)$, it satisfies $i(t) \leq \log_\gamma N_{t,a}$. We get that if $\theta_{N_{t,a},a} \in \theta_a + A_a(N_{t,a}, \delta)$, then

$$N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) \leq d \overline{W}_{-1} \left(1 + \frac{1}{d} \log \frac{\zeta(s)}{\delta} + \frac{s}{d} \log(1 + \log_\gamma N_{t,a}) + \log\left(\gamma \frac{\lambda_{+,a}(n_{i(t)}, \delta)}{\lambda_{-,a}(n_{i(t)}, \delta)}\right) \right).$$

■

To obtain Theorem 42, it remains to extend this argument to a sum of $\mathcal{S} \subseteq [K]$ arms.

F.2. Sum Over Arms

In this section we consider a subset of arms $\mathcal{S} \subseteq [K]$. Lemma 39 extends Corollary 31 to the sum over \mathcal{S} arms.

Lemma 39 *Let $(\rho_{0,a})_{a \in \mathcal{S}}$ be distributions each supported on $\Theta_{D,a}$. With probability $1 - \delta$, for all $t \in \mathbb{N}$, either there exists $a \in \mathcal{S}$ such that $F_{N_{t,a},a} \notin \nabla \phi_a(\Theta_{D,a})$ or*

$$\sum_{a \in \mathcal{S}} N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) \leq - \sum_{a \in \mathcal{S}} \log \mathbb{E}_{y \sim \rho_{0,a}} \exp(-N_{t,a} d_{\phi_a}(y, \theta_{N_{t,a},a})) + \log \frac{1}{\delta}.$$

Proof The proof is the same as the one of Lemma 30 and Corollary 31, except that the likelihood ratio martingale considered contains observations of all arms in \mathcal{S} . ■

Let $\gamma > 1$. We define $n_i = \gamma^i$ for $i \in \mathbb{N}$. Lemma 40 extends Lemma 36 the sum over \mathcal{S} arms.

Lemma 40 *Let $(A_a)_{a \in \mathcal{S}}$ be convex sets containing 0, $\lambda_{+,a}$ and $\lambda_{-,a}$ be the maximum and minimum of the eigenvalues of $\nabla^2 \phi_a(\lambda)$ on A_a and let $\eta > 0$, $(i^a)_{a \in \mathcal{S}} \in \mathbb{N}^{|\mathcal{S}|}$. With probability $1 - \delta$, for all $t \in \mathbb{N}$, if for all $a \in \mathcal{S}$, $F_{N_{t,a},a} \in \nabla \phi_a(\Theta_{D,a})$, $N_{t,a} \in [n_{i^a}, n_{i^a+1})$ and $\theta_{N_{t,a},a} \in \theta_a + A_a$ then*

$$\sum_{a \in \mathcal{S}} N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) \leq (1 + \eta) \left(d|\mathcal{S}| \log\left(1 + \frac{1}{\eta}\right) + \log \frac{1}{\delta} + d \sum_{a \in \mathcal{S}} \log\left(\gamma \frac{\lambda_{+,a}}{\lambda_{-,a}}\right) \right).$$

Proof For each arm, we choose a truncated Gaussian prior, as in Lemma 32, and consider upper and lower bound multiples of the identity, as in Lemma 33. Then, we choose the variance of that prior according to the interval, as in Lemma 36. ■

Lemma 41 extends Lemma 37 the sum over \mathcal{S} arms.

Lemma 41 *Let $(A_a)_{a \in \mathcal{S}}$ be convex sets containing 0, $\lambda_{+,a}$ and $\lambda_{-,a}$ be the maximum and minimum of the eigenvalues of $\nabla^2 \phi_a(\lambda)$ on $\theta_a + A_a$ and let $(i^a)_{a \in \mathcal{S}} \in \mathbb{N}^{|\mathcal{S}|}$. With probability $1 - \delta$, for all $t \in \mathbb{N}$, if for all $a \in \mathcal{S}$, $F_{N_{t,a},a} \in \nabla \phi_a(\Theta_{D,a})$, $N_{t,a} \in [n_{i^a}, n_{i^a+1})$ and $\theta_{N_{t,a},a} \in \theta_a + A_a$ then*

$$\sum_{a \in \mathcal{S}} N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) \leq d|\mathcal{S}| \overline{W}_{-1} \left(1 + \frac{1}{d|\mathcal{S}|} \log \frac{1}{\delta} + \frac{1}{|\mathcal{S}|} \sum_{a \in \mathcal{S}} \log\left(\gamma \frac{\lambda_{+,a}}{\lambda_{-,a}}\right) \right).$$

Proof Optimize η in Lemma 40 by using Lemma 76. ■

Theorem 42 extends Lemma 38 to the sum of \mathcal{S} arms. Therefore, we have obtained a time-uniform upper tail concentration of the sum of \mathcal{S} KL divergence of d -dimensional exponential family with support $\Theta_{D,a} \subseteq \mathbb{R}^d$.

Theorem 42 For all $a \in \mathcal{S}$, let $(A_a(n, \delta))_{n \in \mathbb{N}}$ be a sequence of non-increasing convex sets containing 0, $\lambda_{+,a}(n, \delta)$ and $\lambda_{-,a}(n, \delta)$ be the maximum and minimum of the eigenvalues of $\nabla^2 \phi_a(\lambda)$ on $\theta_a + A_a(n, \delta)$. Let $s > 1$ and let ζ be the Riemann ζ function. For all t , let $i^a(t) = \lfloor \log_\gamma N_{t,a} \rfloor$. With probability $1 - \delta$, for all $t \in \mathbb{N}$, if for all $a \in \mathcal{S}$, $F_{N_{t,a},a} \in \nabla \phi_a(\Theta_{D,a})$ and $\theta_{N_{t,a},a} \in \theta_a + A_a(N_{t,a}, \delta)$, then

$$\begin{aligned} & \sum_{a \in \mathcal{S}} N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) \\ & \leq d|\mathcal{S}| \overline{W}_{-1} \left(1 + \frac{\log\left(\frac{\zeta(s)^{|\mathcal{S}|}}{\delta}\right)}{d|\mathcal{S}|} + \frac{s}{d|\mathcal{S}|} \sum_{a \in \mathcal{S}} \log(1 + \log_\gamma N_{t,a}) + \frac{1}{|\mathcal{S}|} \sum_{a \in \mathcal{S}} \log\left(\gamma \frac{\lambda_{+,a}(n_{i^a(t)}, \delta)}{\lambda_{-,a}(n_{i^a(t)}, \delta)}\right) \right). \end{aligned}$$

Proof Suppose that for all $a \in \mathcal{S}$, $F_{N_{t,a},a} \in \nabla \phi_a(\Theta_{D,a})$. For all $(i^a)_{a \in \mathcal{S}} \in \mathbb{N}^{|\mathcal{S}|}$, Lemma 41 gives that with probability $1 - \frac{\delta}{\zeta(s)^{|\mathcal{S}|} \prod_{a \in \mathcal{S}} (i^a + 1)^s}$, if for all $a \in \mathcal{S}$, $N_{t,a} \in [n_{i^a}, n_{i^a+1})$ and $\theta_{N_{t,a},a} \in \theta_a + A_a(n_{i^a}, \delta)$, then

$$\begin{aligned} & \sum_{a \in \mathcal{S}} N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) \\ & \leq d|\mathcal{S}| \overline{W}_{-1} \left(1 + \frac{1}{d|\mathcal{S}|} \log \frac{\zeta(s)^{|\mathcal{S}|} \prod_{a \in \mathcal{S}} (i^a + 1)^s}{\delta} + \frac{1}{|\mathcal{S}|} \sum_{a \in \mathcal{S}} \log\left(\gamma \frac{\lambda_{+,a}(n_{i^a}, \delta)}{\lambda_{-,a}(n_{i^a}, \delta)}\right) \right). \end{aligned}$$

With probability $1 - \delta$, this is true for all $(i^a)_{a \in \mathcal{S}} \in \mathbb{N}^{|\mathcal{S}|}$. In particular we have the inequality for $(i^a(t))_{a \in \mathcal{S}}$: if for all $a \in \mathcal{S}$, $\theta_{N_{t,a},a} \in \theta_a + A_a(n_{i^a(t)}, \delta)$, then

$$\begin{aligned} & \sum_{a \in \mathcal{S}} N_{t,a} d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) \\ & \leq d|\mathcal{S}| \overline{W}_{-1} \left(1 + \frac{1}{d|\mathcal{S}|} \log \frac{\zeta(s)^{|\mathcal{S}|} \prod_{a \in \mathcal{S}} (i^a(t) + 1)^s}{\delta} + \frac{1}{|\mathcal{S}|} \sum_{a \in \mathcal{S}} \log\left(\gamma \frac{\lambda_{+,a}(n_{i^a(t)}, \delta)}{\lambda_{-,a}(n_{i^a(t)}, \delta)}\right) \right). \end{aligned}$$

Since the sets $A_a(n, \delta)$ are non-increasing with respect to n , a sufficient condition for $\theta_{N_{t,a},a} \in \theta_a + A_a(n_{i^a(t)}, \delta)$ is $\theta_{N_{t,a},a} \in \theta_a + A_a(N_{t,a}, \delta)$. By definition of $i^a(t)$, it satisfies $i^a(t) \leq \log_\gamma N_{t,a}$. We get the inequality of the theorem. \blacksquare

F.3. Univariate Gaussian

Theorem 43 gives a time-uniform upper tail concentration of the sum of \mathcal{S} KL divergence of Gaussian with unknown variance. It is obtained by using Theorem 42 with adequate pre-concentration results.

Theorem 43 For $i \in \{0, -1\}$, let $\bar{W}_i(x) = -W_i(-e^{-x})$ for $x \geq 1$, $\delta \in (0, 1)$, $\mathcal{S} \subseteq [K]$, $\eta_0 > 0$, $\eta_1 > 0$, $\gamma > 1$, $s > 1$ and ζ be the Riemann ζ function. Let

$$\begin{aligned}\varepsilon_\mu(t, \delta) &= \frac{1}{t} \bar{W}_{-1} \left(1 + 2 \log \left(\frac{6|\mathcal{S}|\zeta(s)}{\delta} \right) + 2s + 2s \log \left(1 + \frac{\log t}{2s} \right) \right), \\ 1 + \varepsilon_{+, \sigma}(t, \delta) &= \bar{W}_{-1} \left(1 + \frac{2(1 + \eta_1)}{t} \left(\log \left(\frac{6|\mathcal{S}|\zeta(s)}{\delta} \right) + s \log \left(1 + \frac{\log(t)}{\log(1 + \eta_1)} \right) \right) \right) - \frac{1}{t}, \\ 1 - \varepsilon_{-, \sigma}(t, \delta) &= \bar{W}_0 \left(1 + \frac{2(1 + \eta_0)}{t} \left(\log \left(\frac{6|\mathcal{S}|\zeta(s)}{\delta} \right) + s \log \left(1 + \frac{\log(t)}{\log(1 + \eta_0)} \right) \right) \right) - \frac{1}{t}.\end{aligned}$$

For all $a \in [K]$ and $t \geq t_a(\delta)$, let $i_a(t) = \lfloor \log_\gamma N_{t,a} \rfloor$ and $\bar{t}_a = \inf \{t \mid N_{t,a} \geq i_a(t)\}$,

$$\begin{aligned}\mu_{\pm, t, a} &= \mu_{\bar{t}_a, a} \pm 2\sigma_{\bar{t}_a, a} \sqrt{\frac{\varepsilon_\mu(N_{\bar{t}_a, a}, \delta)}{1 - \varepsilon_{-, \sigma}(N_{\bar{t}_a, a} - 1, \delta)}}, \quad \sigma_{\pm, t, a}^2 = \sigma_{\bar{t}_a, a}^2 \frac{1 \pm \varepsilon_{\pm, \sigma}(N_{\bar{t}_a, a} - 1, \delta)}{1 \mp \varepsilon_{\mp, \sigma}(N_{\bar{t}_a, a} - 1, \delta)}, \\ \mu_{+, t, a}^2 &= \max \mu_{\pm, t, a}^2, \quad R_{t, a}(\delta) = \frac{\sigma_{+, t, a}^2 f_+ (g(\sigma_{+, t, a}^2, \mu_{+, t, a}^2))}{\sigma_{-, t, a}^2 f_- (g(\sigma_{-, t, a}^2, \mu_{+, t, a}^2))},\end{aligned}$$

where $f_\pm(x) = \frac{1 \pm \sqrt{1-x}}{\sqrt{x}}$ and $g(x, y) = \frac{2x}{(x+2y+\frac{1}{2})^2}$. Then, with probability $1 - \delta$, for all $t \geq \max_{a \in \mathcal{S}} t_a(\delta)$,

$$\begin{aligned}& \sum_{a \in \mathcal{S}} N_{t, a} \text{KL}((\mu_{t, a}, \sigma_{t, a}^2), (\mu_a, \sigma_a^2)) \\ & \leq 2|\mathcal{S}| \bar{W}_{-1} \left(1 + \frac{\log \frac{2\zeta(s)|\mathcal{S}|}{\delta}}{2|\mathcal{S}|} + \frac{s}{2|\mathcal{S}|} \sum_{a \in \mathcal{S}} \log(1 + \log_\gamma N_{t, a}) + \frac{1}{|\mathcal{S}|} \sum_{a \in \mathcal{S}} \log(\gamma R_{t, a}(\delta)) \right).\end{aligned}$$

For all $a \in \mathcal{S}$, the stochastic initial time is defined as

$$t_a(\delta) = \inf \left\{ t \mid N_{t, a} > 1 + \max \left\{ t_\pm(\delta), e^{1+W_0 \left(\frac{2(1+\eta_0)}{e} \left(\log \left(\frac{6|\mathcal{S}|\zeta(s)}{\delta} \right) + s \log \left(1 + \frac{\log(N_{t, a} - 1)}{\log(1 + \eta_0)} \right) \right) - e^{-1} \right)} \right\} \right\}. \quad (16)$$

where $\log(t_-(\delta)) = \frac{s}{\log \left(\frac{6|\mathcal{S}|\zeta(s)}{\delta} \right)} - \log(1 + \eta_0)$ and $\log(t_+(\delta)) = \frac{s}{\log \left(\frac{6|\mathcal{S}|\zeta(s)}{\delta} \right) - \frac{1}{2(1+\eta_1)}} - \log(1 + \eta_1)$.

As $W_0(x) \in [-1, +\infty)$, $t \geq t_a(\delta)$ implies that $N_{t, a} > 2$. Numerically, we always observed $\max t_\pm(\delta) < 2$.

The distribution family \mathcal{D} of Gaussian distributions with positive variance is an exponential family, with natural parameter domain $\Theta_D = \mathbb{R} \times \mathbb{R}_+^*$ and log-partition function $\phi(\theta) = -\frac{\theta_1^2}{4\theta_2} - \frac{1}{2} \log(-2\theta_2)$. The distribution with mean x and variance v corresponds to natural parameters $(\frac{x}{v}, -\frac{1}{2v})$. The distribution with natural parameter θ_a has mean $\mu_a = -\frac{\theta_{a,1}}{2\theta_{a,2}}$ and variance $\sigma_a^2 = -\frac{1}{2\theta_{a,2}}$. With those correspondences, we get $\phi(\theta_a) = \frac{\mu_a^2}{2\sigma_a^2} + \log \sigma_a$ and $\nabla \phi(\theta_a) = (\mu_a, \mu_a^2 + \sigma_a^2)$. The domain $\nabla \phi(\Theta_D)$ is the set $\{(x, y) \in \mathbb{R}^2 \mid y > x^2\}$. Finally, the sufficient statistic for that exponential family is $F(X) = (X, X^2)$. Indeed we have

$$d\nu_{\mu_a, \sigma_a^2}(x) = \frac{1}{\sqrt{2\pi}} e^{-(x-\mu_a)^2/2\sigma_a^2} = \frac{1}{\sqrt{2\pi}} e^{\theta_a^\top (x, x^2) - \phi(\theta_a)}.$$

We define the estimator $F_{t,a} = \frac{1}{N_{t,a}} \sum_{s=1}^t F(X_s) \mathbf{1}\{a_s = a\} = (\mu_{t,a}, \mu_{t,a}^2 + \sigma_{t,a}^2)$. When $\sigma_{t,a}^2 > 0$, $F_{t,a}$ belongs to $\nabla\phi(\Theta_D)$ and the maximum likelihood estimator for θ_a in $\mathbb{R} \times \mathbb{R}_+^*$ is well defined: it is $\theta_{t,a} \stackrel{\text{def}}{=} \nabla\phi^{-1}(F_{t,a})$. With probability 1, we have $\sigma_{t,a}^2 > 0$ if and only if $N_{t,a} > 1$. That maximum likelihood estimator (MLE) has mean $\mu_{t,a}$ and variance $\sigma_{t,a}^2$. Therefore, the condition $F_{N_{t,a},a} \in \nabla\phi_a(\Theta_{D,a})$ will be satisfied almost surely if and only if $N_{t,a} > 1$.

Preliminary concentration We will need non-increasing sequence of convex sets $A_a(N_{t,a}, \delta)$ containing 0 such that with probability $1 - \delta$, for all t greater than an initial time $t_0(\delta)$, the MLE $\theta_{t,a}$ belongs to $\theta_a + A_a(N_{t,a}, \delta)$. Thanks to Lemma 44, the problem is reduced to the task of finding intervals for the mean and variance. The non-increasingness of the sequence of convex sets will be a direct consequence of the monotonicity of the bounds defining the rectangles $[\mu_-, \mu_+] \times [\sigma_-^2, \sigma_+^2]$ for all $a \in \mathcal{S}$ (Lemma 46).

Lemma 44 *A rectangle in mean and variance space maps to a (convex) trapeze in natural parameter space. That is, for $\sigma_-^2 > 0$ the set $[\mu_-, \mu_+] \times [\sigma_-^2, \sigma_+^2]$ represents the same distributions as the natural parameters $\{\tilde{\theta} \in \Theta_D \mid \tilde{\theta}_2 \in [-\frac{1}{2\sigma_-^2}, -\frac{1}{2\sigma_+^2}], \frac{\tilde{\theta}_1}{\tilde{\theta}_2} \in [-2\tilde{\mu}_+, -2\tilde{\mu}_-]\}$. If $(\mu, \sigma^2) \in [\mu_-, \mu_+] \times [\sigma_-^2, \sigma_+^2]$, it rewrites as $\theta + A$, where A is a (convex) trapeze containing 0_2 .*

Proof Given $(\tilde{\mu}, \tilde{\sigma}^2)$, the associated natural parameter is $\tilde{\theta} = \left(\frac{\tilde{\mu}}{\tilde{\sigma}^2}, -\frac{1}{2\tilde{\sigma}^2}\right)$. In particular, for $\theta = (\mu, \sigma^2)$, we define A to be the following trapeze in the natural parameter space

$$\begin{aligned} \theta + A &\stackrel{\text{def}}{=} \left\{ \tilde{\theta} \mid (\tilde{\mu}, \tilde{\sigma}^2) \in [\mu_-, \mu_+] \times [\sigma_-^2, \sigma_+^2] \right\} \\ &= \left\{ \tilde{\theta} \mid \tilde{\theta}_2 \in \left[-\frac{1}{2\sigma_-^2}, -\frac{1}{2\sigma_+^2}\right], \tilde{\theta}_1 \in [-2\mu_- \tilde{\theta}_2, -2\mu_+ \tilde{\theta}_2] \right\}. \end{aligned}$$

where the second equality uses that $\tilde{\theta}_2 < 0$. Since convexity is preserved by translation, A is convex if and only if $\theta + A$ is convex, which is direct by considering the last expression. Likewise, A is a trapeze since $\theta + A$ and geometry is preserved by translation. By assumption, we have $(\mu, \sigma^2) \in [\mu_-, \mu_+] \times [\sigma_-^2, \sigma_+^2]$. Therefore, $\theta \in \theta + A$, in other words A contains 0_2 . Note that $A \subset \left[\frac{\mu_-}{\sigma_+^2}, \frac{\mu_+}{\sigma_-^2}\right] \times \left[-\frac{1}{2\sigma_-^2}, -\frac{1}{2\sigma_+^2}\right]$. \blacksquare

Combining the time-uniform concentration inequalities for the mean (Lemma 28) and the variance (Corollary 26) derived in Appendix E, a direct union bound yields that with probability at least $1 - \frac{\delta}{2}$, for all $a \in \mathcal{S}$ and all $t \geq t_a(\delta)$

$$\begin{aligned} \sigma_a^2 (1 - \varepsilon_{-, \sigma}(N_{t,a} - 1, \delta)) &\leq \sigma_{t,a}^2 \leq \sigma_a^2 (1 + \varepsilon_{+, \sigma}(N_{t,a} - 1, \delta)) , \\ |\mu_{t,a} - \mu_a| &\leq \sigma_a \sqrt{\varepsilon_{\mu}(N_{t,a}, \delta)} , \end{aligned}$$

where, abusing notations compared to Lemma 5 (re-scaling of δ),

$$\begin{aligned} t_a(\delta) &= \inf \left\{ t \mid N_{t,a} > 1 + e^{1+W_0\left(\frac{2(1+\eta_0)}{e}\left(\log\left(\frac{6|\mathcal{S}|\zeta(s)}{\delta}\right)+s\log\left(1+\frac{\log(N_{t,a}-1)}{\log(1+\eta_0)}\right)\right)-e^{-1}\right)} \right\}, \\ \varepsilon_\mu(t, \delta) &= \frac{1}{t} \overline{W}_{-1} \left(1 + 2 \log \left(\frac{6|\mathcal{S}|\zeta(s)}{\delta} \right) + 2s + 2s \log \left(1 + \frac{\log t}{2s} \right) \right), \\ 1 + \varepsilon_{+, \sigma}(t, \delta) &= \overline{W}_{-1} \left(1 + \frac{2(1+\eta_1)}{t} \left(\log \left(\frac{6|\mathcal{S}|\zeta(s)}{\delta} \right) + s \log \left(1 + \frac{\log(t)}{\log(1+\eta_1)} \right) \right) \right) - \frac{1}{t}, \\ 1 - \varepsilon_{-, \sigma}(t, \delta) &= \overline{W}_0 \left(1 + \frac{2(1+\eta_0)}{t} \left(\log \left(\frac{6|\mathcal{S}|\zeta(s)}{\delta} \right) + s \log \left(1 + \frac{\log(t)}{\log(1+\eta_0)} \right) \right) \right) - \frac{1}{t}. \end{aligned}$$

Note that $t \geq t_a(\delta)$ implies that $\sigma_{t,a}^2 > 0$ and $N_{t,a} > 2$ hence $F_{N_{t,a},a} \in \nabla \phi_a(\Theta_{D,a})$. Defining

$$\mu_{\pm,t,a} = \mu_a \pm \sigma_a \sqrt{\varepsilon_\mu(N_{t,a}, \delta)} \quad \text{and} \quad \sigma_{\pm,t,a}^2 = \sigma_a^2 (1 \pm \varepsilon_{\pm, \sigma}(N_{t,a} - 1, \delta)),$$

we have obtain non-decreasing sequences $\mu_{-,t,a}$ and $\sigma_{-,t,a}^2$, and non-increasing sequences $\mu_{+,t,a}$ and $\sigma_{+,t,a}^2$, such that with probability $1 - \frac{\delta}{2}$, for all $a \in \mathcal{S}$ and all $t \geq t_a(\delta)$,

$$(\mu_a, \sigma_a^2) \in [\mu_{-,t,a}, \mu_{+,t,a}] \times [\sigma_{-,t,a}^2, \sigma_{+,t,a}^2] \quad \text{and} \quad (\mu_{t,a}, \sigma_{t,a}^2) \in [\mu_{-,t,a}, \mu_{+,t,a}] \times [\sigma_{-,t,a}^2, \sigma_{+,t,a}^2].$$

For all $a \in \mathcal{S}$, applying Lemma 44 on this non-increasing sequence of rectangles yield a sequence of non-increasing convex sets containing 0_2 , denoted by $(A_a(n, \delta))_{n \in \mathbb{N}}$, such that $\theta_a \in \theta_a + A_a(n, \delta)$. By construction, we have with probability $1 - \frac{\delta}{2}$, for all $a \in \mathcal{S}$ and all $t \geq t_a(\delta)$,

$$\theta_{t,a} \in \theta_a + A_a(N_{t,a}, \delta).$$

This concludes the construction of the preliminary concentration sets used to apply Theorem 42. Doing a union bound (splitting δ in two) and restricting to $t \geq \max_{a \in \mathcal{S}} t_a(\delta)$, we obtain that with probability $1 - \delta$, for all $t \geq \max_{a \in \mathcal{S}} t_a(\delta)$,

$$\begin{aligned} & \sum_{a \in \mathcal{S}} N_{t,a} \text{KL}((\mu_{t,a}, \sigma_{t,a}^2), (\mu_a, \sigma_a^2)) \\ & \leq 2|\mathcal{S}| \overline{W}_{-1} \left(1 + \frac{\log \frac{2\zeta(s)^{|\mathcal{S}|}}{\delta}}{2|\mathcal{S}|} + \frac{s}{2|\mathcal{S}|} \sum_{a \in \mathcal{S}} \log(1 + \log_\gamma N_{t,a}) + \frac{1}{|\mathcal{S}|} \sum_{a \in \mathcal{S}} \log \left(\gamma \frac{\lambda_{+,a}(n_{i^a(t)}, \delta)}{\lambda_{-,a}(n_{i^a(t)}, \delta)} \right) \right). \end{aligned}$$

where we used that $d_{\phi_a}(\theta_a, \theta_{N_{t,a},a}) = \text{KL}((\mu_{t,a}, \sigma_{t,a}^2), (\mu_a, \sigma_a^2))$.

Controlling the eigenvalues While the above is enough to obtain a concentration result, the ratio $\frac{\lambda_{+,a}(n_{i^a(t)}, \delta)}{\lambda_{-,a}(n_{i^a(t)}, \delta)}$ cannot be computed since it depends on (μ, σ^2) . To circumvent this issue, we derive an upper bound on this unknown quantity, which will yield a valid concentration inequality.

Upper bounding this quantity might be done by considering a larger rectangle containing the one used to apply Theorem 42, which can be done by chaining the concentration to replace (μ_a, σ_a^2) by $(\mu_{t,a}, \sigma_{t,a}^2)$. Defining

$$\bar{\mu}_{\pm,t,a} = \mu_{t,a} \pm 2\sigma_{t,a} \sqrt{\frac{\varepsilon_\mu(N_{t,a}, \delta)}{1 - \varepsilon_{-, \sigma}(N_{t,a} - 1, \delta)}} \quad \text{and} \quad \bar{\sigma}_{\pm,t,a}^2 = \sigma_{t,a}^2 \frac{1 \pm \varepsilon_{\pm, \sigma}(N_{t,a} - 1, \delta)}{1 \mp \varepsilon_{\mp, \sigma}(N_{t,a} - 1, \delta)},$$

we have by direct manipulations that

$$[\mu_{-,t,a}, \mu_{+,t,a}] \times [\sigma_{-,t,a}^2, \sigma_{+,t,a}^2] \subset [\bar{\mu}_{-,t,a}, \bar{\mu}_{+,t,a}] \times [\bar{\sigma}_{-,t,a}^2, \bar{\sigma}_{+,t,a}^2].$$

By Lemma 44, there is also inclusion of the associated trapeze in the natural space parameters, hence implying the following ordering of the eigenvalues

$$\lambda_{-,a}(n_{i^a}(t), \delta) \geq \bar{\lambda}_{-,a}(n_{i^a}(t), \delta) \quad \text{and} \quad \lambda_{+,a}(n_{i^a}(t), \delta) \leq \bar{\lambda}_{+,a}(n_{i^a}(t), \delta).$$

Lemma 45 provides a control on the eigenvalues λ_+ and λ_- in the case of Gaussian distributions based on the trapeze obtained in Lemma 44.

Lemma 45 *Let μ_- and μ_+ be real values such that $\mu_- < \mu_+$. Let σ_-^2 and σ_+^2 be real values such that $0 < \sigma_-^2 < \sigma_+^2$. Let Θ be the trapeze corresponding to that set in natural parameter space (see Lemma 44). The minimal and maximal eigenvalues of $\nabla^2 \phi(\lambda)$ on Θ are*

$$\lambda_- \geq \sqrt{2} \sigma_-^3 f_- (g(\sigma_-^2, \mu_{++}^2)), \quad \lambda_+ \leq \sqrt{2} \sigma_+^3 f_+ (g(\sigma_+^2, \mu_{++}^2)),$$

where $\mu_{++}^2 = \max\{\mu_-^2, \mu_+^2\}$, $f_{\pm}(x) = \frac{1 \pm \sqrt{1-x}}{\sqrt{x}}$ and $g(x, y) = \frac{2x}{(x+2y+\frac{1}{2})^2}$.

Proof Recall that $\phi(\theta) = -\frac{\theta_1^2}{4\theta_2} - \frac{1}{2} \log(-2\theta_2)$, therefore we obtain

$$\nabla \phi(\theta) = \begin{bmatrix} -\frac{\theta_1}{2\theta_2} \\ \left(\frac{\theta_1}{2\theta_2}\right) - \frac{1}{2\theta_2} \end{bmatrix} \quad \text{and} \quad \nabla^2 \phi(\theta) = \frac{1}{2\theta_2^2} \begin{bmatrix} -\theta_2 & \theta_1 \\ \theta_1 & 1 - \frac{\theta_1^2}{\theta_2} \end{bmatrix}.$$

Computing the eigenvalues of $\nabla \phi(\theta)$, we have $\lambda_-(\theta)I_2 \leq \nabla^2 \phi(\theta) \leq \lambda_+(\theta)I_2$ for all $\theta \in \Theta_D$, where the values of the eigenvalues expressed with the mean parameters are

$$\lambda_{\pm}(\theta) = \sigma^2 \left(\sigma^2 + 2\mu^2 + \frac{1}{2} \pm \sqrt{\left(\sigma^2 + 2\mu^2 + \frac{1}{2}\right)^2 - 2\sigma^2} \right).$$

Note that $(\sigma^2 + 2\mu^2 + \frac{1}{2})^2 - 2\sigma^2 \geq 0 \iff (\sigma^2)^2 + (4\mu^2 - 1)\sigma^2 + (2\mu^2 + \frac{1}{2})^2$. Since $(4\mu^2 - 1)^2 - 4(2\mu^2 + \frac{1}{2})^2 = -16\mu^2 \leq 0$, we have $(\sigma^2 + 2\mu^2 + \frac{1}{2})^2 - 2\sigma^2 \geq 0$.

Defining $f_{\pm}(x) = \frac{1 \pm \sqrt{1-x}}{\sqrt{x}}$ and $g(x, y) = \frac{2x}{(x+2y+\frac{1}{2})^2}$, we have, for all $\tilde{\theta} \in \theta + A$,

$$\lambda_{\pm}(\tilde{\theta}) = \tilde{\sigma}^2 \left(\tilde{\sigma}^2 + 2\tilde{\mu}^2 + \frac{1}{2} \pm \sqrt{\left(\tilde{\sigma}^2 + 2\tilde{\mu}^2 + \frac{1}{2}\right)^2 - 2\tilde{\sigma}^2} \right) = \sqrt{2} \tilde{\sigma}^3 f_{\pm} (g(\tilde{\sigma}^2, \tilde{\mu}^2)).$$

Direct computations yield, for all $x \in (0, 1)$

$$f'_+(x) = -\frac{1 + \sqrt{1-x}}{2x\sqrt{x(1-x)}} < 0 \quad \text{and} \quad f'_-(x) = \frac{1 - \sqrt{1-x}}{2x\sqrt{x(1-x)}} > 0,$$

hence f_+ is decreasing on $(0, 1)$ and f_- is increasing on $(0, 1)$. For all $(x, y) \in \mathbb{R}^+ \times \mathbb{R}^+$

$$\partial_y g(x, y) = -\frac{8x}{(x + 2y + \frac{1}{2})^3} \leq 0,$$

hence $y \mapsto g(x, y)$ is decreasing for all $x \in \mathbb{R}^+$. Let $\mu_{++}^2 = \max\{\mu_-^2, \mu_+^2\}$. By composition rule and $\lambda_{\pm}(\tilde{\theta}) = \sqrt{2}\tilde{\sigma}^3 f_{\pm}(g(\tilde{\sigma}^2, \tilde{\mu}^2))$, we obtain that: for all $\tilde{\theta} \in \theta + A(n, \delta)$

$$\lambda_+(\tilde{\theta}) \leq \sqrt{2}\tilde{\sigma}^3 f_+(g(\tilde{\sigma}^2, \mu_{++}^2)) \quad \text{and} \quad \lambda_-(\tilde{\theta}) \geq \sqrt{2}\tilde{\sigma}^3 f_-(g(\tilde{\sigma}^2, \mu_{++}^2)).$$

To show that $\tilde{\sigma}^2 \mapsto \tilde{\sigma}^3 f_+(g(\tilde{\sigma}^2, \mu_{++}^2))$ and $\tilde{\sigma}^2 \mapsto \tilde{\sigma}^3 f_-(g(\tilde{\sigma}^2, \mu_{++}^2))$ are increasing on \mathbb{R}^+ , we use $\tilde{\sigma}^3 f_{\pm}(g(\tilde{\sigma}^2, \mu_{++}^2)) = \tilde{\sigma}^2 f_{2\mu_{++}^2 + \frac{1}{2}, \pm}(\tilde{\sigma}^2) = h_{a, \pm}(\tilde{\sigma}^2)$ where $f_{a, \pm}(x) = x + a \pm \sqrt{(x+a)^2 - 2x} \geq 0$ and $h_{a, \pm}(x) = x f_{a, \pm}(x)$. Since $h'_{a, \pm}(x) = f_{a, \pm}(x) + x f'_{a, \pm}(x)$ and $f_{a, \pm}(x) \geq 0$, having $f'_{a, \pm}(x) \geq 0$ on \mathbb{R}^+ is sufficient to conclude that $h'_{a, \pm}(x) \geq 0$ on \mathbb{R}^+ . Therefore, a sufficient condition to conclude is to show that $f_{a, \pm}$ is increasing on \mathbb{R}^+ .

$$f'_{a, \pm}(x) = 1 \pm \frac{x + a - 1}{\sqrt{(x+a)^2 - 2x}}$$

When $x + a - 1 \geq 0$, we have directly $f'_{a, +}(x) \geq 0$. Moreover,

$$f'_{a, -}(x) \geq 0 \iff 1 \geq \frac{x + a - 1}{\sqrt{(x+a)^2 - 2x}} \iff 0 \geq 1 - 2a$$

Since $1 - 2a = -2\mu_{++}^2$, we can conclude that $f'_{a, -}(x) \geq 0$.

When $x + a - 1 < 0$, we have directly $f'_{a, -}(x) \geq 0$. Moreover,

$$f'_{a, +}(x) \geq 0 \iff 1 \geq -\frac{x + a - 1}{\sqrt{(x+a)^2 - 2x}} \iff 0 \geq 1 - 2a$$

Since $1 - 2a = -2\mu_{++}^2$, we can conclude that $f'_{a, -}(x) \geq 0$. ■

Lemma 45 and using the above arguments concludes the proof of Theorem 43.

Monotonicity of the preliminary concentration The non-increasingness of the sequence of convex sets is obtained by monotonicity of the bounds of the rectangles $[\mu_{-, a}, \mu_{+, a}] \times [\sigma_{-, a}^2, \sigma_{+, a}^2]$ for all $a \in \mathcal{S}$. Given their definitions, it is sufficient to show that the functions $t \mapsto \varepsilon_{\mu}(t, \delta)$ and $t \mapsto \varepsilon_{\pm, \sigma}(t, \delta)$ are decreasing. This is shown in Lemma 46. The conditions on the initial time to have monotonicity are quite mild since $t_{\pm}(\delta) < 1$ after a relatively mild condition on δ . Depending on the choice of $(|\mathcal{S}|, s, \eta_0, \eta_1)$, it can even hold for all $\delta \in (0, 1]$. Numerically, for any practical choices of parameters we always had $t_{\pm}(\delta) < 2$. Since $t \geq t_a(\delta)$ implies that $N_{t, a} \geq 2$, the condition $N_{t, a} \geq t_{\pm}(\delta)$ is milder.

Lemma 46 *Let $\delta \in (0, 1]$. The function $t \mapsto \varepsilon_{\mu}(t, \delta)$ is decreasing on $[1, +\infty)$. There exists $t_+(\delta)$ such that the function $t \mapsto \varepsilon_{\pm, \sigma}(t, \delta)$ is decreasing on $[t_+(\delta), +\infty)$. In particular, $t_-(\delta) \leq 1$ if and only if $\delta \leq 6|\mathcal{S}|\zeta(s)e^{-\frac{s}{\log(1+\eta_0)}}$ and $t_+(\delta) \leq 1$ if and only if $\delta \leq 6|\mathcal{S}|\zeta(s)e^{-\frac{s}{\log(1+\eta_1)} - \frac{1}{2(1+\eta_1)}}$.*

Proof Let $\delta \in (0, 1]$. Recall that $\overline{W}'_i(x) = \frac{1}{1 - \frac{1}{\overline{W}_i(x)}}$ for all $x > 1$ and $i \in \{-1, 0\}$ (Lemma 75).

Decreasing ε_μ . Let $f(t) = 2s \log\left(1 + \frac{\log t}{2s}\right)$ and $c = 1 + 2 \log\left(\frac{6|S|\zeta(s)}{\delta}\right) + 2s$. Since $s > 1$ and $\frac{6|S|\zeta(s)}{\delta} > 1$, we have $c > 3$. Directly, we have $f'(t) = \frac{1}{t} \frac{1}{1 + \frac{\log t}{2s}}$. Then, by composition of the derivatives, we obtain

$$t^2 \frac{\partial \varepsilon_\mu(t, \delta)}{\partial t} = \frac{1 - \left(1 + \frac{\log t}{2s}\right) (\overline{W}_{-1}(c + f(t)) - 1)}{\left(1 + \frac{\log t}{2s}\right) \left(1 - \frac{1}{\overline{W}_{-1}(c + f(t))}\right)}.$$

Since $t \geq 1$ and $\overline{W}_{-1}(c + f(t)) > 1$, we have

$$\begin{aligned} \frac{\partial \varepsilon_\mu(t, \delta)}{\partial t} < 0 &\iff \left(1 + \frac{\log t}{2s}\right) (\overline{W}_{-1}(c + f(t)) - 1) > 1 \\ &\iff \frac{\log t}{2s} > \frac{1}{\overline{W}_{-1}(c + f(t)) - 1} - 1. \end{aligned}$$

Using that $f(t) \geq 0$ for all $t \geq 1$, $c > 3$ and $\overline{W}_{-1}(x) > x$, we obtain

$$\frac{1}{\overline{W}_{-1}(c + f(t)) - 1} - 1 \leq -\frac{1}{2} < 0 \leq \frac{\log t}{2s}.$$

Therefore, we have shown that $t \mapsto \varepsilon_\mu(t, \delta)$ is decreasing on $[1, +\infty)$ for all $\delta \in (0, 1]$.

Decreasing $\varepsilon_{\pm, \sigma}$. Let $(i_+, i_-) = (-1, 0)$ and $(\eta_+, \eta_-) = (\eta_1, \eta_0)$. Let $a_\pm = 2(1 + \eta_\pm)$, $b = \log\left(\frac{6|S|\zeta(s)}{\delta}\right)$, $f_\pm(t) = s \log\left(1 + \frac{\log t}{\log(1 + \eta_\pm)}\right)$ and $g_\pm(t) = \frac{b + f_\pm(t)}{t}$. Directly, we have

$$f'_\pm(t) = \frac{1}{t} \frac{s}{\log(1 + \eta_\pm) + \log t} \quad \text{and} \quad g'_\pm(t) = \frac{t f'_\pm(t) - (b + f_\pm(t))}{t^2}.$$

Then, by composition of the derivatives, we obtain

$$\pm t^2 \frac{\partial \varepsilon_{\pm, \sigma}(t, \delta)}{\partial t} = \frac{\overline{W}_{i_\pm}(1 + a_\pm g_\pm(t)) (1 - a_\pm (b + f_\pm(t) - t f'_\pm(t))) - 1}{1 - \overline{W}_{i_\pm}(1 + a_\pm g_\pm(t))}.$$

Since $t \geq 1$, $\overline{W}_{-1}(1 + a_+ g_+(t)) > 1$ and $\overline{W}_0(1 + a_- g_-(t)) < 1$, we have

$$\begin{aligned} \frac{\partial \varepsilon_{\pm, \sigma}(t, \delta)}{\partial t} < 0 &\iff a_\pm (b + f_\pm(t) - t f'_\pm(t)) > 1 - \frac{1}{\overline{W}_{i_\pm}(1 + a_\pm g_\pm(t))} \\ &\iff a_\pm \left(b - \frac{s}{\log(1 + \eta_\pm) + \log t}\right) \geq 1 - \frac{1}{\overline{W}_{i_\pm}(1 + a_\pm g_\pm(t))}, \end{aligned}$$

where the sufficient condition is obtained by noting that $f_{\pm}(t) > 0$ for $t > 1$. Using that $\overline{W}_0(1+x) \in (0, 1)$ and that $\overline{W}_{-1}(1+x) > 1$, we can obtain further sufficient conditions

$$\begin{aligned} \frac{\partial \varepsilon_{+,\sigma}(t, \delta)}{\partial t} < 0 &\iff \log\left(\frac{6|\mathcal{S}|\zeta(s)}{\delta}\right) \geq \frac{1}{2(1+\eta_1)} + \frac{s}{\log(1+\eta_1) + \log t} \\ &\iff t \geq t_+(\delta) = \exp\left(\frac{s}{\log\left(\frac{6|\mathcal{S}|\zeta(s)}{\delta}\right) - \frac{1}{2(1+\eta_1)}} - \log(1+\eta_1)\right), \\ \frac{\partial \varepsilon_{-,\sigma}(t, \delta)}{\partial t} < 0 &\iff t \geq t_-(\delta) = \exp\left(\frac{s}{\log\left(\frac{6|\mathcal{S}|\zeta(s)}{\delta}\right) - \log(1+\eta_0)}\right). \end{aligned}$$

■

Appendix G. Thresholds

After leveraging the link between the GLR and the EV-GLR statistics (Lemma 47) and showing how to calibrate δ -correct thresholds (Lemma 48), we derive several δ -correct family of thresholds with increasing complexities both theoretically and numerically: Student thresholds (Appendix G.1), box thresholds (Appendix G.2), KL thresholds (Appendix G.3) and BoB thresholds (Appendix G.4). In Appendix G.5, we study whether the derived family of thresholds is asymptotically tight.

Thresholds Relationship Lemma 47 shows that the relationship between the GLR and the EV-GLR statistics (Lemma 16) allows to obtain δ -correct thresholds for the EV-GLR stopping rule by using the ones obtained for GLR stopping rule, and vice-versa.

Lemma 47 *Let $(c_{a,b})_{(a,b) \in [K]^2}$ be a family of thresholds.*

If $(c_{a,b})_{(a,b) \in [K]^2}$ ensures δ -correctness of the EV-GLR stopping rule, then it ensures δ -correctness of the GLR stopping rule.

Let $C_b(\mu_t, \sigma_t^2) = \frac{(\mu_t, \hat{a}_t - \mu_{t,b})^2}{\min\{\sigma_{t,\hat{a}_t}^2, \sigma_{t,b}^2\}}$. If $(c_{a,b})_{(a,b) \in [K]^2}$ ensures δ -correctness of the GLR stopping rule, then $(\tilde{c}_{a,b})_{(a,b) \in [K]^2}$ ensures δ -correctness of the EV-GLR stopping rule, where $\tilde{c}_{a,b}(N_t, \delta) = \frac{C_b(\mu_t, \sigma_t^2)}{\log(1+C_b(\mu_t, \sigma_t^2))} c_{a,b}(N_t, \delta)$.

Proof Using Lemma 16, we have the following inequalities between the statistics involved in the GLR and EV-GLR stopping rules

$$Z_a^{\text{EV}}(t) \geq Z_a(t) \geq \frac{\log(1 + C_a(\mu_t, \sigma_t^2))}{C_a(\mu_t, \sigma_t^2)} Z_a^{\text{EV}}(t).$$

Let $(c_{a,b})_{(a,b) \in [K]^2}$ a family of thresholds ensuring δ -correctness of $\tau_{\delta}^{\text{EV}}$. We show by inclusion of event that $(c_{a,b})_{(a,b) \in [K]^2}$ is a family of thresholds ensuring δ -correctness of τ_{δ} .

$$\begin{aligned} \{\tau_{\delta} < +\infty, \hat{a}_{\tau_{\delta}} \neq a^*\} &= \{\exists t \in \mathbb{N}, \forall a \neq \hat{a}_t, Z_a(t) > c_{\hat{a}_t, a}(N_t, \delta), \hat{a}_t \neq a^*\} \\ &\subseteq \{\exists t \in \mathbb{N}, \forall a \neq \hat{a}_t, Z_a^{\text{EV}}(t) > c_{\hat{a}_t, a}(N_t, \delta), \hat{a}_t \neq a^*\} \end{aligned}$$

The exact same argument can be used to show the second statement. ■

Calibration by concentration Lemma 48 gives the terms to concentrate to ensure δ -correctness. Due to the structure of the identification problem which doesn't involve the variance, two distance metrics can be used as starting point for the GLR stopping rule. They both rely on a weighted sum of the per-arm KL divergences between the current estimator and an unknown parameter $(\mu, \tilde{\sigma}_t^2)$. While the choice $\tilde{\sigma}_t^2 = \sigma^2$ seems natural in (9), it doesn't fully leverage the BAI structure. This can be done in (6) by choosing $\tilde{\sigma}_t^2 = \sigma_t^2 + (\mu_t - \mu)^2$, which yields smaller thresholds.

Lemma 48 *If with probability $1 - \delta$, for all $t \in \mathbb{N}$ and for all $a \neq a^*(\mu)$,*

$$\sum_{b \in \{a, a^*(\mu)\}} \frac{N_{t,b}}{2} \log \left(1 + \frac{(\mu_{t,b} - \mu_b)^2}{\sigma_{t,b}^2} \right) \leq c_{a, a^*(\mu)}(N_t, \delta),$$

then the GLR stopping rule using the family of thresholds $(c_{b,a})_{(b,a) \in [K]^2}$ is δ -correct on \mathcal{D}^K .

If with probability $1 - \delta$, for all $t \in \mathbb{N}$ and for all $a \neq a^(\mu)$,*

$$\sum_{b \in \{a, a^*(\mu)\}} N_{t,b} \text{KL}((\mu_{t,b}, \sigma_{t,b}^2), (\mu_b, \sigma_b^2)) \leq c_{a, a^*(\mu)}(N_t, \delta),$$

then the GLR stopping rule using the family of thresholds $(c_{b,a})_{(b,a) \in [K]^2}$ is δ -correct on \mathcal{D}^K .

If with probability $1 - \delta$, for all $t \in \mathbb{N}$ and for all $a \neq a^(\mu)$,*

$$\sum_{b \in \{a, a^*(\mu)\}} N_{t,b} \frac{(\mu_{t,b} - \mu_b)^2}{2\sigma_{t,b}^2} \leq c_{a, a^*(\mu)}(N_t, \delta), \quad (17)$$

then the EV-GLR stopping rule using the family of thresholds $(c_{b,a})_{(b,a) \in [K]^2}$ is δ -correct on \mathcal{D}^K .

Proof Let $\hat{a}_t = \hat{a}_t$ and $a^* = a^*(\mu)$. Using Lemma 15, the stopping time (4) for the GLR stopping rule involves the statistics

$$\begin{aligned} Z_a(t) &= \inf_{(\lambda, \kappa^2): \lambda_a \geq \lambda_{\hat{a}_t}} \sum_{b \in \{a, \hat{a}_t\}} N_{t,b} \text{KL}((\mu_{t,b}, \sigma_{t,b}^2), (\lambda_b, \kappa_b^2)) \\ &= \inf_{(\lambda, \kappa^2): \lambda_a \geq \lambda_{\hat{a}_t}} \sum_{b \in \{a, \hat{a}_t\}} \frac{N_{t,b}}{2} \log \left(1 + \frac{(\mu_{t,b} - \lambda_b)^2}{\sigma_{t,b}^2} \right). \end{aligned}$$

Then, by definition,

$$\mathbb{P}(\tau_\delta < +\infty, \hat{a}_{\tau_\delta} \neq a^*) \leq \mathbb{P}(\exists t \in \mathbb{N}, \exists a \neq a^*, a = \hat{a}_t, \forall c \neq a, Z_c(t) > c_{a,c}(N_t, \delta)).$$

Since a valid choice is $c = a^*$ and $\lambda = \mu$, we obtain

$$\begin{aligned} &\mathbb{P}(\tau_\delta < +\infty, \hat{a}_{\tau_\delta} \neq a^*) \\ &\leq \left\{ \begin{array}{l} \mathbb{P} \left(\exists t \in \mathbb{N}, \exists a \neq a^*, \sum_{b \in \{a^*, a\}} N_{t,b} \text{KL}((\mu_{t,b}, \sigma_{t,b}^2), (\mu_b, \sigma_b^2)) > c_{a, a^*}(N_t, \delta) \right) \\ \mathbb{P} \left(\exists t \in \mathbb{N}, \exists a \neq a^*, \sum_{b \in \{a^*, a\}} \frac{N_{t,b}}{2} \log \left(1 + \frac{(\mu_{t,b} - \mu_b)^2}{\sigma_{t,b}^2} \right) > c_{a, a^*}(N_t, \delta) \right) \end{array} \right\}. \end{aligned}$$

The concentration assumptions yield δ -correctness of the first two family of thresholds. The proof for the EV-GLR stopping rule is identical to the one for the GLR stopping rule except that it uses the statistics $Z_a^{\text{EV}}(t)$. ■

G.1. Student Thresholds

Lemma 49 gives the family of Student thresholds.

Lemma 49 *Let $s > 1$ and ζ be the Riemann ζ function. Let a family of thresholds $c_{a,b}(N_t, \delta)$ with value $+\infty$ if $t < \max_{c \in \{a,b\}} t_c^S(\delta)$ and otherwise $c_{a,b}^S(N_t, \delta) = \max \{\beta^S(N_{t,a}, \delta), \beta^S(N_{t,b}, \delta)\}$. Taking*

$$\beta^S(t, \delta) = t \log \left(1 + \frac{1}{t-1} Q \left(1 - \frac{\delta}{4(K-1)\zeta(s)t^s}; \mathcal{T}_{t-1} \right) \right)^2 \quad (18)$$

yields a δ -correct family of thresholds for the GLR stopping rule. The stochastic initial times are

$$\forall a \in [K], \quad t_a^S(\delta) \stackrel{\text{def}}{=} \inf \left\{ t \in \mathbb{N} \mid N_{t,a} \geq \max \left\{ 2, \left(\frac{\delta}{4(K-1)\zeta(s)} \right)^{1/s} \right\} \right\}. \quad (19)$$

Proof Let $s > 1$ and ζ be the Riemann ζ function. Let \mathcal{T}_n denotes the Student distribution with n degrees of freedom and Q its quantile function. We define a threshold $c_{a,b}(N_t, \delta)$ with value $+\infty$ if $t < \max_{c \in \{a,b\}} t_c(\delta)$ and otherwise $c_{a,b}(N_t, \delta) = \max \{c(N_{t,a}, \delta), c(N_{t,b}, \delta)\}$.

Using Lemma 48, δ -correctness of the family of thresholds can be obtained directly by upper bounding (6). An initial time condition $t \geq \max_{c \in \{a,b\}} t_c(\delta)$ is necessary for the threshold $(c_{a,b})_{a,b}$ to be defined since they involve the quantiles of Student distribution. We obtain

$$t_a(\delta) \stackrel{\text{def}}{=} \inf \left\{ t \in \mathbb{N} \mid N_{t,a} \geq \max \left\{ 2, \left(\frac{\delta}{4(K-1)\zeta(s)} \right)^{1/s} \right\} \right\}.$$

Let $a^* = a^*(\mu)$. A simple approach to control the sum of two terms is to control each term individually. Each individual term is a function of $\frac{\mu_{t,a} - \mu_a}{\sqrt{\tilde{\sigma}_{t,a}^2/N_{t,a}}}$ which has a Student distribution, where $\tilde{\sigma}_{t,a}^2 = \frac{N_{t,a}}{N_{t,a}-1} \sigma_{t,a}^2$ is the unbiased variance. Let $\hat{\mu}_t$ be the empirical mean of t standard Gaussian, $\hat{\sigma}_t^2 = \frac{1}{t} \sum_{s=1}^t (X_s - \hat{\mu}_t)^2$ is the empirical variance and $\tilde{\sigma}_t^2 = \frac{t}{t-1} \hat{\sigma}_t^2$ is its unbiased version. Using a union bound and the fact that $\frac{\hat{\mu}_t}{\sqrt{\tilde{\sigma}_t^2/t}} \sim \mathcal{T}_{t-1}$ we obtain

$$\mathbb{P} \left(\exists t \geq \tilde{t}_0(\delta) : \frac{\hat{\mu}_t}{\sqrt{\tilde{\sigma}_t^2/t}} > \tilde{c}(t, \delta) \right) \leq \sum_{t \geq \tilde{t}_0(\delta)} \mathbb{P} \left(\frac{\hat{\mu}_t}{\sqrt{\tilde{\sigma}_t^2/t}} > \tilde{c}(t, \delta) \right) \leq \frac{\delta}{4(K-1)}$$

where $\tilde{c}(t, \delta) = Q \left(1 - \frac{\delta}{4(K-1)\zeta(s)t^s}; \mathcal{T}_{t-1} \right)$. Using this result, direct computations yield that

$$\begin{aligned} & \mathbb{P} \left(\exists t \geq t_0(\delta) : \exists a \neq a^*, \sum_{b \in \{a, a^*\}} \frac{N_{t,b}}{2} \log \left(1 + \frac{(\mu_{t,b} - \mu_b)^2}{\sigma_{t,b}^2} \right) > \max_{b \in \{a, a^*\}} c(N_{t,b}, \delta) \right) \\ & \leq \mathbb{P} \left(\exists t \geq t_0(\delta) : \exists a \neq a^*, \exists b \in \{a, a^*\}, N_{t,b} \log \left(1 + \frac{(\mu_{t,b} - \mu_b)^2}{\sigma_{t,b}^2} \right) > c(N_{t,b}, \delta) \right) \\ & \leq 2(K-1) \mathbb{P} \left(\exists t \geq \tilde{t}_0(\delta) : t \log \left(1 + \frac{t}{t-1} \frac{\hat{\mu}_t^2}{\tilde{\sigma}_t^2} \right) > c(t, \delta) \right) \\ & \leq 4(K-1) \mathbb{P} \left(\exists t \geq \tilde{t}_0(\delta) : \frac{\hat{\mu}_t}{\sqrt{\tilde{\sigma}_t^2/t}} > \sqrt{(t-1) \left(\exp \left(\frac{c(t, \delta)}{t} \right) - 1 \right)} \right) \leq \delta \end{aligned}$$

where the last equation is obtained by choice of the stopping threshold

$$c(t, \delta) = t \log \left(1 + \frac{1}{t-1} Q \left(1 - \frac{\delta}{4(K-1)\zeta(s)t^s}; \mathcal{T}_{t-1} \right)^2 \right). \quad (20)$$

Since it satisfies the hypothesis of Lemma 48, this yields the desired result. \blacksquare

EV-GLR stopping rule Up to a log-transform the same arguments yield a family of thresholds for the EV-GLR stopping rule. The proof of Lemma 50 is omitted since it is almost identical to the above.

Lemma 50 *Let $s > 1$ and ζ be the Riemann ζ function. Let $(t_a)_{a \in [K]}$ as in (19). We define a threshold $c_{a,b}(N_t, \delta)$ with value $+\infty$ if $t < \max_{c \in \{a,b\}} t_c(\delta)$ and otherwise $c_{a,b}(N_t, \delta) = \max \{c(N_{t,a}, \delta), c(N_{t,b}, \delta)\}$. Taking*

$$c(t, \delta) = \frac{t}{t-1} Q \left(1 - \frac{\delta}{4(K-1)\zeta(s)t^s}; \mathcal{T}_{t-1} \right)^2 \quad (21)$$

yields a δ -correct family of thresholds for the EV-GLR stopping rule.

G.2. Box Thresholds

Before presenting the counterparts for the EV-GLR stopping rule, we first present the proof of Lemma 5.

Proof Let $a \neq a^* = a^*(\mu)$. Using Lemma 48, we only need to exhibit threshold ensuring the required concentration behavior. One way of obtaining such an upper bound is to maximize the above quantities under constraints obtained by our concentration results. The form of the optimization is independent of the considered pair of arms $\{a, a^*(\mu)\}$ and of the time t (omitted in the following).

As we will see it only depends on $y_b = \frac{(\mu_{t,b} - \mu_b)^2}{\sigma_{t,b}^2}$ and $x_b = \frac{\sigma_{t,b}^2}{\sigma_b^2}$.

We can show that the family of thresholds in (7) is the solutions of an optimization problem. Let $C, D \in (\mathbb{R}_+^*)^2$ and $N \in (\mathbb{N}^*)^2$,

$$\begin{aligned} & \text{maximize} \quad \sum_{b \in \{1,2\}} \frac{N_b}{2} \log(1 + y_b) \\ & \text{such that} \quad \forall b \in \{1,2\}, \quad y_b \geq 0, \quad x_b y_b \leq C_b, \quad x_b \geq D_b. \end{aligned}$$

Since $y \mapsto \log(1 + y)$ is concave and increasing, $y \mapsto \sum_{b \in \{1,2\}} N_b \log(1 + y_b)$ is concave and increasing in each of its coordinates. Since the constraints and the objective are separate between each coordinate, the maximum is achieved at $\frac{C_b}{D_b}$ and as value

$$\sum_{b \in \{1,2\}} \frac{N_b}{2} \log \left(1 + \frac{C_b}{D_b} \right).$$

To obtain the family of thresholds as in (7), we simply need to use concentration results to specify the constraints C_b and D_b , which differ depending on the considered pair of arms $\{a, a^*(\mu)\}$ and of

the time t . This can be done by combining the lower tail concentration on the empirical variance (Corollary 26) and the upper and lower tail concentration of the empirical mean (Lemma 28). By direct union bound, we have with probability greater than $1 - \frac{\delta}{K-1}$, for all $b \in \{a, a^*\}$ and all $t \geq \max_{b \in \{a, a^*\}} t_b(\delta)$,

$$\begin{aligned} (\mu_{t,b} - \mu_b)^2 &\leq \sigma_b^2 \varepsilon_\mu(N_{t,b}, \delta), \\ \sigma_{t,b}^2 &\geq \sigma_b^2 (1 - \varepsilon_{-, \sigma}(N_{t,b} - 1, \delta)). \end{aligned}$$

where ε_μ , $\varepsilon_{-, \sigma}$ and t_b are defined as in Lemma 5. Using Lemma 75, this initial time condition ensures that $1 - \varepsilon_{-, \sigma}(N_{t,b} - 1, \delta) > 0$. Since \overline{W}_0 has values in $(0, 1)$, we obtain that $\varepsilon_{-, \sigma}(t, \delta) \in (0, 1)$. Taking a union bound over $a \neq a^*$ concludes the proof. \blacksquare

EV-GLR stopping rule Up to a log-transform the same arguments yield a family of thresholds for the EV-GLR stopping rule (Lemma 50). The proof of Lemma 51 is omitted since it is almost identical to the above.

Lemma 51 *Let ε_μ , $\varepsilon_{-, \sigma}$ and $(t_a)_{a \in [K]}$ as in Lemma 5. We define a threshold $c_{a,b}(N_t, \delta)$ with value $+\infty$ if $t < \max_{c \in \{a,b\}} t_c(\delta)$ and otherwise*

$$c_{a,b}(N_t, \delta) = \sum_{c \in \{a,b\}} \frac{N_{t,c} \varepsilon_\mu(N_{t,c}, \delta)}{2(1 - \varepsilon_{-, \sigma}(N_{t,c} - 1, \delta))}. \quad (22)$$

This yields a δ -correct family of thresholds for the EV-GLR stopping rule.

G.3. KL Thresholds

Thanks to Lemma 48, it is sufficient to concentrate the summed KL divergence. Let $a^* = a^*(\mu)$. This can be done by Theorem 43 with $\mathcal{S} = \{a, a^*\}$ for all $a \neq a^*$ (and taking a bound over those $K - 1$ terms). For all $a \in [K]$, the additional initial time condition to obtain monotonicity rewrites

$$t_a(\delta) = \inf \left\{ t \mid N_{t,a} > 1 + \max \left\{ \frac{e^{s/\log\left(\frac{12(K-1)\zeta(s)}{\delta}\right)}}{1 + \eta_0}, \frac{e^{s/\left(\log\left(\frac{12(K-1)\zeta(s)}{\delta}\right) - \frac{1}{2(1+\eta_1)}\right)}}{1 + \eta_1} \right\} \right\}.$$

Numerically, we always observed that this was satisfied after initialization. Therefore, those terms have no impact.

G.4. BoB Thresholds

In Appendix G.2, we saw how to calibrate stopping threshold based on an optimization problem using concentration constraints. Modifying the optimization problem therein, we can leverage Theorem 6 by adding a constraint. While this argument requires doing the union bound over two concentration results (hence considering $\delta/2$), the added constraint can result in a smaller stopping threshold (hence faster stopping).

With our notations, for $b \in \{1, 2\}$, the KL divergence rewrites as $\text{KL}((\mu_{t,b}, \sigma_{t,b}^2), (\mu_b, \sigma_b^2)) = \frac{1}{2} f(x_b, y_b)$ with $f(x, y) = (1 + y)x - 1 - \log(x)$. Since $\nabla^2 f(x, y) = \begin{bmatrix} x^{-2} & 1 \\ 1 & 0 \end{bmatrix}$ is a positive

semi-definite matrix on $(\mathbb{R}_+^*)^2$, f is a convex function of $(\mathbb{R}_+^*)^2$. Let $E > 0$. The KL constraint from Theorem 6 is convex and can be expressed as

$$\sum_{b \in \{1,2\}} \frac{N_b}{2} f(x_b, y_b) \leq E.$$

Since this constraint mixes the two coordinates, we can't conclude by using the separation arguments. To our knowledge, there is no closed form solution for the resulting optimization problem

$$\begin{aligned} & \text{maximize} \quad \sum_{b \in \{1,2\}} \frac{N_b}{2} \log(1 + y_b) \\ & \text{such that} \quad \forall b \in \{1,2\}, \quad y_b \geq 0, \quad x_b y_b \leq C_b, \quad x_b \geq D_b, \\ & \text{and} \quad \sum_{b \in \{1,2\}} \frac{N_b}{2} f(x_b, y_b) \leq E. \end{aligned}$$

However, as a maximization of a concave function under linear and convex inequalities, we can solve it numerically. Corollary 7 is a direct consequence of the above manipulations and an union bound over concentration result.

EV-GLR stopping rule The same ideas lead to Corollary 52, which optimizes a different function under the same constraints. The resulting optimization problem is computationally faster since the objective is linear. The proof of Corollary 52 is omitted since it is almost identical to the above.

Corollary 52 *Let $f(x, y) = (1 + y)x - 1 - \log(x)$ for all $(x, y) \in (\mathbb{R}_+^*)^2$. Let $(t_a)_{a \in [K]}$ in (8), $\varepsilon_\mu, \varepsilon_{-, \sigma}$ as in Lemma 5 and $(c_{b,a}^{\text{KL}})_{b,a \in [K]}$ in (10). The family of thresholds $c_{a,b}(N_t, \delta)$ with value $+\infty$ if $t < \max_{c \in \{a,b\}} t_c(\delta/6)$ and otherwise solution of the optimization problem*

$$\begin{aligned} & \text{maximize} \quad \frac{1}{2} \sum_{c \in \{a,b\}} N_{t,c} y_c \\ & \text{such that} \quad \forall c \in \{a,b\}, \quad y_c \geq 0, \quad x_c y_c \leq \varepsilon_\mu(N_{t,c}, \delta/2), \quad x_c \geq 1 - \varepsilon_{-, \sigma}(N_{t,c} - 1, \delta/2), \\ & \text{and} \quad \frac{1}{2} \sum_{c \in \{a,b\}} N_{t,c} f(x_c, y_c) \leq c_{b,a}^{\text{KL}}(N_t, \delta/2) \end{aligned}$$

yields a δ -correct family of thresholds for the EV-GLR stopping rule.

G.5. Asymptotically Tight Thresholds

Among the class of δ -correct family of thresholds for the GLR stopping rule, Theorem 9 suggests we should select the asymptotically tight ones since they yield an asymptotically optimal algorithm.

As observed empirically in Figure 1(a), the Student and the box thresholds are not asymptotically tight since their slope in $\log(\frac{1}{\delta})$ is higher than the ones of the KL and BoB thresholds. Theoretically, the arguments used in the proofs also justify why they couldn't reach the $\log(\frac{1}{\delta})$. For the Student threshold, it comes from the fact that we consider the inequality $\{(I) + (II) > 2\beta\} \subset \{(I) > \beta\} \cup \{(II) > \beta\}$, hence we could at most reach $2 \log(\frac{1}{\delta})$. For the box threshold, it comes from the fact that we concentrate the term of each arm individually as well, hence we could also reach at most $2 \log(\frac{1}{\delta})$ (higher in practice).

EV-GLR stopping rule Similar arguments could be used to study the $\log\left(\frac{1}{\delta}\right)$ -dependency of the families of thresholds derived for the EV-GLR stopping rule. However, since they are δ -correct, they can't be asymptotically tight. Otherwise, the Theorem 9 will contradict the asymptotic lower bound on the expected sample complexity as $T_{\sigma^2}^*(\mu) < T^*(\mu, \sigma^2)$ (Lemma 4).

KL and BoB thresholds In order to hope to achieve the $\log\left(\frac{1}{\delta}\right)$ dependency one should concentrate the whole sum, as done for the KL threshold. Since $\log(1+y) = \inf_{x>0} f(x,y)$, the KL constraint is also an upper bound on the solution of the optimization problem, i.e. the BoB thresholds are smaller than the KL thresholds $c_{b,a}^{\text{KL}}(N_t, \delta/2)$, hence they inherit their $\log\left(\frac{1}{\delta}\right)$ dependency. Therefore, to show the asymptotically tightness of both the KL and the BoB thresholds, we only need to study the KL family of thresholds. This result is established in Lemma 53.

Lemma 53 *The KL thresholds defined in (10) is an asymptotically tight family of thresholds.*

Proof We analyze our threshold under a sampling rule that starts by pulling all arms $t(\delta) = \max\{t_0(\delta), t_1(\delta)\}$ times, where

$$t_1(\delta) \stackrel{\text{def}}{=} \min \left\{ t \in \mathbb{N} \mid t > e^{1+W_0\left(\frac{2(1+\eta_0)}{e}\left(\log\left(\frac{12(K-1)\zeta(s)}{\delta}\right)+s\log\left(1+\frac{\log(t)}{\log(1+\eta_0)}\right)\right)-\frac{1}{e}}\right)} \right\},$$

$$t_0(\delta) \stackrel{\text{def}}{=} \max \left\{ \frac{e^{s/\log\left(\frac{12(K-1)\zeta(s)}{\delta}\right)}}{1+\eta_0}, \frac{e^{s/\left(\log\left(\frac{12(K-1)\zeta(s)}{\delta}\right)-\frac{1}{2(1+\eta_1)}\right)}}{1+\eta_1} \right\}.$$

After $Kt(\delta)$ pulls, our threshold is finite for all arms. We have $t_0(\delta) \rightarrow_{\delta \rightarrow 0} \frac{1}{1+\min\{\eta_0, \eta_1\}}$. Using that $W_0(x) \approx \log(x) - \log \log(x)$ (Appendix I), $t_1(\delta)$ is asymptotically equivalent to $\frac{2(1+\eta_0)\log(1/\delta)}{\log \log(1/\delta)}$. Therefore, this initial pulling count $t(\delta)$ satisfies $\limsup_{\delta \rightarrow 0} t(\delta)/\log(1/\delta) = 0$.

Our threshold is in fact a family of threshold functions $(c_{a,b}(N_t, \delta))_{a,b \in [K]}$, which after that initialization, are defined by

$$c_{a,b}(N_t, \delta) = 4\overline{W}_{-1} \left(1 + \frac{\log \frac{2\zeta(s)^2}{\delta}}{4} + \frac{s}{4} \sum_{c \in \{a,b\}} \log(1 + \log_{\gamma} N_{t,c}) + \frac{1}{2} \sum_{c \in \{a,b\}} \log(\gamma R_{t,c}(\delta)) \right).$$

Using concavity of $x \mapsto \log(1 + \log_{\gamma}(x))$ and $\sum_{c \in \{a,b\}} N_{t,c} \leq t$, we have

$$c_{a,b}(N_t, \delta) \leq 4\overline{W}_{-1} \left(1 + \frac{\log \frac{2\zeta(s)^2}{\delta}}{4} + \frac{s}{2} \log \left(1 + \frac{\log t/2}{\log \gamma} \right) + \log(\gamma R(t, \delta)) \right),$$

where $R(t, \delta) = \max_{a \in [K]} R_{t,a}(\delta)$. Using the above, the concavity of \overline{W}_{-1} (Lemma 75) yields for $t \geq Kt_0(\delta)$

$$c_{a,b}(N_t, \delta) \leq 4\overline{W}_{-1} \left(1 + \frac{1}{4} \log \left(\frac{2\zeta(s)^2}{\delta} \right) \right) + 4 \left(\frac{s}{2} \log \left(1 + \frac{\log t/2}{\log \gamma} \right) + \log(\gamma R(t, \delta)) \right) \overline{W}'_{-1} \left(1 + \frac{1}{4} \log \left(\frac{2\zeta(s)^2}{\delta} \right) \right).$$

For all x for which it is defined, $\bar{W}'_{-1}(x) = \frac{1}{1-1/\bar{W}_{-1}(x)}$. Let $c > 1$, we have $\bar{W}'_{-1}(x) \leq c \iff \bar{W}_{-1}(x) \geq \frac{1}{1-\frac{1}{c}}$. Taking x_c such that $\bar{W}_{-1}(x_c) = \frac{1}{1-\frac{1}{c}}$, which is possible since it is a strictly increasing function with values in $[1, +\infty)$, Lemma 75 yields that $x_c \leq \frac{1}{1-\frac{1}{c}} + \log(1 - \frac{1}{c})$. For $c = 2$, we obtain $\bar{W}'_{-1}(x) \leq 2$ for all $x \geq 2 - \log(2)$, hence $\bar{W}'_{-1}\left(1 + \frac{1}{4} \log\left(\frac{2\zeta(s)^2}{\delta}\right)\right) \leq 2$ for all $\delta < \min\{1, \delta_0\}$ where $\delta_0 = 2\zeta(s)^2 e^{-4}$. Then, for all $\delta < \min\{1, \delta_0\}$,

$$c_{a,b}(N_t, \delta) \leq 4\bar{W}_{-1}\left(1 + \frac{1}{4} \log\left(\frac{2\zeta(s)^2}{\delta}\right)\right) + 8 \log(\gamma R(t, \delta)) + 4s \log\left(1 + \frac{\log t/2}{\log \gamma}\right).$$

We know that $4s \log\left(1 + \frac{\log t}{\log \gamma}\right) = O(t^\alpha)$ for some $\alpha \in (0, 1)$ and, using $\bar{W}_{-1}(x) =_{+\infty} x + \log(x) + o(1)$ (Lemma 75), that $4\bar{W}_{-1}\left(1 + \frac{1}{4} \log\left(\frac{2\zeta(s)^2}{\delta}\right)\right) = \log\left(\frac{1}{\delta}\right) + o\left(\log\left(\frac{1}{\delta}\right)\right)$. To conclude the proof, we need to show that $\limsup_{\delta \rightarrow 0} \frac{8 \log(R(t, \delta))}{\log(1/\delta)} \leq 1$.

The term $R(t, \delta)$ is a data-dependent term quantifying the goodness of the local quadratic approximation for the KL (as a function of the natural parameters). Taking the notations from Theorem 6 and setting $\mu_{+,t,a}^2 = \max \mu_{\pm,t,a}^2$, we have

$$R(t, \delta) = \max_{a \in [K]} \frac{\sigma_{-,t,a}^3 f_-(g(\sigma_{-,t,a}^2, \mu_{+,t,a}^2))}{\sigma_{+,t,a}^3 f_+(g(\sigma_{+,t,a}^2, \mu_{+,t,a}^2))}$$

where $f_{\pm}(x) = \frac{1 \pm \sqrt{1-x}}{\sqrt{x}}$ and $g(x, y) = \frac{2x}{(x+2y+\frac{1}{2})^2}$. Using the definition of ε_{μ} and $\varepsilon_{\pm, \sigma}$, we obtain directly that $\lim_{t \rightarrow +\infty} \varepsilon_{\mu}(t, \delta) = 0$ and $\lim_{t \rightarrow +\infty} \varepsilon_{\pm, \sigma}(t, \delta) = 0$.

Let $a \in [K]$. For all t , define $i_{t,a} = \lfloor \log_{\gamma} N_{t,a} \rfloor$, $n_{t,a} = \gamma^{i_{t,a}}$, $\bar{t}_a = \inf\{t \mid N_{t,a} = n_{t,a}\}$. Using the law of large number, we also have that $\lim_{N_{t,a} \rightarrow +\infty} \mu_{\bar{t}_a, a} = \mu_a$ and $\lim_{N_{t,a} \rightarrow +\infty} \sigma_{\bar{t}_a, a}^2 = \sigma_a^2$. Chaining the limits, we obtain $\lim_{N_{t,a} \rightarrow +\infty} \mu_{+,t,a}^2 = \mu_a^2$ and $\lim_{N_{t,a} \rightarrow +\infty} \sigma_{\pm, t, a} = \sigma_a$. Since the functions f_{\pm} and g are continuous, we have shown for all $a \in [K]$

$$\lim_{N_{t,a} \rightarrow +\infty} \frac{\sigma_{-,t,a}^3 f_-(g(\sigma_{-,t,a}^2, \mu_{+,t,a}^2))}{\sigma_{+,t,a}^3 f_+(g(\sigma_{+,t,a}^2, \mu_{+,t,a}^2))} = \frac{f_-(g(\sigma_a^2, \mu_a^2))}{f_+(g(\sigma_a^2, \mu_a^2))}$$

When $\delta \rightarrow 0$, the initialization yields $t_0(\delta) \rightarrow \infty$. Therefore, by using the above convergence and the fact that $N_{t,a} \geq t_0(\delta)$, there exists $\delta_1 \in (0, \min\{1, \delta_0\})$ such that for all $\delta \leq \delta_1$ and all $t \geq K t_0(\delta)$,

$$R(t, \delta) \leq 2 \max_{a \in [K]} \frac{f_-(g(\sigma_a^2, \mu_a^2))}{f_+(g(\sigma_a^2, \mu_a^2))},$$

which is a constant independent of δ , hence $\limsup_{\delta \rightarrow 0} \frac{8 \log(R(t, \delta))}{\log(1/\delta)} \leq 1$. ■

Appendix H. Expected Sample Complexity

Theorem 9 gives asymptotic upper bound on the expected sample complexity of the algorithms using the GLR stopping rule. Theorem 10 gives an impossibility results for the algorithms using

the EV-GLR stopping rule, which is based on first deriving an upper bound on the expected sample complexity. In Appendix H.1, we prove the upper bound for TaS with the GLR stopping rule and for EV-TaS with the EV-GLR stopping rule. In Appendix H.2, we prove it for β -EB-TCI with the GLR stopping rule and for β -EB-EVTCI with the EV-GLR stopping rule.

H.1. Wrapped Track-and-Stop

Showing an asymptotic upper bound on the expected sample complexity of TaS and EV-TaS can be done with similar asymptotic arguments as when the variance is known (Garivier and Kaufmann, 2016). The proof relies on two main ingredients. First, a concentration result for (μ_t, σ_t^2) is obtained thanks to the forced exploration (Lemma 54 proved in Appendix H.1.1). Second, we use the Lemma 20 in Garivier and Kaufmann (2016), which ensures that the empirical allocation $\frac{N_t}{t}$ converges towards the optimal allocation being targeted. While it corresponds to $w^*(\mu, \sigma^2)$ for TaS, it is $w_{\sigma^2}^*(\mu)$ for EV-TaS (where the σ^2 dependency is hidden by notation).

Without loss of generality and for the sake of simpler notations, we assume that the Gaussian bandit model with parameter $(\mu, \sigma^2) \in (\mathbb{R} \times \mathbb{R}_+^*)^K$ is such that $a^*(\mu) = 1$.

Track-and-Stop Since the proofs share the same structure as in Garivier and Kaufmann (2016), we detail the one for TaS and highlight the differences for EV-TaS later.

Proof Let $\varepsilon > 0$. From the continuity of $(\mu, \sigma^2) \mapsto w^*(\mu, \sigma^2)$ (Degenne and Koolen, 2019), there exists $\xi_1 = \xi_1(\varepsilon) \leq \frac{\min_{a \neq 1}(\mu_1 - \mu_a)}{4}$, $\xi_2 = \xi_2(\varepsilon) > 1$ and $\xi_3 = \xi_3(\varepsilon) \in (0, 1)$ such that

$$\mathcal{I}_\varepsilon \stackrel{\text{def}}{=} \left(\prod_{a \in [K]} [\mu_a - \xi_1, \mu_a + \xi_1] \right) \times \left(\prod_{a \in [K]} [\sigma_a^2 \xi_3, \sigma_a^2 \xi_2] \right)$$

satisfies $\max_a |w_a^*(\tilde{\mu}, \tilde{\sigma}^2) - w_a^*(\mu, \sigma^2)| \leq \varepsilon$ for all $(\tilde{\mu}, \tilde{\sigma}^2) \in \mathcal{I}_\varepsilon$. Since $a^*(\tilde{\mu}) = 1$ for all $(\tilde{\mu}, \tilde{\sigma}^2) \in \mathcal{I}_\varepsilon$, the empirical best arm is $\hat{a}_t = 1$ whenever $(\mu_t, \sigma_t^2) \in \mathcal{I}_\varepsilon$. Let $T \in \mathbb{N}$, $h(T) \stackrel{\text{def}}{=} T^{1/4}$ and define the concentration event $\mathcal{E}_T = \bigcap_{t=h(T)}^T \{(\mu_t, \sigma_t^2) \in \mathcal{I}_\varepsilon\}$.

The forced exploration ensures that each arm is drawn at least of order \sqrt{t} times at round t . Thanks to concentration results on both (μ_t, σ_t^2) , Lemma 54 upper bounds $\mathbb{P}_\nu(\mathcal{E}_T^c)$.

Lemma 54 *Let T such that $h(T) > \left(K + \frac{1}{1-\xi_3}\right)^2$. There exist two constants B, C (that depend on (μ, σ^2) and ε) such that $\mathbb{P}_\nu(\mathcal{E}_T^c) \leq BT \exp(-CT^{1/8})$.*

Lemma 55 is exactly Lemma 20 in Garivier and Kaufmann (2016), hence the proof is omitted.

Lemma 55 (Lemma 20 in Garivier and Kaufmann (2016)) *There exists a constant T_ε such that for $T \geq T_\varepsilon$, it holds that on \mathcal{E}_T , for either C Tracking or D-Tracking,*

$$\forall t \geq \sqrt{T}, \quad \max_{a \in [K]} \left| \frac{N_{t,a}}{t} - w_a^*(\mu, \sigma^2) \right| \leq 3(K-1)\varepsilon$$

On the event \mathcal{E}_T , it holds for $t \geq h(T)$ that $\hat{a}_t = 1$ and the GLR rewrites

$$\min_{a \neq 1} Z_a(t) = \min_{a \neq 1} \inf_{\lambda \in [\mu_{t,a}, \mu_{t,1}]} \sum_{b \in \{1, a\}} \frac{N_{t,b}}{2} \log \left(1 + \frac{(\mu_{t,b} - \lambda)^2}{\sigma_{t,b}^2} \right) = tg \left(\mu_t, \sigma_t^2, \frac{N_t}{t} \right)$$

where, for $(\tilde{\mu}, \tilde{\sigma}^2) \in (\mathbb{R} \times \mathbb{R}_+^*)^K$ such that $a^*(\tilde{\mu}) = 1$ and $\tilde{w} \in \Delta_K$, we introduced the function

$$g(\tilde{\mu}, \tilde{\sigma}^2, \tilde{w}) = \min_{a \neq 1} \inf_{\lambda \in [\tilde{\mu}^a, \tilde{\mu}^1]} \sum_{b \in \{1, a\}} \frac{\tilde{w}_b}{2} \log \left(1 + \frac{(\tilde{\mu}_b - \lambda)^2}{\tilde{\sigma}_b^2} \right)$$

Recall that $a^*(\tilde{\mu}) = 1$ on \mathcal{I}_ε . We introduce $C_\varepsilon^*(\mu, \sigma^2) = \inf_{(\tilde{\mu}, \tilde{\sigma}^2, \tilde{w}) \in \mathcal{H}_\varepsilon(\mu, \sigma^2)} g(\tilde{\mu}, \tilde{\sigma}^2, \tilde{w})$, where

$$\mathcal{H}_\varepsilon(\mu, \sigma^2) = \mathcal{I}_\varepsilon \times \{ \tilde{w} \in \Delta_K \mid \|\tilde{w} - w^*(\mu, \sigma^2)\|_\infty \leq 2(K-1)\varepsilon \} .$$

Using Lemma 55, for $T \geq T_\varepsilon$, on the event \mathcal{E}_T it holds that for every $t \geq \sqrt{T}$, $\min_{a \neq 1} Z_a(t) \geq tC_\varepsilon^*(\mu, \sigma^2)$.

Let $\alpha \in [0, 1)$, $\delta_0 \in (0, 1]$, functions $f, \bar{T} : (0, 1] \rightarrow \mathbb{R}_+$ and C as in the definition of an asymptotically tight family of thresholds. In the following, we consider $\delta \leq \delta_0$ and $T \geq \{\bar{T}(\delta), T_\varepsilon\}$. Using that for all $a \neq 1$, $Z_a(t) \geq \min_{a \neq 1} Z_a(t) \geq tC_\varepsilon^*(\mu, \sigma^2)$ and $c_{1,a}(N_t, \delta) \leq f(\delta) + CT^\alpha \leq f(\delta) + CT^\alpha$, we obtain on \mathcal{E}_T ,

$$\begin{aligned} \min \{ \tau_\delta, T \} &\leq \sqrt{T} + \sum_{t=\sqrt{T}}^T \mathbf{1}\{ \tau_\delta > t \} \leq \sqrt{T} + \sum_{t=\sqrt{T}}^T \mathbf{1}\{ \exists a \neq 1, Z_a(t) \leq c_{1,a}(N_t, \delta) \} \\ &\leq \sqrt{T} + \sum_{t=\sqrt{T}}^T \mathbf{1}\{ tC_\varepsilon^*(\mu, \sigma^2) \leq f(\delta) + CT^\alpha \} \leq \sqrt{T} + \frac{f(\delta) + CT^\alpha}{C_\varepsilon^*(\mu, \sigma^2)} \end{aligned}$$

Introducing

$$T_0(\delta) = \inf \left\{ T \geq \bar{T}(\delta) : \sqrt{T} + \frac{f(\delta) + CT^\alpha}{C_\varepsilon^*(\mu, \sigma^2)} \leq T \right\},$$

for every $T \geq \max \{ T_0(\delta), \bar{T}(\delta), T_\varepsilon \}$, one has $\mathcal{E}_T \subseteq \{ \tau_\delta \leq T \}$, therefore, by using Lemma 54,

$$\mathbb{P}_\nu(\tau_\delta > T) \leq \mathbb{P}_\nu(\mathcal{E}_T^c) \leq BT \exp(-CT^{1/8})$$

and

$$\mathbb{E}_\nu[\tau_\delta] \leq T_0(\delta) + \bar{T}(\delta) + T_\varepsilon + \sum_{T=1}^{\infty} BT \exp(-CT^{1/8}).$$

We now provide an upper bound on $T_0(\delta)$. Letting $\eta > 0$ and introducing the constant

$$D(\eta) \stackrel{\text{def}}{=} \inf \{ T \in \mathbb{N} : T - \sqrt{T} \geq T/(1+\eta) \} \leq 1 + \left(1 + \frac{1}{\eta} \right)^2$$

one has

$$T_0(\delta) \leq D(\eta) + \inf \left\{ T \geq \bar{T}(\delta) \mid f(\delta) + CT^\alpha \leq T \frac{C_\varepsilon^*(\mu, \sigma^2)}{1+\eta} \right\}.$$

For all $\gamma > 0$, there exists $T_{\alpha, \gamma}$ (depending on μ, σ^2) such that for all $T \geq T_{\alpha, \gamma}$,

$$T \frac{C_\varepsilon^*(\mu, \sigma^2)}{1+\eta} - CT^\alpha \geq T \frac{C_\varepsilon^*(\mu, \sigma^2)}{(1+\eta)(1+\gamma)}.$$

Then,

$$\begin{aligned} T_0(\delta) &\leq D(\eta) + \bar{T}(\delta) + T_{\alpha,\gamma} + \inf \left\{ T \mid f(\delta) \leq T \frac{C_\varepsilon^*(\mu, \sigma^2)}{(1+\eta)(1+\gamma)} \right\} \\ &\leq 1 + D(\eta) + \bar{T}(\delta) + T_{\alpha,\gamma} + \frac{(1+\eta)(1+\gamma)}{C_\varepsilon^*(\mu, \sigma^2)} f(\delta). \end{aligned}$$

Dividing by $\log(1/\delta)$ and taking limits, we get

$$\limsup_{\delta \rightarrow 0} \frac{T_0(\delta)}{\log(1/\delta)} \leq \frac{(1+\eta)(1+\gamma)}{C_\varepsilon^*(\mu, \sigma^2)}.$$

We obtain that for every $\eta, \varepsilon, \gamma > 0$,

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\nu[\tau_\delta]}{\log(1/\delta)} \leq \limsup_{\delta \rightarrow 0} \frac{T_0(\delta)}{\log(1/\delta)} \leq \frac{(1+\eta)(1+\gamma)}{C_\varepsilon^*(\mu, \sigma^2)}.$$

By continuity of g and by definition of $w^*(\mu, \sigma^2)$, we obtain $\lim_{\varepsilon \rightarrow 0} C_\varepsilon^*(\mu, \sigma^2) = T^*(\mu, \sigma^2)^{-1}$. Letting η and ε go to zero yields

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\nu[\tau_\delta]}{\log(1/\delta)} \leq T^*(\mu, \sigma^2).$$

■

Empirical Variance Track-and-Stop We only highlight some differences since the proofs are similar. We emphasize here that the notations $w_{\sigma^2}^*(\mu)$, $T_{\sigma^2}^*(\mu)$ and all the ones we will use below hide the dependency in σ^2 to distinguish this from the quantity defined when the variance is assumed to be unknown. While this might be unfortunate, we believe it eases greatly the notations and the highlight the difference between the two complexities.

Lemma 56 *Using the EV-GLR stopping rule with an asymptotically tight family of thresholds, EV-TaS satisfies that, for all ν with $|a^*(\mu)| = 1$,*

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\nu[\tau_\delta^{EV}]}{\log(1/\delta)} \leq T_{\sigma^2}^*(\mu).$$

Proof Let $\varepsilon > 0$. Similarly, there exists $\xi_1 \leq \frac{\min_{a \neq 1}(\mu_1 - \mu^a)}{4}$, $\xi_2 > 1$, $\xi_3 \in (0, 1)$ such that \mathcal{I}_ε defined as above satisfies $\max_a |w_a^*(\tilde{\mu}) - w_a^*(\mu)| \leq \varepsilon$ for all $(\tilde{\mu}, \tilde{\sigma}^2) \in \mathcal{I}_\varepsilon$. Let $T \in \mathbb{N}$, $h(T) := T^{1/4}$ and define the concentration event \mathcal{E}_T as above. Since Lemma 54 relies solely on forced exploration, its result still hold in that case. A tracking result similar to Lemma 55 gives the existence of T_ε such that for $T \geq T_\varepsilon$, it holds that on \mathcal{E}_T ,

$$\forall t \geq \sqrt{T}, \quad \max_{a \in [K]} \left| \frac{N_{t,a}}{t} - w_a^*(\mu) \right| \leq 3(K-1)\varepsilon$$

On the event \mathcal{E}_T , it holds for $t \geq h(T)$ that $\hat{a}_t = 1$ and the EV-GLR rewrites

$$\min_{a \neq 1} Z_a^{EV}(t) = \min_{a \neq 1} \inf_{\lambda \in [\mu_{t,a}, \mu_{t,1}]} \sum_{b \in \{1, a\}} N_{t,b} \frac{(\mu_{t,b} - \lambda)^2}{2\sigma_{t,b}^2} = tg \left(\mu_t, \sigma_t^2, \frac{N_t}{t} \right)$$

where, for $(\tilde{\mu}, \tilde{\sigma}^2) \in (\mathbb{R} \times \mathbb{R}_+^*)^K$ such that $a^*(\tilde{\mu}) = 1$ and $\tilde{w} \in \Delta_K$, we introduced the function

$$g(\tilde{\mu}, \tilde{\sigma}^2, \tilde{w}) = \min_{a \neq 1} \inf_{\lambda \in [\tilde{\mu}^a, \tilde{\mu}^1]} \sum_{b \in \{1, a\}} \tilde{w}_b \frac{(\tilde{\mu}_b - \lambda)^2}{2\tilde{\sigma}_b^2}$$

We introduce $C_\varepsilon^*(\mu, \sigma^2) = \inf_{(\tilde{\mu}, \tilde{\sigma}^2, \tilde{w}) \in \mathcal{H}_\varepsilon(\mu, \sigma^2)} g(\tilde{\mu}, \tilde{\sigma}^2, \tilde{w})$, where

$$\mathcal{H}_\varepsilon(\mu, \sigma^2) = \mathcal{I}_\varepsilon \times \left\{ \tilde{w} \in \Delta_K \mid \|\tilde{w} - w_{\sigma^2}^*(\mu)\|_\infty \leq 2(K-1)\varepsilon \right\}.$$

By tracking, for $T \geq T_\varepsilon$, on the event \mathcal{E}_T it holds that for every $t \geq \sqrt{T}$, $\min_{a \neq 1} Z_a^{\text{EV}}(t) \geq tC_\varepsilon^*(\mu, \sigma^2)$. Let $T \geq \max\{\bar{T}(\delta), T_\varepsilon\}$. As above, on \mathcal{E}_T , we have $\min\{\tau_\delta^{\text{EV}}, T\} \leq \sqrt{T} + \frac{f(\delta) + CT^\alpha}{C_\varepsilon^*(\mu, \sigma^2)}$ and we introduce $T_0(\delta) = \inf\left\{T \geq \bar{T}(\delta) : \sqrt{T} + \frac{f(\delta) + CT^\alpha}{C_\varepsilon^*(\mu, \sigma^2)} \leq T\right\}$. For every $T \geq \{T_0(\delta), \bar{T}(\delta), T_\varepsilon\}$, one has $\mathcal{E}_T \subseteq \{\tau_\delta^{\text{EV}} \leq T\}$, therefore, by using Lemma 54,

$$\mathbb{E}_\nu[\tau_\delta^{\text{EV}}] \leq T_0(\delta) + \bar{T}(\delta) + T_\varepsilon + \sum_{T=1}^{\infty} BT \exp(-CT^{1/8}).$$

Manipulations similar as above yield an upper bound on $T_0(\delta)$. By continuity of g and by definition of $w_{\sigma^2}^*(\mu)$, we obtain $\lim_{\varepsilon \rightarrow 0} C_\varepsilon^*(\mu, \sigma^2) = T_{\sigma^2}^*(\mu)^{-1}$. Letting η and ε go to zero yields

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\nu[\tau_\delta^{\text{EV}}]}{\log(1/\delta)} \leq T_{\sigma^2}^*(\mu).$$

■

We are now ready to prove the impossibility result for EV-TaS (Theorem 10).

Proof Let $(c_{a,b})_{(a,b) \in [K]^2}$ be an asymptotically tight family of thresholds and a problem independent constant $c_0 > 0$. Combining EV-TaS with the EV-GLR stopping rule using $(c_0 c_{a,b})_{(a,b) \in [K]^2}$ yields an algorithm such that, for all ν with $|a^*(\mu)| = 1$,

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\nu[\tau_\delta^{\text{EV}}]}{\log(1/\delta)} \leq c_0 T_{\sigma^2}^*(\mu).$$

Lemma 56 shows the above result for $c_0 = 1$, and generalizing to $c_0 > 0$ is direct. Suppose towards contradiction that the obtained algorithm is δ -correct. Therefore, using Lemmas 2 and 4, we have shown that

$$T^*(\mu, \sigma^2) \leq \liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\nu[\tau_\delta^{\text{EV}}]}{\log(1/\delta)} \leq \limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\nu[\tau_\delta^{\text{EV}}]}{\log(1/\delta)} \leq c_0 T_{\sigma^2}^*(\mu) < c_0 T_\beta^*(\mu, \sigma^2).$$

For $c_0 \in (0, 1)$, the contradiction is direct. For $c_0 \geq 1$, we have shown that there is a problem independent constant $c_0 > 0$ such that $T^*(\mu, \sigma^2)/T_{\sigma^2}^*(\mu) \leq c_0$. This is a direct contradiction with Lemma 12 showing that there exists a sequence of instances $(\nu_n)_{n \in \mathbb{N}}$ with $|a^*(\nu_n)| = 1$ such that $\lim_{n \rightarrow +\infty} T^*(\mu_n, \sigma_n^2)/T_{\sigma_n^2}^*(\mu_n) = +\infty$. Therefore, the obtained algorithm is not δ -correct. ■

H.1.1. PROOF OF LEMMA 54

Since T is such that $h(T) > \left(K + \frac{1}{1-\xi_3}\right)^2$, we have $h(T) \geq K^2$ and $\xi_3 < 1 - \frac{1}{\sqrt{h(T)}-K}$. In particular, for all $t \in \{h(T), \dots, T\}$ and all $s \in \{\sqrt{t}-K, \dots, t\}$, $\xi_3 + \frac{1}{s} < 1$. Then,

$$\mathbb{P}(\mathcal{E}_T^c) \leq \sum_{t=h(T)}^T \sum_{a=1}^K \left[\mathbb{P}(|\mu_{t,a} - \mu_a| \geq \xi_1) + \mathbb{P}(\sigma_{a,t}^2 \geq \sigma_a^2 \xi_2) + \mathbb{P}\left(\sigma_{a,t}^2 \leq \frac{\sigma_a^2}{\xi_2}\right) \right]$$

By forced exploration, for $t \geq h(T)$ one has $N_{t,a} \geq (\sqrt{t}-K/2)_+ - 1 \geq \sqrt{t}-K$ for every arm a . In the proof of Lemma 19 in [Garivier and Kaufmann \(2016\)](#) it was shown using a union bound over time and Lemma 29 that, for all $t \in [h(T), T]$,

$$\mathbb{P}(|\mu_{t,a} - \mu_a| \geq \xi_1) \leq \frac{2e^{-(\sqrt{t}-K)\frac{\xi_1^2}{2\sigma_a^2}}}{1 - e^{-\frac{\xi_1^2}{2\sigma_a^2}}}.$$

Since we use the same method to show our result on the variance, the proof for the mean is omitted.

Let $\hat{\sigma}_{a,s}^2$ be the empirical variance of the first s reward from arm a (such that $\hat{\sigma}_{a,N_{t,a}}^2 = \sigma_{a,t}^2$). We adopt their proof strategy to derive the equivalent upper bound on the concentration of the variance.

$$\begin{aligned} \mathbb{P}(\sigma_{a,t}^2 \geq \sigma_a^2 \xi_2) &= \mathbb{P}\left(\sigma_{a,t}^2 \geq \sigma_a^2 \xi_2, N_{t,a} \geq \sqrt{t}-K\right) \\ &\leq \sum_{s=\sqrt{t}-K-1}^{t-1} \mathbb{P}(\hat{\sigma}_{a,s+1}^2 \geq \sigma_a^2 \xi_2) \leq \sum_{s=\sqrt{t}-K-1}^{t-1} \exp\left(-\frac{s}{2} \left(h\left(\xi_2 + \frac{1}{s}\right) - 1\right)\right) \end{aligned}$$

The first inequality is obtained by taking a union bound over the values of $N_{t,a} \in [\sqrt{t}-K, t]$. The second inequality is obtained by Corollary 27 with $x = \xi_2 + \frac{1}{s} > 1$. Using that $h(x) = x - \log(x)$ and $\log(1+x) \leq x$, direct computations yield that

$$\begin{aligned} sh\left(\xi_2 + \frac{1}{s}\right) &= s\xi_2 + 1 - s \log\left(\xi_2 + \frac{1}{s}\right) = sh(\xi_2) + 1 - s \log\left(1 + \frac{1}{s\xi_2}\right) \geq sh(\xi_2) + 1 - \frac{1}{\xi_2}, \\ \sum_{s=\sqrt{t}-K-1}^{t-1} \left(e^{-\frac{1}{2}(h(\xi_2)-1)}\right)^s &\leq \frac{1}{1 - e^{-\frac{1}{2}(h(\xi_2)-1)}} e^{-\frac{(\sqrt{t}-K-1)}{2}(h(\xi_2)-1)}. \end{aligned}$$

Putting those together, we obtain, for all $t \in [h(T), T]$,

$$\mathbb{P}(\sigma_{a,t}^2 \geq \sigma_a^2 \xi_2) \leq \frac{e^{-\frac{1}{2}\left(1-\frac{1}{\xi_2}\right)}}{1 - e^{-\frac{1}{2}(h(\xi_2)-1)}} e^{-\frac{(\sqrt{t}-K-1)}{2}(h(\xi_2)-1)}.$$

The same manipulations using Corollary 27 with $x = \xi_3 + \frac{1}{s} \in (0, 1)$ for all $s \in [\sqrt{t}-K, t]$ and all $t \in [h(T), T]$ (see above by choice of T), yield that, for all $t \in [h(T), T]$,

$$\mathbb{P}(\sigma_{a,t}^2 \leq \sigma_a^2 \xi_3) \leq \frac{e^{-\frac{1}{2}\left(1-\frac{1}{\xi_3}\right)}}{1 - e^{-\frac{1}{2}(h(\xi_3)-1)}} e^{-\frac{(\sqrt{t}-K-1)}{2}(h(\xi_3)-1)}.$$

Finally, letting

$$C = \frac{1}{2} \min \left\{ \frac{\xi_1^2}{\max_{a \in [K]} \sigma_a^2}, h(\xi_2) - 1, h(\xi_3) - 1 \right\} \quad \text{and}$$

$$B = \sum_{a=1}^K \left(\frac{2e^{K \frac{\xi_1^2}{2\sigma_a^2}}}{1 - e^{-\frac{\xi_1^2}{2\sigma_a^2}}} + \frac{e^{\frac{K-1}{2}(h(\xi_2)-1) - \frac{1}{2}(1-\frac{1}{\xi_2})}}{1 - e^{-\frac{1}{2}(h(\xi_2)-1)}} + \frac{e^{\frac{K-1}{2}(h(\xi_3)-1) - \frac{1}{2}(1-\frac{1}{\xi_3})}}{1 - e^{-\frac{1}{2}(h(\xi_3)-1)}} \right),$$

one obtains

$$\mathbb{P}(\mathcal{E}_T^c) \leq \sum_{t=h(T)}^T B \exp(-\sqrt{t}C) \leq BT \exp(-\sqrt{h(T)}C) = BT \exp(-CT^{1/8}).$$

H.2. Wrapped β -EB-TCI

Showing an asymptotic upper bound on the expected sample complexity of β -EB-TCI and β -EB-EVTCI can be done with similar asymptotic arguments as when the variance is known (Shang et al., 2020). We will use the unified analysis of Top Two algorithms introduced in Jourdan et al. (2022), which highlights simple properties that the leader and the challenger should satisfy to obtain the desired upper bound. While they introduced it for single-parameter exponential families and bounded distributions, it also allows to cope for our two-parameters setting. The proof is composed of three steps: showing sufficient exploration of all arms, proving that the expectation of the convergence time towards the β -optimal allocation is finite and then concluding on the asymptotic upper bound.

To ensure sufficient exploration, we assume that $\min_{a \neq b} |\mu_a - \mu_b| > 0$, i.e. all the arms have distinct means. To our knowledge, this assumption is shared by all Top Two algorithms analysis.

H.2.1. β -EB-TCI

Let $a^* = a^*(\mu)$. Since the proofs share the same structure as in Jourdan et al. (2022), we detail the one for β -EB-TCI and highlight the differences for β -EB-EVTCI. The β -optimal allocation for Gaussian with unknown variance is defined as

$$w_\beta^*(\mu, \sigma^2) = \arg \max_{w \in \Delta_K, w_{a^*} = \beta} \min_{a \neq a^*} \inf_{u \in \mathbb{R}} \sum_{b \in \{a^*, a\}} w_b \log \left(1 + \frac{(\mu_b - u)^2}{\sigma_b^2} \right).$$

In Russo (2016) and Jourdan et al. (2022), they show the β -optimal allocation was unique for any single-parameter exponential families and for bounded distributions. Since the proof only relies on strict convexity already shown in Appendix C.2, it is straightforward to see that $w_\beta^*(\mu, \sigma^2)$ is a singleton (Property 1 in Jourdan et al. (2022)), denoted by $\{w^\beta\}$, such that $\min_{a \in [K]} w_a^\beta > 0$.

Before delving into the specifics of the proof, Lemma 57 gathers concentration results on which the subsequent analysis heavily relies on.

Lemma 57 *There exists a sub-Gaussian random variable W_μ and a sub-exponential random variable W_σ , which are independent, such that almost surely for all $a \in [K]$ and all t such that $N_{t,a} \geq 2$,*

$$N_{t,a} |\mu_{t,a} - \mu_a| \leq W_\mu \log(e + N_{t,a}) \quad \text{and} \quad \left| N_{t,a} \left(\frac{\sigma_{t,a}^2}{\sigma_a^2} - 1 \right) + 1 \right| \leq W_\sigma \log(e + N_{t,a}).$$

There exists a random variable $W_0 = \max_{a \in [K]} W_a^{-1}$, independent of W_μ , where $W_a/\sigma_a^2 \sim \chi^2(n_0 - 1)$ such that almost surely, for all $a \in [K]$ and all t such that $N_{t,a} \geq n_0$, $N_{t,a}\sigma_{t,a}^2 \geq W_0^{-1}$. In particular, W_0 admits a finite mean for $n_0 \geq 4$ and a finite variance for $n_0 \geq 6$, and any random variable which is polynomial in both W_σ and W_μ has a finite expectation.

Proof The existence of W_μ is given by Lemma 73 in Jourdan et al. (2022).

Using Lemma 25, we have for all $a \in [K]$ and t such that $N_{t,a} \geq 2$, $N_{t,a} \left(\frac{\sigma_{t,a}^2}{\sigma_a^2} - 1 \right) + 1 = S_{N_{t,a}-1,a}$, where $S_{n,a}$ is sub-exponential with $c = 2$ and $-S_{n,a}$ is sub-exponential with $c = -2$ and $S_{n-1,a} + n - 1 = \sum_{s \in [n-1]} Y_{s,a}^2$ is such that $Y_{s,a}$ are i.i.d. $\mathcal{N}(0, 1)$. Using Lemma 72 in Jourdan et al. (2022), $\sup_{N_{t,a} \geq 2} \frac{|S_{N_{t,a}-1,a}|}{\log(e + N_{t,a})}$ is also sub-exponential. By defining W_σ as the maximum over the finitely many arms of the above quantity, we obtain that W_σ is sub-exponential.

Since the increments $Y_{s,a}^2$ are positive, we have for all $a \in [K]$ and t such that $N_{t,a} \geq n_0$

$$\frac{\sigma_{t,a}^2}{\sigma_a^2} = \frac{S_{N_{t,a}-1,a} + N_{t,a} - 1}{N_{t,a}} = \frac{\sum_{s \in [N_{t,a}-1]} Y_{s,a}^2}{N_{t,a}} \geq \frac{Z_a}{N_{t,a}},$$

where $Z_a = \sum_{s \in [n_0-1]} Y_{s,a}^2 \sim \chi^2(n_0 - 1)$. Taking $W_a = \sigma_a^2 Z_a$ and the maximum over $[K]$ yields the result. As maximum of finitely many inverse-chi-squared distributions, W_0 admits a finite mean for $n_0 \geq 4$ and a finite variance for $n_0 \geq 6$.

For the independence property between W_σ and W_μ and between W_0 and W_μ , this is a direct consequence of the fact that the empirical mean and empirical variance are independent. Therefore, any random variable which is polynomial in both W_σ and W_μ has a finite expectation. \blacksquare

Using Lemma 57, it is direct to obtain that, for all $\varepsilon > 0$, there exists N_ε with $\mathbb{E}_\nu[N_\varepsilon] < +\infty$ such that for all $a \in [K]$ and t such that $N_{t,a} \geq N_\varepsilon$, $|\mu_{t,a} - \mu_a| \leq \varepsilon$ and $\left| \frac{\sigma_{t,a}^2}{\sigma_a^2} - 1 \right| \leq \varepsilon$.

As in Qin et al. (2017); Shang et al. (2020); Jourdan et al. (2022), our goal is to upper bound the expectation of the *convergence time*. For $\varepsilon > 0$, the random variable T_β^ε quantifies the number of samples required for the empirical allocations $\frac{N_t}{t}$ to be ε -close to w^β :

$$T_\beta^\varepsilon \stackrel{\text{def}}{=} \inf \left\{ T \geq 1 \mid \forall t \geq T, \left\| \frac{N_t}{t} - w^\beta \right\|_\infty \leq \varepsilon \right\}. \quad (23)$$

Lemma 58 shows that a sufficient condition for asymptotic β -optimality is to show $\mathbb{E}_\nu[T_\beta^\varepsilon] < +\infty$ for all ε small enough.

Lemma 58 *Let $(\delta, \beta) \in (0, 1)^2$. Assume that there exists $\varepsilon_1(\mu, \sigma^2) > 0$ such that for all $\varepsilon \in (0, \varepsilon_1(\mu, \sigma^2)]$, $\mathbb{E}_\nu[T_\beta^\varepsilon] < +\infty$. Combining the GLR stopping rule (4) with an asymptotically tight family of threshold yields an algorithm such that, for all $(\mu, \sigma^2) \in \mathbb{R}^K \times (\mathbb{R}_+^*)^K$ with $|a^*(\mu)| = 1$,*

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\nu[\tau_\delta]}{\log(1/\delta)} \leq T_\beta^*(\mu, \sigma^2).$$

Proof Let $\varepsilon_1 = \varepsilon_1(\mu, \sigma^2)$. Let $c_\beta = \frac{1}{2} \min_{a \in [K]} w_a^\beta > 0$ and $\Delta = \min_{a \neq a^*} |\mu_{a^*} - \mu_a| > 0$. Let $\zeta > 0$. By continuity of

$$(\mu, \sigma^2, w) \mapsto \inf_{x \in \mathbb{R}} \sum_{c \in \{a, b\}} w_c \log \left(1 + \frac{(\mu_c - x)^2}{\sigma_c^2} \right) \mathbf{1}\{\mu_a > \mu_b\}$$

on $\mathbb{R}^K \times (\mathbb{R}_+^*)^K \times \Delta_K$, there exists $\varepsilon_2 > 0$ such that

$$\begin{aligned} & \max_{a \in [K]} \left| \frac{N_{t,a}}{t} - w_a^\beta \right| \leq \varepsilon_2 \quad , \quad \max_{a \in [K]} |\mu_{t,a} - \mu_a| \leq \varepsilon_2 \quad \text{and} \quad \max_{a \in [K]} \left| \frac{\sigma_{t,a}^2}{\sigma_a^2} - 1 \right| \leq \varepsilon_2 \\ \implies & \max_{a \in [K]} |\mu_{t,a} - \mu_a| \leq \frac{\Delta}{4} \quad \text{and} \quad \frac{1}{t} \min_{a \neq a^*} C_t(a^*, a) \geq \frac{1 - \zeta}{T_\beta^*(\mu, \sigma^2)}. \end{aligned}$$

Choosing such a ε_2 , we take $\varepsilon \in (0, \min\{\varepsilon_1, \varepsilon_2, c_\beta\})$. By assumption, we have $\mathbb{E}_\nu[T_\beta^\varepsilon] < +\infty$, hence $\frac{N_{t,a}}{t} \geq w_a^\beta - \varepsilon \geq c_\beta$ for all $a \in [K]$.

Let N_ε as described above (obtained with Lemma 57). For all $t \geq c_\beta^{-1} N_\varepsilon$, we have $N_{t,a} \geq N_\varepsilon$ for all $a \in [K]$, hence $\max_{a \in [K]} |\mu_{t,a} - \mu_a| \leq \varepsilon \leq \varepsilon_2$ and $\max_{a \in [K]} \left| \frac{\sigma_{t,a}^2}{\sigma_a^2} - 1 \right| \leq \varepsilon \leq \varepsilon_2$. Therefore, we have $\hat{a}_t \in \arg \max_{a \in [K]} \mu_{t,a} = \arg \max_{a \in [K]} \mu_a = a^*$ as $\max_{a \in [K]} |\mu_{t,a} - \mu_a| \leq \frac{\Delta}{4}$.

Let $\alpha \in [0, 1)$, $\delta_0 \in (0, 1]$, functions $f, \bar{T} : (0, 1] \rightarrow \mathbb{R}_+$ and C as in the definition of an asymptotically tight family of thresholds. In the following, we consider $\delta \leq \delta_0$. Let $\kappa > 0$. Let $T \geq \frac{1}{\kappa} \max\{T_\beta^\varepsilon, c_\beta^{-1} N_\varepsilon, \bar{T}(\delta)\}$. Using the definition of the GLR stopping rule (4) with a family of asymptotically tight threshold, we have

$$\begin{aligned} \min\{\tau_\delta, T\} & \leq \kappa T + \sum_{t=\kappa T}^T \mathbf{1}\{\tau_\delta > t\} \leq \kappa T + \sum_{n=\kappa T}^T \mathbf{1}\{\exists a \neq a^*, C_t(a^*, a) \leq c_{a^*, a}(N_t, \delta)\} \\ & \leq \kappa T + \sum_{n=\kappa T}^T \mathbf{1}\left\{t \frac{1 - \zeta}{T_\beta^*(\mu, \sigma^2)} \leq f(\delta) + CT^\alpha\right\} \\ & \leq \kappa T + \frac{T_\beta^*(\mu, \sigma^2)}{1 - \zeta} (f(\delta) + CT^\alpha). \end{aligned}$$

Let $T_\zeta(\delta)$ as

$$T_\zeta(\delta) \stackrel{\text{def}}{=} \inf \left\{ T \geq 1 \mid \frac{T_\beta^*(\mu, \sigma^2)}{(1 - \zeta)(1 - \kappa)} (f(\delta) + CT^\alpha) \leq T \right\}.$$

For every $T \geq \max\{T_\zeta(\delta), \frac{1}{\kappa} \max\{T_\beta^\varepsilon, c_\beta^{-1} N_\varepsilon, \bar{T}(\delta)\}\}$, we have $\tau_\delta \leq T$, hence

$$\mathbb{E}_\nu[\tau_\delta] \leq \frac{1}{\kappa} \mathbb{E}_\nu[T_\beta^\varepsilon] + \frac{1}{\kappa c_\beta} \mathbb{E}_\nu[N_\varepsilon] + \frac{1}{\kappa} \bar{T}(\delta) + T_\zeta(\delta).$$

As $\mathbb{E}_\nu[T_\beta^\varepsilon] + c_\beta^{-1} \mathbb{E}_\nu[N_\varepsilon] < +\infty$ and $\lim_{\delta \rightarrow 0} \frac{\bar{T}(\delta)}{\log(1/\delta)} = 0$, we obtain for all $\zeta, \kappa > 0$

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\nu[\tau_\delta]}{\log(1/\delta)} \leq \limsup_{\delta \rightarrow 0} \frac{T_\zeta(\delta)}{\log(1/\delta)} \leq \frac{T_\beta^*(\mu, \sigma^2)}{(1 - \zeta)(1 - \kappa)},$$

where the last inequality uses an inversion result. Letting ζ and κ go to zero yields the result. \blacksquare

To upper bound the expected convergence time, as prior work we first establish sufficient exploration. Given an arbitrary threshold $L \in \mathbb{R}_+^*$, we define the sampled enough set and its arms with highest mean (when not empty) as

$$S_t^L \stackrel{\text{def}}{=} \{a \in [K] \mid N_{t,a} \geq L\} \quad \text{and} \quad \mathcal{I}_t^* \stackrel{\text{def}}{=} \arg \max_{a \in S_t^L} \mu_a. \quad (24)$$

We define the highly and the mildly under-sampled sets

$$U_t^L \stackrel{\text{def}}{=} \{a \in [K] \mid N_{t,a} < L^{1/6}\} \quad \text{and} \quad V_t^L \stackrel{\text{def}}{=} \{a \in [K] \mid N_{t,a} < L^{5/6}\}, \quad (25)$$

where we used $L^{1/6}$ and $L^{5/6}$ instead of \sqrt{L} and $L^{3/4}$ as done in [Shang et al. \(2020\)](#). However, [Jourdan et al. \(2022\)](#) noted the results would hold with L^{α_1} and L^{α_2} with $0 < \alpha_1 < \alpha_2 < 1$.

Lemma 59 shows that the transportation cost is strictly positive and increases linearly.

Lemma 59 *Let $\mu \in \mathbb{R}^K$ with $\mu_b < \mu_a$. There exists L with $\mathbb{E}[|L|^\alpha] < +\infty$ for all $\alpha > 0$ and $D_\nu > 0$ such that for $N_{t,a} \geq L$ and $N_{t,b} \geq L$, $C_t(a, b) > LD_\nu$, where $D_\nu > 0$ is a problem dependent constant.*

Proof Suppose that $N_{t,a} \geq L$ and $N_{t,b} \geq L$, for some L to be determined. First we get

$$C_t(a, b) \geq \frac{L}{2} \mathbf{1}\{\mu_{t,a} > \mu_{t,b}\} \inf_{x \in \mathbb{R}} \sum_{c \in \{a,b\}} \log \left(1 + \frac{(\mu_{t,c} - x)^2}{\sigma_{t,c}^2} \right).$$

For any compact interval $\mathcal{I}_C \subseteq \mathbb{R}$, the function defined by

$$(\mu, \sigma^2) \mapsto \mathbf{1}\{\mu_a > \mu_b\} \inf_{x \in \mathcal{I}_C} \sum_{c \in \{a,b\}} \log \left(1 + \frac{(\mu_c - x)^2}{\sigma_c^2} \right)$$

is continuous on $\mathbb{R}^K \times (\mathbb{R}_+^*)^K$. For L greater than some L_1 with finite moments, we have $\mu_{t,a} > \mu_{t,b}$, $[\mu_b - \varepsilon, \mu_a + \varepsilon] \subseteq \mathcal{I}_C$ and $(\mu_{t,a}, \sigma_{t,a}^2)$ is ε -close to (μ_a, σ_a^2) (and same thing for (μ_b, σ_b^2)). The continuity then gives that there exists L with finite moments such that

$$\inf_{x \in \mathbb{R}} \sum_{c \in \{a,b\}} \log \left(1 + \frac{(\mu_{t,c} - x)^2}{\sigma_{t,c}^2} \right) \geq \frac{1}{2} \inf_{x \in \mathcal{I}_C} \sum_{c \in \{a,b\}} \log \left(1 + \frac{(\mu_c - x)^2}{\sigma_c^2} \right).$$

This is strictly positive since $\mu_b < \mu_a$ due to the continuity and strict convexity properties holding for Gaussian with unknown variance. ■

Lemma 60 *Let $D_\nu > 0$ as in Lemma 59. Let S_t^L and \mathcal{I}_t^* as in (24). There exists L_4 with $\mathbb{E}_\nu[(L_4)^\alpha] < +\infty$ for all $\alpha > 0$ such that if $L \geq L_4$, for all t such that $S_t^L \neq \emptyset$, for all $(a, b) \in \mathcal{I}_t^* \times (S_t^L \setminus \mathcal{I}_t^*)$, we have $C_t(a, b) \geq LD_\nu$.*

Proof Applying Lemma 59 yields the result directly. ■

Lemma 61 gives an upper bound on the transportation costs between a sampled enough arm and an under-sampled one.

Lemma 61 *Let S_t^L as in (24). There exists L_5 with $\mathbb{E}_\nu[(L_5)^\alpha] < +\infty$ for all $\alpha > 0$ such that for all $L \geq L_5$ and all $t \in \mathbb{N}$,*

$$\forall (a, b) \in S_t^L \times \overline{S_t^L}, \quad C_t(a, b) \leq \frac{L}{2} \log \left(1 + LW_0 (D_1 + 4W_\mu)^2 \right),$$

where $D_1 > 0$ is a problem dependent constant and W_0, W_μ are the random variables defined in Lemma 57.

Proof Let $(a, b) \in S_t^L \times \overline{S_t^L}$ (a is sampled more than L times, b is not). Taking $x = \mu_{t,a}$ yields

$$C_t(a, b) \leq \frac{1}{2} N_{t,b} \log \left(1 + \frac{(\mu_{t,b} - \mu_{t,a})^2}{\sigma_{t,b}^2} \right) \leq \frac{L}{2} \log \left(1 + \frac{(\mu_{t,b} - \mu_{t,a})^2}{\sigma_{t,b}^2} \right),$$

where we used that $b \in \overline{S_t^L}$ and that $C_t(a, b) = 0$ when $\mu_{t,b} \geq \mu_{t,a}$.

By definition of W_0 and W_μ , we have for all $k \in [K]$

$$\mu_{t,k} \leq \mu_k + W_\mu \log(e + N_{t,k})/N_{t,k} \quad \text{and} \quad \sigma_{t,k}^2 \geq W_k/N_{t,k}.$$

Then, we obtain

$$\frac{(\mu_{t,b} - \mu_{t,a})^2}{\sigma_{t,b}^2} \leq \frac{N_{t,b}}{W_b} \left(|\mu_a - \mu_b| + W_\mu \left(\frac{\log(e + N_{t,a})}{N_{t,a}} + \frac{\log(e + N_{t,b})}{N_{t,b}} \right) \right)^2.$$

Since $x \mapsto \frac{\log(e+x)}{x}$ is decreasing on \mathbb{R}_+^* , we have $\frac{\log(e+N_{t,k})}{N_{t,k}} \leq 2$ for $N_{t,k} \geq 1$. Then, since $W_0 = \max_{a \in [K]} W_a^{-1}$, we have

$$\frac{(\mu_{t,b} - \mu_{t,a})^2}{\sigma_{t,b}^2} \leq LW_0 (|\mu_a - \mu_b| + 4W_\mu)^2,$$

which yields the result by taking $D_1 = \max_{a \neq b} |\mu_a - \mu_b|$. ■

Lemma 62 shows that the EB leader satisfies the required condition on a leader to obtain sufficient exploration, i.e. Property 2 in Jourdan et al. (2022).

Lemma 62 (Lemma 17 in Jourdan et al. (2022)) *Let S_t^L and \mathcal{I}_t^* as in (24). Let L_4 in Lemma 60. Then, for all $L \geq L_4$, for all n such that $S_t^L \neq \emptyset$, $B_{t+1}^{EB} \in S_t^L$ implies $B_{t+1}^{EB} \in \mathcal{I}_t^*$.*

Lemma 63 shows that the TCI challenger satisfies the required condition on a challenger to obtain sufficient exploration, i.e. Property 3 in Jourdan et al. (2022). The proof of Lemma 63 resembles the proof of Lemma 21 in Jourdan et al. (2022). The sole technical difference lies in the necessities of having an finite mean for W_0 , which explains why we consider $n_0 = 4$ for β -EB-TCI.

Lemma 63 *Let U_t^L and V_t^L as in (25) and $\mathcal{J}_t^* = \arg \max_{a \in \overline{V_t^L}} \mu_a$. There exists L_6 with $\mathbb{E}_\nu[L_6] < +\infty$ such that if $L \geq L_6$, for all n such that $U_t^L \neq \emptyset$, $B_{t+1}^{EB} \notin V_t^L$ implies $A_{t+1}^{TCI} \in V_t^L \cup (\mathcal{J}_t^* \setminus \{B_{t+1}^{EB}\})$.*

Proof Using Lemmas 60, 61 and 62, for all L larger than a random variable L_7 with finite expectation, $B_{t+1}^{EB} \in \mathcal{J}_t^*$ and

$$\forall (a, b) \in \mathcal{J}_t^* \times (\overline{V_t^L} \setminus \mathcal{J}_t^*), \quad C_t(a, b) + \log N_{t,b} \geq L^{5/6} D_\nu + \frac{5}{6} \log L,$$

$$\forall (a, b) \in \overline{U_t^L} \times U_t^L, \quad C_t(a, b) + \log N_{t,b} \leq \frac{L^{1/6}}{2} \log \left(1 + L^{1/6} W_0 (D_1 + 4W_\mu)^2 \right) + \frac{1}{6} \log L.$$

Using that $\log(1+x) \leq \sqrt{x}$, we obtain

$$\frac{L^{1/6}}{2} \log \left(1 + L^{1/6} W_0 (D_1 + 4W_\mu)^2 \right) \leq \frac{L^{1/4}}{2} \sqrt{W_0} (D_1 + 4W_\mu)$$

Therefore, for $L \geq L_8 := W_0^{6/7} \left(\frac{D_1 + 4W_\mu}{2D_\nu} \right)^{12/7} + 1$, we have that

$$L^{5/6} D_\nu + \frac{5}{6} \log L > \frac{L^{1/6}}{2} \log \left(1 + L^{1/6} W_0 (D_1 + 4W_\mu)^2 \right) + \frac{1}{6} \log L.$$

Let $L_6 = \max\{L_7, L_8\}$. Therefore, at least one under-sampled arm has transportation cost lower than all the ones that are much sampled. This implies that $A_{t+1}^{\text{TCI}} \in V_t^L \cup (\mathcal{J}_t^* \setminus \{B_{t+1}^{\text{EB}}\})$. Then, we have

$$\mathbb{E}_\nu[L_6] \leq \mathbb{E}_\nu[L_7] + 1 + \mathbb{E}_\nu \left[W_0^{6/7} \right] \mathbb{E}_\nu \left[\left(\frac{D_1 + 4W_\mu}{2D_\nu} \right)^{12/7} \right] < +\infty,$$

The first inequality and the last strict inequality are obtained by Lemma 57, thanks to the independence of W_0 and W_μ , the fact that W_0 has finite mean since $n_0 \geq 4$ and that polynomial of W_μ have finite expectation. \blacksquare

Since Properties 2 and 3 of Jourdan et al. (2022) are satisfied (Lemmas 62 and 63), Lemma 64 holds for β -EB-TCI on instances such that $\min_{a \neq b} |\mu_a - \mu_b| > 0$.

Lemma 64 (Lemma 7 in Jourdan et al. (2022)) *There exist N_0 with $\mathbb{E}_\nu[N_0] < +\infty$ such that for all $t \geq N_0$ and all $a \in [K]$, $N_{t,a} \geq (t/K)^{1/6}$.*

Now that we have proved sufficient exploration (Lemma 64), we will show convergence towards the β -optimal allocation. Lemma 65 shows that the EB leader satisfies the required condition on a leader to obtain convergence, i.e. Property 5 in Jourdan et al. (2022).

Lemma 65 (Lemma 18 in Jourdan et al. (2022)) *There exists N_6 with $\mathbb{E}_\nu[N_6] < +\infty$ such that for all $t \geq N_6$, $B_{t+1}^{\text{EB}} = a^*$.*

For all $a \in [K]$ and all $t > t_0$, let $\psi_{t,a} \stackrel{\text{def}}{=} \mathbb{P}_{|(t-1)}[a_t = a]$ be the probability of sampling arm a at time t and $\Psi_{t,a} \stackrel{\text{def}}{=} \sum_{s \in [t]} \psi_{s,a}$ its cumulative sum.

Using the EB leader, Lemma 66 shows that the TCI challenger satisfies the required condition on a challenger to obtain convergence towards the β -optimal allocation, i.e. Property 6 in Jourdan et al. (2022). The proof of Lemma 66 is very similar to the proof of Lemma 22 in Jourdan et al. (2022), hence we omit certain details.

Lemma 66 *Let $\varepsilon > 0$. There exists N_7 with $\mathbb{E}_\nu[N_7] < +\infty$ such that for all $n \geq N_7$ and all $a \neq a^*(\mu)$, $\Psi_{t,a}/t \geq w_a^\beta + \varepsilon$ implies that $A_{t+1}^{\text{TCI}} \neq a$.*

Proof Using Lemma 65, we know that $B_{t+1}^{\text{EB}} = a^*$ for all $t \geq N_6$. By definition of A_{t+1}^{TCI} and algebraic manipulation, $A_{t+1}^{\text{TCI}} \neq a$ is implied by

$$\frac{1}{t} \left(C_t(a^*, a) - \min_{b \neq a^*} C_t(a^*, b) \right) > \frac{\log(tK)}{2t}.$$

In Appendix C.2, we obtain the desired regularity properties: joint continuity, strict convexity, equality at the equilibrium. Therefore, by using the same proof as the one of Lemma 20 in Jourdan et al. (2022), there exists $C_\nu > 0$ and N_7 with $\mathbb{E}_\nu[N_7] < +\infty$ such that for all $t \geq N_7$ and all $a \neq a^*$,

$$\frac{\Psi_{t,a}}{t} \geq w_a^\beta + \varepsilon \implies \frac{1}{t} \left(C_t(a^*, a) - \min_{b \neq a^*} C_t(a^*, b) \right) \geq C_\nu.$$

Since $\frac{\log(tK)}{2t} \rightarrow_{\infty} 0$, there exists a deterministic N_8 such that for all $n \geq N_8$, $\frac{\log(tK)}{2t} < C_\nu$. Therefore, for all $n \geq \tilde{N}_7 \stackrel{\text{def}}{=} \max\{N_8, N_7\}$ and all $a \neq a^*$, $\Psi_{t,a}/t \geq w_a^\beta + \varepsilon$ implies that $A_{t+1}^{\text{TCI}} \neq a$. Since $\mathbb{E}_\nu[\tilde{N}_7] = N_8 + \mathbb{E}_\nu[N_7] < +\infty$, this concludes the proof. \blacksquare

Since there is sufficient exploration (Lemma 64) and Properties 5 and 6 of Jourdan et al. (2022) are satisfied (Lemmas 65 and 66), Lemma 67 holds for β -EB-TCI on instances such that $\min_{a \neq b} |\mu_a - \mu_b| > 0$.

Lemma 67 (Lemma 10 in Jourdan et al. (2022)) *Let $\varepsilon > 0$ and T_β^ε as in (23). Then, β -EB-TCI satisfies $\mathbb{E}_\nu[T_\beta^\varepsilon] < +\infty$.*

Combining Lemmas 67 and 58 concludes the proof for β -EB-TCI.

H.2.2. β -EB-EVTCI

Let $a^* = a^*(\mu)$. Similarly, the β -optimal allocation for Gaussian with known variance is defined as

$$w_{\sigma^2, \beta}^*(\mu) = \arg \max_{w \in \Delta_K, w_{a^*} = \beta} \min_{a \neq a^*} \inf_{u \in \mathbb{R}} \sum_{b \in \{a^*, a\}} w_b \frac{(\mu_b - u)^2}{\sigma_b^2},$$

is a singleton, denoted by $\{w^\beta\}$, such that $\min_{a \in [K]} w_a^\beta > 0$.

Lemma 68 *Using the EV-GLR stopping rule with an asymptotically tight family of thresholds, β -EB-EVTCI satisfies that, for all ν with $\min_{a \neq b} |\mu_a - \mu_b| > 0$,*

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\nu[\tau_\delta^{\text{EV}}]}{\log(1/\delta)} \leq T_\beta^*(\mu, \sigma^2).$$

For $\varepsilon > 0$, the random variable $\bar{T}_\beta^\varepsilon$ quantifies the number of samples required for the empirical allocations $\frac{N_t}{t}$ to be ε -close to w^β :

$$\bar{T}_\beta^\varepsilon \stackrel{\text{def}}{=} \inf \left\{ T \geq 1 \mid \forall t \geq T, \left\| \frac{N_t}{t} - w^\beta \right\|_\infty \leq \varepsilon \right\}. \quad (26)$$

Lemma 69 shows that a sufficient condition to obtain an upper bound on the asymptotic expected sample complexity is to show $\mathbb{E}_\nu[\bar{T}_\beta^\varepsilon] < +\infty$ for ε small enough. The proof of Lemma 69 is omitted since it is almost identical to the one of Lemma 58.

Lemma 69 *Let $\delta, \beta \in (0, 1)$. Assume that there exists $\varepsilon_1 > 0$ such that for all $\varepsilon \in (0, \varepsilon_1]$, $\mathbb{E}_\nu[\bar{T}_\beta^\varepsilon] < +\infty$. Combining the EV-GLR stopping rule (5) with an asymptotically tight family of threshold yields an algorithm such that, for all $(\mu, \sigma^2) \in \mathbb{R}^K \times (\mathbb{R}_+^*)^K$ with $|a^*(\mu)| = 1$,*

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\nu[\tau_\delta^{\text{EV}}]}{\log(1/\delta)} \leq T_{\sigma^2, \beta}^*(\mu).$$

Lemma 70 is obtained similarly as Lemma 60, hence we omit the proof.

Lemma 70 *Let S_t^L and \mathcal{I}_t^* as in (24). There exists L_4 with $\mathbb{E}_\nu[(L_4)^\alpha] < +\infty$ for all $\alpha > 0$ such that if $L \geq L_4$, for all t such that $S_t^L \neq \emptyset$, for all $(a, b) \in \mathcal{I}_t^* \times (S_t^L \setminus \mathcal{I}_t^*)$, we have $C_t^{EV}(a, b) \geq LD_\nu$, where $D_\nu > 0$ is a problem dependent constant.*

Lemma 71 is obtained similarly as Lemma 61, hence we omit the proof.

Lemma 71 *Let S_t^L as in (24). There exists L_5 with $\mathbb{E}_\nu[(L_5)^\alpha] < +\infty$ for all $\alpha > 0$ such that for all $L \geq L_5$ and all $t \in \mathbb{N}$,*

$$\forall (a, b) \in S_t^L \times \overline{S_t^L}, \quad C_t^{EV}(a, b) \leq \frac{L^2}{2} W_0 (D_1 + 4W_\mu)^2,$$

where $D_1 > 0$ is a problem dependent constant and W_0, W_μ are the random variables defined in Lemma 57.

Lemma 72 is obtained similarly as Lemma 63. The sole technical difference lies in the necessities of having an finite variance for W_0 , which explains why we consider $n_0 = 6$ for β -EB-EVTCI.

Lemma 72 *Let U_t^L and V_t^L as in (25) and $\mathcal{J}_t^* = \arg \max_{a \in \overline{V_t^L}} \mu_a$. There exists L_6 with $\mathbb{E}_\nu[L_6] < +\infty$ such that if $L \geq L_6$, for all n such that $U_t^L \neq \emptyset$, $B_{t+1}^{EB} \notin V_t^L$ implies $A_{t+1}^{EVTCI} \in V_t^L \cup (\mathcal{J}_t^* \setminus \{B_{t+1}^{EB}\})$.*

Proof Using Lemmas 70, 71 and 62, for all L larger than a random variable L_7 with finite expectation, $B_{t+1}^{EB} \in \mathcal{J}_t^*$ and

$$\begin{aligned} \forall (a, b) \in \mathcal{J}_t^* \times (\overline{V_t^L} \setminus \mathcal{J}_t^*), \quad C_t(a, b) + \log N_{t,b} &\geq L^{5/6} D_\mu + \frac{5}{6} \log L, \\ \forall (a, b) \in \overline{U_t^L} \times U_t^L, \quad C_t(a, b) + \log N_{t,b} &\leq \frac{L^{1/3}}{2} W_0 (D_1 + 4W_\mu)^2 + \frac{1}{6} \log L. \end{aligned}$$

Therefore, for $L \geq L_8 := W_0^2 \left(\frac{D_1 + 4W_\mu}{2D_\nu} \right)^4 + 1$, we have that

$$L^{5/6} D_\mu + \frac{5}{6} \log L > \frac{L^{1/3}}{2} W_0 (D_1 + 4W_\mu)^2 + \frac{1}{6} \log L.$$

Let $L_6 = \max\{L_7, L_8\}$. Therefore, at least one under-sampled arm has transportation cost lower than all the ones that are much sampled. This implies that $A_{t+1}^{EVTCI} \in V_t^L \cup (\mathcal{J}_t^* \setminus \{B_{t+1}^{EB}\})$. Then, we have

$$\mathbb{E}_\nu[L_6] \leq \mathbb{E}_\nu[L_7] + 1 + \mathbb{E}_\nu[W_0^2] \mathbb{E}_\nu \left[\left(\frac{D_1 + 4W_\mu}{2D_\nu} \right)^4 \right] < +\infty.$$

The first inequality and the last strict inequality are obtained by Lemma 57, thanks to the independence of W_0 and W_μ , the fact that W_0 has finite variance since $n_0 \geq 6$ and that polynomial of W_μ have finite expectation. \blacksquare

Since Properties 2 and 3 of Jourdan et al. (2022) are satisfied (Lemmas 62 and 72), Lemma 64 holds for β -EB-EVTCI on instances such that $\min_{a \neq b} |\mu_a - \mu_b| > 0$.

Lemma 73 is obtained similarly as Lemma 66, hence we omit the proof.

Lemma 73 *Let $\varepsilon > 0$. There exists N_7 with $\mathbb{E}_\nu[N_7] < +\infty$ such that for all $n \geq N_7$ and all $a \neq a^*(\mu)$, $\Psi_{t,a}/t \geq w_a^\beta + \varepsilon$ implies that $A_{t+1}^{EVTCl} \neq a$.*

Since there is sufficient exploration (Lemma 64) and Properties 5 and 6 of Jourdan et al. (2022) are satisfied (Lemmas 65 and 73), Lemma 74 holds for β -EB-EVTCl on instances such that $\min_{a \neq b} |\mu_a - \mu_b| > 0$.

Lemma 74 (Lemma 10 in Jourdan et al. (2022)) *Let $\varepsilon > 0$ and $\bar{T}_\beta^\varepsilon$ as in (26). Then, β -EB-EVTCl satisfies $\mathbb{E}_\nu[\bar{T}_\beta^\varepsilon] < +\infty$.*

Combining Lemmas 74 and 69 concludes the proof of the asymptotic upper bound on the expected sample complexity of β -EB-EVTCl, i.e. Lemma 68.

In Appendix C.1, we explained that $T_{\sigma^2, \beta}^*(\mu)$ and $T_\beta^*(\mu, \sigma^2)$ satisfied inequalities like the ones between $T_{\sigma^2}^*(\mu)$ and $T^*(\mu, \sigma^2)$ in (2). Similar to Lemma 12, we can show that there exists a sequence of instances $(\nu_n)_{n \in \mathbb{N}}$ with $\min_{b \neq a} |\mu_{n,a} - \mu_{n,b}| > 0$ such that $\lim_{n \rightarrow +\infty} T_\beta^*(\mu_n, \sigma_n^2) / T_{\sigma_n^2, \beta}^*(\mu_n) = +\infty$. Therefore, we can conclude the impossibility result for β -EB-EVTCl (Theorem 10) with the same arguments as for EV-TaS in Appendix H.1.

Appendix I. The Lambert W Function

The Lambert W function is implicitly defined by the equation $W(x)e^{W(x)} = x$. It defines two main branches W_{-1} (negative) and W_0 (positive).

- W_{-1} , defined on $[-e^{-1}, 0)$, is decreasing and $W_{-1}(-e^{-1}) = -1$.
- W_0 , defined on $[-e^{-1}, +\infty)$, is increasing and $W_0(-e^{-1}) = -1$.

The function W_0 satisfies for all $x \geq e$, $W_0(e^x) \leq x$ and

$$\frac{\log \log(x)}{2 \log(x)} \leq W_0(x) - (\log(x) - \log \log(x)) \leq \frac{e}{e-1} \frac{\log \log(x)}{\log(x)}.$$

Lambert's branches are involved in the inversion of $h(x) = x - \log(x)$. When $x \geq 1$, it involves the negative branch. When $x \leq 1$, it involves the negative part of the positive branch. To make the notations clearer, we define for all $x \geq 1$

$$\bar{W}_{-1}(x) = -W_{-1}(-e^{-x}) \quad \text{and} \quad \bar{W}_0(x) = -W_0(-e^{-x}). \quad (27)$$

Lemma 75 gather useful properties on \bar{W}_{-1} and \bar{W}_0 that we will use.

Lemma 75 (1) For $x \geq 1$, let $h(x) = x - \log(x)$. Then,

$$\begin{aligned} \forall y \geq 1, \quad y \leq h(x) &\iff \begin{cases} \bar{W}_{-1}(y) \leq x & \text{if } x \geq 1 \\ \bar{W}_0(y) \geq x & \text{if } x \in (0, 1] \end{cases}, \\ \forall \delta > 0, \forall c > 0, \quad e^{-c(h(x)-1)} \leq \delta &\iff \begin{cases} \bar{W}_{-1}\left(1 + \frac{1}{c} \log \frac{1}{\delta}\right) \leq x & \text{if } x > 1 \\ \bar{W}_0\left(1 + \frac{1}{c} \log \frac{1}{\delta}\right) \geq x & \text{if } x \in (0, 1) \end{cases}, \\ \forall x > 1, \quad \exp(-x + e^{-x}) \leq \bar{W}_0(x) \leq \exp(-x + e^{1-x}), \\ \forall x > 1, \quad x + \log(x) \leq \bar{W}_{-1}(x) \leq x + \log(x) + \min\left\{\frac{1}{2}, \frac{1}{\sqrt{x}}\right\}, \\ \forall u > 1, \forall t > 1, \quad \bar{W}_0\left(1 + \frac{u}{t}\right) \geq \frac{1}{t} &\iff t \geq \exp\left(1 + W_0\left(\frac{u-1}{e}\right)\right). \end{aligned}$$

(2) The function \bar{W}_{-1} is increasing and strictly concave on $(1, +\infty)$. The function \bar{W}_0 is decreasing and strictly convex on $(1, +\infty)$. In particular,

$$\forall x > 1, \quad (\bar{W}_0)'(x) = \left(1 - \frac{1}{\bar{W}_0(x)}\right)^{-1} \quad \text{and} \quad (\bar{W}_{-1})'(x) = \left(1 - \frac{1}{\bar{W}_{-1}(x)}\right)^{-1}.$$

Proof (1) Let $y \geq 1$ and $x \in (0, 1]$. We obtain

$$\bar{W}_0(y) \geq x \iff W_0(-e^{-y}) \leq -x \iff -e^{-y} \leq -xe^{-x} \iff y \leq x - \log(x)$$

where the second equivalence uses that $-e^{-y} = W_0(-e^{-y})e^{W_0(-e^{-y})}$, $y \mapsto ye^y$ is increasing on $[-1, +\infty)$ and $W_0(x)$ has values on $[-1, 0)$ for $x \in [-e^{-1}, 0)$. Let $x \in (0, 1)$, $\delta, c > 0$. Then,

$$\bar{W}_0\left(1 + \frac{1}{c} \log \frac{1}{\delta}\right) \geq x \iff 1 + \frac{1}{c} \log \frac{1}{\delta} \leq h(x) \iff \exp(-c(h(x) - 1)) \leq \delta$$

Let $x > 1$ and $f(x) \in (0, 1)$. Then, we obtain

$$\bar{W}_0(x) \geq f(x) \iff x \leq f(x) - \log(f(x))$$

For $f(x) = e^{-x+e^{-x}}$, we have $x \leq f(x) - \log(f(x)) \iff e^{-x} \geq 0$, hence this condition holds and $\bar{W}_0(x) \geq f(x)$. For $f(x) = e^{-x+e^{1-x}}$, we have $x \leq f(x) - \log(f(x)) \iff x \leq 1$, hence this condition doesn't hold for $x > 1$, hence $\bar{W}_0(x) \leq f(x)$.

For $\bar{W}_{-1}(y)$, the same arguments yield the three results, which were first proven in Lemma A.1 and A.2 of [Degenne \(2019\)](#).

We denote $v = \frac{u-1}{t} > 0$. Since $t > 1$, direct manipulations show that

$$\begin{aligned} \bar{W}_0\left(1 + \frac{u}{t}\right) \geq \frac{1}{t} &\iff 1 + \frac{u}{t} \leq \frac{1}{t} - \log\left(\frac{1}{t}\right) \iff v + \log(v) \leq \log\left(\frac{u-1}{e}\right) \\ &\iff ve^v \leq \frac{u-1}{e} \iff v \leq W_0\left(\frac{u-1}{e}\right) \iff t \geq \frac{u-1}{W_0\left(\frac{u-1}{e}\right)} = e^{1+W_0\left(\frac{u-1}{e}\right)} \end{aligned}$$

The equivalence introducing W_0 uses that for $\alpha = \frac{u-1}{e} > 0$, $W_0(\alpha)e^{W_0(\alpha)} = \alpha$, $y \mapsto ye^y$ is increasing on $[-1, +\infty)$ and $v > 0$. The last equality uses that $e^{W_0(x)} = \frac{x}{W_0(x)}$.

(2) Let W denote W_0 or W_{-1} and $\bar{W}(x) = -W(-e^{-x})$. It is known (by implicit derivation) that $W'(z) = \frac{1}{z+e^{W(z)}}$ for $z \neq -e^{-1}$. Using that $e^{W(z)} = \frac{z}{\bar{W}(z)}$, this yields that $zW'(z) = \left(1 + \frac{e^{W(z)}}{z}\right)^{-1} = \left(1 + \frac{1}{\bar{W}(z)}\right)^{-1}$. For $x \neq 1$, using the above with $z = -e^{-x}$, we obtain

$$\bar{W}'(x) = -\frac{d}{dx} (W(-e^{-x})) = -e^{-x}W'(-e^{-x}) = \left(1 + \frac{1}{\bar{W}(-e^{-x})}\right)^{-1} = \left(1 - \frac{1}{\bar{W}(x)}\right)^{-1}$$

Since $W_0(-e^{-x}) \in (-1, 0)$ for all $x > 1$ (positive branch on $(-e^{-1}, 0)$), we have $\bar{W}_0(x) \in (0, 1)$, hence $\bar{W}'_0(x) < 0$ for $x > 1$. Therefore, \bar{W}_0 is decreasing on $(1, +\infty]$. Using that $\bar{W}'(x) = \left(1 - \frac{1}{\bar{W}(x)}\right)^{-1}$ for $x \neq 1$, we obtain that \bar{W}'_0 is increasing on $(1, +\infty]$, hence strictly convex. The same arguments yield that \bar{W}_{-1} is increasing and strictly concave on $(1, +\infty]$. ■

Lemma 76 was proven in [Degenne \(2019\)](#). It is needed when using the peeling method.

Lemma 76 (Lemma A.3 in [Degenne \(2019\)](#)) *For $a, b \geq 1$, the minimal value of $f(\eta) = (1 + \eta)(a + \log(b + \frac{1}{\eta}))$ is attained at η^* such that $f(\eta^*) \leq 1 - b + \bar{W}_{-1}(a + b)$. If $b = 1$, then there is equality.*

Appendix J. Implementation Details and Additional Experiments

Implementations details are given in [Appendix J.1](#). Supplementary experiments are provided in [Appendix J.2](#).

J.1. Implementation Details

We detail below the most relevant implementation details, both regarding the sampling rules and the considered family of thresholds.

Stopping thresholds The implementation of the different families of thresholds require to specify the hyper-parameters $(s, \eta_0, \eta_1, \gamma)$. Based on the obtained formulas, we see that there is a trade-off in their choice. Smaller values of (η_0, η_1) yield close to ideal dependency in δ , at the cost of increasing the dependency in t . Smaller values of s yield close to ideal dependency in t , at the cost of increasing the constant term. Smaller values of γ yield more frequent update of the approximation term error (i.e. the ratio of eigenvalues), at the cost of increasing the dependency in t . Those updates are responsible for stair-step shaped curves of the KL and (EV-)BoB thresholds. For reasonable choices of the parameters, the relative performance of the thresholds are not changed. We conduct our experiments with $s = 2$, $\gamma = 1.2$, $\eta_0 = \eta_1 = \log(1/\delta)^{-1}$ (as $\delta < 1$).

Optimal allocation oracles To compute the optimal allocation in TaS, we perform nested binary searches to solve the optimization problem described in [Theorem 13](#). The outer binary search is done on $y \in [0, \min_{a \neq a^*} d_{a^*}(\mu_a)]$. The inner binary searches are done to compute $x_a(y)$ for all $a \in [K]$. To obtain $\lambda(x)$, we compute the (at most three) real solutions of the third order polynomial equation, and then return the one minimizing the function of interest. For EV-TaS, the procedure is similar as the one implemented in [Garivier and Kaufmann \(2016\)](#), and it is significantly faster as $\lambda(x)$ has a closed-form solution.

Sampling rules As recommended in Jourdan et al. (2022), we implemented the β -EB-TCI algorithm with $\beta = 0.5$. This is a deterministic instance of Top Two algorithm with an efficient implementation since it can re-use computations done in the stopping and recommendation rules.

Based on Degenne et al. (2019), we consider DKM with one learner on Δ_K instead of K learners. For unknown variances, the wrapped DKM algorithm is obtained by considering adapted transportation costs (and the corresponding alternatives). To cope for an additional uncertainty due to the unknown variance, the bonus $\frac{\log(t)}{N_{t,a}}$ is multiplied by $1 + 2\frac{\log(t)}{N_{t,a}}$. This heuristic is legitimated by our box concentration. The implementation of EV-DKM is direct by plugging in the empirical variance. Based on Wang et al. (2021), the wrapped algorithms FWS and EV-FWS are obtained similarly.

FHN2 refers to the procedure 2 in Fan et al. (2016), whose pseudo-code (see Algorithm 1) was given in Hong et al. (2021). FHN2 is an elimination strategy which repeatedly samples all arms until only one arm is left. For all $a \in [K]$, let $\bar{X}_{n,a} = \frac{1}{n} \sum_{t \in [n]} X_{t,a}$ where $(X_{t,a})_{t \in [n]}$ are i.i.d. observations from ν_{μ_a, σ_a^2} . Its elimination mechanism compares the pairwise statistics $t_{b,a}(n)(\bar{X}_{n,b} - \bar{X}_{n,a})$ to an elimination threshold $g_{b,a}(t_{b,a}(n), \delta)$, where the effective time $t_{b,a}(n) = n/S_{b,a}(n)$ is defined with the empirical variances of pairwise comparison between arms b and a

$$S_{b,a}(n) = \frac{1}{n-1} \sum_{t \in [n]} (X_{t,b} - X_{t,a} - \bar{X}_{n,b} + \bar{X}_{n,a})^2.$$

The calibration of the elimination threshold, which is done by simulations arguments and continuous-time approximations, yields

$$g_{b,a}(x, \delta) = \sqrt{(x+1) \left(2 \log \left(\frac{K-1}{2\delta} \right) + \log(x+1) \right)}.$$

The choice of the initial time n_0 impacts the empirical performance of the algorithm, yet few practical guidelines were given. In our experiments we took $n_0 = \max\{2, 10 \log(1/\delta)\}$.

Algorithm 1: FHN2

Let $I = [K]$ and $n = n_0$ (user-specified)

Sample n_0 observations from each arm

while $|I| > 1$ **do**

 Compute $S_{b,a}(n)$ and set $t_{b,a}(n) = n/S_{b,a}(n)$ for all $b \neq a$

 Update active arms, $I \leftarrow I \setminus \{b \mid \forall a \neq b, t_{b,a}(n)(\bar{X}_{n,b} - \bar{X}_{n,a}) < -g_{b,a}(t_{b,a}(n), \delta)\}$

 Pull each arm once, and set $n \leftarrow n + 1$

end

Return the unique arm in I

In Fan et al. (2016), they also propose the procedure 1, which we refer to as FHN1. The main difference is that FHN1 uses the empirical variances obtained after initialization, meaning it uses $t_{b,a}(n) = n/S_{b,a}(n_0)$ instead of $t_{b,a}(n) = n/S_{b,a}(n)$. While one can argue that this wastes precious information, FHN1 enjoys better theoretical guaranty. The impact of then choice of n_0 is more important for FHN1 since $S_{b,a}(n_0)$ is used for elimination.

Reproducibility Our code is implemented in Julia 1.7.2, and the plots are generated with the StatsPlots.jl package. Optimizations are performed based on the JuMP.jl optimization

package. In particular, we use the Ipopt non-linear solver (Wächter and Biegler, 2006), available in `Ipopt.jl`, and the `LambertW.jl` package to compute the Lambert W function. Other dependencies are listed in the `Readme.md`. The `Readme.md` file also provides detailed julia instructions to reproduce our experiments, as well as a `script.sh` to run them all at once. The general structure of the code (and some functions) is taken from the `tidnabbil` library.²

J.2. Supplementary Experiments

In Appendix J.2.1, we perform simulations on the characteristic times $T^*(\mu, \sigma^2)$ and $T_{\sigma^2}^*(\mu)$. In Appendix J.2.2, we compare family of stopping thresholds for the EV-GLR stopping rule.

J.2.1. CHARACTERISTIC TIMES

We perform numerical simulations to compare $T^*(\mu, \sigma^2)$ and $T_{\sigma^2}^*(\mu)$. Based on the inequalities (2), we will compare $\frac{T^*(\mu, \sigma^2)}{T_{\sigma^2}^*(\mu)}$ and its upper bound $\frac{d(\mu, \sigma^2)}{\log(1+d(\mu, \sigma^2))}$ on $\frac{T^*(\mu, \sigma^2)}{T_{\sigma^2}^*(\mu)}$, where $d(\mu, \sigma^2) = \max_{a \neq a^*(\mu)} \frac{(\mu_{a^*(\mu)} - \mu_a)^2}{\min\{\sigma_a^2, \sigma_{a^*(\mu)}^2\}}$. Taking $K = 2$, we consider the canonical instance $\mu = (0, -\Delta)$ and $\sigma^2 = (1, r)$, hence $d(\mu, \sigma^2) = \frac{\Delta^2}{\min\{1, r\}}$, and instantiate it by default with $\Delta = 0.2$ and $r = 0.5$. We perform 10000 simulations for varying Δ and r .




Figure 3: Evolution of $\frac{T^*(\mu, \sigma^2)}{T_{\sigma^2}^*(\mu)}$ as a function of (a) Δ for $r = 0.5$ and (b) r for $\Delta = 0.2$.

Figure 3 empirically confirms that $\frac{d(\mu, \sigma^2)}{\log(1+d(\mu, \sigma^2))}$ is an upper bound on $\frac{T^*(\mu, \sigma^2)}{T_{\sigma^2}^*(\mu)}$, which is always above 1. It shows that $\frac{T^*(\mu, \sigma^2)}{T_{\sigma^2}^*(\mu)}$ is close to one as long as the gap is not too large. Moreover, the upper bound $\frac{d(\mu, \sigma^2)}{\log(1+d(\mu, \sigma^2))}$ has the same behavior, except when the variance of the optimal arm is significantly larger than the one of the sub-optimal arm. In this particular regime, our upper bound appears to be loose.

2. This library was created by Degenne et al. (2019), see <https://bitbucket.org/wmkoolen/tidnabbil>. No license were available on the repository, but we obtained the authorization from the authors.

J.2.2. THRESHOLDS

Using the same experimental setup as in Section 3.3, we perform numerical simulations to compare the family of thresholds derived for the EV-GLR stopping rule: EV-Student (Lemma 50), EV-Box (Lemma 51) and EV-BoB (Corollary 52).




Figure 4: Threshold for (5) as a function of (a) $\log(1/\delta)$ for $t = 5000$ and (b) t for $\delta = 0.01$.

Figure 4 is the equivalent of Figure 1 and similar messages can be extracted from it. However, for the EV-GLR stopping rule, the δ -correct family of thresholds can be asymptotically tight, otherwise Theorem 9 would contradict the asymptotic lower bound. The Student threshold in (21) has poor performance in t . While the box threshold in (22) yields better performance in t , it comes at a worse dependency in $\log(1/\delta)$ (higher slope). The BoB threshold in Corollary 52 has the good performance in t of the box threshold and the same dependency in $\log(1/\delta)$ as the Student threshold.




Figure 5: Thresholds as a function of (a) $\log(1/\delta)$ for $t = 5000$ and (b) t for $\delta = 0.01$.

Figure 5 groups Figures 1 and 4 together. By comparing the thresholds obtained by using the same method, all the thresholds for the EV-GLR stopping rule are above the ones obtained for the GLR stopping rule. Given the ordering of the statistic (Lemma 16), this was expected.