
On Explore-Then-Commit Strategies

Aurélien Garivier*

Institut de Mathématiques de Toulouse; UMR5219
Université de Toulouse; CNRS
UPS IMT, F-31062 Toulouse Cedex 9, France
aurelien.garivier@math.univ-toulouse.fr

Emilie Kaufmann

Univ. Lille, CNRS, Centrale Lille, Inria SequeL
UMR 9189, CRISTAL - Centre de Recherche en Informatique Signal et Automatique de Lille
F-59000 Lille, France
emilie.kaufmann@univ-lille1.fr

Tor Lattimore

University of Alberta
116 St & 85 Ave, Edmonton, AB T6G 2R3, Canada
tor.lattimore@gmail.com

Abstract

We study the problem of minimising regret in two-armed bandit problems with Gaussian rewards. Our objective is to use this simple setting to illustrate that strategies based on an exploration phase (up to a stopping time) followed by exploitation are necessarily suboptimal. The results hold regardless of whether or not the difference in means between the two arms is known. Besides the main message, we also refine existing deviation inequalities, which allow us to design fully sequential strategies with finite-time regret guarantees that are (a) asymptotically optimal as the horizon grows and (b) order-optimal in the minimax sense. Furthermore we provide empirical evidence that the theory also holds in practice and discuss extensions to non-gaussian and multiple-armed case.

1 Introduction

It is now a very frequent issue for companies to optimise their daily profits by choosing between one of two possible website layouts. A natural approach is to start with a period of A/B Testing (exploration) during which the two versions are uniformly presented to users. Once the testing is complete, the company displays the version believed to generate the most profit for the rest of the month (exploitation). The time spent exploring may be chosen adaptively based on past observations, but could also be fixed in advance. Our contribution is to show that strategies of this form are much worse than if the company is allowed to dynamically select which website to display without restrictions for the whole month.

Our analysis focusses on a simple sequential decision problem played over T time-steps. In time-step $t \in 1, 2, \dots, T$ the agent chooses an action $A_t \in \{1, 2\}$ and receives a normally distributed reward

*This work was partially supported by the CIMI (Centre International de Mathématiques et d'Informatique) Excellence program while Emilie Kaufmann visited Toulouse in November 2015. The authors acknowledge the support of the French Agence Nationale de la Recherche (ANR), under grants ANR-13-BS01-0005 (project SPADRO) and ANR-13-CORD-0020 (project ALICIA).

$Z_t \sim \mathcal{N}(\mu_{A_t}, 1)$ where $\mu_1, \mu_2 \in \mathbb{R}$ are the unknown mean rewards for actions 1 and 2 respectively. The goal is to find a strategy π (a way of choosing each action A_t based on past observation) that maximises the cumulative reward over T steps in expectation, or equivalently minimises the regret

$$R_\mu^\pi(T) = T \max\{\mu_1, \mu_2\} - \mathbb{E}_\mu \left[\sum_{t=1}^T \mu_{A_t} \right]. \quad (1)$$

This framework is known as the multi-armed bandit problem, which has many applications and has been studied for almost a century [Thompson, 1933]. Although this setting is now quite well understood, the purpose of this article is to show that strategies based on distinct phases of exploration and exploitation are necessarily suboptimal. This is an important message because exploration followed by exploitation is the most natural approach and is often implemented in applications (including the website optimisation problem described above). Moreover, strategies of this kind have been proposed in the literature for more complicated settings [Auer and Ortner, 2010, Perchet and Rigollet, 2013, Perchet et al., 2015]. Recent progress on optimal exploration policies (e.g., by Garivier and Kaufmann [2016]) could have suggested that well-tuned variants of two-phase strategies might be near-optimal. We show, on the contrary, that optimal strategies for multi-armed bandit problems *must* be fully-sequential, and in particular should mix exploration and exploitation. It is known since the work of Wald [1945] on simple hypothesis testing that sequential procedures can lead to significant gains. Here, the superiority of fully sequential procedures is consistent with intuition: if one arm first appears to be better, but if subsequent observations are disappointing, the obligation to commit at some point can be restrictive. In this paper, we give a crisp and precise description of how restrictive it is: it leads to regret asymptotically twice as large on average. The proof of this result combines some classical techniques of sequential analysis and of the bandit literature.

We study two settings, one when the gap $\Delta = |\mu_1 - \mu_2|$ is known and the other when it is not. The most straight-forward strategy in the former case is to explore each action a fixed number of times n and subsequently exploit by choosing the action that appeared best while exploring. It is easy to calculate the optimal n and consequently show that this strategy suffers a regret of $R_\mu^\pi(T) \sim 4 \log(T)/\Delta$. A more general approach is to use a so-called *Explore-Then-Commit* (ETC) strategy, following a nomenclature introduced by Perchet et al. [2015]. An ETC strategy explores each action alternately until some data-dependent stopping time and subsequently commits to a single action for the remaining time-steps. We show in Theorem 2 that by using a sequential probability ratio test (SPRT) it is possible to design an ETC strategy for which $R_\mu^\pi(T) \sim \log(T)/\Delta$, which improves on the above result by a factor of 4. We also prove a lower bound showing that no ETC strategy can improve on this result. Surprisingly it is possible to do even better by using a fully sequential strategy inspired by the UCB algorithm for multi-armed bandits [Katehakis and Robbins, 1995]. We design a new strategy for which $R_\mu^\pi(T) \sim \log(T)/(2\Delta)$, which improves on the fixed-design strategy by a factor of 8 and on SPRT by a factor of 2. Again we prove a lower bound showing that no strategy can improve on this result.

For the case where Δ is unknown, fixed-design strategies are hopeless because there is no reasonable tuning for the exploration budget n . However, it is possible to design an ETC strategy for unknown gaps. Our approach uses a modified fixed-budget best arm identification (BAI) algorithm in its exploration phase (see e.g., Even-Dar et al. [2006], Garivier and Kaufmann [2016]) and chooses the recommended arm for the remaining time-steps. In Theorem 5 we show that a strategy based on this idea satisfies $R_\mu^\pi(T) \sim 4 \log(T)/\Delta$, which again we show is optimal within the class of ETC strategies. As before, strategies based on ETC are suboptimal by a factor of 2 relative to the optimal rates achieved by fully sequential strategies such as UCB, which satisfies $R_\mu^\pi(T) \sim 2 \log(T)/\Delta$ [Katehakis and Robbins, 1995].

In a nutshell, strategies based on fixed-design or ETC are necessarily suboptimal. That this failure occurs even in the simple setting considered here is a strong indicator that they are suboptimal in more complicated settings. Our main contribution, presented in more details in Section 2, is to fully characterise the achievable asymptotic regret when Δ is either known or unknown and the strategies are either fixed-design, ETC or fully sequential. All upper bounds have explicit finite-time forms, which allow us to derive optimal minimax guarantees. For the lower bounds we give a novel and generic proof of all results. All proofs contain new, original ideas that we believe are fundamental to the understanding of sequential analysis.

2 Notation and Summary of Results

We assume that the horizon T is known to the agent. The optimal action is $a^* = \arg \max(\mu_1, \mu_2)$, its mean reward is $\mu^* = \mu_{a^*}$, and the gap between the means is $\Delta = |\mu_1 - \mu_2|$. Let $\mathcal{H} = \mathbb{R}^2$ be the set of all possible pairs of means, and $\mathcal{H}_\Delta = \{\mu \in \mathbb{R}^2 : |\mu_1 - \mu_2| = \Delta\}$. For $i \in \{1, 2\}$ and $n \in \mathbb{N}$ let $\hat{\mu}_{i,n}$ be the empirical mean of the i th action based on the first n samples. Let A_t be the action chosen in time-step t and $N_i(t) = \sum_{s=1}^t \mathbb{1}\{A_s = i\}$ be the number of times the i th action has been chosen after time-step t . We denote by $\hat{\mu}_i(t) = \hat{\mu}_{i,N_i(t)}$ the empirical mean of the i th arm after time-step t .

A strategy is denoted by π , which is a function from past actions/rewards to a distribution over the next actions. An ETC strategy is governed by a sampling rule (which determines which arm to sample at each step), a stopping rule (which specifies when to stop the exploration phase) and a decision rule indicating which arm is chosen in the exploitation phase. As we consider two-armed, Gaussian bandits with equal variances, we focus here on uniform sampling rules, which have been shown in Kaufmann et al. [2014] to be optimal in that setting. For this reason, we define an ETC strategy as a pair (τ, \hat{a}) , where τ is an even stopping time with respect to the filtration $(\mathcal{F}_t = \sigma(Z_1, \dots, Z_t))_t$ and $\hat{a} \in \{1, 2\}$ is \mathcal{F}_τ -measurable. In all the ETC strategies presented in this paper, the stopping time τ depends on the horizon T (although this is not reflected in the notation). At time t , the action

$$\text{picked by the ETC strategy is } A_t = \begin{cases} 1 & \text{if } t \leq \tau \text{ and } t \text{ is odd,} \\ 2 & \text{if } t \leq \tau \text{ and } t \text{ is even,} \\ \hat{a} & \text{otherwise.} \end{cases}$$

The regret for strategy π , given in Eq. (1), depends on T and μ . Assuming, for example that $\mu_1 = \mu_2 + \Delta$, then an ETC strategy π chooses the suboptimal arm $N_2(T) = \frac{\tau \wedge T}{2} + (T - \tau)_+ \mathbb{1}\{\hat{a} = 2\}$ times, and the regret $R_\mu^\pi(T) = \Delta \mathbb{E}_\mu[N_2(T)]$ thus satisfies

$$\Delta \mathbb{E}_\mu[(\tau \wedge T)/2] \leq R_\mu^\pi(T) \leq (\Delta/2) \mathbb{E}_\mu[\tau \wedge T] + \Delta T \mathbb{P}_\mu(\tau \leq T, \hat{a} \neq a^*). \quad (2)$$

We denote the set of all ETC strategies by Π_{ETC} . A fixed-design strategy is an ETC strategy for which there exists an integer n such that $\tau = 2n$ almost surely, and the set of all such strategies is denoted by Π_{DETC} . The set of all strategies is denoted by Π_{ALL} . For $\mathcal{S} \in \{\mathcal{H}, \mathcal{H}_\Delta\}$, we are interested in strategies π that are uniformly efficient on \mathcal{S} , in the sense that

$$\forall \mu \in \mathcal{S}, \forall \alpha > 0, R_\mu^\pi(T) = o(T^\alpha). \quad (3)$$

We show in this paper that any uniformly efficient strategy in Π has a regret at least equal to $C_S^\Pi \log(T)/|\mu_1 - \mu_2|(1 - o_T(1))$ for every parameter $\mu \in \mathcal{S}$, where C_S^Π is given in the adjacent table. Furthermore, we prove that these results are tight. In each case, we propose a uniformly efficient strategy matching this bound. In addition, we prove a tight and non-asymptotic regret bound which also implies, in particular, minimax rate-optimality.

	Π_{ALL}	Π_{ETC}	Π_{DETC}
\mathcal{H}	2	4	NA
\mathcal{H}_Δ	1/2	1	4

The paper is organised as follows. First we consider ETC and fixed-design strategies when Δ known and unknown (Section 3). We then analyse fully sequential strategies that interleave exploration and exploitation in an optimal way (Section 4). For known Δ we present a novel algorithm that exploits the additional information to improve the regret. For unknown Δ we briefly recall the well-known results, but also propose a new regret analysis of the UCB* algorithm, a variant of UCB that can be traced back to Lai [1987], for which we also obtain order-optimal minimax regret. Numerical experiments illustrate and empirically support our results in Section 5. We conclude with a short discussion on non-uniform exploration, and on models with more than 2 arms, possibly non Gaussian. All the proofs are given in the supplementary material. In particular, our simple, unified proof for all the lower bounds is given in Appendix A.

3 Explore-Then-Commit Strategies

Fixed Design Strategies for Known Gaps. As a warm-up we start with the fixed-design ETC setting where Δ is known and where the agent chooses each action n times before committing for the remainder.

The optimal decision rule is obviously $\hat{a} = \arg \max_i \hat{\mu}_{i,n}$ with ties broken arbitrarily. The formal description of the strategy is given in Algorithm 1, where W denotes the Lambert function implicitly defined for $y > 0$ by $W(y) \exp(W(y)) = y$. We denote the regret associated to the choice of n by $R_\mu^n(T)$. The following theorem is not especially remarkable except that the bound is sufficiently refined to show certain negative lower-order terms that would otherwise not be apparent.

```

input:  $T$  and  $\Delta$ 
 $n := \left\lceil 2W(T^2 \Delta^4 / (32\pi)) / \Delta^2 \right\rceil$ 
for  $k \in \{1, \dots, n\}$  do
    choose  $A_{2k-1} = 1$  and  $A_{2k} = 2$ 
end for
 $\hat{a} := \arg \max_i \hat{\mu}_{i,n}$ 
for  $t \in \{2n+1, \dots, T\}$  do
    choose  $A_t = \hat{a}$ 
end for

```

Algorithm 1: FB-ETC algorithm

Theorem 1. Let $\mu \in \mathcal{H}_\Delta$, and let

$$\bar{n} = \left\lceil \frac{2}{\Delta^2} W\left(\frac{T^2 \Delta^4}{32\pi}\right) \right\rceil. \quad \text{Then} \quad R_\mu^{\bar{n}}(T) \leq \frac{4}{\Delta} \log\left(\frac{T \Delta^2}{4.46}\right) - \frac{2}{\Delta} \log \log\left(\frac{T \Delta^2}{4\sqrt{2\pi}}\right) + \Delta$$

whenever $T \Delta^2 > 4\sqrt{2\pi}e$, and $R_\mu^{\bar{n}}(T) \leq T \Delta / 2 + \Delta$ otherwise. In all cases, $R_\mu^{\bar{n}}(T) \leq 2.04\sqrt{T} + \Delta$. Furthermore, for all $\varepsilon > 0$, $T \geq 1$ and $n \leq 4(1 - \varepsilon) \log(T) / \Delta^2$,

$$R_\mu^n(T) \geq \left(1 - \frac{2}{n \Delta^2}\right) \left(1 - \frac{8 \log(T)}{\Delta^2 T}\right) \frac{\Delta T^\varepsilon}{2\sqrt{\pi \log(T)}}.$$

As $R_\mu^n(T) \geq n \Delta$, this entails that $\inf_{1 \leq n \leq T} R_\mu^n(T) \sim 4 \log(T) / \Delta$.

The proof of Theorem 1 is in Appendix B. Note that the "asymptotic lower bound" $4 \log(T) / \Delta$ is actually not a lower bound, even up to an additive constant: $R_\mu^{\bar{n}}(T) - 4 \log(T) / \Delta \rightarrow -\infty$ when $T \rightarrow \infty$. Actually, the same phenomenon applies many other cases, and it should be no surprise that, in numerical experiments, some algorithm reach a regret smaller than Lai and Robbins asymptotic lower bound, as was already observed in several articles (see e.g. Garivier et al. [2016]). Also note that the term Δ at the end of the upper bound is necessary: if Δ is large, the problem is statistically so simple that one single observation is sufficient to identify the best arm; but that observation cannot be avoided.

Explore-Then-Commit Strategies for Known Gaps. We now show the existence of ETC strategies that improve on the optimal fixed-design strategy. Surprisingly, the gain is significant. We describe an algorithm inspired by ideas from hypothesis testing and prove an upper bound on its regret that is minimax optimal and that asymptotically matches our lower bound.

Let P be the law of $X - Y$, where X (resp. Y) is a reward from arm 1 (resp. arm 2). As Δ is known, the exploration phase of an ETC algorithm can be viewed as a statistical test of the hypothesis $H_1 : (P = \mathcal{N}(\Delta, 2))$ against $H_2 : (P = \mathcal{N}(-\Delta, 2))$. The work of Wald [1945] shows that a significant gain in terms of expected number of samples can be obtained by using a sequential rather than a batch test. Indeed, for a batch test, a sample size of $n \sim (4/\Delta^2) \log(1/\delta)$ is necessary to guarantee that both type I and type II errors are upper bounded by δ . In contrast, when a random number of samples is permitted, there exists a sequential probability ratio test (SPRT) with the same guarantees that stops after a random number N of samples with expectation $\mathbb{E}[N] \sim \log(1/\delta) / \Delta^2$ under both H_1 and H_2 . The SPRT stops when the absolute value of the log-likelihood ratio between H_1 and H_2 exceeds some threshold. Asymptotic upper bound on the expected number of samples used by a SPRT, as well as the (asymptotic) optimality of such procedures among the class of all sequential tests can be found in [Wald, 1945, Siegmund, 1985].

Algorithm 2 is an ETC strategy that explores each action alternately, halting when sufficient confidence is reached according to a SPRT. The threshold depends on the gap Δ and the horizon T corresponding to a risk of $\delta = 1/(T \Delta^2)$. The exploration phase ends at the stopping time

$$\tau = \inf \left\{ t = 2n : |\hat{\mu}_{1,n} - \hat{\mu}_{2,n}| \geq \frac{\log(T \Delta^2)}{n \Delta} \right\}.$$

If $\tau < T$ then the empirical best arm \hat{a} at time τ is played until time T . If $T \Delta^2 \leq 1$, then $\tau = 1$

```

input:  $T$  and  $\Delta$ 
 $A_1 = 1, A_2 = 2, t := 2$ 
while  $(t/2)\Delta |\hat{\mu}_1(t) - \hat{\mu}_2(t)| < \log(T \Delta^2)$  do
    choose  $A_{t+1} = 1$  and  $A_{t+2} = 2$ ,
     $t := t + 2$ 
end while
 $\hat{a} := \arg \max_i \hat{\mu}_i(t)$ 
while  $t \leq T$  do
    choose  $A_t = \hat{a}$ ,
     $t := t + 1$ 
end while

```

Algorithm 2: SPRT ETC algorithm

(one could even define $\tau = 0$ and pick a random arm). The following theorem gives a non-asymptotic upper bound on the regret of the algorithm. The results rely on non-asymptotic upper bounds on the expectation of τ , which are interesting in their own right.

Theorem 2. *If $T\Delta^2 \geq 1$, then the regret of the SPRT-ETC algorithm is upper-bounded as*

$$R_\mu^{\text{SPRT-ETC}}(T) \leq \frac{\log(eT\Delta^2)}{\Delta} + \frac{4\sqrt{\log(T\Delta^2)} + 4}{\Delta} + \Delta.$$

Otherwise it is upper bounded by $T\Delta/2 + \Delta$, and for all T and Δ the regret is less than $10\sqrt{T/e} + \Delta$.

The proof of Theorem 2 is given in Appendix C. The following lower bound shows that no uniformly efficient ETC strategy can improve on the asymptotic regret of Algorithm 2. The proof is given in Section A together with the other lower bounds.

Theorem 3. *Let π be an ETC strategy that is uniformly efficient on \mathcal{H}_Δ . Then for all $\mu \in \mathcal{H}_\Delta$,*

$$\liminf_{T \rightarrow \infty} \frac{R_\mu^\pi(T)}{\log(T)} \geq \frac{1}{\Delta}.$$

Explore-Then-Commit Strategies for Unknown Gaps. When the gap is unknown it is not possible to tune a fixed-design strategy that achieves logarithmic regret. ETC strategies can enjoy logarithmic regret and these are now analysed. We start with the asymptotic lower bound.

Theorem 4. *Let π be a uniformly efficient ETC strategy on \mathcal{H} . For all $\mu \in \mathcal{H}$, if $\Delta = |\mu_1 - \mu_2|$ then*

$$\liminf_{T \rightarrow \infty} \frac{R_\mu^\pi(T)}{\log(T)} \geq \frac{4}{\Delta}.$$

A simple idea for constructing an algorithm that matches the lower bound is to use a (fixed-confidence) best arm identification algorithm for the exploration phase. Given a risk parameter δ , a δ -PAC BAI algorithm consists of a sampling rule (A_t) , a stopping rule τ and a recommendation rule \hat{a} which is \mathcal{F}_τ measurable and satisfies, for all $\mu \in \mathcal{H}$ such that $\mu_1 \neq \mu_2$, $\mathbb{P}_\mu(\hat{a} = a^*) \geq 1 - \delta$. In a bandit model with two Gaussian arms, Kaufmann et al. [2014] propose a δ -PAC algorithm using a uniform sampling rule and a stopping rule τ_δ that asymptotically attains the minimal sample complexity $\mathbb{E}_\mu[\tau_\delta] \sim (8/\Delta^2) \log(1/\delta)$. Using the regret decomposition (2), it is easy to show that the ETC algorithm using the stopping rule τ_δ for $\delta = 1/T$ matches the lower bound of Theorem 4.

Algorithm 3 is a slight variant of this optimal BAI algorithm, based on the stopping time

$$\tau = \inf \left\{ t = 2n : |\hat{\mu}_{1,n} - \hat{\mu}_{2,n}| > \sqrt{\frac{4 \log(T/(2n))}{n}} \right\}.$$

The motivation for the difference (which comes from a more carefully tuned threshold featuring $\log(T/2n)$ in place of $\log(T)$) is that the confidence level should depend on the unknown gap Δ , which determines the regret when a mis-identification occurs. The improvement only appears in the non-asymptotic regime where we are able to prove both asymptotic optimality and order-optimal minimax regret. The latter would not be possible using a fixed-confidence BAI strategy. The proof of this result can be found in Appendix D. The main difficulty is developing a sufficiently strong deviation bound, which we do in Appendix G, and that may be of independent interest. Note that a similar strategy was proposed and analysed by Lai et al. [1983], but in the continuous time framework and with asymptotic analysis only.

Theorem 5. *If $T\Delta^2 > 4e^2$, the regret of the BAI-ETC algorithm is upper bounded as*

$$R_\mu^{\text{BAI-ETC}}(T) \leq \frac{4 \log\left(\frac{T\Delta^2}{4}\right)}{\Delta} + \frac{334\sqrt{\log\left(\frac{T\Delta^2}{4}\right)}}{\Delta} + \frac{178}{\Delta} + 2\Delta.$$

It is upper bounded by $T\Delta$ otherwise, and by $32\sqrt{T} + 2\Delta$ in any case.

```

input:  $T(\geq 3)$ 
 $A_1 = 1, A_2 = 2, t := 2$ 
while  $|\hat{\mu}_1(t) - \hat{\mu}_2(t)| < \sqrt{\frac{8 \log(T/t)}{t}}$  do
    choose  $A_{t+1} = 1$  and  $A_{t+2} = 2$ 
     $t := t + 2$ 
end while
 $\hat{a} := \arg \max_i \hat{\mu}_i(t)$ 
while  $t \leq T$  do
    choose  $A_t = \hat{a}$ 
     $t := t + 1$ 
end while

```

Algorithm 3: BAI-ETC algorithm

4 Fully Sequential Strategies for Known and Unknown Gaps

In the previous section we saw that allowing a random stopping time leads to a factor of 4 improvement in terms of the asymptotic regret relative to the naive fixed-design strategy. We now turn our attention to fully sequential strategies when Δ is known and unknown. The latter case is the classic 2-armed bandit problem and is now quite well understood. Our modest contribution in that case is the first algorithm that is simultaneously asymptotically optimal and order optimal in the minimax sense. For the former case, we are not aware of any previous research where the gap is known except the line of work by [Bubeck et al. \[2013\]](#), [Bubeck and Liu \[2013\]](#), where different questions are treated. In both cases we see that fully sequential strategies improve on the best ETC strategies by a factor of 2.

Known Gaps. We start by stating the lower bound (proved in Section A), which is a straightforward generalisation of Lai and Robbins' lower bound.

Theorem 6. *Let π be a strategy that is uniformly efficient on \mathcal{H}_Δ . Then for all $\mu \in \mathcal{H}_\Delta$,*

$$\liminf_{T \rightarrow \infty} \frac{R_\mu^\pi(T)}{\log T} \geq \frac{1}{2\Delta}$$

We are not aware of any existing algorithm matching this lower bound, which motivates us to introduce a new strategy called Δ -UCB that exploits the knowledge of Δ to improve the performance of UCB. In each round the algorithm chooses the arm that has been played most often so far unless the other arm has an upper confidence bound that is close to Δ larger than the empirical estimate of the most played arm. Like ETC strategies, Δ -UCB is not anytime in the sense that it requires the knowledge of both the horizon T and the gap Δ .

```

1: input:  $T$  and  $\Delta$ 
2:  $\varepsilon_T = \Delta \log^{-\frac{1}{8}}(e + T\Delta^2)/4$ 
3: for  $t \in \{1, \dots, T\}$  do
4:   let  $A_{t,\min} := \arg \min_{i \in \{1,2\}} N_i(t-1)$  and  $A_{t,\max} = 3 - A_{t,\min}$ 
5:   if  $\hat{\mu}_{A_{t,\min}}(t-1) + \sqrt{\frac{2 \log \left( \frac{T}{N_{A_{t,\min}}(t-1)} \right)}{N_{A_{t,\min}}(t-1)}} \geq \hat{\mu}_{A_{t,\max}}(t-1) + \Delta - 2\varepsilon_T$  then
6:     choose  $A_t = A_{t,\min}$ 
7:   else
8:     choose  $A_t = A_{t,\max}$ 
9:   end if
10: end for

```

Algorithm 4: Δ -UCB

Theorem 7. *If $T(2\Delta - 3\varepsilon_T)^2 \geq 2$ and $T\varepsilon_T^2 \geq e^2$, the regret of the Δ -UCB algorithm is upper bounded as*

$$\begin{aligned}
R_\mu^{\Delta\text{-UCB}}(T) &\leq \frac{\log(2T\Delta^2)}{2\Delta(1 - 3\varepsilon_T/(2\Delta))^2} + \frac{\sqrt{\pi \log(2T\Delta^2)}}{2\Delta(1 - 3\varepsilon_T/\Delta)^2} \\
&\quad + \Delta \left[\frac{30e\sqrt{\log(\varepsilon_T^2 T)}}{\varepsilon_T^2} + \frac{80}{\varepsilon_T^2} + \frac{2}{(2\Delta - 3\varepsilon_T)^2} \right] + 5\Delta.
\end{aligned}$$

Moreover $\limsup_{T \rightarrow \infty} R_\mu^{\Delta\text{-UCB}}(T)/\log(T) \leq (2\Delta)^{-1}$ and $\forall \mu \in \mathcal{H}_\Delta$, $R_\mu^{\Delta\text{-UCB}}(T) \leq 328\sqrt{T} + 5\Delta$.

The proof may be found in Appendix E.

Unknown Gaps. In the classical bandit setting where Δ is unknown, UCB by [Kathakis and Robbins \[1995\]](#) is known to be asymptotically optimal: $R_\mu^{\text{UCB}}(T) \sim 2\log(T)/\Delta$, which matches the lower bound of [Lai and Robbins \[1985\]](#). Non-asymptotic regret bounds are given for example by [Auer et al. \[2002\]](#), [Cappé et al. \[2013\]](#). Unfortunately, UCB is not optimal in the minimax sense, which is so far only achieved by algorithms that are not asymptotically optimal [[Audibert and Bubeck, 2009](#), [Lattimore, 2015](#)]. Here, with only two arms, we are able to show that Algorithm 5 below is

simultaneously minimax order-optimal and asymptotically optimal. The strategy is essentially the same as suggested by Lai [1987], but with a fractionally smaller confidence bound. The proof of Theorem 8 is given in Appendix F. Empirically the smaller confidence bonus used by UCB* leads to a significant improvement relative to UCB.

```

1: input:  $T$ 
2: for  $t \in \{1, \dots, T\}$  do
3:    $A_t = \arg \max_{i \in \{1, 2\}} \hat{\mu}_i(t-1) + \sqrt{\frac{2}{N_i(t-1)} \log \left( \frac{T}{N_i(t-1)} \right)}$ 
4: end for

```

Algorithm 5: UCB*

Theorem 8. For all $\varepsilon \in (0, \Delta)$, if $T(\Delta - \varepsilon)^2 \geq 2$ and $T\varepsilon^2 \geq e^2$, the regret of the UCB* strategy is upper bounded as


$$R_\mu^{UCB^*}(T) \leq \frac{2 \log \left(\frac{T\Delta^2}{2} \right)}{\Delta \left(1 - \frac{\varepsilon}{\Delta} \right)^2} + \frac{2\sqrt{\pi \log \left(\frac{T\Delta^2}{2} \right)}}{\Delta \left(1 - \frac{\varepsilon}{\Delta} \right)^2} + \Delta \left(\frac{30e\sqrt{\log(\varepsilon^2 T)} + 16e}{\varepsilon^2} \right) + \frac{2}{\Delta \left(1 - \frac{\varepsilon}{\Delta} \right)^2} + \Delta.$$

Moreover, $\limsup_{T \rightarrow \infty} R_\mu^\pi(T)/\log(T) = 2/\Delta$ and for all $\mu \in \mathcal{H}$, $R_\mu^\pi(T) \leq 33\sqrt{T} + \Delta$.

Note that if there are $K > 2$ arms, then the strategy above is still asymptotically optimal, but suffers a minimax regret of $\Omega(\sqrt{TK \log(K)})$, which is a factor of $\sqrt{\log(K)}$ suboptimal.

5 Numerical Experiments

We represent here the regret of the five strategies presented in this article on a bandit problem with $\Delta = 1/5$, for different values of the horizon. The regret is estimated by $4 \cdot 10^5$ Monte-Carlo replications. In the legend, the estimated slopes of $\Delta R^\pi(T)$ (in logarithmic scale) are indicated after the policy names.



The experimental behavior of the algorithms reflects the theoretical results presented above: the regret asymptotically grows as the logarithm of the horizon, the experimental coefficients correspond approximately to theory, and the relative ordering of the policies is respected. However, it should be noted that for short horizons the hierarchy is not quite the same, and the growth rate is not logarithmic; this question is raised in Garivier et al. [2016]. In particular, on short horizons the Best-Arm Identification procedure performs very well with respect to the others, and starts to be beaten (even by the gap-aware strategies) only when $T\Delta^2$ is much larger than 10.

6 Conclusion: Beyond Uniform Exploration, Two Arms and Gaussian distributions

It is worth emphasising the impossibility of non-trivial lower bounds on the regret of ETC strategies using any possible (non-uniform) sampling rule. Indeed, using UCB as a sampling rule together with an a.s. infinite stopping rule defines an artificial but formally valid ETC strategy that achieves the best possible rate for general strategies. This strategy is not a faithful counter-example to our claim that ETC strategies are sub-optimal, because UCB is not a satisfying exploration rule. If exploration is the objective, then uniform sampling is known to be optimal in the two-armed Gaussian case [Kaufmann et al., 2014], which justifies the uniform sampling assumption.

The use of ETC strategies for regret minimisation (e.g., as presented by Perchet and Rigollet [2013]) is certainly not limited to bandit models with 2 arms. The extension to multiple arms is based on the successive elimination idea in which a set of active arms is maintained with arms chosen according to a round robin within the active set. Arms are eliminated from the active set once their optimality becomes implausible and the exploration phase terminates when the active set contains only a single arm (an example is by Auer and Ortner [2010]). The Successive Elimination algorithm has been introduced by Even-Dar et al. [2006] for best-arm identification in the fixed-confidence setting. It was shown to be rate-optimal, and thus a good compromise for both minimizing regret and finding the best arm. If one looks more precisely at multiplicative constants, however, Garivier and Kaufmann [2016] showed that it is suboptimal for the best arm identification task in almost all settings except two-armed Gaussian bandits. Regarding regret minimization, the present paper shows that it is sub-optimal by a factor 2 on every two-armed Gaussian problem.

It is therefore interesting to investigate the performance in terms of regret of an ETC algorithm using an optimal BAI algorithm. This is actually possible not only for Gaussian distributions, but more generally for one-parameter exponential families, for which Garivier and Kaufmann [2016] propose the asymptotically optimal Track-and-Stop strategy. Denoting $d(\mu, \mu') = \text{KL}(\nu_\mu, \nu_{\mu'})$ the Kullback-Leibler divergence between two distributions parameterised by μ and μ' , they provide results which can be adapted to obtain the following bound.

Proposition 1. *For μ such that $\mu_1 > \max_{a \neq 1} \mu_a$, the regret of the ETC strategy using Track-and-Stop exploration with risk $1/T$ satisfies*

$$\limsup_{T \rightarrow \infty} \frac{R_\mu^{\text{TaS}}(T)}{\log T} \leq T^*(\mu) \left(\sum_{a=2}^K w_a^*(\mu) (\mu_1 - \mu_a) \right),$$

where $T^*(\mu)$ (resp. $w^*(\mu)$) is the the maximum (resp. maximiser) of the optimisation problem

$$\max_{w \in \Sigma_K} \inf_{a \neq 1} \left[w_1 d \left(\mu_1, \frac{w_1 \mu_1 + w_a \mu_a}{w_1 + w_a} \right) + w_a d \left(\mu_a, \frac{w_a \mu_1 + w_a \mu_a}{w_1 + w_a} \right) \right],$$

where Σ_K is the set of probability distributions on $\{1, \dots, K\}$.

In general, it is not easy to quantify the difference to the lower bound of Lai and Robbins

$$\liminf_{T \rightarrow \infty} \frac{R_\mu^\pi(T)}{\log T} \geq \sum_{a=2}^K \frac{\mu_1 - \mu_a}{d(\mu_a, \mu_1)}.$$

Even for Gaussian distributions, there is no general closed-form formula for $T^*(\mu)$ and $w^*(\mu)$ except when $K = 2$. However, we conjecture that the worst case is when μ_1 and μ_2 are much larger than the other means: then, the regret is almost the same as in the 2-arm case, and ETC strategies are suboptimal by a factor 2. On the other hand, the most favourable case (in terms of relative efficiency) seems to be when $\mu_2 = \dots = \mu_K$: then

$$w_1^*(\mu) = \frac{\sqrt{K-1}}{K-1 + \sqrt{K-1}}, \quad w_2^*(\mu) = \dots = w_K^*(\mu) = \frac{1}{K-1 + \sqrt{K-1}}$$

and $T^* = 2(\sqrt{K-1} + 1)^2 / \Delta^2$, leading to

$$\limsup_{T \rightarrow \infty} \frac{R_\mu^{\text{TaS}}(T)}{\log(T)} \leq \left(1 + \frac{1}{\sqrt{K-1}} \right) \frac{2(K-1)}{\Delta},$$

while Lai and Robbins' lower bound yields $2(K-1)/\Delta$. Thus, the difference grows with K as $2\sqrt{K-1} \log(T)/\Delta$, but the relative difference decreases.

References

- Jean-Yves Audibert and Sébastien Bubeck. Minimax policies for adversarial and stochastic bandits. In *Proceedings of Conference on Learning Theory (COLT)*, pages 217–226, 2009.
- Peter Auer and Ronald Ortner. UCB revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61(1-2):55–65, 2010.
- Peter Auer, Nicolás Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002.
- Sébastien Bubeck and Che-Yu Liu. Prior-free and prior-dependent regret bounds for thompson sampling. In *Advances in Neural Information Processing Systems*, pages 638–646, 2013.
- Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet. Bounded regret in stochastic multi-armed bandits. In *Proceedings of the 26th Conference On Learning Theory*, pages 122–134, 2013.
- Olivier Cappé, Aurélien Garivier, Odalric-Ambrym Maillard, Rémi Munos, and Gilles Stoltz. Kullback–Leibler upper confidence bounds for optimal sequential allocation. *The Annals of Statistics*, 41(3):1516–1541, 2013.
- 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.
- Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *Proceedings of the 29th Conference On Learning Theory (to appear)*, 2016.
- Aurélien Garivier, Pierre Ménard, and Gilles Stoltz. Explore first, exploit next: The true shape of regret in bandit problems. *arXiv preprint arXiv:1602.07182*, 2016.
- Abdolhossein Hoorfar and Mehdi Hassani. Inequalities on the lambert w function and hyperpower function. *J. Inequal. Pure and Appl. Math*, 9(2):5–9, 2008.
- Michael N Katehakis and Herbert Robbins. Sequential choice from several populations. *Proceedings of the National Academy of Sciences of the United States of America*, 92(19):8584, 1995.
- Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the Complexity of A/B Testing. In *Proceedings of the 27th Conference On Learning Theory*, 2014.
- Tze Leung Lai. Adaptive treatment allocation and the multi-armed bandit problem. *The Annals of Statistics*, pages 1091–1114, 1987.
- Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- Tze Leung Lai, Herbert Robbins, and David Siegmund. Sequential design of comparative clinical trials. In M. Haseeb Rizvi, Jagdish Rustagi, and David Siegmund, editors, *Recent advances in statistics: papers in honor of Herman Chernoff on his sixtieth birthday*, pages 51–68. Academic Press, 1983.
- Tor Lattimore. Optimally confident UCB: Improved regret for finite-armed bandits. Technical report, 2015. URL <http://arxiv.org/abs/1507.07880>.
- Peter Mörters and Yuval Peres. *Brownian motion*, volume 30. Cambridge University Press, 2010.
- Vianney Perchet and Philippe Rigollet. The multi-armed bandit with covariates. *The Annals of Statistics*, 2013.
- Vianney Perchet, Philippe Rigollet, Sylvain Chassang, and Eric Snowberg. Batched bandit problems. In *Proceedings of the 28th Conference On Learning Theory*, 2015.
- David Siegmund. *Sequential Analysis*. Springer-Verlag, 1985.
- William Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- Abraham Wald. Sequential Tests of Statistical Hypotheses. *Annals of Mathematical Statistics*, 16(2): 117–186, 1945.

Notation for the Proofs

We denote by (X_s) and (Y_s) the sequence of successive observations from arm 1 and arm 2, so that

$$\hat{\mu}_{1,s} = \frac{1}{s} \sum_{i=1}^s X_i \quad \text{and} \quad \hat{\mu}_{2,s} = \frac{1}{s} \sum_{i=1}^s Y_i$$

In the proofs of all regret upper bounds given in Appendix B to F we assume without loss of generality that μ is such that $\mu_1 > \mu_2$ and let $\Delta = \mu_1 - \mu_2$.

A Proof of the Lower Bounds (Theorems 3, 4, 6 and Lai&Robbins)

Let π be a uniformly efficient strategy on some class \mathcal{S} , as defined in (3), and let $\lambda \in \mathcal{H}$. If $m(\lambda) = \arg \min\{\lambda_1, \lambda_2\}$, as $\mathbb{E}_\lambda[R_\mu^\pi(T)] = |\lambda_1 - \lambda_2| \mathbb{E}_\lambda[N_{m(\lambda)}(T)]$ this implies in particular that

$$\forall \alpha \in]0, 1], \quad \mathbb{E}_\lambda[N_{m(\lambda)}(T)] = o(T^\alpha).$$

Without loss of generality we assume that $\mu_1 = \mu_2 + \Delta$ with $\Delta > 0$. All the lower bounds are based on a change of measure argument, which involves considering an alternative reward vector (μ'_1, μ'_2) that is “not too far” from (μ_1, μ_2) , but for which the expected behaviour of the algorithm is very different. This is the same approach used by Lai and Robbins [1985], but rewritten and generalised in a more powerful way (in particular regarding the ETC strategies). The improvements come thanks to Inequality 4 in [Garivier et al., 2016], which states that for every $(\mu'_1, \mu'_2) \in \mathcal{H}$ and for every stopping time σ such that $N_2(T)$ is \mathcal{F}_σ -measurable,

$$\mathbb{E}_\mu[N_1(\sigma)] \frac{(\mu'_1 - \mu_1)^2}{2} + \mathbb{E}_\mu[N_2(\sigma)] \frac{(\mu'_2 - \mu_2)^2}{2} \geq \text{kl} \left(\mathbb{E}_\mu \left[\frac{N_2(T)}{T} \right], \mathbb{E}_{\mu'} \left[\frac{N_2(T)}{T} \right] \right),$$

where $\text{kl}(p, q)$ is the relative entropy between Bernoulli distributions with parameters $p, q \in [0, 1]$ respectively. Since $\text{kl}(p, q) \geq (1 - p) \log(1/(1 - q)) - \log(2)$ for all $p, q \in (0, 1)$, one obtains

$$\begin{aligned} \mathbb{E}_\mu[N_1(\sigma)] \frac{(\mu'_1 - \mu_1)^2}{2} + \mathbb{E}_\mu[N_2(\sigma)] \frac{(\mu'_2 - \mu_2)^2}{2} \\ \geq \left(1 - \frac{\mathbb{E}_\mu[N_2(T)]}{T} \right) \log \left(\frac{T}{\mathbb{E}_{\mu'}[N_1(T)]} \right) - \log(2). \end{aligned} \quad (4)$$

For $\mu' \in \mathcal{S}$ such that $\mu'_1 < \mu'_2$, $\mathbb{E}_\mu[N_2(T)] = o(T^\alpha)$ and $\mathbb{E}_{\mu'}[N_1(T)] = o(T^\alpha)$ for all $\alpha \in]0, 1]$, thus

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\mu[N_1(\sigma)] (\mu'_1 - \mu_1)^2 / 2 + \mathbb{E}_\mu[N_2(\sigma)] (\mu'_2 - \mu_2)^2 / 2}{\log T} \geq 1.$$

Let us now draw the conclusions in each setting. Observe that while this argument is now routine for general policies, we show here how to apply it very nicely to ETC strategies as well.

Known gap, General strategy: $\mathcal{S} = \mathcal{H}_\Delta$. By choosing $\sigma = T$, $\mu'_1 = \mu_1$ and $\mu'_2 = \mu_1 + \Delta = \mu_2 + 2\Delta$, we obtain

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\mu[N_2(T)]}{\log T} \geq \frac{1}{(2\Delta)^2/2}.$$

Known gap, ETC strategy: $\mathcal{S} = \mathcal{H}_\Delta$. For an ETC strategy π with a stopping rule τ , one has $N_1(\tau \wedge T) = N_2(\tau \wedge T) = (\tau \wedge T)/2$. Besides, $N_2(T)$ is indeed $\mathcal{F}_{\tau \wedge T}$ -measurable: after $\tau \wedge T$ draws, the agent knows whether she will draw arm 2 for the last $T - \tau \wedge T$ steps or not. With $\mu'_1 = \mu_2$, $\mu'_2 = \mu_1$, Inequality (4) thus yields:

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\mu \left[\frac{\tau \wedge T}{2} \right]}{\log T} \geq \frac{1}{2\Delta^2/2}.$$

Unknown gap, General strategy: $\mathcal{S} = \mathcal{H}$. We use the same choices, except $\mu'_2 = \mu_1 + \varepsilon$ for some $\varepsilon > 0$:

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\mu[N_2(T)]}{\log T} \geq \frac{1}{(\Delta + \varepsilon)^2/2}.$$

Unknown gap, ETC strategy: $\mathcal{S} = \mathcal{H}$. Choosing this time $\mu'_1 = (\mu_1 + \mu_2 - \varepsilon)/2$ and $\mu'_2 = (\mu_1 + \mu_2 + \varepsilon)/2$, for some $\varepsilon > 0$ yields

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\mu \left[\frac{\tau \wedge T}{2} \right]}{\log T} \geq \frac{1}{\left(\frac{\Delta + \varepsilon}{2} \right)^2}.$$

$R_\mu^\pi(T) = \Delta \mathbb{E}_\mu[N_2(T)]$ for general strategies, while Equation (2) shows that $R_\mu^\pi(T) \geq \Delta \mathbb{E}_\mu[(\tau \wedge T)/2]$ for ETC strategies. Therefore letting ε go to zero when needed shows that

$$\liminf_{T \rightarrow \infty} \frac{R_\mu^\pi(T)}{\log(T)} \geq \frac{C_S^\Pi}{\Delta}, \quad \text{for the value } C_S^\Pi \text{ given at the end of Section 2.}$$

Note that the asymptotic lower bound of Theorem 1 can also be proved by similar arguments, if one really wants to bring in an elephant to kill a mouse. Moreover, note that this proof may also lead to (not-so-simple) non-asymptotic lower-bounds, as shown in Garivier et al. [2016] for example.

B Proof of Theorem 1

Let $n \leq T/2$. The number of draws of the suboptimal arm 2 is $N_2 = n + (T - 2n)\mathbb{1}\{S_n \leq 0\}$, where $S_n = (X_1 - Y_1) + \dots + (X_n - Y_n) \sim \mathcal{N}(n\Delta, 2n)$. The expected regret of the strategy using $2n$ exploration steps is

$$R_\mu^n(T) = \Delta \mathbb{E}_\mu[N_2] = \Delta(n + (T - 2n)\mathbb{P}_\mu(S_n \leq 0)). \quad (5)$$

But

$$\mathbb{P}_\mu(S_n \leq 0) = \mathbb{P}_\mu\left(\frac{S_n - n\Delta}{\sqrt{2n}} \leq \frac{-n\Delta}{\sqrt{2n}} = -\Delta\sqrt{\frac{n}{2}}\right).$$

Denote by Φ (resp. ϕ) the pdf (resp. cdf) of the standard Gaussian distribution, and recall that W is the Lambert function defined for all $y > 0$ by $W(y) \exp(W(y)) = y$. The regret is thus upper-bounded as $R_\mu^n(T) \leq \Delta g(n)$ where, for all $x > 0$, $g(x) = x + T\Phi(-\Delta\sqrt{x/2})$. By differentiating g , one can see that its maximum is reached at x^* such that

$$\phi\left(\Delta\sqrt{\frac{x^*}{2}}\right) = \frac{\Delta\sqrt{x^*/2}}{T\Delta^2/4}, \quad \text{and thus} \quad x^* = \frac{2}{\Delta^2}W\left(\frac{T^2\Delta^4}{32\pi}\right).$$

By choosing $\bar{n} = \lceil x^* \rceil$, we obtain

$$R_\mu^{\bar{n}}(T) \leq \frac{2}{\Delta}W\left(\frac{T^2\Delta^4}{32\pi}\right) + \Delta + \Delta T\Phi\left(-\Delta\sqrt{\frac{x^*}{2}}\right).$$

As $g(\bar{n}) \leq g(x^*) + 1 \leq g(0) + 1 = T/2 + 1$, $R_\mu^{\bar{n}}(T) \leq (T/2 + 1)\Delta \leq 2.04\sqrt{T} + \Delta$ for $T\Delta^2 \leq 4\sqrt{2\pi}e$. If $T\Delta^2 \geq 4\sqrt{2\pi}e$, the inequality $W(y) \leq \log((1 + e^{-1})y/\log(y))$ valid for all $y \geq e$ (see Hoofar and Hassani [2008]) entails

$$W\left(\frac{T^2\Delta^4}{32\pi}\right) \leq \log\left(\frac{(1 + e^{-1})\frac{T^2\Delta^4}{32\pi}}{\log\left(\frac{T^2\Delta^4}{32\pi}\right)}\right) = 2\log\left(\frac{1}{8}\sqrt{\frac{1 + e^{-1}}{\pi}}\frac{T\Delta^2}{\sqrt{\log\left(\frac{T\Delta^2}{4\sqrt{2\pi}}\right)}}\right)$$

In addition, the classic bound on the Gaussian tail $\Phi(-y) \leq \phi(y)/y$ yields by definition of x^* :

$$\Phi\left(-\Delta\sqrt{\frac{x^*}{2}}\right) \leq \frac{\phi\left(-\Delta\sqrt{\frac{x^*}{2}}\right)}{\Delta\sqrt{\frac{x^*}{2}}} = \frac{4}{T\Delta^2}.$$

Hence, for $T\Delta^2 \geq 4\sqrt{2\pi}e$,

$$R_\mu^{\bar{n}}(T) \leq \frac{4}{\Delta}\log\left(\frac{e}{8}\sqrt{\frac{1 + e^{-1}}{\pi}}\frac{T\Delta^2}{\sqrt{\log\left(\frac{T\Delta^2}{4\sqrt{2\pi}}\right)}}\right) + \Delta < \frac{4}{\Delta}\log\left(\frac{T\Delta^2}{4.46\sqrt{\log\left(\frac{T\Delta^2}{4\sqrt{2\pi}}\right)}}\right) + \Delta.$$

To complete the proof of the uniform upper-bound, we start from

$$R_\mu^{\bar{n}}(T) \leq \frac{2}{\Delta}W\left(\frac{T^2\Delta^4}{32\pi}\right) + \frac{4}{\Delta} + \Delta.$$




Figure 1: Regret of the Fixed-Budget ETC algorithm with optimal n , for $\Delta = 1/5$ (solid line). The linear upper bound $T\Delta/2$ and the logarithmic bound of Theorem 1 are dashed, bold lines. The thin, dotted line is $2.04\sqrt{T} + \Delta$. The bold, dotted line is $\Delta g(\bar{n})$, with g and \bar{n} as in the proof.

Denoting by $r \approx 1.09$ the root of $(4r - 1)W(r) = 2$, the maximum of $2W\left(\frac{T^2\Delta^4}{32\pi}\right)/\Delta + 4/\Delta$ is reached at $\underline{\Delta} = (32\pi r/T^2)^{1/4}$, and is equal to

$$\frac{2}{\underline{\Delta}} W\left(\frac{T^2\Delta^4}{32\pi}\right) + \frac{4}{\underline{\Delta}} = \frac{8r^{3/4}}{(4r - 1)(2\pi)^{1/4}} \sqrt{T} < 2\sqrt{T}.$$

Let us now prove that choosing n too small leads to catastrophic regret. If $n \leq 4(1 - \epsilon)\log(T)/\Delta^2$, then Equation (5) yields

$$\begin{aligned} \mathbb{P}_\mu(S_n \leq 0) &\geq \frac{1}{\Delta\sqrt{\pi n}} \left(1 - \frac{2}{n\Delta^2}\right) \exp\left(-\frac{\Delta^2 n}{4}\right) \\ &\geq \left(1 - \frac{2}{n\Delta^2}\right) \frac{1}{2\sqrt{\pi(1 - \epsilon)\log(T)}} \exp(-(1 - \epsilon)\log(T)) \\ &\geq \left(1 - \frac{2}{n\Delta^2}\right) \frac{T^{\epsilon-1}}{2\sqrt{\pi\log(T)}}. \end{aligned}$$

Thus, the expected regret is lower-bounded as

$$R_\mu^n(T) \geq \Delta(T - 2n)\mathbb{P}_\mu(S_n \leq 0) \geq \Delta \left(1 - \frac{2}{n\Delta^2}\right) \left(1 - \frac{8\log(T)}{\Delta^2 T}\right) \frac{T^\epsilon}{2\sqrt{\pi\log(T)}}.$$

Let us now turn to the last statement of the theorem. The previous inequality shows that for all $\epsilon > 0$,

$$\liminf_T \inf_{\frac{3}{\Delta^2} < n \leq \frac{4(1-\epsilon)\log(T)}{\Delta^2}} \frac{R_\mu^n(T)}{\log(T)} = +\infty.$$

For $n \leq 3/\Delta^2$, we have that

$$\mathbb{P}_\mu(S_n \leq 0) \geq \mathbb{P}_\mu\left(\frac{S_n - n\Delta}{\sqrt{2n}} \leq -\sqrt{\frac{3}{2}}\right) > 0,$$

and hence that $\liminf_T \inf_{n \leq 3/\Delta^2} R_\mu^n(T)/T > 0$. As $R_\mu^n(T) \geq n\Delta$, the result follows.

C Proof of Theorem 2

Recall that $\mu_1 > \mu_2$. Using (2), one has

$$R_\mu^\pi(T) \leq \Delta \mathbb{E}_\mu\left[\frac{\tau}{2}\right] + T\Delta \mathbb{P}_\mu(\tau < T, \hat{a} = 2).$$

If $T\Delta^2 \leq 1$, then $\tau = 2$ and \hat{a} is based on a single sample from each action. Therefore $\hat{a} = 2$ with probability less than $1/2$ and the regret is upper-bounded by $\Delta + T\Delta/2$. Otherwise, let $S_0 = 0$, $S_n = (X_1 - Y_1) + \dots + (X_n - Y_n)$ for every $n \geq 1$. For every $u > 0$, let $n_u = (\log(T\Delta^2) + u)/\Delta^2$. Observe that

$$\left\{\frac{\tau}{2} \geq \left\lceil \frac{\log(T\Delta^2) + u}{\Delta^2} \right\rceil\right\} \subset \left\{S_{\lceil n_u \rceil} \leq \frac{\log(T\Delta^2)}{\Delta}\right\}.$$

Moreover, if $n_u \in \mathbb{N}$ then $S_{n_u} \sim \mathcal{N}(n_u\Delta, 2n_u)$ and

$$\begin{aligned} \mathbb{P}_\mu\left(S_{n_u} \leq \frac{\log(T\Delta^2)}{\Delta}\right) &= \mathbb{P}_\mu\left(\frac{S_{n_u} - n_u\Delta}{\sqrt{2n_u}} \leq \frac{\log(T\Delta^2)/\Delta - \Delta(\log(T\Delta^2) + u)/\Delta^2}{\sqrt{2(\log(T\Delta^2) + u)/\Delta^2}}\right) \\ &= \mathbb{P}_\mu\left(\frac{S_{n_u} - n_u\Delta}{\sqrt{2n_u}} \leq \frac{-u}{\sqrt{2(\log(T\Delta^2) + u)}}\right) \leq \exp\left(-\frac{u^2}{4(\log(T\Delta^2) + u)}\right). \end{aligned}$$

Hence, for $a = 2\sqrt{\log(T\Delta^2)}$,

$$\begin{aligned} \int_{n_a}^\infty \mathbb{P}_\mu\left(\frac{\tau}{2} - 1 \geq v\right) dv &= \int_a^\infty \mathbb{P}_\mu\left(\frac{\tau}{2} - 1 \geq \frac{\log(T\Delta^2) + u}{\Delta^2}\right) \frac{du}{\Delta^2} \\ &\leq \int_a^\infty \mathbb{P}_\mu\left(\frac{\tau}{2} \geq \left\lceil \frac{\log(T\Delta^2) + u}{\Delta^2} \right\rceil\right) \frac{du}{\Delta^2} \\ &\leq \frac{1}{\Delta^2} \int_a^\infty \exp\left(-\frac{u^2}{4(\log(T\Delta^2) + u)}\right) du \\ &\leq \frac{1}{\Delta^2} \int_a^\infty \exp\left(-\frac{u}{2\sqrt{\log(T\Delta^2)} + 4}\right) du \quad \text{as } \log(T\Delta^2) \leq u\sqrt{\log(\Delta^2 T)}/2 \\ &\leq \frac{1}{\Delta^2} \int_0^\infty \exp\left(-\frac{u}{2\sqrt{\log(T\Delta^2)} + 4}\right) du \\ &= \frac{2\sqrt{\log(T\Delta^2)} + 4}{\Delta^2}, \end{aligned}$$

and

$$\mathbb{E}_\mu\left[\frac{\tau}{2}\right] \leq 1 + n_a + \int_{n_a}^\infty \mathbb{P}_\mu\left(\frac{\tau}{2} - 1 \geq v\right) dv \leq 1 + \frac{\log(T\Delta^2) + 4\sqrt{\log(T\Delta^2)} + 4}{\Delta^2}.$$

To conclude the proof of the first statement, it remains to show that $\mathbb{P}(\tau < T, \hat{a} = 2) \leq 1/(T\Delta^2)$. Since $X_1 - Y_1 \sim \mathcal{N}(\Delta, 2)$, $\mathbb{E}_\mu[\exp(-\Delta(X_1 - Y_1))] = \exp(-\Delta^2 + 2\Delta^2/2) = 1$ and $M_n = \exp(-\Delta S_n)$ is a martingale. Let $\tau_2 = T \wedge \inf\{n \geq 1 : S_n \leq -\log(T\Delta^2)/\Delta\}$, and observe that

$$\{\tau < T, \hat{a} = 2\} \subset \left\{\exists n < T : S_n \leq -\frac{\log(T\Delta^2)}{\Delta}\right\} = \{\tau_2 < T\}.$$

Doob's optional stopping theorem yields $\mathbb{E}_\mu[M_{\tau_2}] = \mathbb{E}_\mu[M_0] = 1$. But as $M_{\tau_2} = \exp(-\Delta S_{\tau_2}) \geq \exp(\Delta \log(T\Delta^2)/\Delta) = T\Delta^2$ on the event $\{\tau_2 < T\}$,

$$\mathbb{P}_\mu(\tau < T, \hat{a} = 2) \leq \mathbb{P}_\mu(\tau_2 < T) \leq \mathbb{E}_\mu \left[\mathbb{1}_{\{\tau_2 < T\}} \frac{M_{\tau_2}}{T\Delta^2} \right] \leq \frac{\mathbb{E}_\mu[M_{\tau_2}]}{T\Delta^2} = \frac{1}{T\Delta^2}.$$

The last statement Theorem 2 is obtained by maximising the bound (except for the Δ summand). Let $u = 1/(\Delta\sqrt{T})$ and $f(u) = -u \log(u/\sqrt{e}) + 2u\sqrt{-\log(u)}$. Denoting $\ell = -\log(u)$, $f'(u) = \ell - 1/2 + 2\sqrt{\ell} - 1/\sqrt{\ell} = (\sqrt{\ell} + 2)(\ell - 1/2)/\sqrt{\ell}$, thus the maximum is reached at $\ell = \log(\Delta\sqrt{T}) = 1/2$ and $\Delta = \sqrt{e/T}$. Re-injecting this value into the bound, we obtain that for every $\Delta > 0$,

$$R_\mu^{\text{SPRT-ETC}}(T) \leq \frac{2 + 4\sqrt{1} + 4}{\sqrt{e/T}} = 10\sqrt{\frac{T}{e}}.$$

D Proof of Theorem 5

Recall that $\mu_1 = \mu_2 + \Delta$ with $\Delta > 0$. Let $W_s = (X_s - Y_s - \Delta)/\sqrt{2}$, which means that W_1, W_2, \dots are i.i.d. standard Gaussian random variables. Introducing $\mathcal{F} = (\hat{a} \neq 1, \tau < T)$, one has by (2) that

$$R_\mu^{\text{BAI-ETC}}(T) \leq T\Delta\mathbb{P}_\mu(\mathcal{F}) + (\Delta/2)\mathbb{E}_\mu[\tau \wedge T].$$

From the definition of τ and Lemma 1.(c) in Appendix G, assuming that $T\Delta^2 \geq 4e^2$, one obtains

$$\begin{aligned} \mathbb{P}_\mu(\mathcal{F}) &\leq \mathbb{P}_\mu \left(\exists s : 2s \leq T, \hat{\mu}_{1,s} - \hat{\mu}_{2,s} \leq -\sqrt{\frac{4}{s} \log\left(\frac{T}{2s}\right)} \right) \\ &= \mathbb{P}_\mu \left(\exists s \leq T/2 : \frac{\sum_{i=1}^s W_i}{s} \leq -\sqrt{\frac{2}{s} \log\left(\frac{T/2}{s}\right)} - \frac{\Delta}{\sqrt{2}} \right) \\ &\leq \frac{120e\sqrt{\log(\Delta^2 T/4)}}{\Delta^2 T} + \frac{64e}{\Delta^2 T}. \end{aligned}$$

The last step is bounding $\mathbb{E}_\mu[\tau \wedge T]$ for which as $T\Delta^2 \geq 4$, Lemma 1.(b) yields

$$\begin{aligned} \mathbb{E}_\mu[\tau \wedge T] &= \sum_{t=1}^T \mathbb{P}_\mu(\tau \geq t) \leq 2 + 2 \sum_{s=1}^{T/2} \mathbb{P}_\mu(\tau \geq 2s + 1) \\ &\leq 2 + 2 \sum_{s=1}^{T/2} \mathbb{P}_\mu \left(\frac{\sum_{i=1}^s W_i}{s} \leq \sqrt{\frac{2}{s} \log\left(\frac{T/2}{s}\right)} - \frac{\Delta}{\sqrt{2}} \right) \\ &\leq \frac{8 \log\left(\frac{T\Delta^2}{4}\right)}{\Delta^2} + \frac{8\sqrt{\pi \log\left(\frac{T\Delta^2}{4}\right)}}{\Delta^2} + \frac{8}{\Delta^2} + 4, \end{aligned}$$

Therefore, if $T\Delta^2 \geq 4e^2$,

$$\begin{aligned} R_\mu^{\text{BAI-ETC}}(T) &\leq \frac{4 \log\left(\frac{T\Delta^2}{4}\right)}{\Delta} + \frac{4\sqrt{\pi \log\left(\frac{T\Delta^2}{4}\right)}}{\Delta} + \frac{4}{\Delta} + 2\Delta + \frac{120e\sqrt{\log\left(\frac{\Delta^2 T}{4}\right)}}{\Delta} + \frac{64e}{\Delta} \\ &\leq \frac{4 \log\left(\frac{T\Delta^2}{4}\right)}{\Delta} + \frac{334\sqrt{\log\left(\frac{T\Delta^2}{4}\right)}}{\Delta} + \frac{178}{\Delta} + 2\Delta. \end{aligned}$$

Taking the limit as $T \rightarrow \infty$ shows that $\limsup_{T \rightarrow \infty} R_\mu^{\text{BAI-ETC}}(T)/\log(T) \leq 4$. Noting that $R_\mu^{\text{BAI-ETC}}(T) \leq \Delta T$ and taking the minimum of this bound and the finite-time bound given above leads arduously to $R_\mu^{\text{BAI-ETC}}(T) \leq 2\Delta + 32\sqrt{T}$ for all μ .

E Proof of Theorem 7

Define random time $\tau = \max\{\tau_1, \tau_2\}$ where τ_i is given by

$$\tau_i = \min \left\{ t \leq T : \sup_{s \geq t} |\hat{\mu}_{i,s} - \mu_i| < \varepsilon_T \right\}.$$

By the concentration Lemma 1.(a) in Appendix G we have $\mathbb{E}_\mu[\tau_i] \leq 1 + 9/\varepsilon_T^2$ and so $\mathbb{E}_\mu[\tau] \leq \mathbb{E}_\mu[\tau_1 + \tau_2] \leq 2 + 18/\varepsilon_T^2$. For $t > 2\tau$ we have $|\hat{\mu}_{A_{t,\max}}(t-1) - \mu_{A_{t,\max}}| < \varepsilon_T$. Therefore the expected number of draws of the suboptimal arm may be bounded by

$$\begin{aligned} \mathbb{E}_\mu[N_2(T)] &= \mathbb{E}_\mu \left[\sum_{t=1}^T \mathbb{1}\{I_t = 2\} \right] \leq \mathbb{E}_\mu[2\tau] + \mathbb{E}_\mu \left[\sum_{t=2\tau+1}^T \mathbb{1}\{I_t = 2\} \right] \\ &\leq \mathbb{E}_\mu[2\tau] + \mathbb{E}_\mu \left[\sum_{t=1}^T \mathbb{1} \left\{ \hat{\mu}_2(t-1) + \sqrt{\frac{2 \log(T/N_2(t-1))}{N_2(t-1)}} \geq \mu_1 + \Delta - 3\varepsilon_T \text{ and } I_t = 2 \right\} \right] \\ &\quad + \mathbb{E}_\mu \left[\sum_{t=1}^T \mathbb{1} \left\{ \hat{\mu}_1(t-1) + \sqrt{\frac{2 \log(T/N_1(t-1))}{N_1(t-1)}} \leq \mu_2 + \Delta - \varepsilon_T = \mu_1 - \varepsilon_T \right\} \right] \end{aligned} \quad (6)$$

By Lemma 1.(b), whenever $T(2\Delta - 3\varepsilon_T)^2 \geq 2$,

$$\begin{aligned} \mathbb{E}_\mu \left[\sum_{t=1}^T \mathbb{1} \left\{ \hat{\mu}_2(t-1) + \sqrt{\frac{2 \log(T/N_2(t-1))}{N_2(t-1)}} \geq \mu_1 + \Delta - 3\varepsilon_T = \mu_2 + 2\Delta - 3\varepsilon_T \text{ and } I_t = 2 \right\} \right] \\ \leq \sum_{s=1}^T \mathbb{P}_\mu \left(\hat{\mu}_{2,s} - \mu_2 + \sqrt{\frac{2 \log(T/s)}{s}} \geq 2\Delta - 3\varepsilon_T \right) \\ \leq \frac{2 \log \left(\frac{T(2\Delta - 3\varepsilon_T)^2}{2} \right)}{(2\Delta - 3\varepsilon_T)^2} + \frac{2 \sqrt{\pi \log \left(\frac{T(2\Delta - 3\varepsilon_T)^2}{2} \right)}}{(2\Delta - 3\varepsilon_T)^2} + \frac{2}{(2\Delta - 3\varepsilon_T)^2} + 1. \end{aligned}$$

For the second term in (6) we apply Lemma 1.(c) to obtain, whenever $T\varepsilon_T^2 \geq e^2$,

$$\begin{aligned} \mathbb{E}_\mu \left[\sum_{t=1}^T \mathbb{1} \left\{ \hat{\mu}_1(t-1) + \sqrt{\frac{2 \log(T/N_1(t-1))}{N_1(t-1)}} \leq \mu_1 - \varepsilon_T \right\} \right] \\ \leq T \mathbb{P}_\mu \left(\exists s \leq T : \hat{\mu}_{1,s} + \sqrt{\frac{2}{s} \log(T/s)} \leq \mu_1 - \varepsilon_T \right) \leq \frac{30e \sqrt{\log(\varepsilon_T^2 T)}}{\varepsilon_T^2} + \frac{16e}{\varepsilon_T^2}. \end{aligned}$$

Therefore, if $T(2\Delta - 3\varepsilon_T)^2 \geq 2$ and $T\varepsilon_T^2 \geq e^2$,

$$\begin{aligned} \mathbb{E}_\mu[N_2(T)] &\leq \frac{2 \log \left(\frac{T(2\Delta - 3\varepsilon_T)^2}{2} \right)}{(2\Delta - 3\varepsilon_T)^2} + \frac{2 \sqrt{\pi \log \left(\frac{T(2\Delta - 3\varepsilon_T)^2}{2} \right)}}{(2\Delta - 3\varepsilon_T)^2} + \frac{30e \sqrt{\log(\varepsilon_T^2 T)}}{\varepsilon_T^2} \\ &\quad + \frac{80}{\varepsilon_T^2} + \frac{2}{(2\Delta - 3\varepsilon_T)^2} + 5. \end{aligned}$$

The first result follows since the regret is $R_\mu^{\Delta\text{-UCB}}(T) = \Delta \mathbb{E}_\mu[N_2(T)]$. For the second it easily noted that for the choice of ε_T given in the definition of the algorithm that $\limsup_{T \rightarrow \infty} \mathbb{E}_\mu[N_2(T)]/\log(T) \leq 1/(2\Delta^2)$. Therefore $\limsup_{T \rightarrow \infty} R_\mu^{\Delta\text{-UCB}}(T)/\log(T) \leq 1/(2\Delta)$. The third result follows from a laborious optimisation step to upper-bound the minimum of $T\Delta$ and the finite-time regret bound above.

F Proof of Theorem 8

For any $\varepsilon \in (0, \Delta)$ we have

$$\begin{aligned} \mathbb{E}_\mu[N_2(T)] &= \mathbb{E}_\mu \left[\sum_{t=1}^T \mathbb{1} \{A_t = 2\} \right] \\ &\leq \mathbb{E}_\mu \left[\sum_{t=1}^T \mathbb{1} \left\{ A_t = 2 \text{ and } \hat{\mu}_2(t-1) + \sqrt{\frac{2}{N_2(t-1)} \log \left(\frac{T}{N_2(t-1)} \right)} \geq \mu_1 - \varepsilon \right\} \right] \\ &\quad + T\mathbb{P}_\mu \left(\exists s \leq T : \hat{\mu}_{1,s} + \sqrt{\frac{2}{s} \log \left(\frac{T}{s} \right)} \leq \mu_1 - \varepsilon \right). \end{aligned}$$

By the concentration Lemma 1.(b) in Appendix G we have, whenever $T(\Delta - \varepsilon)^2 \geq 2$,

$$\begin{aligned} &\mathbb{E}_\mu \left[\sum_{t=1}^T \mathbb{1} \left\{ A_t = 2 \text{ and } \hat{\mu}_2(t-1) + \sqrt{\frac{2}{N_2(t-1)} \log \left(\frac{T}{N_2(t-1)} \right)} \geq \mu_1 - \varepsilon \right\} \right] \\ &\leq \sum_{s=1}^T \mathbb{P}_\mu \left(\hat{\mu}_{2,s} - \mu_2 + \sqrt{\frac{2}{s} \log \left(\frac{T}{s} \right)} \geq \Delta - \varepsilon \right) \\ &\leq \frac{2 \log \left(\frac{T(\Delta - \varepsilon)^2}{2} \right)}{(\Delta - \varepsilon)^2} + \frac{2 \sqrt{\pi \log \left(\frac{T(\Delta - \varepsilon)^2}{2} \right)}}{(\Delta - \varepsilon)^2} + \frac{2}{(\Delta - \varepsilon)^2} + 1 \end{aligned}$$

For the second term we apply Lemma 1.(c) to obtain, whenever $T\varepsilon^2 \geq e^2$,

$$T\mathbb{P}_\mu \left(\exists s \leq T : \hat{\mu}_{1,s} + \sqrt{\frac{2}{s} \log \left(\frac{T}{s} \right)} \leq \mu_1 - \varepsilon \right) \leq \frac{30e\sqrt{\log(\varepsilon^2 T)}}{\varepsilon^2} + \frac{16e}{\varepsilon^2}.$$

Finally, if $T(\Delta - \varepsilon)^2 \geq 2$ and $T\varepsilon^2 \geq e^2$,

$$\begin{aligned} R_\mu^{\text{UCB}}(T) &\leq \Delta \mathbb{E}_\mu[N_2(T)] \\ &\leq \frac{2 \log \left(\frac{T(\Delta - \varepsilon)^2}{2} \right)}{\Delta(1 - \varepsilon/\Delta)^2} + \frac{2 \sqrt{\pi \log \left(\frac{T(\Delta - \varepsilon)^2}{2} \right)}}{\Delta(1 - \varepsilon/\Delta)^2} + \frac{2}{\Delta(1 - \varepsilon/\Delta)^2} + \Delta + \Delta \left(\frac{30e\sqrt{\log(\varepsilon^2 T)}}{\varepsilon^2} + \frac{16e}{\varepsilon^2} \right) \end{aligned}$$

The asymptotic result follows by taking the limit as T tends to infinity and choosing $\varepsilon = \log^{-\frac{1}{8}}(T)$ while the minimax result follows by finding the minimum of the finite-time bound given above and the naive $R_\mu^{\text{UCB}}(T) \leq T\Delta$.

G Deviation Inequalities

As was already the case in the proof of Theorem 1, we heavily rely on the following well-known inequality on the tail of a Gaussian distribution: if $X \sim \mathcal{N}(0, 1)$, then for all $x > 0$

$$\mathbb{P}(X \geq x) \leq \min \left\{ 1, \frac{1}{x\sqrt{2\pi}} \right\} \exp(-x^2/2).$$

Lemma 1 gathers some more specific results that are useful in our regret analyses, and that we believe to be of a certain interest on their own.

Lemma 1. Let $\varepsilon > 0$ and $\Delta > 0$ and W_1, W_2, \dots be standard i.i.d. Gaussian random variables and $\hat{\mu}_t = \sum_{s=1}^t W_s/t$. Then the following hold:

- (a). $\mathbb{E} \left[\min \left\{ t : \sup_{s \geq t} |\hat{\mu}_s| \geq \varepsilon \right\} \right] \leq 1 + 9/\varepsilon^2$
- (b). if $T\Delta^2 \geq 2$ then $\sum_{n=1}^T \mathbb{P} \left(\hat{\mu}_n + \sqrt{\frac{2}{n} \log \left(\frac{T}{n} \right)} \geq \Delta \right) \leq \frac{2 \log \left(\frac{T\Delta^2}{2} \right)}{\Delta^2} + \frac{2\sqrt{\pi \log \left(\frac{T\Delta^2}{2} \right)}}{\Delta^2} + \frac{2}{\Delta^2} + 1$
- (c). if $T\varepsilon^2 \geq e^2$ then $\mathbb{P} \left(\exists s \leq T : \hat{\mu}_s + \sqrt{\frac{2}{s} \log \left(\frac{T}{s} \right)} + \varepsilon \leq 0 \right) \leq \frac{30e\sqrt{\log(\varepsilon^2 T)}}{\varepsilon^2 T} + \frac{16e}{\varepsilon^2 T}$

The proof of Lemma 1 follows from standard peeling techniques and inequalities for Gaussian sums of random variables. So far we do not know if this statement holds for subgaussian random variables where slightly weaker results may be shown by adding log log terms to the confidence interval, but unfortunately by doing this one sacrifices minimax optimality.

Proof of Lemma 1.(a). We use a standard peeling argument and the maximal inequality.

$$\begin{aligned} \mathbb{P}(\exists s \geq t : |\hat{\mu}_s| \geq \varepsilon) &\leq \sum_{k=1}^{\infty} \mathbb{P}(\exists s \in [kt, (k+1)t] : |\hat{\mu}_s| \geq \varepsilon) \\ &\leq \sum_{k=1}^{\infty} \mathbb{P}(\exists s \leq (k+1)t : |s\hat{\mu}_s| \geq kt\varepsilon) \\ &\leq \sum_{k=1}^{\infty} 2 \exp \left(-\frac{(kt\varepsilon)^2}{2(k+1)t} \right) = \sum_{k=1}^{\infty} 2 \exp \left(-\frac{kt\varepsilon^2}{2(1+1/k)} \right) \\ &\leq \sum_{k=1}^{\infty} 2 \exp \left(-\frac{kt\varepsilon^2}{4} \right) \end{aligned}$$

Therefore

$$\begin{aligned} \mathbb{E}[\tau] &\leq 1 + \sum_{t=1}^{\infty} \mathbb{P}(\tau \geq t) \leq 1 + \sum_{t=1}^{\infty} \min \left\{ 1, 2 \sum_{k=1}^{\infty} \exp \left(-\frac{kt\varepsilon^2}{4} \right) \right\} \\ &= 1 + \sum_{t=1}^{\infty} \min \left\{ 1, \frac{2}{\exp(t\varepsilon^2/4) - 1} \right\} \leq 1 + \int_0^{\infty} \min \left\{ 1, \frac{2}{\exp(t\varepsilon^2/4) - 1} \right\} dt \\ &\leq 1 + \frac{4 \log(4)}{\varepsilon^2} + \int_{4/\varepsilon^2 \log(4)}^{\infty} \frac{8}{3 \exp(t\varepsilon^2/4)} dt = 1 + \frac{4 \log(4)}{\varepsilon^2} + \frac{8}{3\varepsilon^2} \leq 1 + \frac{9}{\varepsilon^2}. \end{aligned}$$

□

Proof of Lemma 1.(b). Let ν be the solution of $\sqrt{2 \log(T/n)/n} = \Delta$, that is $\nu = 2W(\Delta^2 T/2)/\Delta^2$. Then

$$\sum_{n=1}^T \mathbb{P} \left(\hat{\mu}_n + \sqrt{\frac{2}{n} \log \left(\frac{T}{n} \right)} \geq \Delta \right) \leq \nu + \sum_{n=\lceil \nu \rceil}^T \mathbb{P} \left(\hat{\mu}_n \geq \Delta - \sqrt{\frac{2}{n} \log \left(\frac{T}{n} \right)} \right).$$

As for all $n \geq \nu$

$$\frac{2}{n} \log \frac{T}{n} \leq \frac{2}{\nu} \log \left(\frac{T}{\nu} \right) \frac{\nu}{n} = \Delta^2 \frac{\nu}{n},$$

$$\begin{aligned}
\sum_{n=\lceil \nu \rceil}^T \mathbb{P} \left(\hat{\mu}_n \geq \Delta - \sqrt{\frac{2}{n} \log \left(\frac{T}{n} \right)} \right) &\leq \sum_{n=\lceil \nu \rceil}^{\infty} \mathbb{P} \left(\hat{\mu}_n \geq \Delta \left(1 - \sqrt{\frac{\nu}{n}} \right) \right) \\
&\leq \sum_{n=\lceil \nu \rceil}^{\infty} \exp \left(-\frac{n\Delta^2}{2} \left(1 - \sqrt{\frac{\nu}{n}} \right)^2 \right) \\
&= \sum_{n=\lceil \nu \rceil}^{\infty} \exp \left(-\frac{\Delta^2}{2} (\sqrt{n} - \sqrt{\nu})^2 \right) \\
&\leq 1 + \int_{\nu}^{\infty} \exp \left(-\frac{\Delta^2}{2} (\sqrt{x} - \sqrt{\nu})^2 \right) dx \\
&= 1 + \frac{2}{\Delta} \int_0^{\infty} \left(\frac{y}{\Delta} + \sqrt{\nu} \right) \exp \left(-\frac{y^2}{2} \right) dy \\
&= 1 + \frac{2}{\Delta^2} + \frac{\sqrt{2\pi\nu}}{\Delta}.
\end{aligned}$$

Hence, as $\nu \leq 2 \log(\Delta^2 T/2)/\Delta^2$ whenever $T\Delta^2 \geq 2$,

$$\sum_{n=1}^T \mathbb{P} \left(\hat{\mu}_n + \sqrt{\frac{2}{n} \log \left(\frac{T}{n} \right)} \geq \Delta \right) \leq \frac{2}{\Delta^2} \left(\log \left(\frac{T\Delta^2}{2} \right) + 1 + \sqrt{\pi \log \left(\frac{T\Delta^2}{2} \right)} \right) + 1.$$

□

Proof of Lemma 1.(c). Let $x > 0$ and $n \in \mathbb{N}$, then by the reflection principle (eg., [Mörters and Peres \[2010\]](#)) it holds that

$$\mathbb{P}(\exists s \leq n : s\hat{\mu}_s + x \leq 0) = 2\mathbb{P}(n\hat{\mu}_n + x \leq 0) \leq 2 \min \left\{ \frac{1}{x} \sqrt{\frac{n}{2\pi}}, 1 \right\} \exp \left(-\frac{x^2}{2n} \right) \quad (7)$$

We prepare to use the peeling technique with a carefully chosen grid. Let

$$\eta = \frac{\log(\varepsilon^2 T)}{\log(\varepsilon^2 T) - 1} \quad \text{and} \quad G_k = [\eta^k, \eta^{k+1}[.$$

As $T\varepsilon^2 > e^2$, one has $\eta \in]1, 2[$. Moreover, our choice of η leads to the following inequality, that will be useful in the sequel

$$\forall x \geq \varepsilon^{-2}, \quad (x/T)^{\frac{1}{\eta}} \leq e(x/T) \quad (8)$$

Using a union bound and then Eq. (7), one can write

$$\begin{aligned}
&\mathbb{P} \left(\exists s \leq T : \hat{\mu}_s + \sqrt{\frac{2}{s} \log \left(\frac{T}{s} \right)} + \varepsilon \leq 0 \right) \\
&\leq \sum_{k=0}^{\infty} \mathbb{P} \left(\exists s \in G_k : s\hat{\mu}_s + \sqrt{2\eta^k \log \left(1 \vee \frac{T}{\eta^{k+1}} \right)} + \eta^k \varepsilon \leq 0 \right) \\
&\leq \sum_{k=0}^{\infty} 2 \min \left\{ 1, \sqrt{\frac{\eta}{4\pi}} \frac{1}{\sqrt{\log \left(1 \vee \frac{T}{\eta^{k+1}} \right)} + \varepsilon \sqrt{\eta^k/2}} \right\} \left(\frac{\eta^{k+1}}{T} \right)^{\frac{1}{\eta}} \exp \left(-\frac{\eta^{k-1} \varepsilon^2}{2} \right) \\
&\leq \sum_{k=0}^{\infty} f(k) \leq 2 \max_k f(k) + \int_0^{\infty} f(u) du,
\end{aligned}$$

where the function f is defined on $[0, +\infty[$ by

$$f(u) = 2 \min \left\{ 1, \sqrt{\frac{\eta}{4\pi \log \left(\frac{T}{\eta^{u+1}} \right)}} \right\} \left(\frac{\eta^{u+1}}{T} \right)^{\frac{1}{\eta}} \exp \left(-\frac{\eta^{u-1} \varepsilon^2}{2} \right),$$

with the convention that $\sqrt{\frac{\eta}{4\pi \log\left(\frac{T}{\eta^{u+1}}\right)}} = +\infty$ for u such that $T/\eta^{u+1} < 1$.

The last inequality relies on the fact that f can be checked to be unimodal, which permits to upper bound the sum with an integral. The maximum of f is easily upper bounded as follows, using notably (8):

$$\max_k f(k) \leq 2 \sup_{k \geq 0} \left(\frac{\eta^{k+1}}{T} \right)^{\frac{1}{\eta}} \exp\left(-\frac{\eta^{k-1}\varepsilon^2}{2}\right) = 2 \exp(-1/\eta) \left(\frac{2\eta}{\varepsilon^2 T} \right)^{\frac{1}{\eta}} \leq \frac{8e}{\varepsilon^2 T}.$$

The remainder of the proof is spent bounding the integral, which will be split into three disjoint intervals with boundaries at constants $k_1 < k_2$ given by

$$k_1 = \log(1/\varepsilon^2)/\log(\eta) \quad k_2 = 1 + \frac{\log\left(\frac{-\log \log(\eta)}{\varepsilon^2}\right)}{\log(\eta)}.$$

These are chosen such that

$$\eta^{k_1} = \varepsilon^{-2} \quad \exp\left(-\frac{\eta^{k_2-1}\varepsilon^2}{2}\right) = \sqrt{\log(\eta)}.$$

First, one has

$$\begin{aligned} I_1 &:= \int_0^{k_1} f(u) du \leq \int_0^{k_1} \frac{\sqrt{2}}{\sqrt{\pi \log\left(\frac{T}{\eta^{u+1}}\right)}} \left(\frac{\eta^{u+1}}{T} \right)^{\frac{1}{\eta}} \exp\left(-\frac{\eta^{u-1}\varepsilon^2}{2}\right) du \\ &\leq \frac{\sqrt{2}}{\sqrt{\pi \log\left(\frac{T}{\eta^{k_1+1}}\right)}} \int_0^{k_1} \left(\frac{\eta^{u+1}}{T} \right)^{\frac{1}{\eta}} du \leq \frac{\sqrt{2}}{\sqrt{\pi \log\left(\frac{\varepsilon^2 T}{\eta}\right)}} \frac{\eta}{\log(\eta)} \left(\frac{\eta}{\varepsilon^2 T} \right)^{\frac{1}{\eta}} \\ &\leq \frac{e\sqrt{2}\eta^2}{\sqrt{\pi \log\left(\frac{\varepsilon^2 T}{\eta}\right)} \log(\eta) \varepsilon^2 T} \leq \frac{e\sqrt{2}\eta^2}{\sqrt{\pi \log\left(\frac{\varepsilon^2 T}{\eta}\right)}} \frac{\log(\varepsilon^2 T)}{\varepsilon^2 T}, \end{aligned}$$

where the last inequality follows from the fact that $(\log(\eta))^{-1} \leq \eta/(\eta-1) = \log(\varepsilon^2 T)$.

Secondly,

$$\begin{aligned} I_2 &:= \int_{k_1}^{k_2} f(u) du \leq \int_{k_1}^{k_2} \frac{\sqrt{2}}{\sqrt{\pi \log\left(\frac{T}{\eta^{u+1}}\right)}} \left(\frac{\eta^{u+1}}{T} \right)^{\frac{1}{\eta}} \exp\left(-\frac{\eta^{u-1}\varepsilon^2}{2}\right) du \\ &\leq \frac{e\sqrt{2}}{\sqrt{\pi \log\left(\frac{T}{\eta^{k_2+1}}\right)}} \int_{k_1}^{\infty} \frac{\eta^{u+1}}{T} \exp\left(-\frac{\eta^{u-1}\varepsilon^2}{2}\right) du = \frac{e\sqrt{2}}{\sqrt{\pi \log\left(\frac{T}{\eta^{k_2+1}}\right)}} \frac{2\eta^2 \exp\left(-\frac{\eta^{k_1-1}\varepsilon^2}{2}\right)}{\varepsilon^2 T \log(\eta)} \\ &\leq \frac{2e\sqrt{2}}{\sqrt{\pi \log\left(\frac{T}{\eta^{k_2+1}}\right)}} \frac{\eta^2}{\varepsilon^2 T \log(\eta)} \leq \frac{2e\sqrt{2}\eta^2}{\sqrt{\pi \log\left(\frac{\varepsilon^2 T}{-\eta^2 \log \log(\eta)}\right)}} \frac{\log(\varepsilon^2 T)}{\varepsilon^2 T} \\ &\leq \frac{2e\sqrt{2}\eta^2}{\sqrt{\pi \log\left(\frac{\varepsilon^2 T}{\eta^2 \log \log(\varepsilon^2 T)}\right)}} \frac{\log(\varepsilon^2 T)}{\varepsilon^2 T} \end{aligned}$$

The second inequality follows from (8), since by definition of k_1 , one has $\eta^{u+1} \geq \varepsilon^{-2}$ for $u \geq k_1$, whereas the last inequalities use again that $(\log(\eta))^{-1} \leq \log(\varepsilon^2 T)$. Using similar arguments for the

third term, one has

$$\begin{aligned}
I_3 &:= \int_{k_2}^{\infty} f(u) du \leq 2 \int_{k_2}^{\infty} \left(\frac{\eta^{u+1}}{T} \right)^{\frac{1}{\eta}} \exp \left(-\frac{\eta^{u-1} \varepsilon^2}{2} \right) du \\
&\leq 2e \int_{k_2}^{\infty} \frac{\eta^{u+1}}{T} \exp \left(-\frac{\eta^{u-1} \varepsilon^2}{2} \right) du = \frac{4e\eta^2 \exp \left(-\frac{\eta^{k_2-1} \varepsilon^2}{2} \right)}{\varepsilon^2 T \log(\eta)} \\
&= \frac{4e\eta^2}{\varepsilon^2 T \sqrt{\log(\eta)}} \leq \frac{4e\eta^2}{\varepsilon^2 T} \sqrt{\log(\varepsilon^2 T)}
\end{aligned}$$

Combining the three upper bounds yield

$$\begin{aligned}
\int_0^{\infty} f(u) du &\leq I_1 + I_2 + I_3 \\
&\leq \frac{e\eta^2 \log(\varepsilon^2 T)}{\varepsilon^2 T} \left(\frac{\sqrt{2}}{\sqrt{\pi \log \left(\frac{\varepsilon^2 T}{\eta} \right)}} + \frac{2\sqrt{2}}{\sqrt{\pi \log \left(\frac{\varepsilon^2 T}{\eta^2 \log \log(\varepsilon^2 T)} \right)}} + \frac{4}{\sqrt{\log(\varepsilon^2 T)}} \right) \\
&\leq \frac{4e \log(\varepsilon^2 T)}{\varepsilon^2 T} \left(\frac{\sqrt{2}}{\sqrt{\pi \log \left(\frac{\varepsilon^2 T}{2} \right)}} + \frac{2\sqrt{2}}{\sqrt{\pi \log \left(\frac{\varepsilon^2 T}{4 \log \log(\varepsilon^2 T)} \right)}} + \frac{4}{\sqrt{\log(\varepsilon^2 T)}} \right)
\end{aligned}$$

It can be shown, using notably the inequality $\log u \leq u/e$, that for all $\varepsilon^2 T \geq e^2$,

$$\begin{aligned}
\log(\varepsilon^2 T/2) &\geq (1/2) \log(\varepsilon^2 T) \\
\log \left(\frac{\varepsilon^2 T}{4 \log \log(\varepsilon^2 T)} \right) &\geq \left(1 - \frac{4}{e^2} \right) \log(\varepsilon^2 T)
\end{aligned}$$

and one obtains

$$\int_0^{\infty} f(u) du \leq \frac{4e \log(\varepsilon^2 T)}{\varepsilon^2 T} \times \frac{2 + 2e\sqrt{2}/(\sqrt{e^2 - 4}) + 4\sqrt{\pi}}{\sqrt{\pi \log(\varepsilon^2 T)}} \leq \frac{30e}{\varepsilon^2 T} \sqrt{\log(\varepsilon^2 T)}.$$

Finally,

$$\mathbb{P} \left(\exists s \leq T : \hat{\mu}_s + \sqrt{\frac{2}{s} \log \left(\frac{T}{s} \right)} + \varepsilon \leq 0 \right) \leq \int_0^{\infty} f(u) du + 2 \max_k f(k) \leq \frac{30e}{\varepsilon^2 T} \sqrt{\log(\varepsilon^2 T)} + \frac{16e}{\varepsilon^2 T}.$$

□