



Pure Exploration in Infinitely-Armed Bandit Models with Fixed-Confidence

Maryam Aziz, Jesse Anderton, Emilie Kaufmann, Javed Aslam

► To cite this version:

Maryam Aziz, Jesse Anderton, Emilie Kaufmann, Javed Aslam. Pure Exploration in Infinitely-Armed Bandit Models with Fixed-Confidence. ALT 2018 - Algorithmic Learning Theory, Apr 2018, Lanzarote, Spain. hal-01729969

HAL Id: hal-01729969

<https://hal.science/hal-01729969>

Submitted on 12 Mar 2018

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

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

Pure Exploration in Infinitely-Armed Bandit Models with Fixed-Confidence

Maryam Aziz

*Northeastern University
Boston, MA*

AZIZM@CCS.NEU.EDU

Jesse Anderton

*Northeastern University
Boston, MA*

JESSE@CCS.NEU.EDU

Emilie Kaufmann

*CNRS & CRISTAL, UMR 9189, Université de Lille
Inria Lille, SequeL team*

EMILIE.KAUFMANN@UNIV-LILLE1.FR

Javed Aslam

*Northeastern University
Boston, MA*

JAA@CCS.NEU.EDU

Editor:

Abstract

We consider the problem of near-optimal arm identification in the fixed confidence setting of the infinitely armed bandit problem when nothing is known about the arm reservoir distribution. We (1) introduce a PAC-like framework within which to derive and cast results; (2) derive a sample complexity lower bound for near-optimal arm identification; (3) propose an algorithm that identifies a nearly-optimal arm with high probability and derive an upper bound on its sample complexity which is within a log factor of our lower bound; and (4) discuss whether our $\log^2 \frac{1}{\delta}$ dependence is inescapable for “two-phase” (select arms first, identify the best later) algorithms in the infinite setting. This work permits the application of bandit models to a broader class of problems where fewer assumptions hold.

Keywords: Infinitely-Armed Bandit Models, Pure Exploration

1. Introduction

We present an extension of the *stochastic multi-armed bandit* (MAB) model, which is applied to many problems in computer science and beyond. In a bandit model, an agent is confronted with a set of arms that are unknown probability distributions. At each round t , the agent chooses an arm to play, based on past observation, after which a reward drawn from the arm’s distribution is observed. This sequential sampling strategy (“bandit algorithm”) is adjusted to optimize some utility measure. Two measures are typical: cumulative regret minimization and pure exploration. For regret minimization, one attempts minimize *regret*, the difference between the expected cumulative rewards of an optimal strategy and the employed strategy. In the pure-exploration framework, one seeks the arm with largest mean irrespective of the observed rewards. Two dual settings have been studied: the *fixed-budget*

setting, wherein one can use only a given number of arm-pulls, and the *fixed-confidence* setting, wherein one attempts to achieve a utility target with minimal arm-pulls.

While the literature mainly considers bandit models with a known, *finite* number of arms, for many applications the number of arms may be very large and even infinite. In these cases, one can often settle for an arm which is “near” the best in some sense, as such an arm can be identified at significantly less cost. One such application is machine learning: given a large pool of possible classifiers (arms), one wants to find the one with minimal risk (mean reward) by sequentially choosing a classifier, training it and measuring its empirical test error (reward). In text, image, and video classification, one often encounters effectively infinite sets of classifiers which are prohibitively expensive to assess individually. Addressing such cases with bandit models is particularly useful when used within ensemble algorithms such as AdaBoost (Freund and Schapire, 1996), and some variations on this idea have already been explored (Appel et al., 2013; Busa-Fekete and Kégl, 2010; Dubout and Fleuret, 2014; Escudero et al., 2001), though the task of efficiently identifying a near-optimal classifier is at present unsolved. We here approach such problems from a theoretical standpoint.

Two distinct lines of work address a potentially *infinite set of arms*. Let \mathcal{W} be a (potentially uncountable) set of arms and assume that there exists $\mu : \mathcal{W} \rightarrow \mathbb{R}$, a mean-reward mapping such that when some arm w is selected, one observes an independent draw of a random variable with mean $\mu(w)$. One line of research (Kleinberg et al., 2008; Bubeck et al., 2011b; Grill et al., 2015) assumes that \mathcal{W} is some metric space, and that μ has some regularity property with respect to the metric (for example it is locally-Lipschitz). Both regret minimization and fixed-budget pure-exploration problems have been studied in this setting. Another line of research, starting with the work of Berry et al. (1997) assumes no particular structure on \mathcal{W} and no regularity for μ . Rather, there is some *reservoir distribution* on the arms’ means (the set $\mu(\mathcal{W})$ with our notation) such that at each round the learner can decide to query a new arm, whose mean is drawn from the reservoir, and sample it, or to sample an arm that was previously queried. While regret minimization was studied by several authors (Wang et al., 2009; Bonald and Proutière, 2013; David and Shimkin, 2014), the recent work of Carpentier and Valko (2015) is the first to study the pure-exploration problem in the fixed-budget setting.

We present a novel theoretical framework for the fixed-confidence pure-exploration problem in an infinite bandit model with a reservoir distribution. The reservoir setting seems well-suited for machine learning, since it is not clear whether the test error of a parametric classifier is smooth with respect to its parameters. Typically, an assumption is made on the form of the tail of the reservoir which allows estimation of the probability that an independently-drawn arm will be “good;” that is, close to the best possible arm. However, for problems such as that mentioned above such an assumption does not seem warranted. Instead, we employ a parameter, α , indicating the probability of independently drawing a “good” arm. When a tail assumption can be made, α can be computed from this assumption. Otherwise, it can be chosen based on the user’s needs. Note that the problem of identifying a “top- α ” arm in the infinite case corresponds to the finite case problem of finding one of the top m arms from a set of n arms, for $n > m$, with $\alpha = m/n$. The first of two PAC-like frameworks we introduce, the (α, δ) framework, aims to identify an arm in the top- α tail of the reservoir with probability at least $1 - \delta$, using as few samples as possible.

We now motivate our second framework. When no assumptions can be made on the reservoir, one may encounter reservoirs with large probability masses close to the boundary of the top- α tail. Indeed, the distribution of weighted classifier accuracies in later rounds of AdaBoost has this property, as the weights are chosen to drive all classifiers toward random performance. This is a problem for any framework defined purely in terms of α , because such masses make us likely to observe arms which are not in the top- α tail but which are hard to distinguish from top- α arms. However, in practice their similarity to top- α arms makes them reasonable arms to select. For this reason we add an ϵ relaxation, which limits the effort spent on arms near the top- α tail while adding directly to the simple regret a user may observe. Formally, our $(\alpha, \epsilon, \delta)$ framework seeks an arm within ϵ of the top- α fraction of the arms with probability at least $1 - \delta$, using as few samples from the arms as possible.

Although α and ϵ both serve to relax the goal of finding an arm with maximum mean, they have distinct purposes and are both useful. One might wonder, if the inverse CDF G^{-1} for the arm reservoir was available (at least at the tail), why one would not simply compute $\epsilon' = \epsilon + G^{-1}(1 - \alpha)$ and use the established (ϵ, δ) framework. Indeed, α is important precisely when the form of the reservoir tail is unknown. The user of an algorithm will wish to limit the effort spent in finding an optimal arm, and with no assumptions on the reservoir ϵ alone is insufficient to limit an algorithm's sample complexity. Just as there might be large probability close to the α boundary, it may be that there is virtually no probability within ϵ of the top arm. The user applies α to (effectively) specify how hard to work to estimate the reservoir tail, and ϵ to specify how hard to work to differentiate between individual arms.

Our approach differs from the typical reservoir setting in that it does not require any regularity assumption on the tail of the reservoir distribution, although it can take advantage of one when available. Within this framework, we prove a lower bound on the expected number of arm pulls necessary to achieve (α, δ) or $(\alpha, \epsilon, \delta)$ performance by generalizing the information-theoretic tools introduced by Kaufmann et al. (2016) in the finite MAB setting. We also study a simple algorithmic solution to the problem based on the KL-LUCB algorithm of Kaufmann and Kalyanakrishnan (2013), an algorithm for ϵ -best arm identification in bandit models with a finite number of arms, and we compare its performance to our derived lower bound theoretically. Our algorithm is an $(\alpha, \epsilon, \delta)$ algorithm, but we show how to achieve (α, δ) performance when assumptions can be made on the tail of the reservoir.

We introduce the (α, δ) and $(\alpha, \epsilon, \delta)$ frameworks and relate them to existing literature in Section 2. Section 3 proves our sample complexity lower bounds. In Section 4, we present and analyze the (α, ϵ) -KL-LUCB algorithm for one-dimensional exponential family reward distributions. A comparison between our upper and lower bounds can be found in Section 4.4. We defer most proofs to the appendix, along with some numerical experiments.

2. Pure Exploration with Fixed Confidence

Here we formalize our frameworks and connect them to the existing literature.

2.1 Setup, Assumptions, and Notation

Let $(\mathcal{W}, \mathcal{F}_{\mathcal{W}}, M)$ be a probability space over arms with measure M , where each arm $w \in \mathcal{W}$ is some abstract object (e.g. a classifier), and let $(\Theta, \mathcal{F}_{\Theta})$ be a measurable space over expected rewards, where $\Theta \subseteq \mathbb{R}$ is a continuous interval and \mathcal{F}_{Θ} is the Borel σ -algebra over

Θ (i.e. the smallest σ -algebra containing all sub-intervals of Θ). Also let $P_\Theta = \{p_\theta, \theta \in \Theta\}$ be a parametric set of probability distributions such that each distribution is continuously parameterized by its mean. To ease the notation, we shall assume $\mathbb{E}_{X \sim p_\theta}[X] = \theta$. One can think of P_Θ as a one-parameter exponential family (e.g. the family of Bernoulli, Gaussian with fixed and known variance, Poisson or Exponential distributions with means in some interval or other subset of \mathbb{R}), however we do not limit ourselves to such well-behaved reward distributions. We defer our further assumptions on P_Θ to Section 3.1. We will denote by f_θ the density of the element in P_Θ with mean θ .

An *infinite bandit model* is characterized by a probability measure M over $(\mathcal{W}, \mathcal{F}_\mathcal{W})$ together with a measurable mapping $\mu : \mathcal{W} \rightarrow \Theta$ assigning a mean θ (and therefore a reward distribution p_θ) to each arm. The role of the measure M is to define the top- α fraction of arms, as we will show in Eq. 2; it can be used by the algorithm to sample arms. At each time step t , a user selects an arm $W_t \in \mathcal{W}$, based on past observation. He can either query a new arm in \mathcal{W} (which may be sampled $W_t \sim M$, or selected adaptively) or select an arm that has been queried in previous rounds. In any case, when arm W_t is drawn, an independent sample $Z_t \sim p_{\mu(W_t)}$ is observed.

For example, when boosting decision stumps (binary classifiers which test a single feature against a threshold), the set of arms \mathcal{W} consists of all possible decision stumps for the corpus, and the expected reward for each arm is its expected accuracy over the sample space of all possible classification examples. An algorithm may choose to draw the arms at random according to the probability measure M ; this is commonly done by, in effect, placing uniform probability mass over the thresholds placed halfway between the distinct values seen in the training data and placing zero mass over the remaining thresholds. We are particularly interested in the case when the number of arms in the support for M is so large as to be effectively infinite, at least with respect to the available computational resources.

We denote by \mathbb{P}_μ^M and \mathbb{E}_μ^M the probability and expectation under an infinite bandit model with arm probability measure M and mean function μ . The history of the bandit game up to time t is $H_t = ((W_1, Z_1), \dots, (W_t, Z_t))$. By our assumption, the arm selected at round t only depends on H_{t-1} and U_t , which is uniform on $[0, 1]$ and independent of H_{t-1} (used to sample from M if needed). In particular, the conditional density of W_t given H_{t-1} , denoted by $\mathbb{P}_\mu^M(W_t|H_{t-1})$, is independent of the mean mapping μ . Note that this property is satisfied as well if, when querying a new arm, W_t can be chosen arbitrarily in \mathcal{W} (depending on H_{t-1}), and not necessarily at random from M . Under these assumptions, one can compute the likelihood of H_T :

$$\ell(H_T; \mu, M) = \prod_{t=1}^T f_{\mu(W_t)}(Z_t) \mathbb{P}_\mu^M(W_t|H_{t-1}). \quad (1)$$

Note that the arms W_t are not latent objects: they are assumed to be observed, but not their means $\mu(W_t)$. For instance, in our text classification example we know the classifier we are testing but not its true classification accuracy. Treating arms as observed in this way simplifies the likelihood by making the choice of new arms to query independent of their mean mappings. This is key to our approach to dealing with reservoirs about which nothing is known; we can avoid integrating over such reservoirs and so do not require the reservoir to be smooth. For details, see Appendix A.1.

2.2 Objective and Generic Algorithm

Reservoir distribution. The probability space over arms $(\mathcal{W}, \mathcal{F}_{\mathcal{W}}, M)$ and the mapping μ is used to form a pushforward measure over expected rewards $M_{\Theta}(\mathcal{E}) := (\mu_*(M))(\mathcal{E}) = M(\mu^{-1}(\mathcal{E}))$, for $\mathcal{E} \in \mathcal{F}_{\Theta}$, inducing the probability space $(\Theta, \mathcal{F}_{\Theta}, M_{\Theta})$ over expected rewards. We define our *reservoir distribution* CDF $G(\tau) = M_{\Theta}(\{\theta \leq \tau\})$ whose density g is its Radon-Nikodym derivative with respect to M . For convenience, we also define the “inverse” CDF $G^{-1}(p) := \inf\{\theta : G(\theta) \geq p\}$. We assume that G has bounded support and let μ^* be the largest possible mean under the reservoir distribution, $\mu^* := G^{-1}(1) = \inf\{\theta : G(\theta) = 1\}$.

In the general setup introduced above, the reservoir may or may not be useful to query new arms, but it is needed to define the notion of top- α fraction.

Finding an arm in the top- α fraction. In our setting, for some fixed $\alpha \in]0, 1[$ and some $\epsilon \geq 0$, the goal is to identify an arm that belongs to the set

$$\mathcal{G}_{M,\mu}^{\alpha,\epsilon} := \{w \in \mathcal{W} : \mu(w) \geq G^{-1}(1 - \alpha) - \epsilon\} \quad (2)$$

of arms whose expected mean rewards is high, in the sense that their mean is within ϵ of the quantile of order $1 - \alpha$ of the reservoir distribution. For notational convenience, when we set ϵ to zero we write $\mathcal{G}_{M,\mu}^{\alpha} := \mathcal{G}_{M,\mu}^{\alpha,0}$.

Generic algorithm An algorithm is made of a *sampling rule* (W_t) , a *stopping rule* τ (with respect to the filtration generated by H_t) and a *recommendation rule* \hat{s}_{τ} that selects one of the queried arms as a candidate arm from $\mathcal{G}_{M,\mu}^{\alpha,\epsilon}$. This is summarized in Algorithm 1.

Fix $\delta \in]0, 1[$. An algorithm that returns an arm from $\mathcal{G}_{M,\mu}^{\alpha,\epsilon}$ with probability at least $1 - \delta$ is said to be $(\alpha, \epsilon, \delta)$ -correct. Moreover, an $(\alpha, \epsilon, \delta)$ -correct algorithm must perform well on all possible infinite bandit models: $\forall(M, \mu), \mathbb{P}_{\mu}^M(\hat{s}_{\tau} \in \mathcal{G}_{M,\mu}^{\alpha,\epsilon}) \geq 1 - \delta$. Our goal is to build an $(\alpha, \epsilon, \delta)$ -correct algorithm that uses as few samples as possible, i.e. for which $\mathbb{E}_{\mu}^M[\tau]$ is small. We similarly define the notion of (α, δ) -correctness when $\epsilon = 0$.

Input: Arm set \mathcal{W} , target α, ϵ, δ

Output: Some arm \hat{s}

for $t \leftarrow 1, 2, \dots$ **do**
(choose one of:)

1. Pull arm: Choose $W_t \sim \mathbb{P}_{\mu}^M(W_t | H_{t-1})$ and observe reward $Z_t \sim p_{\mu}(W_t)$
2. Stop: Choose $\hat{s} \leftarrow W_s$ for some $s < t$,
return \hat{s}

end for

Algorithm 1: Generic algorithm.

$(\alpha, \epsilon, \delta)$ -correctness When little is known about the reservoir distribution (e.g. it might not even be smooth), an ϵ -relaxed algorithm is appropriate. The choice of ϵ represents a tradeoff between simple regret (defined shortly) and the maximum budget used to differentiate between arms. We provide our lower bound in both ϵ -relaxed and unrelaxed forms. Our algorithm requires an ϵ parameter, but we show how this parameter can be chosen under regularity assumptions on the tail of the reservoir to provide an (α, δ) -correct algorithm.

Simple regret guarantees. In the infinite bandit literature, performance is typically measured in terms of *simple regret*: $r_{\tau} = \mu^* - \mu(\hat{s}_{\tau})$. If the tail of the reservoir distribution is bounded, one can obtain simple regret upper bounds for an algorithm in our framework.

A classic assumption (see, e.g. Carpentier and Valko (2015)) is that there exists $\beta > 0$ and two constants E, E' such that

$$\forall \rho > 0, E\rho^\beta \leq M_\Theta(\{\theta \geq \mu^* - \rho\}) \leq E'\rho^\beta. \quad (3)$$

With $C = E^{-1/\beta}$ and $C' = (E')^{-1/\beta}$, this translates into $\forall \alpha > 0, C\alpha^{1/\beta} \leq \mu^* - G^{-1}(1 - \alpha) \leq C'\alpha^{1/\beta}$, and a (α, δ) -correct algorithm has its simple regret upper bounded as

$$\mathbb{P}(r_\tau \leq C'\alpha^{1/\beta}) \geq 1 - \delta. \quad (4)$$

Similarly, a $(\alpha, \epsilon, \delta)$ -correct algorithm has its simple regret bounded as

$$\mathbb{P}(r_\tau \leq C'\alpha^{1/\beta} + \epsilon) \geq 1 - \delta. \quad (5)$$

If β is known, α and ϵ can be chosen to guarantee a simple regret below an arbitrary bound.

2.3 Related Work

Bandit models were introduced by Thompson (1933). There has been recent interest in pure-exploration problems (Even-Dar et al., 2006; Audibert et al., 2010); for which good algorithms are expected to differ from those for the classic regret minimization objective (Bubeck et al., 2011a; Kaufmann and Garivier, 2017).

For a finite number of arms with means μ_1, \dots, μ_K , the fixed-confidence best arm identification problem was introduced by Even-Dar et al. (2006). The goal is to select an arm $\hat{a} \in \{1, \dots, K\}$ satisfying $\mathbb{P}(\mu_{\hat{a}} \geq \mu^* - \epsilon) \geq 1 - \delta$, where $\mu^* = \max_a \mu_a$. Such an algorithm is called (ϵ, δ) -PAC. In our setting, assuming a uniform reservoir distribution over $\{1, \dots, K\}$ yields an (α, δ) -correct algorithm with α being the fraction of ϵ -good arms. Algorithms are either based on successive eliminations (Even-Dar et al., 2006; Karnin et al., 2013) or on confidence intervals (Kalyanakrishnan et al., 2012; Gabillon et al., 2012). For exponential family reward distributions, the KL-LUCB algorithm of Kaufmann and Kalyanakrishnan (2013) refines the confidence intervals to obtain better performance compared to its Hoeffding-based counterpart, and a sample complexity scaling with the Chernoff information between arms (an information-theoretic measure related to the Kullback-Leibler divergence). We build on this algorithm to define (α, ϵ) -KL-LUCB in Section 4. *Lower bounds* on the sample complexity have also been proposed by Mannor et al. (2004); Kaufmann et al. (2016); Garivier and Kaufmann (2016). In Section 3 we generalize the change of distribution tools used therein to present a lower bound for pure exploration in an infinite bandit model.

Regret minimization has been studied extensively for infinite bandit models (Berry et al., 1997; Wang et al., 2009; Bonald and Proutière, 2013; David and Shimkin, 2014), whereas Carpentier and Valko (2015) is the first work dealing with pure-exploration for general reservoirs. The authors consider the fixed-budget setting, under the tail assumption (3) for the reservoir distribution, already discussed.

Although the fixed-confidence pure-exploration problem for infinitely armed bandits has been rarely addressed for general reservoir distributions, the *most-biased coin problem* studied by Chandrasekaran and Karp (2012); Jamieson et al. (2016) can be viewed as a particular instance, with a specific reservoir distribution that is a mixture of “heavy” coins of mean θ_1 and “light” coins of mean θ_0 : $G = (1 - \alpha)\delta_{\theta_0} + \alpha\delta_{\theta_1}$, where $\theta_1 > \theta_0$ and with δ_θ here denoting the Dirac delta function. The goal is to identify, with probability at least $1 - \delta$, an arm with mean θ_1 . If a lower bound α_0 on α is known, this is equivalent to finding

an (α_0, δ) -correct algorithm by our definition. We suggest in Section 4.4 that the sample complexity of any two-phase algorithm (such as ours) might scale like $\log^2 \frac{1}{\delta}$, while Jamieson et al. achieve a dependence on δ of $\log \frac{1}{\delta}$ for the special case they address.

Finally, the recent work of Chaudhuri and Kalyanakrishnan (2017) studies a framework that is similar to the one introduced in this paper¹. Their first goal of identifying, in a finite bandit model an arm with mean larger than $\mu_{[m]} - \epsilon$ (with $\mu_{[m]}$ the arm with m -th largest mean) is extended to the infinite case, in which the aim is to find an (α, ϵ) -optimal arm. The first algorithm proposed for the infinite case applies the Median Elimination algorithm (Even-Dar et al., 2006) on top of $\frac{1}{\alpha} \log \frac{2}{\delta}$ arms drawn from the reservoir and is proved to have a $O\left(\frac{1}{\alpha \epsilon^2} \log^2 \frac{1}{\delta}\right)$ sample complexity. The dependency in $\log^2 \frac{1}{\delta}$ is the same as the one we obtain for (α, ϵ) -KL-LUCB, however our analysis goes beyond the scaling in $\frac{1}{\epsilon^2}$ and reveals a complexity term based on KL-divergence, that can be significantly smaller. Another algorithm is presented, without sample complexity guarantees, that runs LUCB on successive batches of arms drawn from the reservoir in order to avoid memory storage issues.

3. Lower Bound

We now provide sample complexity lower bounds for our two frameworks.

3.1 Sample complexity lower bound

Our lower bound scales with the Kullback-Leibler divergence between arm distributions p_{θ_1} and p_{θ_2} , denoted by $d(\theta_1, \theta_2) := \text{KL}(p_{\theta_1} \| p_{\theta_2}) = \mathbb{E}_{X \sim p_{\theta_1}} \left[\log \frac{f_{\theta_1}(X)}{f_{\theta_2}(X)} \right]$.

Furthermore, we make the following assumptions on the arm reward distributions, that are typically satisfied for one-dimensional exponential families.

Assumption 1 *The KL divergence, that is the application $(\theta_1, \theta_2) \mapsto d(\theta_1, \theta_2)$ is continuous on $\Theta \times \Theta$, and Θ and d satisfy*

- $\theta_1 \neq \theta_2 \Rightarrow 0 < d(\theta_1, \theta_2) < \infty$
- $\theta_1 < \theta_2 < \theta_3 \Rightarrow d(\theta_1, \theta_3) > d(\theta_2, \theta_3)$ and $d(\theta_1, \theta_2) < d(\theta_1, \theta_3)$

It also relies on the following partition of the arms in \mathcal{W} by their expected rewards. Let $m = \lceil 1/\alpha \rceil$. We partition \mathcal{W} into subsets \mathcal{S}^i for $1 \leq i \leq m$, where $\mathcal{S}^i = \{w \in \mathcal{W} : \mu(w) \in [b_i, b_{i-1}]\}$. The interval boundaries b_i are defined so that each subset has measure α under the reservoir distribution g , with the possible exception of the subset with smallest expected reward. In particular, $b_0 = \mu^*$ and b_i lies at the boundary between subsets i and $i + 1$.

$$b_i = \begin{cases} \mu^* & \text{if } i = 0 \\ G^{-1}(G(b_{i-1}) - \alpha) & \text{if } i \geq 1 \end{cases}, \quad (6)$$

where μ^* is defined in Eq. 2.2. See Figure 1 for an illustration.




Figure 1: Our reservoir partition. Each consecutive α -interval on the CDF defines some subset \mathcal{S}^i .

1. Note that we became aware of their work after submitting our paper.

In the Bernoulli case, Assumption 2 reduces to $\mu^* < 1$ (no arm has perfect performance); when the set of possible means Θ is unbounded it always holds as G has a finite support.

Assumption 2 $\mu^* < \sup_{\theta \in \Theta} \theta$.

Theorem 1 Fix some $\alpha, \delta \in]0, 1[$. Any (α, δ) -correct algorithm needs an expected sample complexity $\mathbb{E}_\mu^M[\tau]$ that is lower bounded as follows.

$$\mathbb{E}_\mu^M[\tau] \geq \left(\frac{1}{d(\mu^*, b_2)} + \sum_{i=2}^{m-1} \frac{1}{d(b_i, \mu^*)} \right) \log \frac{1}{2.4\delta}.$$

Remark 2 When \mathcal{W} is finite s.t. $|\mathcal{W}| = K$, if we choose a uniform reservoir and let $\alpha = 1/K$, then $|\mathcal{S}^i| = 1$ for all i and our lower bound reduces to the bound obtained by Kaufmann et al. (2016) for best arm identification with $\epsilon = 0$. Assuming arm means $\theta_1 > \theta_2 \geq \dots \geq \theta_K$, one has

$$\mathbb{E}_\mu^M[T] \geq \left[\frac{1}{d(\theta_1, \theta_2)} + \sum_{i=2}^K \frac{1}{d(\theta_i, \theta_1)} \right] \log \frac{1}{2.4\delta}.$$

3.2 Proof of Theorem 1

The proof relies on the following lemma that expresses a change of measure in an infinite bandit model. Its proof is detailed in Appendix A.

Lemma 3 Let $\lambda : \mathcal{W} \rightarrow \Theta$ be an alternative mean-mapping. Let $T_i(t) = \sum_{s=1}^t \mathbf{1}_{(w_s \in \mathcal{S}^i)}$ be the number of times an arm in \mathcal{S}^i has been selected. For any stopping time σ and any event $C \in H_\sigma$,

$$\sum_{i=1}^m \mathbb{E}_\mu^M[T_i(\sigma)] \sup_{w \in \mathcal{S}^i} d(\mu(w), \lambda(w)) \geq \text{kl}(\mathbb{P}_\mu^M(C), \mathbb{P}_\lambda^M(C)),$$

where $\text{kl}(x, y) = x \log(x/y) + (1-x) \log((1-x)/(1-y))$ is the Bernoulli relative entropy.

Let $\tau^i = T_i(\tau)$ be the (random) number of draws from arms in \mathcal{S}^i , so $\tau = \sum_{i=1}^m \tau^i$. Our lower bound on $\mathbb{E}_\mu^M[\tau]$ follows from bounds on each of the $\mathbb{E}_\mu^M[\tau^i]$. We omit \mathcal{S}^m because its measure may be less than α . By Assumption 2, there is $\epsilon > 0$ such that $\mu^* + \epsilon < \sup_{\theta \in \Theta} \theta$. Fix i between 2 and $m-1$ and define an alternative arm reward mapping $\lambda^i(w)$ as follows.

$$\lambda^i(w) = \begin{cases} \mu^* + \epsilon & \text{if } w \in \mathcal{S}^i \\ \mu(w) & \text{otherwise} \end{cases} \quad (7)$$

This mapping induces an alternative reservoir distribution g^i under which for all $i < m$, $\mathcal{G}_{M, \lambda^i}^\alpha = \mathcal{S}^i$ because \mathcal{S}^i has measure α (as M is unchanged) and under λ^i the expected rewards of its arms are above all other arms by at least ϵ . Also, by construction $\mathcal{G}_{M, \mu}^\alpha = \mathcal{S}^1$.

Define the event $C_\mu = (\hat{s}_\tau \in \mathcal{S}^1)$. Any (α, δ) -correct algorithm thus satisfies $\mathbb{P}_\mu^M(C_\mu) \geq 1 - \delta$ and $\mathbb{P}_\lambda^M(C_\mu) \leq \delta$. First using some monotonicity properties of the binary relative entropy $\text{kl}(x, y)$, one has $\text{kl}(\mathbb{P}_\mu^M(C_\mu), \mathbb{P}_{\lambda^i}^M(C_\mu)) \geq \text{kl}(1 - \delta, \delta) \geq \log \frac{1}{2.4\delta}$, where the second inequality is due to Kaufmann et al. (2016).

Applying Lemma 3 to event C_μ and using the fact that $\lambda^j(w) = \mu(w)$ for all $j \neq i$, one obtains $\mathbb{E}_\mu^M [T^i] \sup_{w \in \mathcal{S}^i} d(\mu(w), \mu^* + \epsilon) \geq \log \frac{1}{2.4\delta}$. Letting ϵ go to 0 yields, for all $i \neq 1$,

$$\mathbb{E}_\mu^M [\tau^i] \geq \frac{1}{\sup_{w \in \mathcal{S}^i} d(\mu(w), \mu^*)} \log \left(\frac{1}{2.4\delta} \right),$$

and $\sup_{w \in \mathcal{S}^i} d(\mu(w), \mu^*) \leq d(b_i, \mu^*)$ as $\theta \mapsto d(\theta, \mu^*)$ is decreasing when $\theta < \mu^*$.

We now define the alternative mean rewards mapping, for $\epsilon > 0$ small enough

$$\lambda^1(w) = \begin{cases} b_2 - \epsilon & \text{if } w \in \mathcal{S}^1 \\ \mu(w) & \text{otherwise} \end{cases} \quad (8)$$

One has $\mathcal{G}_{M,\mu}^\alpha = \mathcal{S}_1$ whereas $\mathcal{G}_{M,\lambda^1}^\alpha = \mathcal{S}_2$, hence letting $C_\mu = (\hat{s}_\tau \in \mathcal{S}_1)$ satisfies $\mathbb{P}_\mu^M (C_\mu) \geq 1 - \delta$ and $\mathbb{P}_{\lambda^1}^M (C_\mu) \leq \delta$. Using the same reasoning as before yields

$$\log \left(\frac{1}{2.4\delta} \right) \leq \mathbb{E}_\mu^M [\tau^1] \sup_{w \in \mathcal{S}^1} d(\mu(w), b_2 - \epsilon) = d(\mu^*, b_2 - \epsilon)$$

Letting ϵ go to zero yields $\mathbb{E}_\mu^M [\tau^1] \geq \frac{1}{d(\mu^*, b_2)} \log \left(\frac{1}{2.4\delta} \right)$. \square

One can prove an ϵ -relaxed version of this theorem, which provides a lower bound on the number of samples needed to find an arm whose expected reward is within ϵ of the top- α fraction of arms with probability at least $1 - \delta$. When $\epsilon > 0$ multiple subsets may contain such arms, and the proof approach above does not work for these subsets. We instead adopt the strategy of Mannor et al. (2004): at most one such subset can have probability greater than $1/2$ of its arms being chosen by the algorithm, so we exclude this subset from our bound. We arrive at the following, which holds when $\mu^* + \epsilon$ is in Θ (i.e. for ϵ small enough).

Remark 4 Fix some $\alpha, \epsilon, \delta \in]0, 1[$, and let q be the number of subsets containing arms within ϵ of the top α fraction. Any $(\alpha, \epsilon, \delta)$ -correct algorithm needs an expected sample complexity $\mathbb{E}_\mu^M [\tau]$ that is lower bounded as follows.

$$\mathbb{E}_\mu^M [\tau] \geq \left(\frac{q-1}{d(b_1 - \epsilon, \mu^* + \epsilon)} + \sum_{i=q+1}^{m-1} \frac{1}{d(b_i, \mu^* + \epsilon)} \right) \log \frac{1}{4\delta}.$$

4. Algorithm and Upper Bound

In this section we assume that $P_\Theta = \{p_\theta, \theta \in \Theta\}$ is a one-parameter exponential family, meaning that there exists some twice differentiable convex function $b(\theta)$ and some reference measure ν such that p_θ has a density f_θ with respect to ν , where $f_\theta(x) = \exp(\theta x - b(\theta))$. Distributions in an exponential family can indeed be parameterized by their means as $\mu = b^{-1}(\theta)$. We do not make any new assumptions on the reservoir distribution.

Under these assumptions on the arms, we present and analyze a two-phase algorithm called (α, ϵ) -KL-LUCB. We prove the $(\alpha, \epsilon, \delta)$ -correctness of this algorithm and a high probability upper bound on its sample complexity in terms of the complexity of the reservoir distribution induced by the arm measure M and the arm reward mapping function μ . We also show how to obtain (α, δ) -correctness under assumptions on the tail of the reservoir.

4.1 The algorithm

(α, ϵ) -KL-LUCB, presented as Algorithm 2, is a two-phase algorithm. It first queries $n = \frac{1}{\alpha} \log \frac{2}{\delta}$ arms w^1, \dots, w^n from M , the measure over \mathcal{W} , and then runs the KL-LUCB algorithm (Kaufmann and Kalyanakrishnan, 2013) on the queried arms. KL-LUCB identifies the m -best arms in a multi-armed bandit model, up to some $\epsilon > 0$. We use it with $m = 1$. This algorithm adaptively selects pairs of arms to sample from based on confidence intervals on the means of the arms. These confidence intervals rely on some exploration rate

$$\beta(t, \delta) := \log(k_1 n t^\gamma / \delta), \quad (9)$$

for constants $\gamma > 1$ and $k_1 > 2(1 + \frac{1}{\gamma} - 1)$. The upper and lower confidence bounds are

$$U_a(t) := \max \{ \theta \in \Theta : N_a(t) d(\hat{p}_a(t), \theta) \leq \beta(t, \delta) \} \quad (10)$$

$$L_a(t) := \min \{ \theta \in \Theta : N_a(t) d(\hat{p}_a(t), \theta) \leq \beta(t, \delta) \}, \quad (11)$$

where $N_a(t) = \sum_{s=1}^t \mathbb{1}_{(W_s=w^a)}$ is the number of times arm w^a was sampled by round t and $\hat{p}_a(t) = \frac{1}{N_a(t)} \sum_{s=1}^t \mathbb{1}_{(W_s=w^a)} Z_s$ is the empirical mean reward of arm w^a at round t , where Z_s is an i.i.d. draw from arm W_s , with distribution $p_{\mu(W_s)}$. Recall that $d(\mu_1, \mu_2)$ is the KL divergence between arm distributions parameterized by their means μ_1 and μ_2 .

For each queried arm w^a , the algorithm maintains a confidence interval $\mathcal{I}_a(t) = [L_a(t), U_a(t)]$ on $\mu(w^a)$, and at any even round t selects two arm indexes: (1) the empirical best arm $\hat{a}(t) \in \arg\max_{a=1, \dots, n} \hat{p}_a(t)$, and (2) the arm among the empirical worst arms that is most likely to be mistaken with $\hat{a}(t)$, $\hat{b}(t) = \arg\max_{a \neq \hat{a}(t)} U_a(t)$. The two arms are sampled: $W_t = w^{\hat{a}(t)}$ and $W_{t+1} = w^{\hat{b}(t)}$ and the confidence intervals are updated. The algorithm terminates when the overlap between the associated confidence intervals is smaller than some $\epsilon > 0$: $\tau = \inf \{ t \in \mathbb{N}^* : L_{\hat{a}(t)}(t) > \max_{a \neq \hat{a}(t)} U_a(t) + \epsilon \}$. The recommendation rule is $\hat{s}_\tau = w^{\hat{a}(t)}$.

4.2 $(\alpha, \epsilon, \delta)$ -Correctness

For the algorithm to be $(\alpha, \epsilon, \delta)$ -correct, it is sufficient that the following two events occur:

- A is the event that some w^a was drawn from the top- α fraction of the reservoir.
- B is the event that the KL-LUCB algorithm succeeds in identifying an arm within ϵ of the best arm among the n arms drawn in the initialization phase.

Indeed, on $A \cap B$ the recommended arm $\hat{s}_\tau = w^{\hat{a}(\tau)}$ satisfies $\mu(\hat{s}_\tau) > \max_a \mu(w^a) - \epsilon \geq G^{-1}(1 - \alpha) - \epsilon$, hence \hat{s}_τ belongs to the top- α fraction, up to ϵ . We prove in Appendix B that $\mathbb{P}(A^c) \leq \delta/2$ and $\mathbb{P}(B^c) \leq \delta/2$, which yields the following result.

Input: $\alpha, \epsilon, \delta > 0$

$$n = \frac{1}{\alpha} \log \frac{2}{\delta}$$

for $a \leftarrow 1$ **to** n **do**

draw arm $w^a \sim M$

sample arm w^a once

end for

$t = n$ (current number of samples drawn)

$B(n) = \infty$ (stopping index)

Compute confidence bounds $U_a(n)$ and $L_a(n)$

while $B(t) > \epsilon$ **do**

Draw arm $w^{\hat{a}(t)}$ and $w^{\hat{b}(t)}$

$t = t + 2$

Update confidence bounds, compute $\hat{a}(t)$ and $\hat{b}(t)$

$B(t) = U_{\hat{b}(t)}(t) - L_{\hat{a}(t)}(t)$

end while

return $w^{\hat{a}(t)}$

Algorithm 2: (α, ϵ) -KL-LUCB

Lemma 5 *With $\beta(t, \delta)$ defined in (9), (α, ϵ) -KL-LUCB returns an arm from $\mathcal{G}_{M, \mu}^{\alpha, \epsilon}$ with probability at least $1 - \delta$.*

It follows that when the parameter ϵ is chosen small enough, e.g., $\epsilon < G^{-1}(1 - \frac{\alpha}{2}) - G^{-1}(1 - \alpha)$, (α, ϵ) -KL-LUCB is $(\alpha/2, \delta)$ -correct. For example, under the tail assumption 3, ϵ can be chosen of order $c\alpha^{1/\beta}$. However, when nothing is known about the reservoir distribution (e.g. it may not even be smooth) we are not aware of an algorithm to choose ϵ to provide a (α, δ) -correctness guarantee.

4.3 Sample Complexity of (α, ϵ) -KL-LUCB

Recall the partition of \mathcal{W} into subsets \mathcal{S}^i of measure α for $1 \leq i \leq m$, where $\mathcal{S}^i = \{w \in \mathcal{W} : \mu(w) \in [b_i, b_{i-1}]\}$. We define our sample complexity bound in terms of the complexity term:

$$\bar{H}_{\alpha, \epsilon} = \frac{2}{\epsilon^2} + \sum_{i=2}^m \frac{1}{\max(\epsilon^2/2, d^*(b_{i-1}, b_1))}, \quad (12)$$

where $d^*(p, q)$ is the Chernoff information between two reward distributions parameterized by their means p and q . This quantity is closely related to the KL-divergence: it is defined as $d^*(p, q) = d(z^*, p)$ where z^* is the unique solution in z to $d(z, p) = d(z, q)$.

Let the random variable τ be the number of samples used by (α, ϵ) -KL-LUCB. The following upper bound on τ holds.

Theorem 6 *Let α, δ such that $0 < \delta \leq \alpha \leq 1/3$. The (α, ϵ) -KL-LUCB algorithm with exploration rate $\beta(t, \delta)$ defined by (9) and a parameter $\epsilon > 0$ is $(\alpha, \epsilon, \delta)$ -correct and satisfies, with probability at least $1 - 7\delta$,*

$$\tau \leq 12C_0(\gamma)\bar{H}_{\alpha, \epsilon} \log^2 \frac{1}{\delta} + o\left(\log^2 \frac{1}{\delta}\right),$$

with $C_0(\gamma)$ such that $C_0(\gamma) \geq \gamma \log(C_0(\gamma)) + 1 + \frac{\gamma}{e}$.

Theorem 6 only presents the leading term in δ (when δ goes to zero) of the sample complexity upper bound, but an explicit upper bound can be extracted from the proof of Theorem 6, given in Appendix D.

4.4 Comparison and Discussion

Our ϵ -relaxed bounds simplify, for appropriate constants c_1, c_2 and small enough ϵ , to

$$\begin{aligned} \mathbb{E}_{\mu}^M[\tau] &\geq c_1 \left(\frac{1}{d(b_1 - \epsilon, b_0 + \epsilon)} + \sum_{i=3}^{m-1} \frac{1}{d(b_i, b_0 + \epsilon)} \right) \log \frac{1}{\delta}, \\ \tau &\leq c_2 \left(\frac{4}{\epsilon^2} + \sum_{i=2}^{m-1} \frac{1}{d^*(b_i, b_1)} \right) \log^2 \frac{1}{\delta}. \end{aligned}$$

A log factor separates our bounds, and the upper bound complexity term is slightly larger than that in the lower bound. KL-divergence in the lower bound is of comparable scale to Chernoff information in the upper bound: in the Bernoulli case one has $\frac{(\mu^* - x)^2}{2} < d^*(x, \mu^*) <$

$d(x, \mu^*) < \frac{(\mu^* - x)^2}{\mu^*(1 - \mu^*)}$. However, for $i \neq 1$, $d^*(b_i, b_1)$ is slightly smaller than $d(b_i, b_0 + \epsilon)$, while ϵ^2 is smaller than $d(b_1 - \epsilon, b_0 + \epsilon)$. These differences are reduced as α is decreased.

When δ is not too small, the extra $\log \frac{1}{\delta}$ factor is small compared to the constants in our upper bound. It is well-established for finite bandit models that for a wide variety of algorithms the sample complexity scales like $\log \frac{1}{\delta}$. The additional log factor comes from the fact that each phase of our algorithm needs a δ -correctness guarantee. In our first phase, we choose a number of arms to draw from the reservoir without drawing any rewards from those arms. In the second phase, we observe rewards from our arms without drawing any new arms from the reservoir. It is an interesting open question to prove whether in any such two-phase algorithm the $\log^2 \frac{1}{\delta}$ term is avoidable. It is not hard to show that the first phase must draw at least $\frac{c}{\alpha} \log \frac{1}{\delta}$ arms for some constant c in order to obtain a single arm from the top- α fraction with high probability, but the *expected* number of arms in the top- α fraction is already $c \log \frac{1}{\delta}$ in this case. The second phase can be reduced to a problem of finding one of the top m arms, or of finding one arm above the unknown threshold $G^{-1}(1 - \alpha)$, but we are not aware of a lower bound on these problems even for the finite case.

Despite all this, it seems likely that a one-phase algorithm can avoid the quadratic dependence on $\log \frac{1}{\delta}$. Indeed, Jamieson et al. (2016) provides such an algorithm for the special case of reservoirs involving just two expected rewards. They employ a subroutine which returns the target coin with constant probability by drawing a number of coins that does not depend on δ . They wrap this subroutine in a δ -correct algorithm which iteratively considers progressively more challenging reservoirs, terminating when a target coin is identified. We agree with the authors that adapting this approach for general reservoirs is an interesting research direction. However their method relies on the special shape of their reservoir and it is not immediately clear how it might be generalized.

5. Conclusion

In contrast with previous approaches to bandit models, we have limited consideration to changes of distribution which change only the mean mapping μ and not the measure M over arms \mathcal{W} . This allows us to analyze infinite bandit models without a need to integrate over the full reservoir distribution, so we can prove results for reservoirs which are not even smooth. We proved a lower bound on the sample complexity of the problem, and we introduced an algorithm with an upper bound within a log factor of our lower bound.

An interesting future direction is to study improved algorithms, namely one-phase algorithms which alternate between sampling arms to estimate the reservoir and drawing new arms to obtain better arms with higher confidence. These algorithms might be able to have only a $\log \frac{1}{\delta}$ instead of $\log^2 \frac{1}{\delta}$ dependency in the upper bound. In practice, however, the algorithm we present exhibits good empirical performance.

Acknowledgement. We thank Virgil Pavlu for the fruitful discussions and his efforts that made this project better. E. Kaufmann acknowledges the support of the French Agence Nationale de la Recherche (ANR), under grant ANR-16-CE40-0002 (project BADASS).

References

- Ron Appel, Thomas Fuchs, Piotr Dollar, and Pietro Perona. Quickly boosting decision trees – pruning underachieving features early. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, 2013.
- J-Y. Audibert, S. Bubeck, and R. Munos. Best Arm Identification in Multi-armed Bandits. In *Proceedings of the 23rd Conference on Learning Theory*, 2010.
- Donald A. Berry, Robert W. Chen, Alan Zame, David C. Heath, and Larry A. Shepp. Bandit problems with infinitely many arms. *Ann. Statist.*, 25(5):2103–2116, 10 1997. doi: 10.1214/aos/1069362389.
- Thomas Bonald and Alexandre Proutière. Two-target algorithms for infinite-armed bandits with bernoulli rewards. In *Advances in Neural Information Processing Systems (NIPS)*. 2013.
- S. Bubeck, R. Munos, and G. Stoltz. Pure Exploration in Finitely Armed and Continuous Armed Bandits. *Theoretical Computer Science* 412, 1832-1852, 412:1832–1852, 2011a.
- S. Bubeck, R. Munos, G. Stoltz, and C. Szepesvári. X-armed bandits. *Journal of Machine Learning Research*, 12:1587–1627, 2011b.
- A.N Burnetas and M. Katehakis. Optimal adaptive policies for sequential allocation problems. *Advances in Applied Mathematics*, 17(2):122–142, 1996.
- R. Busa-Fekete and B. Kégl. Fast boosting using adversarial bandits. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*, 2010. <http://www.machinelearning.org>.
- Alexandra Carpentier and Michal Valko. Simple regret for infinitely many armed bandits. *CoRR*, abs/1505.04627, 2015.
- Karthekeyan Chandrasekaran and Richard M. Karp. Finding the most biased coin with fewest flips. *CoRR*, abs/1202.3639, 2012.
- Arghya Roy Chaudhuri and Shivaram Kalyanakrishnan. Pac identification of a bandit arm relative to a reward quantile. In *AAAI*, 2017.
- Yahel David and Nahum Shimkin. Infinitely many-armed bandits with unknown value distribution. *European Conference, ECML PKDD*, pages 307–322, 2014.
- Charles Dubout and François Fleuret. Adaptive sampling for large scale boosting. *J. Mach. Learn. Res.*, 15(1):1431–1453, January 2014. ISSN 1532-4435.
- G. Escudero, L. Màrquez, and G. Rigau. Using lazyboosting for word sense disambiguation. In *The Proceedings of the Second International Workshop on Evaluating Word Sense Disambiguation Systems*, 2001.
- E. Even-Dar, S. Mannor, and Y. Mansour. Action Elimination and Stopping Conditions for the Multi-Armed Bandit and Reinforcement Learning Problems. *Journal of Machine Learning Research*, 7: 1079–1105, 2006.
- Yoav Freund and Robert E. Schapire. Experiments with a new boosting algorithm. In *Proceedings of the Thirteenth International Conference on International Conference on Machine Learning, ICML’96*, pages 148–156, San Francisco, CA, USA, 1996. Morgan Kaufmann Publishers Inc. ISBN 1-55860-419-7.

- 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 (NIPS)*. 2012.
- Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *Proceedings of the 29th Conference On Learning Theory*, 2016.
- Jean-Bastien Grill, Michal Valko, and Rémi Munos. Black-box optimization of noisy functions with unknown smoothness. In *Advances on Neural Information Processing Systems (NIPS)*, 2015.
- Kevin Jamieson, Daniel Haas, and Ben Recht. The Power of Adaptivity in Identifying Statistical Alternatives. In *Advances on Neural Information Processing Systems (NIPS)*, 2016.
- Shivaram Kalyanakrishnan, Ambuj Tewari, Peter Auer, and Peter Stone. PAC subset selection in stochastic multi-armed bandits. In *Proceedings of the 29th International Conference on Machine Learning, (ICML)*, 2012.
- Zohar Karnin, Tomer Koren, and Oren Somekh. Almost optimal exploration in multi-armed bandits. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, 2013.
- E. Kaufmann and A. Garivier. Learning the distribution with largest mean: two bandit frameworks. *arXiv:1702.00001*, 2017.
- E. Kaufmann and S. Kalyanakrishnan. Information complexity in bandit subset selection. In *Proceeding of the 26th Conference On Learning Theory.*, 2013.
- E. Kaufmann, O. Cappé, and A. Garivier. On the Complexity of Best Arm Identification in Multi-Armed Bandit Models. *Journal of Machine Learning Research*, 17(1):1–42, 2016.
- R. Kleinberg, A. Slivkins, and E. Upfal. Multi-armed bandit in metric spaces. In *Proceedings of the 40th ACM Symposium on Theory of Computing*, 2008.
- T.L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- S. Magureanu, R. Combes, and A. Proutière. Lipschitz Bandits: Regret lower bounds and optimal algorithms. In *Proceedings on the 27th Conference On Learning Theory*, 2014.
- Shie Mannor, John N. Tsitsiklis, Kristin Bennett, and Nicolò Cesa-bianchi. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research*, 5:2004, 2004.
- William R. 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.
- Yizao Wang, Jean yves Audibert, and Rémi Munos. Algorithms for infinitely many-armed bandits. In *Advances in Neural Information Processing Systems (NIPS)*. 2009.

Appendix A. Changes of distribution in infinite bandit models

A.1 Proof of Lemma 3

We describe in this section the key results that allow us to adapt *changes of distribution* arguments to the infinite bandit setting. Lemma 3 follows easily from Lemma 7 and Lemma 8 that are stated below and proved in the next two sections.

All the regret or sample complexity lower bounds in bandit models rely on change of distributions arguments (see, e.g. Lai and Robbins (1985); Burnetas and Katehakis (1996); Audibert et al. (2010)). A change of distribution relates the probability of an event under a given bandit model to the probability of that event under an alternative bandit model, that is “not too far” from the initial model but under which the performance of the algorithm is supposed to be completely different. Magureanu et al. (2014); Kaufmann et al. (2016) recently found an elegant formulation for such a change of distribution in terms of the expected log-likelihood ratio and we explain below how we generalize these tools to the infinite bandit model.

Given an infinite bandit model (M, μ) , one may consider an alternative bandit model (M, λ) in which the measure M is similar but the mean function is different: $\lambda \neq \mu$. As mentioned in Section 2.1, we consider strategies such that $\mathbb{P}_\mu^M(W_t | H_{T-1})$ is independent from μ . Hence, defining the log-likelihood ratio between (M, μ) and (M, λ) at round t as $L_{\mu, \lambda}(t) = \log(\ell(H_t; \mu, M) / \ell(H_t; \lambda, M))$, where the likelihood is defined in (1), one has

$$L_{\mu, \lambda}(T) = \sum_{t=1}^T \log \frac{f_{\mu(W_t)}(Z_t)}{f_{\lambda(W_t)}(Z_t)}. \quad (13)$$

The following result generalizes Lemma 1 in Kaufmann et al. (2016) to the infinite bandit model. It permits to relate the expected log-likelihood ratio to the probability of any event under the two different models.

Lemma 7 *Let σ be a stopping time and μ and λ be two reward mappings. For any event C in H_σ ,*

$$\mathbb{E}_\mu^M [L_{\mu, \lambda}(\sigma)] \geq \text{kl}(\mathbb{P}_\mu^M(C), \mathbb{P}_\lambda^M(C)),$$

where $\text{kl}(x, y) = x \log(x/y) + (1-x) \log((1-x)/(1-y))$ is the Bernoulli relative entropy.

The next result provides an upper bound on the expected log-likelihood ratio. While in a classic multi-armed bandit model, the log-likelihood can be expressed as a sum that features the expected number of draws of each arm, such a quantity would not be defined in the infinite bandit model. Hence, we need to introduce a partition of \mathcal{W} .

Lemma 8 *Fix $\mathcal{S}_1, \dots, \mathcal{S}_m$ a partition of \mathcal{W} and let $T_i(t) = \sum_{s=1}^t \mathbb{1}_{(w_s \in \mathcal{S}^i)}$ be the number of times an arm in \mathcal{S}^i has been selected. For any stopping time σ ,*

$$\mathbb{E}_\mu^M [L_{\mu, \lambda}(\sigma)] \leq \sum_{i=1}^m \mathbb{E}_\mu^M [T_i(\sigma)] \sup_{w \in \mathcal{S}^i} d(\mu(w), \lambda(w)).$$

A.2 Proof of Lemma 7

The proof for the infinite case follows the argument by Kaufmann et al. (2016) for the finite case. First, the conditional Jensen's inequality is applied, given the convexity of $\exp(-x)$. The expectation derivations hold for our infinite arms case without modification. The only necessary statement for which the finite case proof needs updating, $\mathbb{P}_\lambda^M(C) = \mathbb{E}_\mu^M [\mathbb{1}_{(C)} \exp(-L_{\mu,\lambda}(\sigma))]$, is proven for our infinite case setting as Lemma 9.

Lemma 9 *Let μ and λ be two arm reward mappings, $L_{\mu,\lambda}(T)$ be the log likelihood ratio defined in (13) and C be an (event) subset of histories of length T . Then*

$$\mathbb{P}_\lambda^M(C) = \mathbb{E}_\mu^M [\mathbb{1}_{(C)} \cdot \exp(-L_{\mu,\lambda}(T))]$$

Proof of Lemma 9. Let $\ell(H_T; \mu, M)$ be the likelihood function defined in (1). We introduce furthermore the notation $h_t = (W_1, Z_1, \dots, W_t, Z_t)$, $\mathbf{w}_T = (W_1, \dots, W_T)$ and $\mathbf{z}_T = (Z_1, \dots, Z_T)$. Recall that a strategy is such that conditional density of W_t given H_{t-1} does not depend on the mean reward mapping but only on the reservoir distribution: we denote it by $\mathbb{P}_M(W_t|H_{t-1})$. The proof of Lemma 9 follows from the following inequalities.

$$\begin{aligned} \mathbb{P}_\lambda^M(C) &= \mathbb{E}_\lambda^M [\mathbb{1}_{(C)}] = \int \mathbb{1}_{(C)}(H_T) \ell(H_T; \lambda, M) dH_T \\ &= \int \mathbb{1}_{(C)}(h_T) \prod_{t=1}^T f_{\lambda(W_t)}(Z_t) \mathbb{P}_M(W_t|h_{t-1}) d\mathbf{w}_T d\mathbf{z}_T \\ &= \int \mathbb{1}_{(C)}(h_T) \prod_{t=1}^T \frac{f_{\lambda(W_t)}(Z_t)}{f_{\mu(W_t)}(Z_t)} \prod_{t=1}^T f_{\mu(W_t)}(Z_t) \mathbb{P}_M(W_t|h_{t-1}) d\mathbf{w}_T d\mathbf{z}_T \\ &= \int \mathbb{1}_{(C)}(H_T) \prod_{t=1}^T \frac{f_{\lambda(W_t)}(Z_t)}{f_{\mu(W_t)}(Z_t)} \ell(H_T; \mu, M) dH_T \\ &= \mathbb{E}_\mu^M \left[\mathbb{1}_{(C)} \cdot \prod_{t=1}^T \frac{f_{\lambda(W_t)}(Z_t)}{f_{\mu(W_t)}(Z_t)} \right]. \end{aligned}$$

□

A.3 Proof of Lemma 8

The proof follows from the following inequalities.

$$\begin{aligned} &\mathbb{E}_\mu^M [L_{\mu,\lambda}(\sigma)] \\ &= \mathbb{E}_\mu^M \left[\sum_{t=1}^{\infty} \mathbb{E}_\mu^M \left[\mathbb{1}_{(\sigma \geq t-1)} \log \frac{f_{\mu(W_t)}(Z_t)}{f_{\lambda(W_t)}(Z_t)} \middle| H_{t-1} \right] \right] \\ &= \mathbb{E}_\mu^M \left[\sum_{t=1}^{\infty} \mathbb{1}_{(\sigma \geq t-1)} d(\mu(W_t), \lambda(W_t)) \right] \\ &\leq \mathbb{E}_\mu^M \left[\sum_{i=1}^m \sum_{t=1}^{\sigma} \mathbb{1}_{(W_t \in \mathcal{S}^i)} \sup_{w \in \mathcal{S}^i} d(\mu(w), \lambda(w)) \right] \end{aligned}$$

Appendix B. Proof of Lemma 5

Letting

$$\begin{aligned} A &= (\exists a \leq n : \mu(w^a) > G^{-1}(1 - \alpha)) \\ B &= (\mu_{\hat{a}(\tau)} \geq \max_a \mu_a - \epsilon), \end{aligned}$$

Lemma 5 follows from the fact that $\mathbb{P}(A^c) \leq \frac{\delta}{2}$ and $\mathbb{P}(B^c) \leq \frac{\delta}{2}$, that we now prove.

First, by definition of the reservoir distribution G and the fact that $\mu(W^t)$ are i.i.d. samples from it, one has

$$\begin{aligned} \mathbb{P}(A^c) &= \mathbb{P}\left(\bigcap_{a=1}^n (\mu(w^a) \leq G^{-1}(1 - \alpha))\right) = (1 - \alpha)^n \\ &= \exp\left(-n \log\left(\frac{1}{1 - \alpha}\right)\right) \\ &= \exp\left(-\log\left(\frac{2}{\delta}\right) \frac{1}{\alpha} \log\left(\frac{1}{1 - \alpha}\right)\right) \leq \frac{\delta}{2}, \end{aligned}$$

using that $-\log(1 - x) > x$.

The upper bound on $\mathbb{P}(B^c)$ follows the same lines as that of the correctness of the KL-LUCB algorithm (Kaufmann and Kalyanakrishnan, 2013), however note that we are able to use a smaller exploration rate compared to this work. Abusing notation slightly we let $\mu_a := \mu(w^a)$, and letting $(1) = \operatorname{argmax}_a \mu_a$ we have

$$B^c \subseteq (\exists t \in \mathbb{N} : \mu_{(1)} > U_{(1)}(t)) \bigcup_{a: \mu_a < \mu_{(1)} - \epsilon} (\exists t \in \mathbb{N} : L_a(t) > \mu_a).$$

Let $d^-(x, y) := d(x, y) \mathbf{1}_{(x > y)}$. For each a , $\mathbb{P}(\exists t \in \mathbb{N} : L_a(t) > \mu_a | \mu_a)$ is upper bounded by

$$\begin{aligned} &\mathbb{P}(\exists t \in \mathbb{N} : N_a(t) d^-(\hat{p}_a(t), \mu_a) > \beta(t, \delta)) \\ &\leq \mathbb{P}(\exists t \in \mathbb{N} : N_a(t) d^-(\hat{p}_a(t), \mu_a) > \beta(N_a(t), \delta)) \\ &\leq \mathbb{P}(\exists s \in \mathbb{N} : s d^-(\hat{p}_a(t), \mu_a) > \beta(s, \delta)) \\ &\leq \sum_{s=1}^{\infty} \exp(-\beta(s, \delta)) \leq \frac{\delta}{k_1 n} \sum_{t=1}^{\infty} \frac{1}{t^\gamma} \leq \frac{\delta}{2n}, \end{aligned}$$

where we use a union bound together with Chernoff's inequality and the fact that k_1 is chosen to be larger than $2 \sum_t \frac{1}{t^\gamma}$. Similar reasoning shows that

$$\mathbb{P}(\exists t \in \mathbb{N} : \mu_{(1)} > U_{(1)}(t)) \leq \frac{\delta}{2n},$$

and a union bound yields $\mathbb{P}(B^c) \leq \frac{\delta}{2}$.

Appendix C. Proof of Lemma 11

For every i , let \mathcal{R}_i be the set of arms $a \in \{1, \dots, n\}$ such that $w^a \in \mathcal{S}^i$. Letting $Y_a^i = \mathbb{1}_{(w^a \in \mathcal{S}^i)}$, as $M(\mathcal{S}^i) = \alpha$, $(Y_a^i)_a$ is i.i.d. with a Bernoulli distribution of parameter α and $|\mathcal{R}_i| = \sum_{a=1}^n Y_a^i$.

Using Chernoff's inequality for Bernoulli random variables yields

$$\begin{aligned} \mathbb{P}(|\mathcal{R}_i| > 6 \log(1/\delta)) &= \mathbb{P}(|\mathcal{R}_i| > 3\alpha n) \\ &= \mathbb{P}\left(\frac{1}{n} \sum_{a=1}^n Y_a^i > 3\alpha\right) \leq \exp\{-n \text{kl}(3\alpha, \alpha)\}, \end{aligned}$$

where $\text{kl}(x, y) = x \log(x/y) + (1-x) \log((1-x)/(1-y))$ is the binary relative entropy.

Using Lemma 10 below for $\beta = 2$, if $\alpha \leq 1/3$ it holds that

$$\text{kl}(3\alpha, \alpha) \geq (3 \log 3 - 2)\alpha > \alpha$$

and, using again the definition of n ,

$$\mathbb{P}(|\mathcal{R}_i| > 6 \log(1/\delta)) \leq \exp\{-2 \log(1/\delta)\} = \delta^2.$$

Using a union bound on the m subsets \mathcal{S}^i , and the fact that $m \leq 1/\alpha$ and $\delta \leq \alpha$, one has

$$\mathbb{P}(C^{\mathfrak{L}}) \leq m\delta^2 \leq \frac{\delta^2}{\alpha} \leq \delta,$$

which concludes the proof.

Lemma 10 *Let $\beta > -1$. For all $\alpha \leq \frac{1}{1+\beta}$,*

$$\text{kl}((1+\beta)\alpha, \alpha) \geq ((1+\beta) \log(1+\beta) - \beta)\alpha.$$

This inequality is optimal in the first order in the sense that its two members are equivalent when α goes to zero.

Proof of Lemma 10. By definition

$$\begin{aligned} \text{kl}((1+\beta)\alpha, \alpha) &= (1+\beta)\alpha \log(1+\beta) \\ &\quad + (1 - (1+\beta)\alpha) \log\left(\frac{1 - (1+\beta)\alpha}{1 - \alpha}\right) \\ &= (1+\beta)\alpha \log(1+\beta) \\ &\quad + (1 - (1+\beta)\alpha) \log\left(1 - \frac{\beta\alpha}{1 - \alpha}\right) \end{aligned}$$

Now, using the fact that $1 - (1+\beta)\alpha > 0$ for $\alpha \leq 1/(1+\beta)$ and the following inequality

$$\forall x > -1, \quad \log(1+x) \geq \frac{x}{1+x}$$

one obtains

$$\begin{aligned}
 \text{kl}((1+\beta)\alpha, \alpha) &\geq (1+\beta)\alpha \log(1+\beta) \\
 &\quad + (1 - (1+\beta)\alpha) \frac{-\beta\alpha}{1 - \alpha - \beta\alpha} \\
 &= \alpha((1+\beta)\log(1+\beta) - \beta)
 \end{aligned}$$

□

Appendix D. Proof of Theorem 6

We let $\mu_a := \mu(w^a)$ and denote by $\boldsymbol{\mu} = (\mu_1, \dots, \mu_n)$ the means of the n arms that have been queried, sorted in decreasing order.

In addition to events A and B defined in Section 4.2, we introduce the event C that for every subset $i \leq m$, at most $6 \log \frac{1}{\delta}$ arms belong to \mathcal{S}^i :

$$C = \bigcap_{i=1, \dots, m} \{|\{a \leq n : w^a \in \mathcal{S}^i\}| \leq 6 \log(1/\delta)\} \quad (14)$$

We prove the following in Appendix C.

Lemma 11 *If $\delta \leq \alpha \leq 1/3$, $\mathbb{P}(C) \geq 1 - \delta$.*

For all $t \in \mathbb{N}$ we introduce the event

$$W_t = \bigcap_{1 \leq a \leq n} (L_a(t) \leq \mu_a \leq U_a(t)) \quad (15)$$

and define $W = \bigcap_{t \in \mathbb{N}^*} W_t$. By the same argument as the one used in the proof of Lemma 5 (see Appendix B), one can show that $\mathbb{P}(W) \geq 1 - 2\delta$.

Fix $c \in [\mu_2, \mu_1]$. Our analysis relies on the following crucial statement.

Proposition 12 (Kaufmann and Kalyanakrishnan (2013)) *Let $\tilde{\beta}_a(t) := \sqrt{\frac{\beta(t, \delta)}{2N_a(t)}}$. If W_t holds and $(U_{\hat{b}(t)} - L_{\hat{a}(t)}) > \epsilon$ then there exists $a \in \{\hat{a}(t), \hat{b}(t)\}$ such that*

$$c \in \mathcal{I}_a(t) \text{ and } \tilde{\beta}_a(t) > \epsilon/2. \quad (16)$$

Fixing some integer $T \in \mathbb{N}$, we now upper bound τ on the event $\mathcal{E} := A \cap B \cap C \cap W$.

$$\begin{aligned}
 \min(\tau, T) &\leq \sum_{t=1}^T \mathbf{1}_{(\tau > t)} = n + 2 \sum_{\substack{t \in n+2\mathbb{N} \\ t \leq T}} \mathbf{1}_{(\tau > t)} \\
 &\leq n + 2 \sum_{\substack{t \in n+2\mathbb{N} \\ t \leq T}} \mathbf{1}_{(U_{\hat{b}(t)} - L_{\hat{a}(t)}) > \epsilon} \\
 &\leq n + 2 \sum_{\substack{t \in n+2\mathbb{N} \\ t \leq T}} \mathbf{1}_{(\exists a \in \{\hat{a}(t), \hat{b}(t)\} : c \in \mathcal{I}_a(t), \tilde{\beta}_a(t) > \epsilon/2)},
 \end{aligned}$$

using Proposition 12 and the fact that W holds. To ease the notation, we let $\mathcal{S}_t = \{\hat{a}(t), \hat{b}(t)\}$ be the set of drawn arms at round t . Letting $\mathcal{A}_\epsilon = \{a : d^*(\mu_a, c) < \epsilon^2/2\}$ and noting that $\tilde{\beta}_a(t) > \epsilon/2 \iff N_a(t) < \beta(t, \delta)/(\epsilon^2/2)$,

$$\begin{aligned} \min(\tau, T) &\leq n + 2 \sum_{a \in \mathcal{A}_\epsilon} \sum_{\substack{t \in n+2\mathbb{N} \\ t \leq T}} \mathbb{1}_{(a \in \mathcal{S}_t)} \mathbb{1}_{(N_a(t) < \beta(T, \delta)/(\epsilon^2/2))} \\ &\quad + 2 \sum_{a \in \mathcal{A}_\epsilon^c} \sum_{\substack{t \in n+2\mathbb{N} \\ t \leq T}} \mathbb{1}_{(a \in \mathcal{S}_t)} \mathbb{1}_{(N_a(t) d(\hat{p}_a(t), c) \leq \beta(T, \delta))} \end{aligned}$$

Now defining the event

$$G_T^a = \bigcup_{\substack{t \in n+2\mathbb{N} \\ t \leq T}} \left\{ N_a(t) > \frac{\beta(T, \delta)}{d^*(\mu_a, c)}, N_a(t) d(\hat{p}_a(t), c) \leq \beta(T, \delta) \right\}$$

From Lemma 1 in Kaufmann and Kalyanakrishnan (2013),

$$\mathbb{P}(G_T^a | \boldsymbol{\mu}) \leq \frac{1}{d^*(\mu_a, c)} \exp(-\beta(T, \delta)). \quad (17)$$

Introducing $\mathcal{F}_T = \cap_{a \in \mathcal{A}_\epsilon^c} (G_T^a)^c$, one can further upper bound τ on $\mathcal{E} \cap \mathcal{F}_T$ as

$$\begin{aligned} \min(\tau, T) &\leq n + 2 \sum_{a \in \mathcal{A}_\epsilon} \sum_{\substack{t \in n+2\mathbb{N} \\ t \leq T}} \mathbb{1}_{(a \in \mathcal{S}_t)} \mathbb{1}_{(N_a(t) < \beta(T, \delta)/(\epsilon^2/2))} \\ &\quad + 2 \sum_{a \in \mathcal{A}_\epsilon^c} \sum_{\substack{t \in n+2\mathbb{N} \\ t \leq T}} \mathbb{1}_{(a \in \mathcal{S}_t)} \mathbb{1}_{(N_a(t) < \beta(T, \delta)/d^*(\mu_a, c))} \\ &\leq n + 2 \underbrace{\left[\sum_{a \in \mathcal{A}_\epsilon} \frac{2}{\epsilon^2} + \sum_{a \in \mathcal{A}_\epsilon^c} \frac{1}{d^*(\mu_a, c)} \right]}_{=: H(\boldsymbol{\mu}, c, \epsilon)} \beta(T, \delta). \end{aligned}$$

We now provide a deterministic upper bound on $H(\boldsymbol{\mu}, c, \epsilon)$, by summing over the arms in the different subsets \mathcal{S}^i . Let q be the number of subsets containing arms within \mathcal{A}_ϵ . We choose c such that $c > G^{-1}(1 - \alpha) = b_1$ (such a choice is possible as $\mu_1 > G^{-1}(1 - \alpha)$ as event A holds) and note that $\mathcal{A}_\epsilon \subseteq \cup_{i \leq q} \mathcal{S}^i$. One has

$$\begin{aligned} H(\boldsymbol{\mu}, c, \epsilon) &\leq \sum_{i=1}^q \sum_{a: w_a \in \mathcal{S}^i} \frac{2}{\epsilon^2} + \sum_{i=q+1}^m \sum_{a: w_a \in \mathcal{S}^i} \frac{1}{d^*(\mu_a, G^{-1}(1 - \alpha))} \\ &\leq 6\overline{H}_{\alpha, \epsilon} \log(1/\delta), \end{aligned}$$

using that event C holds and each \mathcal{S}^i contains at least $6 \log(1/\delta)$ arms. Hence on $\mathcal{E} \cap \mathcal{F}_T$,

$$\min(\tau, T) \leq n + 12\overline{H}_{\alpha, \epsilon} \log(1/\delta) \beta(T, \delta).$$

Applying this to $T = T^*$ where

$$T^* := \inf\{T \in \mathbb{N} : n + 12\overline{H}_{\alpha,\epsilon} \log(1/\delta)\beta(T, \delta) \leq T\},$$

one obtains $\min(\tau, T) \leq T$, hence $\tau \leq T$. We proved that

$$\mathbb{P}(\tau \leq T^*) \leq 1 - 4\delta - \mathbb{P}(\mathcal{F}_{T^*}^{\mathcal{C}}).$$

An upper bound on T^* can be extracted from Appendix E of Kaufmann and Kalyanakrishnan (2013):

$$T^* \leq 12C_0(\gamma)\overline{H}_{\alpha,\epsilon} \log \frac{1}{\delta} \log \left(\frac{12k_1n(\log(1/\delta)^\gamma)\overline{H}_{\alpha,\epsilon}^\gamma}{\delta} \right) + n.$$

An upper bound on $\mathbb{P}(\mathcal{F}_{T^*}^{\mathcal{C}})$ concludes the proof:

$$\begin{aligned} \mathbb{P}(\mathcal{F}_{T^*}^{\mathcal{C}}) &\leq 2\delta + \mathbb{P}(\mathcal{F}_{T^*}^{\mathcal{C}} \cap A \cap C) \\ &= 2\delta + \mathbb{E} \left[\mathbb{P}(\mathcal{F}_{T^*}^{\mathcal{C}} | \boldsymbol{\mu}) \mathbf{1}_{(A \cap C)} \right] \end{aligned}$$

Now, on $A \cap C$ (which is used for the second inequality), using (17),

$$\begin{aligned} \mathbb{P}(\mathcal{F}_{T^*}^{\mathcal{C}} | \boldsymbol{\mu}) &\leq \sum_{a \in \mathcal{A}_\epsilon^{\mathcal{C}}} \frac{1}{d^*(\mu_a, c)} \exp(-\beta(T^*, \delta)) \\ &\leq 6\overline{H}_{\alpha,\epsilon} \log(1/\delta) \exp(-\beta(T^*, \delta)) \\ &\leq \left(\frac{6\overline{H}_{\alpha,\epsilon} \log(1/\delta)}{k_1n(T^*)^\gamma} \right) \delta \leq \delta, \end{aligned}$$

where the last inequality follows from the definition of T^* . Thus $\mathbb{P}(\mathcal{F}_{T^*}^{\mathcal{C}}) \leq 3\delta$.

Appendix E. Empirical Results

We exhibit (α, ϵ) -KL-LUCB for infinite models of Bernoulli arms with various parameter values. In order to meet our assumption that $\mu^* < 1$, we truncate our *Beta* distributions to have support on the interval $]0, 0.95]$. We report the fraction of runs in which it fails to find an arm within ϵ of the top- α fraction, the mean simple regret observed, and the mean budget used. We have also tried various minor modifications to this algorithm which appear to reduce the sample complexity without much affecting its success rate or simple regret. These include drawing a sample only from the least-sampled of the two arms chosen at each round and updating confidence intervals only for those arms whose upper bounds overlap with the top arm. However, we provide no performance guarantees for the algorithm with these modifications. Note that our results are for the algorithm as stated in the paper, and not with any of these modifications.

Table 1: Performance of (α, ϵ) -KL-LUCB. Mean of 100 runs. “Effective α ” gives the measure of arms meeting the (α, ϵ) criteria. “Errors” indicates the fraction of runs where the α objective was not achieved; it should be below δ .

The budget T is impacted roughly linearly by $1/\alpha$ and quadratically by $1/\epsilon$; the same regret is achieved with different budgets based on parameter selection.

RESERVOIR	α	ϵ	δ	EFFECTIVE α	ERRORS	SIMPLE REGRET	T
BETA(1,1)	0.025	0.024	0.05	0.049	0.00	0.008	51K
BETA(1,1)	0.025	0.024	0.10	0.049	0.01	0.011	46K
BETA(1,1)	0.050	0.010	0.05	0.060	0.02	0.015	113K
BETA(1,1)	0.050	0.010	0.10	0.060	0.06	0.020	90K
BETA(1,1)	0.050	0.048	0.05	0.098	0.00	0.014	12K
BETA(1,1)	0.050	0.048	0.10	0.098	0.00	0.017	10K
BETA(1,1)	0.050	0.050	0.05	0.100	0.01	0.014	11K
BETA(1,1)	0.050	0.050	0.10	0.100	0.00	0.022	10K
BETA(1,1)	0.100	0.010	0.05	0.110	0.02	0.030	71K
BETA(1,1)	0.100	0.010	0.10	0.110	0.06	0.044	69K
BETA(1,1)	0.100	0.050	0.05	0.150	0.00	0.037	10K
BETA(1,1)	0.100	0.050	0.10	0.150	0.00	0.033	7K
BETA(1,2)	0.025	0.063	0.05	0.221	0.00	0.044	10K
BETA(1,2)	0.025	0.063	0.10	0.221	0.01	0.061	10K
BETA(1,2)	0.050	0.010	0.05	0.234	0.03	0.075	79K
BETA(1,2)	0.050	0.010	0.10	0.234	0.10	0.093	65K
BETA(1,2)	0.050	0.050	0.05	0.274	0.01	0.069	10K
BETA(1,2)	0.050	0.050	0.10	0.274	0.03	0.094	11K
BETA(1,2)	0.050	0.091	0.05	0.314	0.01	0.077	5K
BETA(1,2)	0.050	0.091	0.10	0.314	0.00	0.091	5K
BETA(1,2)	0.100	0.010	0.05	0.326	0.04	0.123	63K
BETA(1,2)	0.100	0.010	0.10	0.326	0.06	0.136	60K
BETA(1,2)	0.100	0.050	0.05	0.366	0.00	0.113	10K
BETA(1,2)	0.100	0.050	0.10	0.366	0.05	0.139	10K
BETA(1,3)	0.025	0.076	0.05	0.368	0.00	0.132	12K
BETA(1,3)	0.025	0.076	0.10	0.368	0.00	0.142	10K
BETA(1,3)	0.050	0.010	0.05	0.378	0.01	0.176	87K
BETA(1,3)	0.050	0.010	0.10	0.378	0.10	0.195	82K
BETA(1,3)	0.050	0.050	0.05	0.418	0.00	0.166	13K
BETA(1,3)	0.050	0.050	0.10	0.418	0.06	0.216	14K
BETA(1,3)	0.050	0.096	0.05	0.464	0.00	0.183	7K
BETA(1,3)	0.050	0.096	0.10	0.464	0.00	0.196	6K
BETA(1,3)	0.100	0.010	0.05	0.474	0.06	0.233	69K
BETA(1,3)	0.100	0.010	0.10	0.474	0.06	0.251	53K
BETA(1,3)	0.100	0.050	0.05	0.514	0.01	0.220	10K
BETA(1,3)	0.100	0.050	0.10	0.514	0.03	0.241	10K