



HAL
open science

Learning the distribution with largest mean: two bandit frameworks

Emilie Kaufmann, Aurélien Garivier

► **To cite this version:**

Emilie Kaufmann, Aurélien Garivier. Learning the distribution with largest mean: two bandit frameworks. *ESAIM: Proceedings and Surveys*, 2017, 60, pp.114 - 131. 10.1051/proc/201760114 . hal-01449822v3

HAL Id: hal-01449822

<https://hal.science/hal-01449822v3>

Submitted on 6 Nov 2017

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.

LEARNING THE DISTRIBUTION WITH LARGEST MEAN: TWO BANDIT FRAMEWORKS *

EMILIE KAUFMANN¹ AND AURÉLIEN GARIVIER²

Abstract. Over the past few years, the multi-armed bandit model has become increasingly popular in the machine learning community, partly because of applications including online content optimization. This paper reviews two different sequential learning tasks that have been considered in the bandit literature ; they can be formulated as (sequentially) learning which distribution has the highest mean among a set of distributions, with some constraints on the learning process. For both of them (*regret minimization* and *best arm identification*) we present recent, asymptotically optimal algorithms. We compare the behaviors of the sampling rule of each algorithm as well as the complexity terms associated to each problem.

Résumé. Le modèle stochastique dit de bandit à plusieurs bras soulève ces dernières années un grand intérêt dans la communauté de l'apprentissage automatique, du fait notamment de ses applications à l'optimisation de contenu sur le web. Cet article présente deux problèmes d'apprentissage séquentiel dans le cadre d'un modèle de bandit qui peuvent être formulés comme la découverte de la distribution ayant la moyenne la plus élevée dans un ensemble de distributions, avec certaines contraintes sur le processus d'apprentissage. Pour ces deux objectifs (*minimisation du regret* d'une part et *identification du meilleur bras* d'autre part), nous présentons des algorithmes optimaux, en un sens asymptotique. Nous comparons les stratégies d'échantillonnage employées par ces deux types d'algorithmes ainsi que les quantités caractérisant la complexité de chacun des problèmes.

INTRODUCTION

Bandit models can be traced back to the 1930s and the work of [Thompson, 1933] in the context of medical trials. It addresses the idealized situation where, for a given symptom, a doctor has several treatments at her disposal, but has no prior knowledge about their efficacies. These efficacies need to be *learnt* by allocating treatments to patients and observing the result. As the doctor aims at healing as many patients as possible, she would like to select the best treatment as often as possible, even though it is unknown to her at the beginning. After each patient, the doctor takes the outcome of the treatment into account in order to decide which treatment to assign to the next patient: the learning process is *sequential*.

This archetypal situation is mathematically captured by the *multi-armed bandit model*. It involves an *agent* (the doctor) interacting with a set of K probability distribution ν^1, \dots, ν^K called *arms* (the treatments), which

* 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).

¹ CNRS & Univ. Lille, UMR 9189 (CRIStAL), Inria Lille Nord-Europe (Sequel), emilie.kaufmann@univ-lille1.fr

² Institut de Mathématiques de Toulouse, UMR 5219, Université de Toulouse, CNRS, aurelien.garivier@math.univ-toulouse.fr

she may sequentially sample. The mean of arm a (which is unknown to the agent) is denoted by μ_a . At round t , the agent selects an arm $A_t \in \{1, \dots, K\}$ and subsequently observes a sample $X_t \sim \nu^{A_t}$ from the associated distribution. The arm A_t is selected according to a *sampling strategy* denoted by $\pi = (\pi_t)_{t \geq 1}$, where π_t maps the history of past arm choices and observations $A_1, X_1, \dots, A_{t-1}, X_{t-1}$ to an arm. In a simplistic model for the clinical trial example, each arm is a Bernoulli distribution that indicates the success or failure of the treatment. After sampling an arm (giving a treatment) at time t , the doctor observes whether the patient was healed ($X_t = 1$) or not ($X_t = 0$). In this example as in many others, the samples gathered can be considered as *rewards*, and a natural goal for the agent is to adjust her sampling strategy so as to maximize the expected sum $\mathbb{E}[\sum_{t=1}^T X_t]$ of the rewards gathered up to some given *horizon* T . This is equivalent to minimizing the *regret*

$$R^\pi(T) = \mu^* T - \mathbb{E} \left[\sum_{t=1}^T X_t \right],$$

which is defined as the gap between the expected efficiency of the strategy π and the expected cumulated reward of an *oracle strategy* always playing the best arm $a^* = \operatorname{argmax}_a \mu_a$ that has mean $\mu^* = \max_a \mu_a$.

A sampling strategy minimizing the regret should not only learn which arm has the highest mean: it should also not incur too big losses during this learning phase. In other words, it has to achieve a good trade-off between *exploration* (experimenting all the arms in order to estimate their means) and *exploitation* (focusing on the arm that appears best so far). Despite its simplicity, the multi-armed bandit model already captures the fundamental dilemma inherent to *reinforcement learning* [Sutton and Barto, 1998], where the goal is to learn how to act optimally in a random environment based on numeric feedback. The fundamental model of reinforcement learning is the Markov Decision Process [Puterman, 1994], which involves the additional notion of system state; a bandit model is simply a Markov Decision Process with a single state.

Sometimes, rewards actually correspond to profits for the agent. In fact, the imaginatively named *multi-armed bandits* refer to casino slot machines: a player sequentially selects one of them (also called a *one-armed bandit*), draws its arm, and possibly collects her wins. While the model was initially motivated by clinical trials, modern applications involve neither bandits, nor casinos, but for example the design of recommender systems [Chu et al., 2011], or more generally content optimization. Indeed, a bandit algorithm may be used by a company for dynamically selecting which version of its website to display to each user, in order to maximize the number of *conversions* (purchase or subscription for example). In the case of two competing options, this problem is known as *A/B testing*. It motivates the consideration of a different optimization problem in a bandit model: rather than continuously changing its website, the company may prefer to experiment during a *testing phase* only, which is aimed at identifying the best version, and then to use that one consistently for a much bigger audience.

In such a testing phase, the objective is different: one aims at learning which arm has highest mean *without constraint* on the cumulative reward. In other words, the company agrees to lose some profit during the testing phase, as long as the length of this phase is as short as possible. In this framework, called *best arm identification*, the sampling rule is designed so as to identify the arm with highest mean as fast and as confidently as possible. Two alternative frameworks are considered in the literature. In the *fixed-budget setting* [Audibert et al., 2010], the length of the trial phase is given and the goal is to minimize the probability of misidentifying the best arm. In the *fixed-confidence setting* [Even-Dar et al., 2006], a risk parameter δ is given and the procedure is allowed to choose when the testing phase stops. It must guarantee that the misidentification probability is smaller than δ while minimizing the *sample complexity*, that is the expected number of samples required before electing the arm. Although the study of best arm identification problems is relatively recent in the bandit literature, similar questions were already addressed in the 1950s under the name *ranking and identification* problems [Bechhofer, 1954, Bechhofer et al., 1968], and they are also related to the sequential adaptive hypothesis testing framework introduced by [Chernoff, 1959].

In this paper, we review a few algorithms for both regret minimization and best arm identification in the fixed-confidence setting. The algorithms and results are presented for simple classes of parametric bandit models, and

we explain along the way how some of them can be extended to more general models. In each case we introduce an asymptotic notion of optimality and present algorithms that are asymptotically optimal. Our optimality notion is *instance-dependent*, in the sense that we characterize the minimal regret or minimal sample complexity achievable on each specific bandit instance. The paper is structured as follows: we introduce in Section 1 the parametric bandit models considered in the paper, and present some useful probabilistic tools for the analysis of bandit algorithms. We discuss the regret minimization problem in Section 2 and the best arm identification problem in Section 3. We comment in Section 4 on the different behaviors of the algorithms aimed at these distinct objectives, and on the different information-theoretic quantities characterizing their complexities.

1. PARAMETRIC BANDIT MODELS AND USEFUL TOOLS

1.1. Some Assumptions on the Arms Distributions

Unless specified otherwise, we assume in the rest of the paper that all the arms belong to a class of distributions parameterized by their means, $\mathcal{D}_I = \{\nu_\mu : \mu \in I\}$, where I is an interval of \mathbb{R} . We assume that for all $\mu \in I$, ν_μ has a density denoted by f_μ with respect to some fixed reference measure, and that $\mathbb{E}_{X \sim \nu_\mu}[X] = \mu$. For all $(\mu, \mu') \in I^2$ we introduce

$$d(\mu, \mu') := \text{KL}(\nu_\mu, \nu_{\mu'}) = \mathbb{E}_{X \sim \nu_\mu} \left[\log \frac{f_\mu(X)}{f_{\mu'}(X)} \right],$$

the Kullback-Leibler divergence between the distribution of mean μ and that of mean μ' . We shall in particular consider examples in which \mathcal{D}_I forms a one-parameter exponential family (e.g. Bernoulli distributions, Gaussian distributions with known variance, exponential distributions), for which there is a closed form formula for the divergence function d (see, e.g. [Cappé et al., 2013]).

Under this assumption, a bandit model is fully described by a vector $\boldsymbol{\mu} = (\mu_1, \dots, \mu_K)$ in I^K such that $\nu^a = \nu_{\mu_a}$ for all $a \in \{1, \dots, K\}$. We denote by $\mathbb{P}_\boldsymbol{\mu}$ and $\mathbb{E}_\boldsymbol{\mu}$ the probability and expectation under the bandit model $\boldsymbol{\mu}$. Under $\mathbb{P}_\boldsymbol{\mu}$, the sequence $(Y_{a,s})_{s \in \mathbb{N}^*}$ of successive observations from arm a is i.i.d. with law ν_{μ_a} , and the families $(Y_{a,s})_a$ are independent. Given a strategy π , we let $N_a^\pi(t) = \sum_{s=1}^t \mathbb{1}_{(A_s=a)}$ be the number of draws of arm a up to and including round $t \geq 1$. Hence, upon selection of the arm A_t , the observation made at round t is $X_t = Y_{A_t, N_{A_t}^\pi(t)}$. When the strategy π is clear from the context, we may remove the superscript π and write simply $N_a(t)$. We define $\hat{\mu}_{a,s} = \frac{1}{s} \sum_{i=1}^s Y_{a,i}$ as the empirical mean of the first s observations from arm a , and $\hat{\mu}_a(t) = \hat{\mu}_{a, N_a(t)}$ as the empirical mean of arm a at round t of the bandit algorithm.

For the two frameworks that we consider, regret minimization and best arm identification, we adopt the same approach. First, we propose a lower bound on the target quantity (regret or sample complexity). Then, we propose strategies whose regret or sample complexity asymptotically matches the lower bound. Two central tools to derive lower bounds and algorithms are *changes of distributions* and *confidence intervals*.

1.2. Changes Of Distribution

Problem-dependent lower bounds in the bandit literature all rely in the end on change of distribution arguments (see e.g. [Lai and Robbins, 1985, Burnetas and Katehakis, 1996, Mannor and Tsitsiklis, 2004, Audibert et al., 2010]). In order to control the probability of some event under the bandit model $\boldsymbol{\mu}$, the idea is to consider an alternative bandit model $\boldsymbol{\lambda}$ under which some assumptions on the strategy make it is easier to control the probability of this event. This alternative model $\boldsymbol{\lambda}$ should be *close* enough to $\boldsymbol{\mu}$, in the sense that the *transportation cost* should not be too high. This transportation cost is related to the log-likelihood ratio of the observations up to time t , that we denote by

$$L_t^\pi(\boldsymbol{\mu}, \boldsymbol{\lambda}) = \sum_{a=1}^K \sum_{s=1}^{N_a^\pi(t)} \log \frac{f_{\mu_a}(Y_{a,s})}{f_{\lambda_a}(Y_{a,s})}.$$

Letting $\mathcal{F}_t = \sigma(X_1, \dots, X_t)$ be the σ -field generated by the observations up to time t , it is indeed well known that for all $\mathcal{E} \in \mathcal{F}_t$, $\mathbb{P}_\mu(\mathcal{E}) = \mathbb{E}_\lambda[\mathbb{1}_\mathcal{E} \exp(L_t^\pi(\mu, \lambda))]$.

The most simple way of writing changes of distribution (see [Kaufmann et al., 2014, Combes and Proutière, 2014] and [Garivier et al., 2016b]) directly relates the *expected* log-likelihood ratio of the observations under two bandit models to the probability of any event under the two models. If S is a stopping time, one can show that for any two bandit models μ, λ and for any event in \mathcal{F}_S ,

$$\mathbb{E}_\mu[L_S^\pi(\mu, \lambda)] \geq \text{kl}(\mathbb{P}_\mu(\mathcal{E}), \mathbb{P}_\lambda(\mathcal{E})),$$

where $\text{kl}(x, y) = x \log(x/y) + (1-x) \log((1-x)/(1-y))$ is the binary relative entropy, i.e. the Kullback-Leibler divergence between two Bernoulli distributions of means x and y . Using Wald's lemma, one can show that in the particular case of bandit models, the expected log-likelihood ratio can be expressed in terms of the expected number of draws of each arms, which yields the following result.

Lemma 1. *Let S be a stopping time. For any event $\mathcal{E} \in \mathcal{F}_S$,*

$$\sum_{a=1}^K \mathbb{E}_\mu[N_a^\pi(S)] d(\mu_a, \lambda_a) \geq \text{kl}(\mathbb{P}_\mu(\mathcal{E}), \mathbb{P}_\lambda(\mathcal{E})).$$

Two different proofs of this result can be found in [Kaufmann et al., 2016] and [Garivier et al., 2016b], in which a slightly more general result is derived based on the entropy contraction principle. As we will see in the next sections, this lemma is particularly powerful to prove lower bounds on the regret or the sample complexity, as both quantities are closely related to the expected number of draws of each arm.

1.3. Confidence Intervals

In both the regret minimization and best arm identification frameworks, the sampling rule has to decide which arm to sample from at a current round, based on the observations gathered at previous rounds. This decision may be based on the set of statistically plausible values for the mean of each arm a , that is materialized by a *confidence interval* on μ_a . Note that in this sequential learning framework, this interval has to be built based on a *random number of observations*.

The line of research leading to the UCB1 algorithm [Auer et al., 2002] worked under the assumption that each arm is a bounded distribution supported in $[0, 1]$. Bounded distributions are particular examples of sub-Gaussian distributions. A random variable X is said to be σ^2 -sub-Gaussian if $\mathbb{E}[e^{\lambda(X - \mathbb{E}[X])}] \leq \exp(\lambda^2 \sigma^2 / 2)$ holds for all $\lambda \in \mathbb{R}$. Hoeffding's lemma states that distributions with a support bounded in $[a, b]$ are $(b-a)^2/4$ -sub-Gaussian. If arm a is σ^2 -sub-Gaussian, Hoeffding's inequality together with a union bound to handle the random number of observations permits to show that

$$\mathbb{P}\left(\hat{\mu}_a(t) + \sqrt{\frac{2\sigma^2\gamma}{N_a(t)}} < \mu_a\right) \leq te^{-\gamma}. \quad (1)$$

Hence one can build an upper-confidence bound on μ_a with probability of coverage $1 - \delta$ by setting $\gamma = \log(t/\delta)$.

There are two levels of improvement here. First, under more specific assumption on the arms (for example if the arms belong to some exponential family of distributions), Chernoff's inequality has an explicit form that can be used directly in place of Hoeffding's inequality. It states that $\mathbb{P}(\hat{\mu}_{a,s} > x) \leq \exp(-sd(x, \mu_a))$, where $d(x, y)$ is the KL-divergence function defined in Section 1.1. Then, to handle the random number of observations, a peeling argument can be used rather than a union bound. This argument, initially developed in the context of Markov order estimation (see [Garivier and Leonardi, 2011]), was used in [Garivier and Moulines, 2011, Bubeck, 2010] under sub-gaussian assumption. Combining these two ideas [Garivier and Cappé, 2011] show that, letting

$$u_a(t) = \sup\{q : N_a(t) d(\hat{\mu}_a(t), q) \leq \gamma\}$$

one has

$$\mathbb{P}(u_a(t) < \mu_a) \leq e^{\lceil \gamma \log(t) \rceil} e^{-\gamma}. \quad (2)$$

The improvement can be measured by specifying this result to Bernoulli distributions, for which the two bounds (1) and (2) hold. By Pinsker's inequality $\text{kl}(x, y) > 2(x-y)^2$, it holds that $u_a(t) \leq \hat{\mu}_a(t) + \sqrt{2\sigma^2\gamma/N_a(t)}$. Hence for γ and t such that $e^{\lceil \gamma \log(t) \rceil} \leq t$, $u_a(t)$ is a smaller upper-confidence bound on μ_a with the same coverage guarantees. As we will see in the next sections, such refined confidence intervals have yield huge improvements in the bandit literature, and lead to simple UCB-type algorithms that are asymptotically optimal for regret minimization.

2. OPTIMAL STRATEGIES FOR REGRET MINIMIZATION

After the initial work of [Thompson, 1933], bandit models were studied again in the 1950s, with for example the paper of [Robbins, 1952], in which the notion of regret is introduced. Interestingly, a large part of the early work on bandit models takes a slightly different *Bayesian* perspective: the goal is also to maximize the expected sum of rewards, but the expectation is also computed over a prior distribution for the arms (see [Berry and Fristedt, 1985] for a survey). It turns out that this Bayesian multi-armed bandit problem can be solved exactly using dynamic programming [Bellman, 1956], but the exact solution is in most cases intractable. Practical solutions may be found when one aims at maximizing the sum of *discounted* rewards over an infinite horizon: the seminal paper of [Gittins, 1979] shows that the Bayesian optimal policy has a simple form where, at each round an index is computed for each arm and the arm with highest index is selected.

Gittins's work motivated the focus on *index policies*, where an index is computed for each arm as a selection procedure. Such index policies have also emerged in the "frequentist" literature on multi-armed bandits. Some asymptotic expansions of the index put forward by Gittins were proposed. They have led to new policies that could be studied directly, forgetting about their Bayesian roots. This line of research includes in particular the seminal work of [Lai and Robbins, 1985].

2.1. A Lower Bound on the Regret

In 1985, Lai and Robbins characterized the optimal regret rate in one-parameter bandit models, by providing an asymptotic lower bound on the regret and a first index policy with a matching regret [Lai and Robbins, 1985]. In order to understand this lower bound, one can first observe that the regret can be expressed in terms of the number of draws of each sub-optimal arm. Indeed, a simple conditioning shows that for any strategy π ,

$$R_{\boldsymbol{\mu}}^{\pi}(T) = \mathbb{E} \left[\sum_{t=1}^T (\mu^* - \mu_{A_t}) \right] = \sum_{a: \mu_a < \mu^*} (\mu^* - \mu_a) \mathbb{E}_{\boldsymbol{\mu}} [N_a^{\pi}(T)], \quad (3)$$

where we recall that $N_a^{\pi}(t) = \sum_{s=1}^t \mathbb{1}_{(A_s=t)}$ is the number of times arm a has been selected up to time t . A strategy is said to be *uniformly efficient* if its regret is small on every bandit model in our class, that is if for all $\boldsymbol{\mu} \in I^K$ and for every $\alpha \in]0, 1]$, $R_{\boldsymbol{\mu}}^{\pi}(T) = o(T^{\alpha})$.

Theorem 2. [Lai and Robbins, 1985] *Any uniformly efficient strategy π satisfies, for all $\boldsymbol{\mu} \in I^K$,*

$$\forall a : \mu_a < \mu^*, \quad \liminf_{T \rightarrow \infty} \frac{\mathbb{E}_{\boldsymbol{\mu}} [N_a^{\pi}(T)]}{\log(T)} \geq \frac{1}{d(\mu_a, \mu^*)}.$$

By Equation (3), this result directly provides a logarithmic regret lower bound on the regret:

$$\liminf_{T \rightarrow \infty} \frac{R_{\boldsymbol{\mu}}^{\pi}(T)}{\log(T)} \geq \sum_{a: \mu_a < \mu^*} \frac{\mu^* - \mu_a}{d(\mu_a, \mu^*)} \stackrel{\text{def}}{=} C(\boldsymbol{\mu}). \quad (4)$$

This lower bound motivates the definition of an *asymptotically optimal algorithm* (on a set of parametric bandit models \mathcal{D}_I) as an algorithm for which for all $\boldsymbol{\mu} \in I^K$, the regret is asymptotically upper bounded by $C(\boldsymbol{\mu}) \log(T)$. This defines an *instance-dependent* notion of optimality, as we want an algorithm that attains the best regret rate for every bandit instance $\boldsymbol{\mu}$. However, for some instances $\boldsymbol{\mu}$ such that some arms are very close to the optimal arm, the constant $C(\boldsymbol{\mu})$ may be really large and the $C(\boldsymbol{\mu}) \log(T)$ bound is not very interesting in finite-time. For such instances, one may prefer having regret upper bounds that scale in \sqrt{KT} and are independent of $\boldsymbol{\mu}$, matching the *minimax* regret lower bound obtained by [Cesa-Bianchi and Lugosi, 2006, Bubeck and Cesa-Bianchi, 2012] for Bernoulli bandits:

$$\inf_{\pi} \sup_{\boldsymbol{\mu} \in I} R_{\boldsymbol{\mu}}^{\pi}(T) \geq \frac{1}{20} \sqrt{KT}.$$

Logarithmic instance-dependent regret lower bound have also been obtained under more general assumptions for the arms distributions [Burnetas and Katehakis, 1996], and even in some examples of structured bandit models, in which they take a less explicit form [Graves and Lai, 1997, Magureanu et al., 2014]. All these lower bounds rely on a change of distribution argument, and we now explain how to easily obtain the lower bound of Theorem 2 by using the tool described in Section 1.2, Lemma 1.

Fixing a suboptimal arm a in the bandit model $\boldsymbol{\mu}$, we define an alternative bandit model $\boldsymbol{\lambda}$ such that $\lambda_i = \mu_i$ for all $i \neq a$ and $\lambda_a = \mu^* + \epsilon$. In $\boldsymbol{\lambda}$, arm a is now the optimal arm, hence a uniformly efficient algorithm will draw this arm very often. As arm a is the only arm that has been modified in $\boldsymbol{\lambda}$, the statement in Lemma 1 takes the simple form:

$$\mathbb{E}_{\boldsymbol{\mu}}[N_a^{\pi}(T)]d(\mu_a, \mu^* + \epsilon) \geq \text{kl}(\mathbb{P}_{\boldsymbol{\mu}}(A_T), \mathbb{P}_{\boldsymbol{\lambda}}(A_T)),$$

for any event $A_T \in \mathcal{F}_T$. Now the event $A_T := (N_a(T) < T/2)$ is very likely under $\boldsymbol{\mu}$ in which a is suboptimal, and very unlikely under $\boldsymbol{\lambda}$ in which a is optimal. More precisely, the uniformly efficient assumption permits to show that $\mathbb{P}_{\boldsymbol{\mu}}(A_T) \rightarrow 1$ and $\mathbb{P}_{\boldsymbol{\lambda}}(A_T) \leq o(T^{\alpha})/T$ for all α when T goes to infinity. This leads to $\text{kl}(\mathbb{P}_{\boldsymbol{\mu}}(A_T), \mathbb{P}_{\boldsymbol{\lambda}}(A_T)) \sim \log(T)$ and proves Theorem 2.

2.2. Asymptotically Optimal Index Policies and Upper Confidence Bounds

Lai and Robbins also proposed the first algorithm whose regret matches the lower bound (4) and this first asymptotically optimal algorithm is actually an index policy, i.e. it is of the form

$$A_{t+1} = \operatorname{argmax}_{a \in \{1, \dots, K\}} U_a(t),$$

but the proposed indices $U_a(t)$ are quite complex to compute. [Agrawal, 1995, Katehakis and Robbins, 1995] later proposed slightly more simple indices and show that they can be interpreted as *Upper Confidence Bounds* (UCB) on the unknown means of the arms. UCB-type algorithms were popularized by [Auer et al., 2002], who introduce the UCB1 algorithm for (non-parametric) bandit models with bounded rewards, and give the first *finite-time* upper bound on its regret. Simple indices like those of UCB1 can be used more generally for σ^2 -sub-Gaussian rewards, and take the form

$$U_a(t) = \hat{\mu}_a(t) + \sqrt{\frac{2\sigma^2 f(t)}{N_a(t)}},$$

for some function f which controls the confidence level. While the original choice of [Auer et al., 2002] is too conservative, one may safely choose $f(t) = \log(t)$ in practice; obtaining finite-time regret bounds is somewhat easier with a slightly larger choice, as in [Garivier and Cappé, 2011]. With such a choice, for Bernoulli

distributions (which are $1/4$ -sub-Gaussian), the regret of this index policy can be shown to be

$$\sum_{a:\mu_a < \mu^*}^K \frac{1}{2(\mu^* - \mu_a)} \log(T) + O(\sqrt{\log(T)}),$$

which is only order-optimal with respect to the lower bound (4), as by Pinsker inequality $d(\mu_a, \mu^*) > 2(\mu^* - \mu_a)^2$. Since the work of [Auer et al., 2002], several improvements of UCB1 have been proposed. They aimed at providing finite-time regret guarantees that would match the asymptotic lower bound (4) (see the review [Bubeck and Cesa-Bianchi, 2012]). Among them, the kl-UCB algorithm studied by [Cappé et al., 2013] is shown to be asymptotically optimal when the arms belong to a one-parameter exponential family. This algorithm is an index policy associated with

$$u_a(t) = \max \{q : N_a(t) d(\hat{\mu}_a(t), q) \leq f(t)\},$$

for the same choice of an exploration function f as mentioned above. The discussion of Section 1.3 explains why this index is actually an upper confidence bound on μ_a : choosing $f(t) = \log(t) + 3 \log \log(t)$, one has $\mathbb{P}_{\boldsymbol{\mu}}(u_a(t) \geq \mu_a) \gtrsim 1 - 1/[t \log^2(t)]$. For this particular choice, [Cappé et al., 2013] give a finite-time analysis of kl-UCB, proving its asymptotic optimality. To conclude on UCB algorithms, let us mention that several improvements have been proposed. A simple but significant one is obtained by replacing $f(t)$ by $\log(t/N_a(t))$ in the definition of $u_a(t)$, leading to a variant sometimes termed kl-UCB⁺ which has a slightly better empirical performance, but also minimax guarantees that plain UCB algorithms do not enjoy (for a discussion and related ideas, see the OCUCB algorithm of [Lattimore, 2016], [Ménard and Garivier, 2017] and the references therein).

2.3. Beyond the Optimism Principle

For simple parametric bandit models, in particular when rewards belong to a one-parameter exponential family, we showed that the regret minimization problem is solved, at least in an asymptotic sense: the kl-UCB algorithm, for example, attains the best possible regret rate on every problem instance. All the UCB-type algorithms described in the previous section are based on the so called principle of “optimism in face of uncertainty”. Indeed, at each round of a UCB algorithm the confidence intervals materialize the set of bandit models that are compatible with our observations (see Figure 1, left), and choosing the arm with largest UCB amounts to acting optimally in an “optimistic” model in which the mean of each arm would be equal to its best possible value. This optimism principle has also been successfully applied in some structured bandit models [Abbasi-Yadkori et al., 2011], as well as in reinforcement learning [Jaksch et al., 2010] and other related problems [Bubeck et al., 2013].

While the Lai and Robbins’ lower bound provides a good guideline to design algorithms, it has sometimes been misunderstood as a justification of the wrong *folk theorem* which is well-known by practitioners mostly interested in using bandit algorithms: “no strategy can have a regret smaller than $C(\boldsymbol{\mu}) \log(t)$, which is reached by good strategies”. But experiments often infirm this claim: it is easy to show settings and algorithms where the regret is much smaller than $C(\boldsymbol{\mu}) \log(t)$ and does not look like a logarithmic curve. The reason is twofold: first, Lai and Robbins’ lower result is asymptotic; a close look at its proof shows that it is relevant only when the horizon T is so large that any reasonable policy has identified the best arm with high probability; second, it only states that the regret *divided by* $\log(t)$ cannot always be smaller than $C(\boldsymbol{\mu})$. In [Garivier et al., 2016a], a more simple but similar bandit model of complexity $C'(\boldsymbol{\mu})$ is given where some strategy is proved to have a regret smaller than $C'(\boldsymbol{\mu}) \log(t) - c \log(\log(t))$ for some positive constant c .

Some recent works try to complement this result and to give a better description of what can be observed in practice. Notably, [Garivier et al., 2016b] focuses mainly on the initial regime: the authors show in particular that all strategies suffer a *linear regret* before T reaches some problem-dependent value. When the problem is very difficult (for example when the number of arms is very large) this initial phase may be the only observable one... They give non-asymptotic inequalities, and above all show a way to prove lower bounds which may lead to further new results (see e.g. [Garivier et al., 2016a]). It would be of great interest (but technically difficult) to exhibit an intermediate regime where, after this first phase, statistical estimation becomes possible but is

still not trivial. This would in particular permit to discriminate from a theoretical perspective between all the bandit algorithms that are now known to be asymptotically optimal, but for which significant differences may be observed in practice.

Indeed, one drawback of the kl-UCB algorithm is the need to construct tight confidence intervals (as explained in Section 1.3), which may not be generalized easily beyond simple parametric models. More flexible, Bayesian algorithms have recently also been shown to be asymptotically optimal, and to have good empirical performances. Given a prior distribution on the arms, Bayesian algorithms are simple procedures exploiting the *posterior distributions* of each arm. In the Bernoulli example, assuming a uniform prior on the mean of each arm, the posterior distribution of μ_a at round t , defined as the conditional distribution of μ_a given past observations, is easily seen to be a Beta distribution with parameters given by the number of ones and zeros observed so far from the arm. A Bayesian algorithm uses the different posterior distributions, that are illustrated in Figure 1 (right), to choose the next arm to sample from.

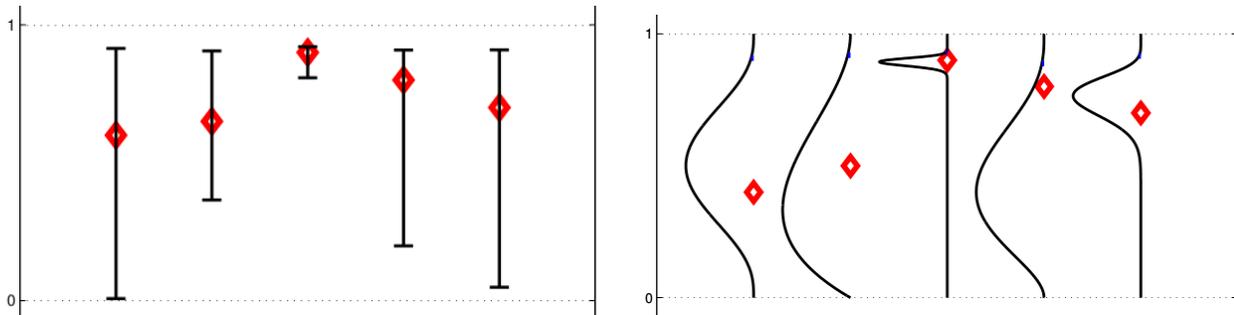


FIGURE 1. Frequentist versus Bayesian algorithms in a five-armed Bernoulli bandit model (the red diamonds are the unknown means). On the left, the confidence intervals used by kl-UCB, on the right the Beta posterior distributions used by Bayesian algorithms.

The Bayes-UCB algorithm of [Kaufmann et al., 2012a] exploits these posterior distributions in an optimistic way: it selects at round t the arm whose posterior on the mean has the largest quantile of order $1 - 1/t$. Another popular algorithm, Thompson Sampling, departs from the optimism principle by selecting arms at random according to their probability of being optimal. This principle was proposed by [Thompson, 1933] as the very first bandit algorithm, and can easily be implemented by drawing one sample from the posterior distribution of the mean of each arm, and selecting the arm with highest sample. This algorithm, also called probability matching, was rediscovered in the 2000s for its good empirical performances in complex bandit models [Scott, 2010, Chapelle and Li, 2011], but its first regret analysis dates back to [Agrawal and Goyal, 2012]. Both Thompson Sampling and Bayes-UCB have been shown recently to be asymptotically optimal in one-parameter models, for some choices of the prior distribution [Kaufmann et al., 2012b, Agrawal and Goyal, 2013, Kaufmann, 2016]. These algorithms are also quite generic, as they can be implemented in any bandit model in which one can define a prior distribution on the arms, and draw samples from the associated posterior. For example, they can be used in (generalized) linear bandit models, that can model recommendation tasks where the features of the items are taken into account (see, e.g. [Agrawal and Goyal, 2013] and the Chapter 4 of [Kaufmann, 2014]).

3. OPTIMAL STRATEGIES FOR BEST ARM IDENTIFICATION

Finding the arm with largest mean (without trying to maximize the *cumulated* rewards) is quite a different task and relates more to classical statistics. It can indeed be cast into the framework of *sequential adaptive hypothesis testing* introduced by [Chernoff, 1959]. In this framework, one has to decide which of the (composite)

hypotheses

$$H_1 : “\mu_1 > \max_{i \neq 1} \mu_i” , H_2 = “\mu_2 > \max_{i \neq 2} \mu_i” , \dots , H_K = “\mu_K > \max_{i \neq K} \mu_i”$$

is true. In order to gain information, one can select at each round one out of K possible *experiments*, each of them consisting in sampling from one of the marginal distributions (arms). Moreover, one has to choose when to stop the trial and decide for one of the hypotheses. Rephrased in a “bandit” terminology, a strategy consists of

- a *sampling rule*, that specifies which arm A_t is selected at time t (A_t is \mathcal{F}_{t-1} -measurable),
- a *stopping rule* τ , that indicates when the trial ends (τ is a stopping time wrt \mathcal{F}_t),
- a *recommendation rule* \hat{a}_τ that provides, upon stopping, a guess for the best arm (\hat{a}_τ is \mathcal{F}_τ -measurable).

However, the objective of the fixed-confidence best arm identification problem differs from that of [Chernoff, 1959], where one aims at minimizing a risk measure of the form

$$R(\boldsymbol{\mu}) = r_{a^*(\boldsymbol{\mu})} \mathbb{P}_{\boldsymbol{\mu}}(\hat{a}_\tau \neq a^*(\boldsymbol{\mu})) + c \mathbb{E}_{\boldsymbol{\mu}}[\tau],$$

where r_i is the cost of wrongly rejecting hypothesis H_i and c is a cost for sampling. Modern bandit literature rather focuses on so-called (ϵ, δ) -PAC *strategies* (for Probably Approximately Correct) which output, with high probability, an arm whose mean is within ϵ of the mean of the best arm:

$$\forall \boldsymbol{\mu} \in I^K, \quad \mathbb{P}_{\boldsymbol{\mu}}(\mu_{\hat{a}_\tau} > \mu^* - \epsilon) \geq 1 - \delta.$$

The goal is to build a (ϵ, δ) -PAC strategy with a sample complexity $\mathbb{E}_{\boldsymbol{\mu}}[\tau]$ that is as small as possible. For simplicity, we focus here on the case $\epsilon = 0$: a strategy is called δ -PAC¹ if

$$\forall \boldsymbol{\mu} \in \mathcal{S}, \quad \mathbb{P}_{\boldsymbol{\mu}}(\hat{a}_\tau = a^*(\boldsymbol{\mu})) \geq 1 - \delta,$$

where $\mathcal{S} = \{\boldsymbol{\mu} \in I^K : \exists i : \mu_i > \max_{j \neq i} \mu_j\}$ is the set of bandit models that have a unique optimal arm.

We show in the next section that, as in the regret minimization framework, there exists an instance-dependent lower bound on the sample complexity of any δ -PAC algorithm. We further present an algorithm whose sample complexity matches the lower bound, at least in the asymptotic regime where δ goes to 0. It is remarkable that this optimal algorithm, described in Section 3.2, is actually a by-product of the lower bound analysis described in Section 3.1, which sheds light on how a good strategy should distribute the draws between the arms.

3.1. The Sample Complexity of δ -PAC Best Arm Identification

The first lower bound on the sample complexity of a (ϵ, δ) -PAC algorithm was given by [Mannor and Tsitsiklis, 2004]. Particularized to the case $\epsilon = 0$, the lower bound says that for Bernoulli bandit models with means in $[0, \alpha]$, there exists a constant C_α and a subset of the sub-optimal arms \mathcal{K}_α such that for any δ -PAC algorithm

$$\mathbb{E}_{\boldsymbol{\mu}}[\tau] \geq C_\alpha \left[\sum_{a \in \mathcal{K}_\alpha} \frac{1}{(\mu^* - \mu_a)^2} \right] \log \left(\frac{1}{8\delta} \right).$$

Following this result, the literature has provided several δ -PAC strategies together with upper bounds on their sample complexity, mostly under the assumption that the rewards are bounded in $[0, 1]$. Existing strategies fall into two categories: those based on successive eliminations [Even-Dar et al., 2006, Karnin et al., 2013], and those based on confidence intervals [Kalyanakrishnan et al., 2012, Gabillon et al., 2012, Jamieson et al., 2014].

¹It would be more correct to call it a δ -PC (*Probably Correct*) strategy.

For all these algorithms, under a bandit instance such that $\mu_1 > \mu_2 \geq \dots \geq \mu_K$, the number of samples used can be shown to be of order

$$C \left[\frac{1}{(\mu_1 - \mu_2)^2} + \sum_{a=2}^K \frac{1}{(\mu_1 - \mu_a)^2} \right] \log \left(\frac{1}{\delta} \right) + o_{\delta \rightarrow 0} \left(\log \frac{1}{\delta} \right),$$

where C is a (large) numerical constant. While explicit finite-time bounds on τ can be extracted from most of the papers listed above, we mostly care here about the first-order term in δ , when δ goes to zero. Both the upper and lower bounds take the form of a sum over the arms of an individual complexity term (involving the inverse squared gap with the best or second best arm), but there is a gap as those sums do not involve the same number of terms; in addition, loose multiplicative constants make it hard to identify the exact minimal sample complexity of the problem.

As for the regret minimization problem, the true sample complexity can be expected to involve information-theoretic quantities (like the Kullback-Leibler divergence between arms distributions), for which the quantities above appear to be only surrogates; for example, for Bernoulli distributions, it holds that $2(\mu_1 - \mu_a)^2 < d(\mu_a, \mu_1) < (\mu_1 - \mu_a)^2 / (\mu_1(1 - \mu_1))$. For exponential families, it has been shown that incorporating the KL-based confidence bounds described in Section 1.3 into existing algorithms lowers the sample complexity [Kaufmann and Kalyanakrishnan, 2013] but the true sample complexity was only recently obtained by [Garivier and Kaufmann, 2016]. The result, and its proof, are remarkably simple:

Theorem 3. *Let $\boldsymbol{\mu} \in \mathcal{S}$, define $\text{Alt}(\boldsymbol{\mu}) := \{\boldsymbol{\lambda} \in \mathcal{S} : a^*(\boldsymbol{\lambda}) \neq a^*(\boldsymbol{\mu})\}$ and let $\Sigma_K = \{w \in [0, 1]^K : \sum_{a=1}^K w_a = 1\}$ be the set of probability vectors. Any δ -PAC algorithm satisfies*

$$\mathbb{E}_{\boldsymbol{\mu}}[\tau] \geq T^*(\boldsymbol{\mu}) \text{kl}(\delta, 1 - \delta),$$

where

$$T^*(\boldsymbol{\mu})^{-1} = \sup_{w \in \Sigma_K} \inf_{\boldsymbol{\lambda} \in \text{Alt}(\boldsymbol{\mu})} \sum_{a=1}^K w_a d(\mu_a, \lambda_a).$$

This lower bound against relies on a change of distribution, but unlike Lai and Robbins' result (and previous results for best arm identification), it is not sufficient to individually lower bound the expected number of draws of each arm using a single alternative model. One needs to consider the set $\text{Alt}(\boldsymbol{\mu}) = \{\boldsymbol{\lambda} : a^*(\boldsymbol{\lambda}) \neq a^*(\boldsymbol{\mu})\}$ of all possible alternative models $\boldsymbol{\lambda}$ in which the optimal arm is different from the optimal arm of $\boldsymbol{\mu}$.

Given a δ -PAC algorithm, let $\mathcal{E} = (\hat{a}_\tau \neq \hat{a}^*(\boldsymbol{\mu}))$. For any $\boldsymbol{\lambda} \in \text{Alt}(\boldsymbol{\mu})$, the δ -PAC property implies that $\mathbb{P}_{\boldsymbol{\mu}}(\mathcal{E}) \leq \delta$ while $\mathbb{P}_{\boldsymbol{\lambda}}(\mathcal{E}) \geq 1 - \delta$. Hence, by Lemma 1,

$$\sum_{a=1}^K \mathbb{E}_{\boldsymbol{\mu}}[N_a(\tau)] d(\mu_a, \lambda_a) \geq \text{kl}(\delta, 1 - \delta).$$

Combining *all* the inequalities thus obtained for the different possible values of $\boldsymbol{\lambda} \in \text{Alt}(\boldsymbol{\mu})$, we conclude that:

$$\begin{aligned} \text{kl}(\delta, 1 - \delta) &\leq \inf_{\boldsymbol{\lambda} \in \text{Alt}(\boldsymbol{\mu})} \sum_{a=1}^K \mathbb{E}_{\boldsymbol{\mu}}[N_a(\tau)] d(\mu_a, \lambda_a) \\ &\leq \mathbb{E}_{\boldsymbol{\mu}}[\tau] \left(\inf_{\boldsymbol{\lambda} \in \text{Alt}(\boldsymbol{\mu})} \sum_{a=1}^K \frac{\mathbb{E}_{\boldsymbol{\mu}}[N_a(\tau)]}{\mathbb{E}_{\boldsymbol{\mu}}[\tau]} d(\mu_a, \lambda_a) \right) \\ &\leq \mathbb{E}_{\boldsymbol{\mu}}[\tau] \left(\sup_{w \in \Sigma_K} \inf_{\boldsymbol{\lambda} \in \text{Alt}(\boldsymbol{\mu})} \sum_{a=1}^K w_a d(\mu_a, \lambda_a) \right). \end{aligned}$$

In the last step, we use the fact that the vector $(\mathbb{E}_{\boldsymbol{\mu}}[N_a(\tau)]/\mathbb{E}_{\boldsymbol{\mu}}[\tau])$ sums to one: upper bounding by the worst probability vector w yields a bound that is independent of the algorithm.

We thus obtain the (not fully explicit, but simple) lower bound of Theorem 3 that holds under the parametric assumption of Section 1.1. Its form involving an optimization problem is reminiscent of the early work of [Agrawal et al., 1989, Graves and Lai, 1997] that provide a lower bound on the regret in general, possibly structured bandit models. For best arm identification, [Vaidhyan and Sundaresan, 2015] consider the particular case of Poisson distribution in which there is only one arm that is different from the others, where a very nice formula can be derived for the sample complexity. For general exponential family bandit models, we now provide a slightly more explicit expression of $T^*(\boldsymbol{\mu})$, that permits to efficiently compute it.

3.2. An Asymptotically Optimal Algorithm

3.2.1. Computing the complexity and the optimal weights

The proof of Theorem 3 reveals that the quantity

$$w^*(\boldsymbol{\mu}) \in \operatorname{argmax}_{w \in \Sigma_K} \inf_{\boldsymbol{\lambda} \in \operatorname{Alt}(\boldsymbol{\mu})} \sum_{a=1}^K w_a d(\mu_a, \lambda_a)$$

can be interpreted as a vector of *optimal proportions*, in the sense that any strategy matching the lower bound should satisfy $\mathbb{E}_{\boldsymbol{\mu}}[N_a(\tau)]/\mathbb{E}_{\boldsymbol{\mu}}[\tau] \simeq w_a^*(\boldsymbol{\mu})$. Some algebra shows that the above optimization problem has a unique solution, and provides an efficient way of computing $w^*(\boldsymbol{\mu})$ for any $\boldsymbol{\mu}$, which boils down to numerically solving a series of scalar equations. In this section, we shall assume that $\boldsymbol{\mu}$ is such that $\mu_1 > \mu_2 \geq \dots \geq \mu_K$.

First, when the distribution belong to a one-dimensional exponential family (which we assume in the rest of this section), one can solve the inner optimization over $\boldsymbol{\lambda}$ in closed form, using Lagrange duality. This yields:

$$w^*(\boldsymbol{\mu}) \in \operatorname{argmax}_{w \in \Sigma_K} \min_{a \neq 1} \left[w_1 d\left(\mu_1, \frac{w_1 \mu_1 + w_a \mu_a}{w_1 + w_a}\right) + w_a d\left(\mu_a, \frac{w_1 \mu_1 + w_a \mu_a}{w_1 + w_a}\right) \right].$$

Then, one can prove that at the optimum the $K - 1$ quantities in the min are equal. Introducing their common value as an auxiliary variable, one can show that the computation of $w^*(\boldsymbol{\mu})$ reduces to solving a one-dimensional optimization problem. For all $a \neq 1$, one introduces the strictly increasing mapping

$$\begin{aligned} g_a : [0, +\infty[&\longrightarrow [0, d(\mu_1, \mu_a)[\\ x &\mapsto d\left(\mu_1, \frac{\mu_1 + x \mu_a}{1 + x}\right) + x d\left(\mu_a, \frac{\mu_1 + x \mu_a}{1 + x}\right), \end{aligned}$$

and defines $x_a : [0, d(\mu_1, \mu_a)[\rightarrow [0, +\infty[$ to be its inverse mapping. With this notation, the following Lemma 4 provides a way to compute $w^*(\boldsymbol{\mu})$. The scalar equation $F_{\boldsymbol{\mu}}(y) = 0$ defined therein may be solved using binary search. At each step of the search, the solution $x_a(y)$ of the equation $g_a(x) = y$ can again be computed by using binary search, or by Newton's method. This algorithm is available as a `julia` code at <https://github.com/jsfunc/best-arm-identification>.

Lemma 4. For every $a \in \{1, \dots, K\}$,

$$w_a^*(\boldsymbol{\mu}) = \frac{x_a(y^*)}{\sum_{a=1}^K x_a(y^*)}, \quad (5)$$

where y^* is the unique solution of the equation $F_{\boldsymbol{\mu}}(y) = 1$, where

$$F_{\boldsymbol{\mu}} : y \mapsto \sum_{a=2}^K \frac{d\left(\mu_1, \frac{\mu_1 + x_a(y) \mu_a}{1 + x_a(y)}\right)}{d\left(\mu_a, \frac{\mu_1 + x_a(y) \mu_a}{1 + x_a(y)}\right)} \quad (6)$$

is a continuous, increasing function on $[0, d(\mu_1, \mu_2)[$ such that $F_{\boldsymbol{\mu}}(0) = 0$ and $F_{\boldsymbol{\mu}}(y) \rightarrow \infty$ when $y \rightarrow d(\mu_1, \mu_2)$.

This result yields an efficient algorithm for computing $T^*(\boldsymbol{\mu})$. But can a closed form formula be derived, at least in some special cases? In the two-armed case, it is easy to see that $T^*(\boldsymbol{\mu})$ is equal to the inverse Chernoff information between the two arms (see [Kaufmann et al., 2014]). However, no closed form is available when $K \geq 3$, even for simple families of distributions. For Gaussian arms with known variance σ^2 , only the following bound is known which captures $T^*(\boldsymbol{\mu})$ up to a factor 2:

$$\frac{2\sigma^2}{(\mu_1 - \mu_2)^2} + \sum_{a=2}^K \frac{2\sigma^2}{(\mu_1 - \mu_a)^2} \leq T^*(\boldsymbol{\mu}) \leq 2 \left[\frac{2\sigma^2}{(\mu_1 - \mu_2)^2} + \sum_{a=2}^K \frac{2\sigma^2}{(\mu_1 - \mu_a)^2} \right].$$

Note that $T^*(\boldsymbol{\mu})$ may be much smaller than $4\sigma^2 K / (\mu_1 - \mu_2)^2$, which is the minimal number of samples required by a strategy using uniform sampling (for which $N_a(t)/t \simeq 1/K$). An optimal strategy actually uses quite unbalanced weights $w^*(\boldsymbol{\mu})$.

3.2.2. An algorithm inspired from the lower bound

Back to general exponential families, building on the lower bound and our ability to compute $w^*(\boldsymbol{\mu})$, we now introduce an efficient algorithm whose sample complexity matches the lower bound, at least for small values of δ . This Track-and-Stop strategy consists of two elements:

- a *tracking sampling rule*, that forces the proportion of draws of each arm a to converge to the associated optimal proportion $w_a^*(\boldsymbol{\mu})$, by using the plug-in estimates $w^*(\hat{\boldsymbol{\mu}}(t))$,
- the *Chernoff stopping rule*, that can be interpreted as the stopping rule of a sequential Generalized Likelihood Ratio Test (GLRT), whose closed form in this particular problem is very similar to our lower bound.

When stopping, our guess \hat{a}_τ is the empirical best arm. We now describe the sampling and stopping rule in details before presenting the theoretical guarantees for Track-and-Stop.

3.2.3. The Tracking sampling rule

Let $\hat{\boldsymbol{\mu}}(t) = (\hat{\mu}_1(t), \dots, \hat{\mu}_K(t))$ be the current maximum likelihood estimate of $\boldsymbol{\mu}$ at time t :

$$\hat{\mu}_a(t) = \frac{1}{N_a(t)} \sum_{s \leq t} X_s \mathbb{1}\{A_s = a\}.$$

A first idea for matching the proportions $w^*(\boldsymbol{\mu})$ is to track the plug-in estimates $w^*(\hat{\boldsymbol{\mu}}(t))$, by drawing at round t the arm a whose empirical proportion of draws lags furthest behind the estimated target $w_a^*(\hat{\boldsymbol{\mu}}(t))$. But a closer inspection shows that (sufficiently fast) convergence of $\hat{\boldsymbol{\mu}}(t)$ towards the true parameter $\boldsymbol{\mu}$ requires some “forced exploration” to make sure each arm has not been under-sampled. More formally, defining $F_t = \{a : N_a(t) < \sqrt{t} - K/2\}$, the Tracking rule is defined as

$$A_{t+1} \in \begin{cases} \underset{a \in F_t}{\operatorname{argmin}} N_a(t) & \text{if } F_t \neq \emptyset & (\text{forced exploration}) \\ \underset{1 \leq a \leq K}{\operatorname{argmax}} \{w_a^*(\hat{\boldsymbol{\mu}}(t)) - N_a(t)/t\} & \text{otherwise} & (\text{tracking the plug-in estimate}) \end{cases}$$

Simple combinatorial arguments prove that the Tracking rule draws each arm at least $(\sqrt{t} - K/2)_+ - 1$ times at round t , and relate the gap between $N_a(t)/t$ and $w_a^*(\boldsymbol{\mu})$ to the gap between $w_a^*(\boldsymbol{\mu})$ and $w_a^*(\hat{\boldsymbol{\mu}}(t))$. This permits in particular to show that the Tracking rule has the following desired behavior:

Proposition 5. *The Tracking sampling rule satisfies*

$$\mathbb{P}_{\boldsymbol{\mu}} \left(\lim_{t \rightarrow \infty} \frac{N_a(t)}{t} = w_a^*(\boldsymbol{\mu}) \right) = 1.$$

3.2.4. The Chernoff stopping rule

Let $\ell(X_1, \dots, X_t; \boldsymbol{\lambda})$ be the log-likelihood of the observations up to time t under a bandit model $\boldsymbol{\lambda}$ and define

$$\hat{Z}(t) = \frac{\sup_{\boldsymbol{\lambda}} \ell(X_1, \dots, X_t; \boldsymbol{\lambda})}{\sup_{\boldsymbol{\lambda} \in \text{Alt}(\hat{\boldsymbol{\mu}}(t))} \ell(X_1, \dots, X_t; \boldsymbol{\lambda})} = \frac{\ell(X_1, \dots, X_t; \hat{\boldsymbol{\mu}}(t))}{\sup_{\boldsymbol{\lambda} \in \text{Alt}(\hat{\boldsymbol{\mu}}(t))} \ell(X_1, \dots, X_t; \boldsymbol{\lambda})}.$$

Intuitively, this generalized likelihood ratio $\hat{Z}(t)$ is large if the current maximum-likelihood estimate is far apart from its “closest alternative” $\hat{\boldsymbol{\mu}}(t)$ defined as the parameter maximizing the likelihood under the constraint that it belongs to $\text{Alt}(\hat{\boldsymbol{\mu}}(t))$, i.e. that its optimal arm is different from that of $\hat{\boldsymbol{\mu}}(t)$. This idea can be traced back to the work of [Chernoff, 1959], in which $\hat{Z}(t)$ is interpreted as the Neyman-Pearson statistic for testing the (data-dependent) pseudo-hypothesis “ $\boldsymbol{\mu} = \hat{\boldsymbol{\mu}}(t)$ ” against “ $\boldsymbol{\mu} = \tilde{\boldsymbol{\mu}}(t)$ ”, based on all samples available up to round t . The analysis of Chernoff, however, only applies to two discrete hypotheses, whereas the best arm identification problem requires to consider K continuous hypotheses.

In the paper [Garivier and Kaufmann, 2016], we provide new insights on this Chernoff stopping rule, formally defined as

$$\tau_\delta = \inf \left\{ t \in \mathbb{N} : \hat{Z}(t) > \beta(t, \delta) \right\}, \quad (7)$$

where $\beta(t, \delta)$ is some threshold function.

The first problem is to set the threshold $\beta(t, \delta)$ such that the probability of error of Track-and-Stop is upper bounded by δ . Our analysis relies on expressing $\hat{Z}(t)$ in terms of pairwise sequential GLRTs of “ $\mu_a < \mu_b$ ” against “ $\mu_a \geq \mu_b$ ”, for which we provide tight bounds on the type I error. Indeed, letting

$$\hat{Z}_{a,b}(t) := \log \frac{\max_{\lambda_a \geq \lambda_b} p_{\lambda_a} \left(\underline{X}_{N_a(t)}^a \right) p_{\lambda_b} \left(\underline{X}_{N_b(t)}^b \right)}{\max_{\lambda_a \leq \lambda_b} p_{\lambda_a} \left(\underline{X}_{N_a(t)}^a \right) p_{\lambda_b} \left(\underline{X}_{N_b(t)}^b \right)},$$

where $\underline{X}_{N_a(t)}^a = (X_s : A_s = a, s \leq t)$ is a vector that contains the observations of arm a available at time t , and where $p_{\lambda}(V_1, \dots, V_n)$ is the likelihood of n i.i.d. observations drawn from ν_{λ} , one can show that

$$\hat{Z}(t) = \min_{b \neq \hat{a}(t)} \hat{Z}_{\hat{a}(t), b}(t), \quad (8)$$

where $\hat{a}(t)$ is the empirical best arm at round t . In other words, one stops when for each arm b that is different from the empirical best arm $\hat{a}(t)$, a GLRT would reject the (data-dependent) pseudo-hypothesis “ $\mu_{\hat{a}(t)} < \mu_b$ ”. This expression also allows for a simple computation of $\hat{Z}(t)$, as $\hat{\mu}_a(t) > \hat{\mu}_b(t)$ implies

$$\hat{Z}_{a,b}(t) = N_a(t) d \left(\hat{\mu}_a(t), \frac{N_a(t) \hat{\mu}_a(t) + N_b(t) \hat{\mu}_b(t)}{N_a(t) + N_b(t)} \right) + N_b(t) d \left(\hat{\mu}_b(t), \frac{N_a(t) \hat{\mu}_a(t) + N_b(t) \hat{\mu}_b(t)}{N_a(t) + N_b(t)} \right).$$

Under the Tracking sampling rule, it is easy to see that the $\hat{Z}(t)$ grows linearly with t , hence with a $\beta(t, \delta)$ that is sub-linear in t , τ_δ is also surely finite. The probability of error of the Chernoff sampling rule is thus upper bounded by

$$\mathbb{P}_{\boldsymbol{\mu}}(\hat{a}_\tau \neq 1) \leq \mathbb{P}_{\boldsymbol{\mu}} \left(\exists t \in \mathbb{N} : \hat{a}(t) \neq 1, \hat{Z}(t) > \beta(t, \delta) \right) \leq \sum_{a=2}^K \mathbb{P}_{\boldsymbol{\mu}} \left(\exists t \in \mathbb{N} : \hat{Z}_{a,1}(t) > \beta(t, \delta) \right).$$

In the Bernoulli case, the following lemma permits to prove that the Chernoff stopping rule is δ -PAC for the choice

$$\beta(t, \delta) = \log \left(\frac{2(K-1)t}{\delta} \right). \quad (9)$$

Lemma 6. *Let $\delta \in (0, 1)$. For any sampling strategy, if $\mu_a < \mu_b$, $\mathbb{P}\left(\exists t \in \mathbb{N} : \hat{Z}_{a,b}(t) > \log(2t/\delta)\right) \leq \delta$.*

Another rewriting of the Chernoff stopping rule permits to understand why it achieves the optimal sample complexity when coupled to the Tracking stopping rule. Indeed, using the particular form of the likelihood in an exponential family yields

$$\hat{Z}(t) = t \left(\inf_{\lambda \in \text{Alt}(\hat{\mu}(t))} \sum_{a=1}^K \frac{N_a(t)}{t} d(\hat{\mu}_a(t), \lambda_a) \right).$$

This expression is reminiscent of the lower bound of Theorem 3. When t is large one expects $\hat{\mu}(t)$ to be close to μ thanks to the forced exploration, and $N_a(t)/t$ to be close to $w_a^*(\mu)$, due to Proposition 5. Hence one has $\hat{Z}(t) \simeq t/T^*(\mu)$ for large values of t . Thus, with the threshold function (9), for small δ , τ_δ is asymptotically upper bounded by the smallest t such that $t \geq T^*(\mu) \log(2(K-1)t/\delta)$, which is of order $T^*(\mu) \log(1/\delta)$ for small values of δ .

3.2.5. Optimality of Track-and-Stop

In the previous section, we sketched an upper bound on the number of samples used by Track-and-Stop that holds with probability one in the Bernoulli case. [Garivier and Kaufmann, 2016] also propose an asymptotic upper bound on the expected sample complexity of this algorithm, beyond the Bernoulli case. The results are summarized below.

Theorem 7. *Let $\alpha > 1$. There exists a constant $C = C_{\alpha, K}$ such that for all $\delta \in]0, 1[$ the Track-and-Stop strategy with threshold*

$$\beta(t, \delta) = \log \left(\frac{Ct^\alpha}{\delta} \right)$$

is δ -PAC and satisfies

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

In the Bernoulli case, one can set $\alpha = 1$ and $C = 2(K-1)$.

Hence, Track-and-Stop can be qualified as asymptotically optimal, in the sense that its sample complexity matches the lower bound of Theorem 3, when δ tends to zero. Inspired by the regret minimization study, an important direction of future work is to obtain finite-time upper bounds on the sample complexity of an algorithm, that would still asymptotically match the lower bound of Theorem 3. A different line of research has studied, for sub-Gaussian rewards, the asymptotic behavior of the sample complexity for fixed values of δ in a regime in which the gap between the best and second best arm goes to zero [Jamieson et al., 2014, Chen et al., 2016], leading to a different notion of optimality. Hence, we should aim for the best of both worlds: an algorithm with a finite-time sample complexity upper bound that would also match the lower bound obtained in this alternative asymptotic regime.

Finally, while Theorem 7 gives asymptotic results for Track-and-Stop, we would like to highlight the practical impact of this algorithm. Experiments in [Garivier and Kaufmann, 2016] reveal that for relatively “large” values of δ (e.g. $\delta = 0.1$) the sample complexity of Track-and-Stop appears to be twice smaller than that of state-of-the-art algorithms in generic scenarios. The sampling rule of Track-and-Stop is slightly more computationally demanding than that of its competitors, as it requires to compute $w^*(\hat{\mu}(t))$ at each round. However, the Tracking sampling rule is the most naive idea, and we will investigate whether other simple heuristics could be used to guarantee that the empirical proportions of draws converge towards the optimal proportions $w^*(\mu)$, while being amenable for finite-time analysis.

4. DISCUSSION

It is known at least since [Bubeck et al., 2011] that good algorithms for regret minimization and pure-exploration are expected to be different: small regrets after t time steps imply a large probability of error $\mathbb{P}_{\boldsymbol{\mu}}(\mu^* \neq \mu_{\hat{a}(t)})$, where $\hat{a}(t)$ is the recommendation for the best arm at time t . In the dual fixed confidence setting that was studied in this paper, we provided other elements to assess the difference of the regret minimization and best arm identification problems.

First, the sampling strategy used by both type of algorithms are very different. Regret minimizing algorithm draw the best arm most of the time ($t - O(\log(t))$ times in t rounds) while each sub-optimal arm gathers a vanishing proportion of draws. On the contrary, identifying the best arm requires the proportions of arm draws to converge to a vector $w^*(\boldsymbol{\mu})$ will all non-zero components. Figure 2 illustrates this different behavior: the number of draws of each arm and associated KL-based confidence intervals are displayed for the kl-UCB (left) and Track-and-Stop (right) strategies. As expected, Track-and-Stop draws more frequently than its competitor the close-to-optimal arms, and has therefore tighter confidence intervals on their means.

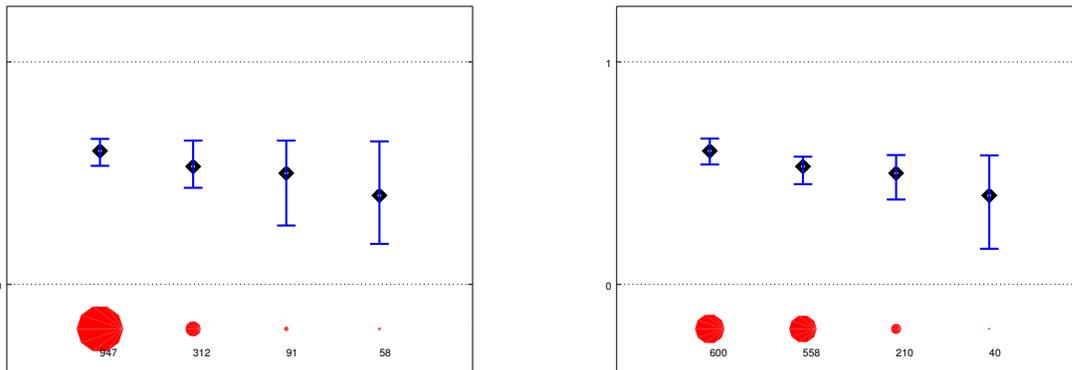


FIGURE 2. Number of draws of the kl-UCB (left) and Track-and-Stop (right) algorithms on one run of length $t = 1408$ (the stopping time of Track-and-Stop). Confidence intervals on the unknown mean of each arm (black diamond) are also displayed.

We also emphasize that the information-theoretic quantities characterizing the complexity of the two problems are different. For regret minimization, we saw that the minimal regret of uniformly efficient (u.e.) strategies satisfies:

$$\inf_{\pi \in \{u.e.\}} \lim_{T \rightarrow \infty} \frac{R_{\boldsymbol{\mu}}^{\pi}(T)}{\log(T)} = \sum_{a: \mu_a < \mu^*} \frac{\mu^* - \mu_a}{d(\mu_a, \mu^*)},$$

where $d(x, y)$ is the Kullback-Leibler divergence between the distribution of mean x and the distribution of mean y in our class. For best arm identification,

$$\inf_{\pi \in \{PAC\}} \lim_{\delta \rightarrow 0} \frac{\mathbb{E}_{\boldsymbol{\mu}}[\tau_{\delta}]}{\log(1/\delta)} = T^*(\boldsymbol{\mu}),$$

where $T^*(\boldsymbol{\mu})$ is the solution of an optimization problem expressed with Kullback-Leibler divergences between arms that has no closed form solution for more than two arms.

Although regret minimization and best arm identification are two very different objectives, both in terms of algorithms and of complexity, best arm identification tools have been used within regret minimization algorithms in so-called *Explore-Then-Commit* strategies [Perchet and Rigollet, 2013, Perchet et al., 2015]. For minimizing

regret up to a horizon T such strategies use a (elimination-based) fixed-confidence best arm identification algorithm with $\delta = 1/T$ to make a guess for the best arm and then commit to play this estimated best arm until the end of the horizon T . In a simple case (two Gaussian arms), we recently quantified the sub-optimality of such approaches: the regret of the best such Explore-Then-Commit strategy is at least twice larger than that of the kl-UCB algorithm [Garivier et al., 2016a]. Even if the article focuses on Gaussian rewards, other cases with possibly more than two arms are also discussed. Coming back to the introductory example of A/B testing, the take-home message is the following: if you prefer to experiment first the two options before using only one of them in production, instead of continuously allocating the two options using a good regret-minimizing strategy, then this will cost you twice larger a regret.

Unlike the asymptotically optimal regret minimizing strategies that we presented, the asymptotically optimal Track-and-Stop strategy for best arm identification has no finite-time guarantees, and its implementation is slightly more complex. An important future work is to see whether useful tools for regret minimization, like the optimism principle or Bayesian methods can be combined with Track-and-Stop to have a simpler algorithm with a finite-time analysis. A starting point may be found in [Russo, 2016], who recently proposed a modified Bayesian Thompson Sampling rule that has some promising properties.

Acknowledgement: The authors are extremely thankful to the reviewers of this paper, who contributed significantly to the clarity of the presentation by their numerous and always relevant comments.

REFERENCES

- [Abbasi-Yadkori et al., 2011] Abbasi-Yadkori, Y., D.Pál, and C.Szepesvári (2011). Improved Algorithms for Linear Stochastic Bandits. In *Advances in Neural Information Processing Systems*.
- [Agrawal, 1995] Agrawal, R. (1995). Sample mean based index policies with $O(\log n)$ regret for the multi-armed bandit problem. *Advances in Applied Probability*, 27(4):1054–1078.
- [Agrawal et al., 1989] Agrawal, R., Teneketzis, D., and Anantharam, V. (1989). Asymptotically Efficient Adaptive Allocation Schemes for Controlled i.i.d. Processes: Finite Parameter Space. *IEEE Transactions on Automatic Control*, 34(3):258–267.
- [Agrawal and Goyal, 2012] Agrawal, S. and Goyal, N. (2012). Analysis of Thompson Sampling for the multi-armed bandit problem. In *Proceedings of the 25th Conference On Learning Theory*.
- [Agrawal and Goyal, 2013] Agrawal, S. and Goyal, N. (2013). Further Optimal Regret Bounds for Thompson Sampling. In *Proceedings of the 16th Conference on Artificial Intelligence and Statistics*.
- [Audibert et al., 2010] Audibert, J.-Y., Bubeck, S., and Munos, R. (2010). Best Arm Identification in Multi-armed Bandits. In *Proceedings of the 23rd Conference on Learning Theory*.
- [Auer et al., 2002] Auer, P., Cesa-Bianchi, N., and Fischer, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2):235–256.
- [Bechhofer, 1954] Bechhofer, R. (1954). A single-sample multiple decision procedure for ranking means of normal populations with known variances. *Annals of Mathematical Statistics*, 25:16–39.
- [Bechhofer et al., 1968] Bechhofer, R., Kiefer, J., and Sobel, M. (1968). *Sequential identification and ranking procedures*. The University of Chicago Press.
- [Bellman, 1956] Bellman, R. (1956). A problem in the sequential design of experiments. *The indian journal of statistics*, 16(3/4):221–229.
- [Berry and Fristedt, 1985] Berry, D. and Fristedt, B. (1985). *Bandit Problems. Sequential allocation of experiments*. Chapman and Hall.
- [Bubeck, 2010] Bubeck, S. (2010). *Jeux de bandits et fondation du clustering*. PhD thesis, Université de Lille 1.
- [Bubeck and Cesa-Bianchi, 2012] Bubeck, S. and Cesa-Bianchi, N. (2012). Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122.
- [Bubeck et al., 2013] Bubeck, S., Ernst, D., and Garivier, A. (Feb. 2013). Optimal discovery with probabilistic expert advice: Finite time analysis and macroscopic optimality. *Journal of Machine Learning Research*, 14:601–623.
- [Bubeck et al., 2011] Bubeck, S., Munos, R., and Stoltz, G. (2011). Pure Exploration in Finitely Armed and Continuous Armed Bandits. *Theoretical Computer Science* 412, 1832–1852, 412:1832–1852.
- [Burnetas and Katehakis, 1996] Burnetas, A. and Katehakis, M. (1996). Optimal adaptive policies for sequential allocation problems. *Advances in Applied Mathematics*, 17(2):122–142.
- [Cappé et al., 2013] Cappé, O., Garivier, A., Maillard, O.-A., Munos, R., and Stoltz, G. (2013). Kullback-Leibler upper confidence bounds for optimal sequential allocation. *Annals of Statistics*, 41(3):1516–1541.

- [Cesa-Bianchi and Lugosi, 2006] Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, Learning and Games*. Cambridge University Press.
- [Chapelle and Li, 2011] Chapelle, O. and Li, L. (2011). An empirical evaluation of Thompson Sampling. In *Advances in Neural Information Processing Systems*.
- [Chen et al., 2016] Chen, L., Li, J., and Qiao, M. (2016). Towards instance optimal bounds for best arm identification. *arXiv:1608.06031*.
- [Chernoff, 1959] Chernoff, H. (1959). Sequential design of Experiments. *The Annals of Mathematical Statistics*, 30(3):755–770.
- [Chu et al., 2011] Chu, W., Li, L., Reyzin, L., and Schapire, R. (2011). Contextual Bandits with Linear Payoff Functions. In *Proceedings of the 14th Conference on Artificial Intelligence and Statistics*.
- [Combes and Proutière, 2014] Combes, R. and Proutière, A. (2014). Unimodal Bandits without Smoothness. Technical report.
- [Even-Dar et al., 2006] Even-Dar, E., Mannor, S., and Mansour, Y. (2006). Action Elimination and Stopping Conditions for the Multi-Armed Bandit and Reinforcement Learning Problems. *Journal of Machine Learning Research*, 7:1079–1105.
- [Gabillon et al., 2012] Gabillon, V., Ghavamzadeh, M., and Lazaric, A. (2012). Best Arm Identification: A Unified Approach to Fixed Budget and Fixed Confidence. In *Advances in Neural Information Processing Systems*.
- [Garivier and Cappé, 2011] Garivier, A. and Cappé, O. (2011). The KL-UCB algorithm for bounded stochastic bandits and beyond. In *Proceedings of the 24th Conference on Learning Theory*.
- [Garivier and Kaufmann, 2016] Garivier, A. and Kaufmann, E. (2016). Optimal best arm identification with fixed confidence. In *Proceedings of the 29th Conference On Learning Theory (to appear)*.
- [Garivier et al., 2016a] Garivier, A., Kaufmann, E., and Lattimore, T. (2016a). On explore-then-commit strategies. In *Advances in Neural Processing Systems (NIPS)*.
- [Garivier and Leonardi, 2011] Garivier, A. and Leonardi, F. (Nov. 2011). Context tree selection: A unifying view. *Stochastic Processes and their Applications*, 121(11):2488–2506.
- [Garivier et al., 2016b] Garivier, A., Ménard, P., and Stoltz, G. (2016b). Explore first, exploit next: The true shape of regret in bandit problems. *arXiv preprint arXiv:1602.07182*.
- [Garivier and Moulines, 2011] Garivier, A. and Moulines, E. (2011). On Upper-Confidence Bound Policies for Switching Bandit Problems. In *Proceedings of the 22nd conference on Algorithmic Learning Theory*.
- [Gittins, 1979] Gittins, J. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society, Series B*, 41(2):148–177.
- [Graves and Lai, 1997] Graves, T. and Lai, T. (1997). Asymptotically Efficient adaptive choice of control laws in controlled markov chains. *SIAM Journal on Control and Optimization*, 35(3):715–743.
- [Jaksch et al., 2010] Jaksch, T., Ortner, R., and Auer, P. (2010). Near-Optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11:1563–1600.
- [Jamieson et al., 2014] Jamieson, K., Malloy, M., Nowak, R., and Bubeck, S. (2014). lil’UCB: an Optimal Exploration Algorithm for Multi-Armed Bandits. In *Proceedings of the 27th Conference on Learning Theory*.
- [Kalyanakrishnan et al., 2012] Kalyanakrishnan, S., Tewari, A., Auer, P., and Stone, P. (2012). PAC subset selection in stochastic multi-armed bandits. In *International Conference on Machine Learning (ICML)*.
- [Karnin et al., 2013] Karnin, Z., Koren, T., and Somekh, O. (2013). Almost optimal Exploration in multi-armed bandits. In *International Conference on Machine Learning (ICML)*.
- [Katehakis and Robbins, 1995] Katehakis, M. and Robbins, H. (1995). Sequential choice from several populations. *Proceedings of the National Academy of Science*, 92:8584–8585.
- [Kaufmann, 2014] Kaufmann, E. (2014). *Analyse de stratégies bayésiennes et fréquentistes pour l’allocation séquentielle de ressources*. PhD thesis.
- [Kaufmann, 2016] Kaufmann, E. (2016). On bayesian index policies for sequential resource allocation. *Preprint arXiv:1601.01190*.
- [Kaufmann et al., 2012a] Kaufmann, E., Cappé, O., and Garivier, A. (2012a). On Bayesian Upper-Confidence Bounds for Bandit Problems. In *Proceedings of the 15th conference on Artificial Intelligence and Statistics*.
- [Kaufmann et al., 2014] Kaufmann, E., Cappé, O., and Garivier, A. (2014). On the Complexity of A/B Testing. In *Proceedings of the 27th Conference On Learning Theory*.
- [Kaufmann et al., 2016] Kaufmann, E., Cappé, O., and Garivier, A. (2016). On the Complexity of Best Arm Identification in Multi-Armed Bandit Models. *Journal of Machine Learning Research*, 17(1):1–42.
- [Kaufmann and Kalyanakrishnan, 2013] Kaufmann, E. and Kalyanakrishnan, S. (2013). Information complexity in bandit subset selection. In *Proceeding of the 26th Conference On Learning Theory*.
- [Kaufmann et al., 2012b] Kaufmann, E., Korda, N., and Munos, R. (2012b). Thompson Sampling : an Asymptotically Optimal Finite-Time Analysis. In *Proceedings of the 23rd conference on Algorithmic Learning Theory*.
- [Lai and Robbins, 1985] Lai, T. and Robbins, H. (1985). Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22.
- [Lattimore, 2016] Lattimore, T. (2016). Regret analysis of the anytime optimally confident UCB algorithm. *CoRR*, abs/1603.08661.
- [Magureanu et al., 2014] Magureanu, S., Combes, R., and Proutière, A. (2014). Lipschitz Bandits: Regret lower bounds and optimal algorithms. In *Proceedings on the 27th Conference On Learning Theory*.

- [Mannor and Tsitsiklis, 2004] Mannor, S. and Tsitsiklis, J. (2004). The Sample Complexity of Exploration in the Multi-Armed Bandit Problem. *Journal of Machine Learning Research*, pages 623–648.
- [Ménard and Garivier, 2017] Ménard, P. and Garivier, A. (Feb. 2017). A minimax and asymptotically optimal algorithm for stochastic bandits.
- [Perchet and Rigollet, 2013] Perchet, V. and Rigollet, P. (2013). The multi-armed bandit with covariates. *The Annals of Statistics*.
- [Perchet et al., 2015] Perchet, V., Rigollet, P., Chassang, S., and Snowberg, E. (2015). Batched bandit problems. In *Proceedings of the 28th Conference On Learning Theory*.
- [Puterman, 1994] Puterman, M. (1994). *Markov Decision Processes. Discrete Stochastic. Dynamic Programming*. Wiley.
- [Robbins, 1952] Robbins, H. (1952). Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535.
- [Russo, 2016] Russo, D. (2016). Simple bayesian algorithms for best arm identification. In *Proceedings of the 29th Conference On Learning Theory*.
- [Scott, 2010] Scott, S. (2010). A modern Bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry*, 26:639–658.
- [Sutton and Barto, 1998] Sutton, R. and Barto, A. (1998). *Reinforcement Learning: an Introduction*. MIT press.
- [Thompson, 1933] Thompson, W. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25:285–294.
- [Vaidhyan and Sundaresan, 2015] Vaidhyan, N. and Sundaresan, R. (2015). Learning to detect an oddball target. *arXiv:1508.05572*.