

Thompson Sampling: An Asymptotically Optimal Finite Time Analysis

Emilie Kaufmann, Nathaniel Korda and Rémi Munos

Telecom Paristech UMR CNRS 5141 & INRIA Lille - Nord Europe

July 20, 2012

Abstract

The question of the optimality of Thompson Sampling for solving the stochastic multi-armed bandit problem had been open since 1933. In this paper we answer it positively for the case of Bernoulli rewards by providing the first finite-time analysis that matches the asymptotic rate given in the Lai and Robbins lower bound for the cumulative regret. The proof is accompanied by a numerical comparison with other optimal policies, experiments that have been lacking in the literature until now for the Bernoulli case.

1 Introduction

In a stochastic bandit problem an agent is repeatedly asked to choose one action from an action set, each of which produces a reward drawn from an underlying, fixed, but unknown distribution associated with each action. Thus he must choose at each time whether to use the observations he has already gathered to gain the greatest immediate reward (exploitation) or whether to choose an action from which few observations have been made and risk immediate loss for greater knowledge and potential future gain (exploration). In this paper we focus on stochastic bandits with Bernoulli rewards, initially proposed by Thompson in his paper of 1933 [14] to model medical allocation problems. Thompson's paper also presented the first bandit algorithm, Thompson Sampling. This algorithm has received much attention in the recent literature, and in this paper we give the first theoretical proof of the asymptotic optimality of this algorithm in the context of cumulative regret minimisation. Furthermore we achieve this result by giving a finite time analysis for the algorithm.

Associated with each action, a , is an unknown Bernoulli distribution $\mathcal{B}(\mu_a)$, whose expectation is μ_a . At each time t the agent chooses to observe an action $A_t \in \{1, \dots, K\}$ and receives a reward R_t drawn from the distribution $\mathcal{B}(\mu_{A_t})$. A policy, or bandit algorithm, is defined to be a (possibly randomised) method for choosing A_t given the past history of observations and actions. The agent's goal is to minimize the expected cumulative regret of his policy, which is defined to be:

$$\mathcal{R}(T) := T\mu^* - \mathbb{E} \left[\sum_{t=1}^T R_t \right] = \sum_{a \in A} (\mu^* - \mu_a) \mathbb{E}[N_{a,t}] \quad (1)$$

where $\mu^* = \max_a \mu_a$ denotes the expectation of the best arm¹, or optimal action, and $N_{a,t}$ the number of draws of arm a at the end of round t . Lai and Robbins proved in [10] that all *strongly consistent* policies

¹The words arms and actions are used interchangeably.

(i.e. policies satisfying $\mathcal{R}(t) = o(t^\alpha)$ for all $\alpha \in (0, 1)$) must satisfy, for any suboptimal arm a

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}[N_{a,T}]}{\ln T} \geq \frac{1}{K(\mu_a, \mu^*)} \quad (2)$$

where $K(p, q)$ denotes the Kullback-Leibler divergence between $\mathcal{B}(p)$ and $\mathcal{B}(q)$:

$$K(p, q) := p \ln \frac{p}{q} + (1-p) \ln \frac{1-p}{1-q}.$$

Their result, which holds for more general classes of reward distributions, leads to the definition of *asymptotically optimal* policies as policies that satisfy (2) with equality.

In the same paper [10] Lai and Robbins were able to describe an asymptotically optimal policy, however no finite-time analysis was provided, nor was it an efficient policy to implement. The UCB1 algorithm by Auer et al. [4] was the first of a series of efficient policies, like UCB-V [3] or MOSS [2], for which good regret bounds in finite time were also provided. These policies all use an upper confidence bound for the empirical mean of past rewards as an optimistic index for each arm, choosing at each time the action with the highest current index. However, for each of these algorithms we only have the result that there exists two constants $K_1 > 2$ and $K > 0$ such that for every suboptimal action a , with $\Delta_a = \mu^* - \mu_a$,

$$\mathbb{E}[N_{a,T}] \leq \frac{K_1}{\Delta_a^2} \ln(T) + K_2. \quad (3)$$

This does not imply (2) with equality since by the Pinsker inequality $2K(\mu_a, \mu^*) > \Delta_a^2$. On the contrary, recently proposed index policies such as DMED [8] and KL-UCB [6, 11], which use indices obtained from KL-based confidence regions, have been shown to be asymptotically optimal.

Unlike most of this family of upper confidence bound algorithms that has been so successful, Thompson Sampling is a policy that uses ideas from Bayesian modelling and yet it solves the fundamentally frequentist problem of regret minimisation. Assume a uniform prior on each parameter μ_a , let $\pi_{a,t}$ denote the posterior distribution for μ_a after the t^{th} round of the algorithm. Let $\theta_{a,t}$ denote a sample from $\pi_{a,t}$; we sometimes refer to $\theta_{a,t}$ as a *Thompson sample*. Thompson sampling is the policy which at time t chooses to observe the action with the highest Thompson sample $\theta_{a,t}$, i.e. it chooses action a with the probability that this action has the highest expected reward under the posterior distribution.

Before Agrawal and Goyal's recent paper [1] Thompson Sampling had been investigated in [7] as the Bayesian Learning Automaton, and in [12] where an optimistic version was also proposed; however these papers only provided weak theoretical guarantees. In [5] extensive numerical experiments were carried out for Thompson Sampling beyond the scope of the Bernoulli bandit setting (to the Generalized Linear Bandit Model) but without any theoretical guarantee at all. Consequently the first finite-time analysis of Thompson Sampling in [1] was a major breakthrough, yet the upper bound for the regret that is shown in this paper scales like (3) and the question of Thompson Sampling's asymptotic optimality was still open.

Meanwhile, there has been a resurgence of interest in Bayesian strategies for bandit problems (see [9] for a review of them). The Bayes-UCB algorithm, an upper confidence bound policy which uses an adaptive quantile of $\pi_{a,t}$ as an optimistic index, was the first Bayesian algorithm to be proved asymptotically optimal. In this paper we are able to show that the same is true for a randomised Bayesian algorithm, Thompson Sampling. Moreover we refer in our analysis to the Bayes-UCB index when introducing the deviation between a Thompson Sample and the corresponding posterior quantile.

Contributions We provide a finite-time regret bound for Thompson Sampling, that follows from (1) and from the result on the expected number of suboptimal draws stated in Theorem 2. More precisely we show the following:

Theorem 1 For every $\epsilon > 0$ there exists a problem-dependent constant $C(\epsilon, \mu_1, \dots, \mu_K)$ such that the regret of Thompson Sampling satisfies:

$$\mathcal{R}(T) \leq (1 + \epsilon) \sum_{a \in A: \mu_a \neq \mu^*} \frac{\Delta_a(\ln(T) + \ln \ln(T))}{K(\mu_a, \mu^*)} + C(\epsilon, \mu_1, \dots, \mu_K).$$

Besides this asymptotically optimal regret bound, we also provide the first numerical experiments that show Thompson Sampling outperforming the current best optimal policies like DMED, KL-UCB or Bayes-UCB. The rest of the paper is structured as follows. Section 2 contains notations or results already used in [1], [6] or [9] that are useful in our finite-time analysis given in Section 3. Numerical experiments are presented in Section 4.

2 Preliminaries

We gather together here some useful preliminaries such as notations not already given in the introduction:

- For the rest of this paper, we assume action 1 is the unique optimal action. Without loss of generality², we can assume that the parameter $\mu = (\mu_1, \dots, \mu_K)$ of the problem is such that $\mu_1 > \mu_2 \geq \dots \geq \mu_K$.
- We shall denote by $S_{a,t}$ the number of successes observed from action a by time t , and denote the empirical mean by:

$$\hat{\mu}_{a,t} := \frac{S_{a,t}}{N_{a,t}}.$$

- In the Bernoulli case, with a uniform prior on the parameters μ_a of the arms, the posterior on arm a at time t is explicitly

$$\pi_{a,t} = \text{Beta}(S_{a,t} + 1, N_{a,t} - S_{a,t} + 1).$$

- Let $F_{a,b}^{\text{Beta}}$ denote the cdf of a Beta distribution with parameters a and b and $F_{j,\mu}^{\text{B}}$ (resp $f_{j,\mu}^{\text{B}}$) the cdf (resp pdf) of a Binomial distribution with parameters j and μ . We recall an important link between Beta and Binomial distribution which was used in both [1] and [9]:

$$F_{a,b}^{\text{Beta}}(y) = 1 - F_{a+b-1,y}^{\text{B}}(a-1)$$

We use this ‘Beta-Binomial trick’ at several stages of our analysis.

- We denote by $u_{a,t}$ (resp. $q_{a,t}$) the KL-UCB (resp. Bayes-UCB) index at time t , and define them by

$$u_{a,t} := \operatorname{argmax}_{x > \frac{S_{a,t}}{N_{a,t}}} \left\{ K \left(\frac{S_{a,t}}{N_{a,t}}, x \right) \leq \frac{\ln(t) + \ln(\ln(T))}{N_{a,t}} \right\}$$

$$q_{a,t} := Q \left(1 - \frac{1}{t \ln(T)}, \pi_{a,t} \right).$$

where $Q(\alpha, \pi)$ denotes the quantile of order α of the distribution π

- A special link between these two indices is shown in [9]: $q_{a,t} < u_{a,t}$.

²In Appendix A of [1] the authors show that adding a second optimal arm can only improve the regret performance of Thompson Sampling.

3 Finite Time Analysis

3.1 Sketch of Analysis

Unlike Agrawal and Goyal’s analysis, which is based on explicit computation of the expectation $\mathbb{E}[N_{2,T}]$, we are more inspired by standard analyses of frequentist index policies. Such policies compute, for each arm a at round t , an index, $l_{a,t}$, from the sequence of rewards observed from a by time t , and choose $A_t = \operatorname{argmax}_a l_{a,t}$. A standard analysis of such a policy aims to bound the number of draws of a suboptimal arm, a , by considering two possible events that might lead to a play of this arm:

- the optimal arm (arm 1) is under-estimated, i.e. $l_{1,t} < \mu_1$;
- the optimal arm is not under-estimated and the suboptimal arm a is drawn.

Taking these to be a good description of when the suboptimal arm is drawn leads to the decomposition

$$\mathbb{E}[N_{a,T}] \leq \sum_{t=1}^T \mathbb{P}(l_{1,t} < \mu_1) + \sum_{t=1}^T \mathbb{P}((l_{a,t} \geq l_{1,t} > \mu_1) \cap (A_t = a))$$

The analysis of an optimistic algorithm then proceeds by showing that the left term (the “under-estimation” term) is $o(\ln(T))$ and the right term is of the form $\frac{1}{K(\mu_a, \mu_1)} \ln(T) + o(\ln(T))$ (or at worst $\frac{2}{\Delta_a^2} \ln(T) + o(\ln(T))$) as in the analysis of UCB1). This style of argument works for example for the KL-UCB algorithm [6] and also for the Bayesian optimistic algorithm Bayes-UCB [9].

However we cannot directly apply this approach to analyse Thompson Sampling, since the sample $\theta_{a,t}$ is not an optimistic estimate of μ_a based on an upper confidence bound. Indeed, even when $\pi_{1,t}$ is well concentrated and therefore close to a Gaussian distribution centred at μ_1 , $\mathbb{P}(\theta_{1,t} < \mu_1)$ is close to $\frac{1}{2}$ and the under-estimation term is not $o(\ln(T))$. Hence we will not compare in our proof the sample $\theta_{a,t}$ to μ_a , but to $\mu_a - \sqrt{6 \ln(t)/N_{a,t}}$ (if $N_{a,t} > 0$) which is the lower bound of an UCB interval. We set the convention that if $N_{a,t} = 0$, $\sqrt{6 \ln(t)/N_{a,t}} = \infty$.

As is observed in [1] the main difficulty in a regret analysis for Thompson Sampling is to control the number of draws of the optimal arm. We provide this control in the form of Proposition 1 whose proof, given in Section 3.3, explores in depth the randomised nature of Thompson Sampling.

Proposition 1 *There exists constants $b = b(\mu_1, \mu_2) \in (0, 1)$ and $C_b = C_b(\mu_1, \mu_2) < \infty$ such that*

$$\sum_{t=1}^{\infty} \mathbb{P}(N_{1,t} \leq t^b) \leq C_b.$$

This proposition tells us that the probability that the algorithm has seen only a small number draws on arm 1 is itself small. As a result we can reduce to analysing the behaviour of the algorithm once it has seen a reasonable number of draws on arm 1, and thus the posterior distribution is well concentrated.

Remark 1 *In general, a result on the regret like $\mathbb{E}[N_{1,t}] \geq t - K \ln(t)$ does not imply a deviation inequality for $N_{1,t}$ (see [13]). Proposition 1 is therefore a strong result, that enables us to adapt the standard analysis mentioned above.*

Using this result, the new decomposition finally yields the following theorem:

Theorem 2 *Let $\epsilon > 0$. With b as in Proposition 1, for every suboptimal arm a , there exist constants $D(\epsilon, \mu_1, \mu_a)$, $N(b, \epsilon, \mu_1, \mu_a)$ and $N_0(b)$ such that:*

$$\mathbb{E}[N_{a,T}] \leq (1 + \epsilon) \frac{\ln(T) + \ln \ln(T)}{K(\mu_a, \mu_1)} + D(\epsilon, \mu_1, \mu_a) + N(b, \epsilon, \mu_1, \mu_a) + N_0(b) + 5 + 2C_b.$$

The constants will be made more explicit in the proofs of Proposition 1 and Theorem 2. The fact that Theorem 2 holds for every $\epsilon > 0$ gives us the asymptotic optimality of Thompson Sampling.

3.2 Proof of Theorem 2

Step 1: Decomposition First we recall the modified decomposition mentioned above:

$$\begin{aligned} \mathbb{E}[N_{a,T}] &\leq \sum_{t=1}^T \mathbb{P} \left(\theta_{1,t} \leq \mu_1 - \sqrt{\frac{6 \ln(t)}{N_{1,t}}} \right) + \sum_{t=1}^T \mathbb{P} \left(\theta_{a,t} > \mu_1 - \sqrt{\frac{6 \ln(t)}{N_{1,t}}}, A_t = a \right) \\ &\leq \sum_{t=1}^T \mathbb{P} \left(\theta_{1,t} \leq \mu_1 - \sqrt{\frac{6 \ln(t)}{N_{1,t}}} \right) \\ &\quad + \sum_{t=1}^T \mathbb{P} \left(\theta_{a,t} > \mu_1 - \sqrt{\frac{6 \ln(t)}{N_{1,t}}}, A_t = a, \theta_{a,t} < q_{a,t} \right) + \sum_{t=1}^T \mathbb{P}(\theta_{a,t} > q_{a,t}) \end{aligned}$$

The sample $\theta_{a,t}$ is not very likely to exceed the quantile of the posterior distribution $q_{a,t}$ we introduced:

$$\sum_{t=1}^T \mathbb{P}(\theta_{a,t} > q_{a,t}) \leq \sum_{t=1}^T \frac{1}{t \ln(T)} \leq \frac{1 + \ln(T)}{\ln(T)} \leq 2$$

where this last inequality follows for $T \geq e$. So finally, using that $u_{a,t} \geq q_{a,t}$,

$$\mathbb{E}[N_{a,t}] \leq \underbrace{\sum_{t=1}^T \mathbb{P} \left(\theta_{1,t} \leq \mu_1 - \sqrt{\frac{6 \ln(t)}{N_{1,t}}} \right)}_A + \underbrace{\sum_{t=1}^T \mathbb{P} \left(u_{a,t} > \mu_1 - \sqrt{\frac{6 \ln(t)}{N_{1,t}}}, A_t = a \right)}_B + 2 \quad (4)$$

Step 2: Bounding term A Dealing with term A boils down to showing a new self-normalized inequality adapted to the randomisation present in each round of the Thompson algorithm.

Lemma 1 *There exists some deterministic constant $N_0(b)$ such that*

$$\sum_{t=1}^{\infty} \mathbb{P} \left(\theta_{1,t} \leq \mu_1 - \sqrt{\frac{6 \ln(t)}{N_{1,t}}} \right) \leq N_0(b) + 3 + C_b < \infty$$

with C_b defined as in Proposition 1.

Proof. Let (U_t) denote a sequence of i.i.d. uniform random variables, and let $\Sigma_{1,s}$ be the sum of the first s rewards from arm 1. In the following, we make the first use of the link between Beta and Binomial distributions:

$$\begin{aligned} \mathbb{P} \left(\theta_{1,t} \leq \mu_1 - \sqrt{\frac{6 \ln(t)}{N_{1,t}}} \right) &= \mathbb{P} \left(U_t \leq F_{S_{1,t}+1, N_{1,t}-S_{1,t}+1}^{\text{Beta}} \left(\mu_1 - \sqrt{\frac{6 \ln(t)}{N_{1,t}}} \right) \right) \\ &= \mathbb{P} \left(\left(U_t \leq 1 - F_{N_{1,t}+1, \mu_1 - \sqrt{\frac{6 \ln(t)}{N_{1,t}}}}^{\text{B}}(S_{1,t}) \right) \cap (N_{1,t} \geq t^b) \right) + \mathbb{P}(N_{1,t} \leq t^b) \\ &= \mathbb{P} \left(\left(F_{N_{1,t}+1, \mu_1 - \sqrt{\frac{6 \ln(t)}{N_{1,t}}}}^{\text{B}}(S_{1,t}) \leq U_t \right) \cap (N_{1,t} \geq t^b) \right) + \mathbb{P}(N_{1,t} \leq t^b) \\ &\leq \mathbb{P} \left(\exists s \in \{t^b \dots t\} : F_{s+1, \mu_1 - \sqrt{\frac{6 \ln(t)}{s}}}^{\text{B}}(\Sigma_{1,s}) \leq U_t \right) + \mathbb{P}(N_{1,t} \leq t^b) \\ &= \sum_{s=\lceil t^b \rceil}^t \mathbb{P} \left(\Sigma_{1,s} \leq (F_{s+1, \mu_1 - \sqrt{\frac{6 \ln(t)}{s}}}^{\text{B}})^{-1}(U_t) \right) + \mathbb{P}(N_{1,t} \leq t^b) \end{aligned}$$

The first term in the final line of this display now deals only with Binomial random variables with large numbers of trials (greater than t^b), and so we can draw on standard concentration techniques to bound this term. Proposition 1 takes care of the second term.

Note that $(F^B)^{-1}_{s+1, \mu_1 - \sqrt{6 \ln(t)/s}}(U_t) \sim \text{Bin}\left(s+1, \mu_1 - \sqrt{6 \ln(t)/s}\right)$ and is independent from $\Sigma_{1,s} \sim \text{Bin}(s, \mu_1)$. For each s , we define two i.i.d. sequences of Bernoulli random variables:

$$(X_{1,l})_l \sim \mathcal{B}\left(\mu_1 - \sqrt{\frac{6 \ln(t)}{s}}\right) \text{ and } (X_{2,l})_l \sim \mathcal{B}(\mu_1),$$

and we let $Z_l := X_{2,l} - X_{1,l}$, another i.i.d. sequence, with mean $\sqrt{\frac{6 \ln(t)}{s}}$. Using these notations,

$$\mathbb{P}\left(\Sigma_{1,s} \leq (F^B)^{-1}_{s+1, \mu_1 - \sqrt{\frac{6 \ln(t)}{s}}}(U_t)\right) \leq \mathbb{P}\left(\sum_{l=1}^s Z_l \leq 1\right) = \mathbb{P}\left(\sum_{l=1}^s \left(Z_l - \sqrt{\frac{6 \ln(t)}{s}}\right) \leq -\left(\sqrt{6s \ln(t)} - 1\right)\right).$$

Let $N_0(b)$ be such that if $t \geq N_0(b)$, $\sqrt{6t^b \ln(t)} - 1 > \sqrt{5t^b \ln(t)}$. For $t \geq N_0(b)$, we can apply Hoeffding's inequality to the bounded martingale difference sequence $Z'_l = Z_l - \sqrt{6 \ln(t)/s}$ to get

$$\mathbb{P}\left(\Sigma_{1,s} < (F^B)^{-1}_{s+1, \mu_1 - \sqrt{\frac{6 \ln(t)}{s}}}(U_t)\right) \leq \exp\left(-2 \frac{(\sqrt{5s \ln(t)})^2}{4s}\right) = e^{-\frac{5}{2} \ln(t)} = \frac{1}{t^{\frac{5}{2}}}.$$

We conclude that

$$\sum_{t=1}^{\infty} \mathbb{P}\left(\theta_1(t) < \mu_1 - \sqrt{\frac{6 \ln(t)}{N_{1,t}}}\right) \leq N_0(b) + \sum_{t=1}^{\infty} \frac{1}{t^{\frac{3}{2}}} + C_b \leq N_0(b) + 3 + C_b.$$

■

Step 3: Bounding Term B We specifically show that:

Lemma 2 For all $a = 2, \dots, K$, for any $\epsilon > 0$ there exist $N(b, \epsilon, \mu_1, \mu_a), D(\epsilon, \mu_1, \mu_a) > 0$ such that for all $T > N(b, \epsilon, \mu_1, \mu_a)$

$$(B) \leq (1 + \epsilon) \frac{\ln(T) + \ln \ln(T)}{K(\mu_a, \mu_1)} + D(\epsilon, \mu_1, \mu_a).$$

Proof. First rewrite term B so that we can apply Proposition 1:

$$\begin{aligned} (B) &\leq \sum_{t=1}^T \mathbb{P}\left(u_{a,t} > \mu_1 - \sqrt{\frac{6 \ln(t)}{N_{1,t}}}, A_t = a, N_{1,t} \geq t^b\right) + \sum_{t=1}^T \mathbb{P}(N_{1,t} \leq t^b) \\ &\leq \sum_{t=1}^T \mathbb{P}\left(u_{a,t} > \mu_1 - \sqrt{\frac{6 \ln(t)}{t^b}}, A_t = a\right) + C_b \end{aligned}$$

For ease of notation we introduce

$$\begin{aligned} K^+(x, y) &:= K(x, y) \mathbf{1}_{(x \leq y)}, \quad f_T(t) := \ln(t) + \ln(\ln(T)) \\ \beta_t &= \sqrt{\frac{6 \ln(t)}{t^b}}, \text{ and } K_{T,a}(\epsilon) = (1 + \epsilon) \frac{\ln(T) + \ln \ln(T)}{K(\mu_a, \mu_1)}. \end{aligned}$$

Now

$$(u_{a,t} \geq \alpha) = (N_{2,t} K^+(\hat{\mu}_{2, N_{2,t}}, \alpha) \leq f_T(t))$$

and so summing over the values of $N_{2,t}$ and inverting the sums we get

$$\begin{aligned} \sum_{t=1}^T \mathbb{P}(u_{a,t} > \mu_1 - \beta_t, A_t = a) &= \mathbb{E} \left[\sum_{s=1}^{\lfloor K_{T,a} \rfloor} \sum_{t=s}^T \mathbf{1}_{(sK^+(\hat{\mu}_{a,s}, \mu_1 - \beta_t) \leq f_T(t))} \mathbf{1}_{(A_t = a, N_{2,t} = s)} \right] \\ &+ \mathbb{E} \left[\sum_{s=\lfloor K_{T,a} \rfloor + 1}^T \sum_{t=s}^T \mathbf{1}_{(sK^+(\hat{\mu}_{a,s}, \mu_1 - \beta_t) \leq f_T(t))} \mathbf{1}_{(A_t = a, N_{2,t} = s)} \right]. \end{aligned}$$

As $y \mapsto K^+(\hat{\mu}_{a,s}, y)$ is increasing and $t \mapsto \beta_t$ is decreasing for $t \geq e^{1/b}$, for T such that

$$K_{T,a}(\epsilon) \geq e^{1/b}, \quad (5)$$

we have that if $t \geq K_{T,a}(\epsilon)$,

$$\mathbf{1}_{(sK^+(\hat{\mu}_{a,s}, \mu_1 - \beta_t) \leq f_T(t))} \leq \mathbf{1}_{(sK^+(\hat{\mu}_{a,s}, \mu_1 - \beta_{K_{T,a}}) \leq f_T(t))} \leq \mathbf{1}_{(sK^+(\hat{\mu}_{a,s}, \mu_1 - \beta_{K_{T,a}}) \leq f_T(T))}.$$

and therefore,

$$\begin{aligned} \sum_{t=1}^T \mathbb{P}(u_{a,t} > \mu_1 - \beta_t, A_t = a) &\leq \mathbb{E} \left[\sum_{s=1}^{\lfloor K_{T,a} \rfloor} \sum_{t=s}^T \mathbf{1}_{(A_t = a, N_{2,t} = s)} \right] \\ &+ \mathbb{E} \left[\sum_{s=\lfloor K_{T,a} \rfloor + 1}^T \mathbf{1}_{(sK^+(\hat{\mu}_{a,s}, \mu_1 - \beta_{K_{T,a}}) \leq f_T(T))} \sum_{t=s}^T \mathbf{1}_{(A_t = a, N_{2,t} = s)} \right]. \end{aligned}$$

Given that $\sum_{t=s}^T \mathbf{1}_{(A_t = a, N_{2,t} = s)} \leq 1$ for all s , the first term is upper bounded by $K_{a,T}$, whereas the second is upper bounded by

$$\mathbb{E} \left[\sum_{s=\lfloor K_{T,a} \rfloor + 1}^T \mathbf{1}_{(K_{T,a} K^+(\hat{\mu}_{a,s}, \mu_1 - \beta_{K_{T,a}}) \leq f_T(T))} \right]$$

So, for T satisfying (5),

$$(B) \leq K_{T,a} + \sum_{s=\lfloor K_{T,a} \rfloor + 1}^T \mathbb{P} \left(K^+(\hat{\mu}_{a,s}, \mu_1 - \beta_{K_{T,a}}) \leq \frac{K(\mu_a, \mu_1)}{1 + \epsilon} \right).$$

Using the convexity of $K^+(\hat{\mu}_{a,s}, \cdot)$, we can show that

$$K^+(\hat{\mu}_{a,s}, \mu_1) \leq K^+(\hat{\mu}_{a,s}, \mu_1 - \beta_{K_{a,T}}) + \frac{2}{\mu_1(1 - \mu_1)} \beta_{K_{a,T}}.$$

If $K^+(\hat{\mu}_{a,s}, \mu_1 - \beta_{K_{T,a}}) \leq K(\mu_a, \mu_1)/(1 + \epsilon)$, then

$$\begin{aligned} K^+(\hat{\mu}_{a,s}, \mu_1) &\leq \frac{K(\mu_a, \mu_1)}{1 + \epsilon} + \frac{2}{\mu_1(1 - \mu_1)} \beta_{K_{a,T}} \\ &\leq \frac{K(\mu_a, \mu_1)}{1 + \epsilon/2} \end{aligned} \quad (6)$$

where the last inequality (6) holds for large enough T . There exists a deterministic constant $N = N(b, \epsilon, \mu_1, \mu_a)$ such that for all $T \geq N$ both (5) and (6) are satisfied. Hence, for all $T \geq N$

$$(B) \leq K_{T,a} + \sum_{s=\lfloor K_{T,a} \rfloor + 1}^T \mathbb{P} \left(K^+(\hat{\mu}_{a,s}, \mu_1) \leq \frac{K(\mu_a, \mu_1)}{1 + \frac{\epsilon}{2}} \right).$$

Since this last sum is bounded above explicitly by some constant $D(\epsilon, \mu_1, \mu_a)$ in [11] we have proved the lemma. To be explicit, $D(\epsilon, \mu_1, \mu_a) = \frac{(1 + \epsilon/2)^2}{\epsilon^2(\min(\mu_a(1 - \mu_a); \mu_1(1 - \mu_1)))^2}$. ■

Conclusion: The result now follows from Lemmas 1, 2 and inequality (4).

3.3 Proof of Proposition 1

Since we focus on the number of draws of the optimal arm, let τ_j be the occurrence of the j^{th} play of the optimal arm (with $\tau_0 := 0$). Let $\xi_j := (\tau_{j+1} - 1) - \tau_j$: this random variable measures the number of time steps between the j^{th} and the $(j+1)^{\text{th}}$ play of the optimal arm, and so $\sum_{a=2}^K N_{a,t} = \sum_{j=0}^{N_{1,t}} \xi_j$.

For each suboptimal arm, a relevant quantity is $C_a = \frac{32}{(\mu_1 - \mu_a)^2}$ and let $C = \max_{a \neq 1} C_a = \frac{32}{(\mu_1 - \mu_2)^2}$. We also introduce $\delta_a = \frac{\mu_1 - \mu_a}{2}$ and let $\delta = \delta_2$.

Step 1: Initial Decomposition of Summands First we use a union bound on the summands to extract the tails of the random variables ξ_j :

$$\begin{aligned} \mathbb{P}(N_{1,t} \leq t^b) &= \mathbb{P}\left(\sum_{a=2}^K N_{a,t} \geq t - t^b\right) \\ &\leq \mathbb{P}(\exists j \in \{0, \dots, \lfloor t^b \rfloor\} : \xi_j \geq t^{1-b} - 1) \\ &\leq \sum_{j=0}^{\lfloor t^b \rfloor} \mathbb{P}(\xi_j \geq t^{1-b} - 1) \end{aligned} \quad (7)$$

This means that there exists a time range of length $t^{1-b} - 1$ during which only suboptimal arms are played. In the case of two arms this implies that the (unique) suboptimal arm is played $\lceil \frac{t^{1-b}-1}{2} \rceil$ times during the first half of this time range. Thus its posterior becomes well concentrated around its mean with high probability, and we can use this fact to show that the probability the suboptimal action is chosen a further $\lceil \frac{t^{1-b}-1}{2} \rceil$ times in a row is very small.

In order to generalise this approach we introduce a notion of a *saturated*, suboptimal action:

Definition 1 Let t be fixed. For any $a \neq 1$, an action a is said to be saturated at time s if it has been chosen at least $C_a \ln(t)$ times. That is $N_{a,s} \geq C_a \ln(t)$. We shall say that it is unsaturated otherwise. Furthermore at any time we call a choice of an unsaturated, suboptimal action an interruption.

We want to study the event $E_j = \{\xi_j \geq t^{1-b} - 1\}$. We introduce the interval $\mathcal{I}_j = \{\tau_j, \tau_j + \lceil t^{1-b} - 1 \rceil\}$ (included in $\{\tau_j, \tau_{j+1}\}$ on E_j) and begin by decomposing it into K subintervals:

$$\mathcal{I}_{j,l} := \left\{ \tau_j + \left\lceil \frac{(l-1)(t^{1-b}-1)}{K} \right\rceil, \tau_j + \left\lceil \frac{l(t^{1-b}-1)}{K} \right\rceil \right\}, \quad l = 1, \dots, K.$$

Now for each interval $\mathcal{I}_{j,l}$, we introduce:

- $F_{j,l}$: the event that by the end of the interval $\mathcal{I}_{j,l}$ at least l suboptimal actions are saturated;
- $n_{j,l}$: the number of interruptions during this interval.

We use the following decomposition to bound the probability of the event E_j :

$$\mathbb{P}(E_j) = \mathbb{P}(E_j \cap F_{j,K-1}) + \mathbb{P}(E_j \cap F_{j,K-1}^c) \quad (8)$$

To bound both probabilities, we will need the fact, stated in Lemma 3, that the probability of $\theta_{1,s}$ being smaller than $\mu_2 + \delta$ during a long subinterval of \mathcal{I}_j is small. This follows from the fact that the posterior on the optimal arm is always Beta($S_{1,\tau_j} + 1, j - S_{1,\tau_j} + 1$) on \mathcal{I}_j : hence, when conditioned on S_{1,τ_j} , $\theta_{1,s}$ is an i.i.d. sequence with non-zero support above $\mu_2 + \delta$, and thus is unlikely to remain below $\mu_2 + \delta$ for a long time period. This idea is also an important tool in the analysis of Thompson Sampling in [1].

Lemma 3 $\exists \lambda_0 = \lambda_0(\mu_1, \mu_2) > 1$ such that for $\lambda \in]1, \lambda_0[$, for every (random) interval \mathcal{J} included in \mathcal{I}_j and for every positive function f , one has

$$\mathbb{P}(\{\forall s \in \mathcal{J}, \theta_{1,s} \leq \mu_2 + \delta\} \cap \{|\mathcal{J}| \geq f(t)\}) \leq (\alpha_{\mu_1, \mu_2})^{f(t)} + C_{\lambda_0, \lambda} \frac{1}{f(t)^\lambda} e^{-j d_{\lambda, \mu_1, \mu_2}}$$

where $C_{\lambda, \mu_1, \mu_2}, d_{\lambda, \mu_1, \mu_2} > 0$, and $\alpha_{\mu_1, \mu_2} = (1/2)^{1-\mu_2-\delta}$.

The proof of this important lemma will be postponed to section 3.4 and all the constants are explicitly defined there. Another keypoint in the proof is the fact that a sample from a saturated suboptimal arm cannot fall too far from its true mean. The following lemma is easily adapted from Lemma 2 in [1].

Lemma 4

$$\mathbb{P}(\exists s \leq t, \exists a \neq 1 : \theta_{a,s} > \mu_a + \delta_a, N_{a,s} > C_a \ln(t)) \leq \frac{2(K-1)}{t^2}.$$

Step 2: Bounding $\mathbb{P}(E_j \cap F_{j,K-1})$ On the event $E_j \cap F_{j,K-1}$, only saturated suboptimal arms are drawn on the interval $\mathcal{I}_{j,K}$. Using the concentration results for samples of these arms in Lemma 4, we get

$$\begin{aligned} \mathbb{P}(E_j \cap F_{j,K-1}) &\leq \mathbb{P}(\{\exists s \in \mathcal{I}_{j,K}, a \neq 1 : \theta_{a,s} > \mu_a + \delta\} \cap E_j \cap F_{j,K-1}) \\ &\quad + \mathbb{P}(\{\forall s \in \mathcal{I}_{j,K}, a \neq 1 : \theta_{a,s} \leq \mu_a + \delta_a\} \cap E_j \cap F_{j,K-1}) \\ &\leq \mathbb{P}(\exists s \leq t, a \neq 1 : \theta_{a,s} > \mu_a + \delta_a, N_{a,t} > C_a \ln(t)) \\ &\quad + \mathbb{P}(\{\forall s \in \mathcal{I}_{j,K}, a \neq 1 : \theta_{a,s} \leq \mu_2 + \delta\} \cap E_j \cap F_{j,K-1}) \\ &\leq \frac{2(K-1)}{t^2} + \mathbb{P}(\theta_{1,s} \leq \mu_2 + \delta, \forall s \in \mathcal{I}_{j,K}). \end{aligned}$$

The last inequality comes from the fact that if arm 1 is not drawn, the sample $\theta_{1,s}$ must be smaller than some sample $\theta_{a,s}$ and therefore smaller than $\mu_2 + \delta$. Since $\mathcal{I}_{j,K}$ is an interval in \mathcal{I}_j of size $\left\lceil \frac{t^{1-b}-1}{K} \right\rceil$ we get using Lemma 3, for some fixed $\lambda \in]1, \lambda_0[$,

$$\mathbb{P}(\theta_{1,s} \leq \mu_2 + \delta, \forall s \in \mathcal{I}_{j,K}) \leq (\alpha_{\mu_1, \mu_2})^{\frac{t^{1-b}-1}{K}} + C_{\lambda, \mu_1, \mu_2} \left(\frac{t^{1-b}-1}{K} \right)^{-\lambda} e^{-j d_{\lambda, \mu_1, \mu_2}} =: g(\mu_1, \mu_2, b, j, t). \quad (9)$$

Hence we have show that

$$\mathbb{P}(E_j \cap F_{j,K-1}) \leq \frac{2(K-1)}{t^2} + g(\mu_1, \mu_2, b, j, t), \quad (10)$$

and choosing b such that $b < 1 - \frac{1}{\lambda}$, the following hypothesis on g holds:

$$\sum_{t \geq 1} \sum_{j \leq t^b} g(\mu_1, \mu_2, b, j, t) < +\infty.$$

Step 3: Bounding $\mathbb{P}(E_j \cap F_{j,K-1}^c)$ We show through an induction that for all $2 \leq l \leq K$, if t is larger than some deterministic constant $N_{\mu_1, \mu_2, b}$ specified in the base case,

$$\mathbb{P}(E_j \cap F_{j,l-1}^c) \leq (l-2) \left(\frac{2(K-1)}{t^2} + f(\mu_1, \mu_2, b, j, t) \right)$$

for some function f such that $\sum_{t \geq 1} \sum_{1 \leq j \leq t^b} f(\mu_1, \mu_2, b, j, t) < \infty$. For $l = K$ we get

$$\mathbb{P}(E_j \cap F_{j,K-1}^c) \leq (K-2) \left(\frac{2(K-1)}{t^2} + f(\mu_1, \mu_2, b, j, t) \right). \quad (11)$$

Step 4: The Base Case of the induction Note that on the event E_j only suboptimal arms are played during $\mathcal{I}_{j,1}$. Hence at least one suboptimal arm must be played $\lceil \frac{t^{1-b}-1}{K^2} \rceil$ times.

There exists some deterministic constant $N_{\mu_1, \mu_2, b}$ such that for $t \geq N_{\mu_1, \mu_2, b}$, $\lceil \frac{t^{1-b}-1}{K^2} \rceil \geq C \ln(t)$ (the constant depends only on μ_1 and μ_2 because $C = C_2$). So when $t \geq N_{\mu_1, \mu_2, b}$, at least one suboptimal arm must be saturated by the end of $\mathcal{I}_{j,1}$. Hence, for $t \geq N_{\mu_1, \mu_2, b}$

$$\mathbb{P}(E_j \cap F_{j,1}^c) = 0.$$

This concludes the base case.

Step 5: The Induction As an inductive hypothesis we assume that for some $2 \leq l \leq K-1$ if $t \geq N_{\mu_1, \mu_2, b}$ then

$$\mathbb{P}(E_j \cap F_{j,l-1}^c) \leq (l-2) \left(\frac{2(K-1)}{t^2} + f(\mu_1, \mu_2, b, j, t) \right).$$

Then, making use of the inductive hypothesis,

$$\begin{aligned} \mathbb{P}(E_j \cap F_{j,l}^c) &\leq \mathbb{P}(E_j \cap F_{j,l-1}^c) + \mathbb{P}(E_j \cap F_{j,l}^c \cap F_{j,l-1}) \\ &\leq (l-2) \left(\frac{2(K-1)}{t^2} + f(\mu_1, \mu_2, b, j, t) \right) + \mathbb{P}(E_j \cap F_{j,l}^c \cap F_{j,l-1}). \end{aligned}$$

To complete the induction we therefore need to show that:

$$\mathbb{P}(E_j \cap F_{j,l}^c \cap F_{j,l-1}) \leq \frac{2(K-1)}{t^2} + f(\mu_1, \mu_2, b, j, t). \quad (12)$$

On the event $(E_j \cap F_{j,l}^c \cap F_{j,l-1})$, there are exactly $l-1$ saturated arms at the beginning of interval $\mathcal{I}_{j,l}$ and no new arm is saturated during this interval. As a result there cannot be more than $KC \ln(t)$ interruptions during this interval, and so we have

$$\mathbb{P}(E_j \cap F_{j,l}^c \cap F_{j,l-1}) \leq \mathbb{P}(E_j \cap F_{j,l-1} \cap \{n_{j,l} \leq KC \ln(t)\}).$$

Let \mathcal{S}_l denote the set of saturated arms at the end of $\mathcal{I}_{j,l}$ and introduce the following decomposition:

$$\begin{aligned} &\mathbb{P}(E_j \cap F_{j,l-1} \cap \{n_{j,l} \leq KC \ln(t)\}) \\ &\leq \underbrace{\mathbb{P}(\{\exists s \in \mathcal{I}_{j,l}, a \in \mathcal{S}_{l-1} : \theta_{a,s} > \mu_a + \delta_a\} \cap E_j \cap F_{j,l-1})}_A \\ &\quad + \underbrace{\mathbb{P}(\{\forall s \in \mathcal{I}_{j,l}, a \in \mathcal{S}_{l-1} : \theta_{a,s} \leq \mu_a + \delta_a\} \cap E_j \cap F_{j,l-1} \cap \{n_{j,l} \leq KC \ln(t)\})}_B. \end{aligned}$$

Clearly, using Lemma 4:

$$(A) \leq \mathbb{P}(\exists s \leq t, \exists a \neq 1 : \theta_{a,s} > \mu_a + \delta_a, N_{a,s} > C_a \ln(t)) \leq \frac{2(K-1)}{t^2}.$$

To deal with term (B), we introduce for k in $\{0, \dots, n_{j,l}-1\}$ the random intervals \mathcal{J}_k as the time range between the k^{th} and $(k+1)^{\text{st}}$ interruption in $\mathcal{I}_{j,l}$. For $k \geq n_{j,l}$ we set $\mathcal{J}_k = \emptyset$. Note that on the event in the probability (B) there is a subinterval of $\mathcal{I}_{j,l}$ of length $\lceil \frac{t^{1-b}-1}{CK^2 \ln(t)} \rceil$ during which there are no interruptions. Moreover on this subinterval of $\mathcal{I}_{j,l}$, for all $a \neq 1$, $\theta_{a,s} \leq \mu_a + \delta_a$. (This holds for unsaturated arms as well as for saturated arms since their samples are smaller than the maximum sample

of a saturated arm.) Therefore,

$$\begin{aligned}
(B) &\leq \mathbb{P} \left(\left\{ \exists k \in \{0, \dots, n_{j,l}\} : |\mathcal{J}_k| \geq \frac{t^{1-b} - 1}{CK^2 \ln(t)} \right\} \cap \{\forall s \in \mathcal{I}_{j,l}, a \in \mathcal{S}_{l-1}1 : \theta_{a,s} \leq \mu_2 + \delta\} \cap E_j \cap F_{j,l-1} \right) \\
&\leq \sum_{k=1}^{KC \ln(t)} \mathbb{P} \left(\left\{ |\mathcal{J}_k| \geq \frac{t^{1-b} - 1}{CK^2 \ln(t)} \right\} \cap \{\forall s \in \mathcal{J}_k, a \neq 1 : \theta_{a,s} \leq \mu_2 + \delta\} \cap E_j \right) \\
&\leq \sum_{k=1}^{KC \ln(t)} \mathbb{P} \left(\left\{ |\mathcal{J}_k| \geq \frac{t^{1-b} - 1}{CK^2 \ln(t)} \right\} \cap \{\forall s \in \mathcal{J}_k, \theta_{1,s} \leq \mu_2 + \delta\} \right) \tag{13}
\end{aligned}$$

Now, we have to bound the probability that $\theta_{1,s} \leq \mu_2 + \delta$ for all s in an interval of size $\frac{t^{1-b}-1}{CK^2 \ln(t)}$ included in \mathcal{I}_j . So we apply Lemma 3 to get:

$$(B) \leq CK \ln(t) (\alpha_{\mu_1, \mu_2})^{\frac{t^{1-b}-1}{CK^2 \ln(t)}} + C_{\lambda, \mu_1, \mu_2} \frac{CK \ln(t)}{\left(\frac{t^{1-b}-1}{CK^2 \ln(t)}\right)^\lambda} e^{-j d_{\lambda, \mu_1, \mu_2}} := f(\mu_1, \mu_2, b, j, t).$$

Choosing the same b as in (9), we get that $\sum_{t \geq 1} \sum_{1 \leq j \leq t^b} f(\mu_1, \mu_2, b, j, t) < +\infty$. It follows that for this value of b , (12) holds and the induction is complete.

Step 8: Conclusion Let b be the constant chosen in Step 2. From the decomposition (8) and the two upper bounds (10) and (11), we get, for $t \geq N_{\mu_1, \mu_2, b}$:

$$\mathbb{P}(E_j) \leq (K-2) \left(\frac{2(K-1)}{t^2} + f(\mu_1, \mu_2, b, j, t) \right) + \frac{2(K-1)}{t^2} + g(\mu_1, \mu_2, b, j, t).$$

Recalling (7), summing over the possible values of j and t we obtain:

$$\sum_{t \geq 1} \mathbb{P}(N_{1,t} \leq t^b) \leq N_{\mu_1, \mu_2, b} + 2(K-1)^2 \sum_{t \geq 1} \frac{1}{t^{2-b}} + \sum_{t \geq 1} \sum_{j=1}^{t^b} [Kf(\mu_1, \mu_2, b, j, t) + g(\mu_1, \mu_2, b, j, t)] < C_{\mu_1, \mu_2, b}$$

for some constant $C_{\mu_1, \mu_2, b} < \infty$.

3.4 Proof of Lemma 3

On the interval the \mathcal{J} (included in \mathcal{I}_j by hypothesis), the posterior distribution $\pi_{1,s} = \pi_{1,\tau_j}$ is fixed and $\theta_{1,s}$ is, when conditioned on S_{1,τ_j} , an i.i.d. sequence with common distribution $\text{Beta}(S_{1,\tau_j} + 1, j - S_{1,\tau_j} + 1)$. Hence,

$$\mathbb{P}(\theta_{1,s} \leq \mu_2 + \delta | s \in \mathcal{J}) = F_{(S_{1,\tau_j} + 1, j - S_{1,\tau_j} + 1)}^{\text{Beta}}(\mu_2 + \delta) = 1 - F_{(j+1, \mu_2 + \delta)}^B(S_{1,\tau_j})$$

where we use the link between the tail of Beta and Bernoulli distribution mentioned above. Using the independence of the $\theta_{1,s}$ gives

$$\mathbb{P}(\forall s \in \mathcal{J}, \theta_{1,s} \leq \mu_2 + \delta | S_{1,\tau_j}) = \left(1 - F_{(j+1, \mu_2 + \delta)}^B(S_{1,\tau_j})\right)^{|\mathcal{J}|} \leq \left(1 - F_{(j+1, \mu_2 + \delta)}^B(S_{1,\tau_j})\right)^{f(t)}$$

Finally

$$\mathbb{P}(\forall s \in \mathcal{J}, \theta_{1,s} \leq \mu_2 + \delta) = \mathbb{E} \left[\mathbb{P}(\forall s \in \mathcal{J}, \theta_{1,s} \leq \mu_2 + \delta | S_{1,\tau_j}) \right] \leq \mathbb{E} \left[\left(1 - F_{(j+1, \mu_2 + \delta)}^B(S_{1,\tau_j})\right)^{f(t)} \right]$$

An exact computation of this expectation leads to

$$\mathbb{E} \left[\left(1 - F_{(j+1, \mu_2 + \delta)}^B(S_{1,\tau_j})\right)^{f(t)} \right] = \sum_{s=0}^j \left(1 - F_{(j+1, \mu_2 + \delta)}^B(s)\right)^{f(t)} f_{j, \mu_1}^B(s)$$

To simplify notation, from now on let $y = \mu_2 + \delta$. Using, as in [1], that $F_{j+1,y}^B(s) = (1-y)F_{j,y}^B(s) + yF_{j,y}^B(s-1) \geq (1-y)F_{j,y}^B(s)$, we get:

$$(1 - F_{(j+1,y)}^B(s))^{f(t)} \leq \exp\left(-f(t)F_{(j+1,y)}^B(s)\right) \leq \exp\left(-f(t)(1-y)F_{(j,y)}^B(s)\right)$$

Therefore,

$$\mathbb{E}\left[(1 - F_{(j+1,\mu_2+\delta)}^B(S_{1,\tau_j}))^{f(t)}\right] \leq \sum_{s=0}^j \exp\left(-f(t)(1-y)F_{(j,y)}^B(s)\right) f_{j,\mu_1}^B(s)$$

Using the fact that for $s \geq \lceil yj \rceil$, $F_{j,y}^B(s) \geq \frac{1}{2}$ (since the median of a binomial distribution with parameters j and y is $\lceil yj \rceil$ or $\lfloor yj \rfloor$), we get

$$\begin{aligned} & \mathbb{E}\left[(1 - F_{(j+1,\mu_2+\delta)}^B(S_{1,\tau_j}))^{f(t)}\right] \\ & \leq \sum_{s=0}^{\lfloor jy \rfloor} \exp\left(-f(t)(1-y)F_{(j,y)}^B(s)\right) f_{j,\mu_1}^B(s) + \sum_{s=\lceil jy \rceil}^j \left(\frac{1}{2}\right)^{(1-y)f(t)} f_{j,\mu_1}^B(s) \\ & \leq \underbrace{\sum_{s=0}^{\lfloor jy \rfloor} \exp\left(-f(t)(1-y)F_{(j,y)}^B(s)\right) f_{j,\mu_1}^B(s)}_E + \left(\frac{1}{2}\right)^{(1-y)f(t)}. \end{aligned}$$

It is easy to show that for every $\lambda > 1, \forall x > 0, x^\lambda \exp(-x) \leq \left(\frac{\lambda}{e}\right)^\lambda$. This allows us to upper-bound the exponential for all $\lambda > 1$, using $C_\lambda = \left(\frac{\lambda}{e}\right)^\lambda$, by:

$$(E) \leq \frac{C_\lambda}{(f(t)(1-y))^\lambda} \sum_{s=0}^{\lfloor jy \rfloor} \frac{f_{j,\mu_1}^B(s)}{\left(F_{(j,y)}^B(s)\right)^\lambda} \leq \frac{C_\lambda}{(f(t)(1-y))^\lambda} \sum_{s=0}^{\lfloor jy \rfloor} \frac{f_{j,\mu_1}^B(s)}{\left(f_{(j,y)}^B(s)\right)^\lambda}$$

Now, inspired by Agrawal and Goyal's work (proof of Lemma 3) we compute:

$$\begin{aligned} \frac{f_{j,\mu_1}^B(s)}{\left(f_{(j,y)}^B(s)\right)^\lambda} &= \frac{\binom{j}{s} \mu_1^s (1-\mu_1)^{j-s}}{\binom{j}{s}^\lambda (y^\lambda)^s ((1-y)^\lambda)^{j-s}} \leq \frac{\mu_1^s (1-\mu_1)^{j-s}}{(y^\lambda)^s ((1-y)^\lambda)^{j-s}} \\ &= \left(\frac{1-\mu_1}{(1-y)^\lambda}\right)^j \left(\frac{\mu_1(1-y)^\lambda}{y^\lambda(1-\mu_1)}\right)^s \end{aligned}$$

Let $R_\lambda(\mu_1, y) = \frac{\mu_1(1-y)^\lambda}{y^\lambda(1-\mu_1)}$. There exists some $\lambda_1 > 1$ such that, if $\lambda < \lambda_1$, $R_\lambda > 1$. More precisely,

$$R_\lambda > 1 \Leftrightarrow \frac{\mu_1}{1-\mu_1} > \left(\frac{y}{1-y}\right)^\lambda \Leftrightarrow \ln\left(\frac{\mu_1}{1-\mu_1}\right) > \lambda \ln\left(\frac{y}{1-y}\right)$$

and so

$$\lambda_1(\mu_1, y) = \begin{cases} \frac{\ln\left(\frac{\mu_1}{1-\mu_1}\right)}{\ln\left(\frac{y}{1-y}\right)} & \text{if } y > \frac{1}{2} \\ +\infty & \text{if } y < \frac{1}{2} \end{cases}$$

For $1 < \lambda < \lambda_1$:

$$\begin{aligned} \sum_{s=0}^{\lfloor jy \rfloor} \frac{f_{j,\mu_1}^B(s)}{(f_{(j,\mu_2+\delta)}^B(s))^\lambda} &\leq \left(\frac{1-\mu_1}{(1-y)^\lambda} \right)^j \sum_{s=0}^{\lfloor jy \rfloor} R_\lambda^s = \left(\frac{1-\mu_1}{(1-y)^\lambda} \right)^j \frac{R_\lambda^{\lfloor jy \rfloor + 1} - 1}{R_\lambda - 1} \\ &\leq \left(\frac{1-\mu_1}{(1-y)^\lambda} \right)^j \frac{R_\lambda}{R_\lambda - 1} R_\lambda^{jy} = \frac{R_\lambda}{R_\lambda - 1} \left(\frac{1-\mu_1}{(1-y)^\lambda} \right)^{j-jy} \left(\frac{\mu_1}{y^\lambda} \right)^{jy} \\ &= \frac{R_\lambda}{R_\lambda - 1} e^{-jd_\lambda(y,\mu_1)} \end{aligned}$$

where $d_\lambda(y, \mu_1) = y \ln \left(\frac{y^\lambda}{\mu_1} \right) + (1-y) \ln \left(\frac{(1-y)^\lambda}{1-\mu_1} \right)$. Rearranging we can write

$$d_\lambda(y, \mu_1) = \lambda [y \ln(y) + (1-y) \ln(1-y)] - [y \ln(\mu_1) + (1-y) \ln(1-\mu_1)]$$

which is an affine function of λ with negative slope ($y \ln(y) + (1-y) \ln(1-y) < 0$ for all $y \in (0, 1)$) and $d_1(y, \mu_1) = K(y, \mu_1) > 0$. Hence, for fixed $0 < y < \mu_1 \leq 1$ this function is positive whenever

$$\lambda < \frac{y \ln(\mu_1) + (1-y) \ln(1-\mu_1)}{y \ln(y) + (1-y) \ln(1-y)} =: \lambda_2(\mu_1, y).$$

Clearly, $\lambda_2(\mu_1, y) > 1$ and we choose $\lambda_0 = \min(\lambda_1, \lambda_2)$. After some calculation one can show that $\lambda_2 \leq \lambda_1$, and therefore that

$$\lambda_0(\mu_1, \mu_2) = \lambda_2(\mu_1, \mu_2 + \delta) = 1 + \frac{K(\mu_2 + \delta, \mu_1)}{(\mu_2 + \delta) \ln \frac{1}{\mu_2 + \delta} + (1 - \mu_2 - \delta) \ln \frac{1}{1 - \mu_2 - \delta}}.$$

To obtain the constants used in the statement of the lemma we define $d_{\lambda, \mu_1, \mu_2} := d_\lambda(y, \mu_1)$

$$C_{\lambda, \mu_1, \mu_2} := C_{\lambda_0} (1 - \mu_2 - \delta)^{-\lambda} \frac{R_\lambda}{1 - R_\lambda}.$$

This concludes the proof.

4 Experiments

We illustrate here the performance of Thompson Sampling on numerical experiments with Bernoulli rewards. First we compare in terms of cumulative regret up to horizon $T = 10000$ Thompson Sampling to UCB, KL-UCB and Bayes-UCB in two different two-arms problem, one with small rewards and the other with high rewards, with different gaps between the parameters of the arms. Figure 1 shows Thompson Sampling always outperforms KL-UCB and also Bayes-UCB for large horizons. The three optimal policies are significantly better than UCB, even for small horizons.

Figure 2 displays for several algorithms an estimation of the distribution of the cumulative regret based on $N = 50000$ trials, for a horizon $T = 20000$ in a 10-armed bandit problem with

$$\mu = (0.1, 0.05, 0.05, 0.05, 0.02, 0.02, 0.02, 0.01, 0.01, 0.01).$$

The first two algorithms are variants of UCB. Of these the UCB-V algorithm is close to the index policy to which Thompson Sampling is compared in [5] in the Bernoulli setting, but this policy is not known to be optimal. This algorithm incorporates an estimation of the variance of the rewards in the index which is defined to be, for an arm that have produced k rewards in n draws,

$$\frac{k}{n} + \sqrt{\frac{2 \log(t)}{n} \frac{k}{n} \left(1 - \frac{k}{n} \right)} + \frac{3 \log(t)}{n}$$

The other algorithms displayed in Figure 2 have a mean regret closer (sometimes smaller) than the lower bound (which is only asymptotic), and among them, Thompson is the best. It is also the easiest optimal policy to implement, since the optimization problem solved in KL-UCB and even the computation of the quantiles in Bayes-UCB are more costly than producing one sample from the posterior for each arm.




Figure 1: Cumulated regret for the two armed-bandit problem with $\mu_1 = 0.2, \mu_2 = 0.25$ (left) and $\mu_1 = 0.8, \mu_2 = 0.9$ (right). Regret is estimated as an average over $N = 20000$ trials.




Figure 2: Regret of the various algorithms as a function of time (on a log scale). On each graph, the red dashed line shows the lower bound, the solid bold curve corresponds to the mean regret while the dark and light shaded regions show respectively the central 99% and the upper 0.05%

5 Discussion

This paper provides the first proof of the asymptotic optimality of Thompson Sampling for Bernoulli bandits. Moreover the proof consists in a finite time analysis comparable with that of other known optimal policies. We also provide here simulations showing that Thompson Sampling outperforms currently known optimal policies.

Our proof of optimality borrows some ideas from Agrawal and Goyal’s paper, such as the notion of saturated arms. However we make use of ideas together with our own to obtain a stronger result, namely control over the tail of $N_{1,t}$ rather than its expectation. This is a valuable result which justifies the complexity of the proof of Proposition 2. Indeed control over these tails allows us to give a simpler finite time analysis for Thompson Sampling which is closer to the arguments for UCB-like algorithms, and also yields the optimal asymptotic rate of Lai and Robbins.

Thanks to the generalisation pointed out in [1], the Bernoulli version of Thompson Sampling can be applied to bandit problems with bounded rewards, and is therefore an excellent alternative to UCB policies. It would also be very natural to generalise Thompson to more complex reward distributions, choosing a prior appropriate for the assumptions on these distributions. Indeed, even in complex settings where the prior is not computable, Thompson Sampling only requires one sample from the posterior, which can be obtained efficiently using MCMC. Encouraging numerical experiments for reward distributions in the exponential family using a conjugate prior suggest that a generalisation of the proof is achievable. However this poses quite a challenge since the proof here is often heavily dependent on specific properties of Beta distributions. A natural generalisation would need a prior-dependent finite-time result controlling the tail probabilities of posterior distributions as the number of samples increases.

Acknowledgments We thank Aurélien Garivier and Olivier Cappé, for many fruitful discussions and for giving us the opportunity to work together.

This work was supported by the French National Research Agency (ANR-08-COSI-004 project EXPLO-RA) and the European Communitys Seventh Framework Programme (FP7/2007-2013) under grant agreements no 216886 (PASCAL2), and no 270327 (CompLACS).

References

- [1] S. Agrawal and N. Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference On Learning Theory (COLT)*, 2012.
- [2] J-Y. Audibert and S. Bubeck. Regret bounds and minimax policies under partial monitoring. *Journal of Learning Resaerch*, 2010.
- [3] J-Y. Audibert, R. Munos, and Cs. Szepesvári. Exploration-exploitation trade-off using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19), 2009.
- [4] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2):235–256, 2002.
- [5] O. Chapelle and L. Li. An empirical evaluation of thompson sampling. In *NIPS*, 2011.
- [6] A. Garivier and O. Cappé. The kl-ucb algorithm for bounded stochastic bandits and beyond. In *Conference On Learning Theory (COLT)*, 2011.
- [7] O.C. Granmo. Solving two-armed bernoulli bandit problems using a bayesian learning automaton. *International Journal of Intelligent Computing and Cybernetics*, 3(2):207–234, 2010.
- [8] J. Honda and A. Takemura. An asymptotically optimal bandit algorithm for bounded support models. In *Conference On Learning Theory (COLT)*, 2010.

- [9] E. Kaufmann, A. Garivier, and O. Cappé. On bayesian upper-confidence bounds for bandit problems. In *AISTATS*, 2012.
- [10] T.L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- [11] O-A. Maillard, R. Munos, and G. Stoltz. A finite-time analysis of multi-armed bandits problems with kullback-leibler divergences. In *Conference On Learning Theory (COLT)*, 2011.
- [12] B.C. May, N. Korda, A. Lee, and D. Leslie. Optimistic bayesian sampling in contextual bandit problems. *Journal of Machine Learning Research*, 13:2069–2106, 2012.
- [13] A. Salomon and J-Y. Audibert. Deviations of stochastic bandit regret. In *Algorithmic Learning Theory (ALT)*, 2011.
- [14] W.R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25:285–294, 1933.