



**HAL**  
open science

## Sub-sampling for Multi-armed Bandits

Akram Baransi, Odalric-Ambrym Maillard, Shie Mannor

► **To cite this version:**

Akram Baransi, Odalric-Ambrym Maillard, Shie Mannor. Sub-sampling for Multi-armed Bandits. Proceedings of the European Conference on Machine Learning, 2014, pp.13. hal-01025651v1

**HAL Id: hal-01025651**

**<https://hal.science/hal-01025651v1>**

Submitted on 18 Jul 2014 (v1), last revised 16 Dec 2014 (v2)

**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.

# Sub-sampling for Multi-armed Bandits

Akram Baransi<sup>1</sup>, Odalric-Ambrym Maillard<sup>1</sup>, and Shie Mannor<sup>1</sup>

Department of Electrical Engineering,  
Technion - Israel Institute of Technology, Haifa, Israel  
abaransi@tx.technion.ac.il,  
odalric-ambrym.maillard@ens-cachan.org, shie@ee.technion.ac.il

**Abstract.** The stochastic multi-armed bandit problem is a popular model of the exploration/exploitation trade-off in sequential decision problems. We introduce a novel algorithm that is based on sub-sampling. Despite its simplicity, we show that the algorithm demonstrates excellent empirical performances against state-of-the-art algorithms, including Thompson sampling and KL-UCB. The algorithm is very flexible, it does not need to know a set of reward distributions in advance nor the range of the rewards. It is not restricted to Bernoulli distributions and is also invariant under rescaling of the rewards. We provide a detailed experimental study comparing the algorithm to the state of the art, the main intuition that explains the striking results, and conclude with a finite-time regret analysis for this algorithm in the simplified two-arm bandit setting.

**Keywords:** Multi-armed Bandits · Sub-sampling · Reinforcement Learning

## 1 Introduction

In sequential decision making under uncertainty, the main dilemma that a decision maker faces is to explore, or not to explore. One of these problems is the popular stochastic multi-armed bandit problem, termed in reference to the 19th century gambling game and introduced by [23, 20]. It illustrates the fundamental trade-off between *exploration*, that is, making decisions that improve the knowledge of the environment, and *exploitation*, that is, choosing the decision that has maximized the previous payoff. Classically, each decision is referred to as an “arm”. There is a finite set of arms and each arm, when pulled, returns a real value, called the *reward*, which is independently and identically drawn from an unknown distribution. At each time step the decision maker chooses an arm based on the sequence of rewards that has been observed so far, pulls this arm and observes a new sample from the corresponding unknown underlying distribution. The objective is to find a policy for choosing the next arm to be pulled, that maximizes the sum of the expected rewards, or equivalently minimize the expected *regret*, that is the loss caused by not pulling the best arm at each time step. If  $\mathcal{A}$  denotes the set of arms and  $\{\mu_a\}_{a \in \mathcal{A}}$  the mean reward of the distribution of each arm, we denote  $\star \in \operatorname{argmax}_{a \in \mathcal{A}} \mu_a$  an optimal arm and the (expected) regret of an algorithm that pulls arms  $a_1, \dots, a_T$  up to time  $T$  is classically defined as

$$\mathfrak{R}_t = \mathbb{E} \left[ \sum_{t=1}^T (\mu_\star - \mu_{a_t}) \right]. \quad (1)$$

**Previous work** Since the formulation of the problem by Robbins (1952), the regret, that measures the cumulative loss resulting from pulling sub-optimal arms, has

been a popular criterion for assessing the quality of a strategy. Gittins index based policies ([13, 14, 12]), which were initially introduced by Gittins in 1979, is a family of Bayesian-optimal policies that based on indices that fully characterize each arm given the current history of the game, and at each time step the arm with the highest index will be pulled. However, the high computational cost of Gittins indices and the fact that they are practically limited to a specific set of distributions, arose the need of modifying the policies and make them more efficient. In [8], extending the seminal work of [18], the authors characterized the achievable performance. They showed that under suitable conditions on the possible distributions associated to each arm, any policy that is “admissible” (that is, not grossly under-performing, see [18] for details) must satisfy the following asymptotic lower-performance bound

$$\liminf_{T \rightarrow \infty} \frac{\mathfrak{R}_T}{\log(T)} \geq \sum_{a: \mu_a < \mu_*} \frac{\mu_* - \mu_a}{\mathcal{K}_{\text{inf}}(\nu_a; \mu_*)}, \quad (2)$$

where  $\mathcal{K}_{\text{inf}}(\nu_a; \mu_*)$  is an information-theoretic quantity which measures the minimal Kullback-Leibler divergence between  $\nu_a$  and distributions in the model that have expectations larger than  $\mu_*$ . In the same papers, [18], [10], [8] suggested that Gittins indices can be approximated by quantities that can be interpreted as upper bounds of confidence intervals.

In [1], the generic class of index policies termed **UCB** (Upper Confidence Bounds) was introduced, together with an asymptotic analysis of their performance. [5] provided the first finite time analysis for a particular variant of **UCB** based on Hoeffding’s inequality, showing that the regret grows logarithmically with the time horizon  $T$ . A few algorithms from the **UCB** family have been recently introduced such as **UCB-V** ([4]), **MOSS** ([3]), **Improved-UCB** ([6]), as well as the recent Kullback-Leibler-based algorithms **DMED** ([15]), **Kinf** ([19]), **k1-UCB** ([11]) and **KL-UCB** ([9]), that were shown to be first-order optimal.

Besides Gittins index and the **UCB**-type algorithms, another important class of algorithms is that introduced by Thompson ([23, 24]), and called **Thompson sampling**. The algorithm assumes that the arms’ distributions belong to a parametric family of distributions  $\mathcal{P} = \{p(\cdot|\theta), \theta \in \Theta\}$  where  $\Theta \subseteq R$ , it starts by putting a prior distributions on each one of the arms’ parameters, and at each time step a posterior distribution is maintained according to the rewards observed so far. In practice each different  $\mathcal{P}$  leads to a different implementation of the algorithm. At each time step, this Bayesian algorithm draws one sample for each arm from its posterior, then pulls the arm that maximizes the expected reward given that parameter. Recently, the analysis developed in [9] enabled to tackle the first frequentist optimal bound for the Thompson-sampling algorithm ([16]) in case of a family of Bernoulli distributions, thus proving that this algorithm also achieves optimality with a the regret that grows logarithmically with  $T$ . See also [2], as well as the recent extension to another class of distributions in [17].

**Contribution** In this paper we introduce a novel algorithm called **BESA** (Best Empirical Sampled Average) for the stochastic multi-armed bandit problem (see Section 2). The algorithm has a different flavor than previously introduced algorithms. It is *not* based on the computation of an empirical confidence bounds but rather on the sampling of some specific quantity. It is for that reason related in spirit to the Thompson sampling algorithm. However, unlike Thompson sampling, **BESA** does not rely on

---

**Algorithm 1** **BESA**(a,b) for a two-arm bandit

---

**Require:** Two arms  $a, b$ , current time  $t$ .

- 1: Sample  $I_t^a \sim \text{Wr}(N_t(a); N_t(b))$  and  $I_t^b \sim \text{Wr}(N_t(b); N_t(a))$ .
- 2: Define  $\tilde{\mu}_{t,a} = \widehat{\mu}(X_{1:N_t(a)}^a(I_t^a))$  and  $\tilde{\mu}_{t,b} = \widehat{\mu}(X_{1:N_t(b)}^b(I_t^b))$ .
- 3: Choose (break ties by choosing the arm with the smaller  $N_t$ )

$$a_t = \underset{a' \in \{a,b\}}{\text{argmax}} \tilde{\mu}_{t,a'}$$

---

a parametric set of distribution (or a prior) and is instead fully *non-parametric*. In Section 3, we compare the performance of the algorithm against state-of-the-art algorithms, including Thompson sampling and KL-UCB, in several scenarios with different types of reward distributions and show that the algorithm demonstrates excellent empirical performances against them. In Section 4, we provide a possible explanation for the strong performance of **BESA**, and then discuss its properties; Perhaps the most important property of **BESA** is its flexibility, since the same implementation can be used for any type of reward distributions, contrary to Thompson sampling or KL-UCB whose implementations differ according to the considered set of distribution. Finally in Section 5, we provide a finite-time regret bound for this algorithm in the two-arm bandit problem. We show with a rough analysis that the expected regret of the algorithm in this case is  $O(\log(T))$  where  $T$  is the time horizon. The focus of the paper is to introduce and report the striking empirical performance of this *simple* and *flexible* algorithm.

**Setup and notations** We consider a multi-armed bandit setting with finitely many arms  $\mathcal{A}$  and respective reward distributions  $\{\nu_a\}_{a \in \mathcal{A}}$ , where  $\nu_a \in \mathcal{P}([0, 1])$  and  $\mathcal{P}([0, 1])$  denotes the set of probability measures with support in  $[0, 1]$ . We denote  $\mu_a \in [0, 1]$  the mean of the distribution  $\nu_a$ , and  $X_{1:n}^a = (X_1^a, \dots, X_n^a)$  a sample of size  $n$ , i.i.d. from  $\nu_a$ . In the sequel, we use the short-hand notation  $[n]$  for the set of integer  $\{1, \dots, n\}$ . For a set of indices  $I \subset [n]$  of size  $m$ , say  $I = (i_1, \dots, i_m)$ , we write  $X_{1:n}^a(I) = (X_{i_1}^a, \dots, X_{i_m}^a)$  for the corresponding sub-sampled set. A sample of size  $m$  drawn *without replacement* from the set  $[n]$  is written  $I(n; m) \sim \text{Wr}(n; m)$ . Here  $\text{Wr}(n; m)$  denotes a distribution over sets of integers (with the convention that  $\text{Wr}(n; m) = \delta_{[n]}$  if  $m > n$ , where  $\delta$  refers to a Dirac distribution). Finally, for a sample  $S$  of real values, we denote  $\widehat{\mu}(S)$  the average of the sample components. For instance we have  $\widehat{\mu}(X_{1:n}^a) = \frac{1}{n} \sum_{i=1}^n X_i^a$ . Let  $\star \in \mathcal{A}$  denote an arm with maximal mean  $\mu_\star$ . The regret of an algorithm that pulls arms  $\{a_t\}_{t \in [T]}$  up to time  $T$  is defined by (1), where the expectation is taken with respect to all sources of randomness. We also denote the number of pulls of an arm  $a \in \mathcal{A}$  up to time  $t$  by  $N_t(a) = \sum_{t'=1}^{t-1} \mathbb{I}\{a_{t'} = a\}$ .

## 2 The **BESA** Algorithm

The main contribution of this paper is to introduce a novel algorithm, called **BESA** (**B**est **E**mpirical **S**ampled **A**verage) that uses a sub-sampling procedure in order to compare between the empirical values of two arms. The pseudo-code of the algorithm, for two arms is provided in Algorithm 1. The version for the more general multi-armed bandit uses a tournament strategy described in Algorithm 2.

The main idea of the algorithm is to make a fair comparison between the arms: Given two arms  $a$  and  $b$  that has been pulled  $n_a = N_t(a)$  and  $n_b = N_t(b) > n_a$  times respectively at time  $t$ , comparing the empirical averages of the arms is not a fair comparison since  $a$  has not gotten the same number of opportunities as  $b$  to show its abilities. **BESA** compensates for this situation by sub-sampling uniformly  $n_a$  rewards out of the

---

**Algorithm 2** **BESA**( $\mathcal{A}$ ) for a multi-armed bandit

---

**Require:** Set of arms  $\mathcal{A}$  of size  $A$ , current time  $t$ .

- 1: **if**  $\mathcal{A} = \{a\}$  **then**
  - 2:   Choose  $a_t = a$ .
  - 3: **else**
  - 4:   Choose  $a_t = \mathbf{BESA}_t(\mathbf{BESA}_t(\{a_i\}_{1 \leq i < \lceil A/2 \rceil}), \mathbf{BESA}_t(\{a_i\}_{\lfloor A/2 \rfloor < i \leq A}))$
  - 5: **end if**
- 

$n_b$  rewards of arm  $b$ . **BESA** then compares the empirical average of the rewards from  $a$ , to the empirical average of the rewards sub-sampled from  $b$ . It finally chooses  $b$  if its computed value is larger than the one of  $a$ . We provide in Algorithm 1 a more formal and unified presentation of this strategy. If  $n_b > n_a$ , the sampled set (line 2) is  $I_t^b \sim \text{Wr}(N_t(b); N_t(a))$  (indeed  $I_t^a \sim \text{Wr}(N_t(a); N_t(b))$  is the full set  $[N_t(a)]$  in this case). Then (line 3,4) the compared values become  $\hat{\mu}(X_{1:N_t(a)}^a)$  and  $\hat{\mu}(X_{1:N_t(b)}^b(I_t^b))$ .

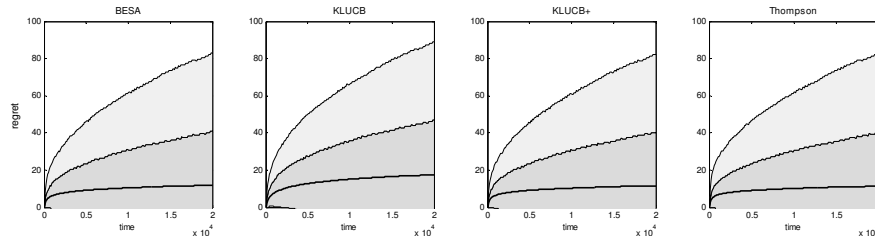
**BESA, FTL and Thompson sampling** At first sight, **BESA** seems close to a version of the standard **Follow The Leader** (FTL) algorithm. This algorithm selects  $\text{argmax}_{a \in \mathcal{A}} \hat{\mu}(X_{1:N_t(a)}^a)$ , that is the best empirical arm (with no sub-sampling). **FTL** is known to be a *bad* strategy in the bandit setting as it can lead to a linear regret in a number of situations. It is thus *a priori* striking that **BESA** can be any reasonable. On the other hand, **BESA** uses a sampling strategy in order to select the subset used to compute the sub-sampled mean. This is in this respect related in spirit to the **Thompson sampling** strategy, that is known to be both Bayesian optimal, and frequentist optimal, achieving the state-of-the-art for the bandit setting with Bernoulli distribution of rewards ([16]), or more recently distributions in the one-dimension exponential family ([17]). Note that **Thompson sampling** actually refers to a collection of algorithms whose implementation depend on the prior we have on reward distributions. Thus **Thompson sampling** for Bernoulli distributions is for instance different than **Thompson sampling** for exponential distributions. In contrast, **BESA** keeps the same form regardless of the distribution on rewards.

**A tournament for many arms** We extend Algorithm 1 written for two arms to the more general case of a finite set  $\mathcal{A}$  of arms by using a divide-and-conquer style algorithm (see Algorithm 2). This intuitively corresponds to organizing a tournament between arms. To avoid relying too much on a specific ordering of the arms that may bias the final result (and look arbitrary), we randomly shuffle the set of arms before each decision. That is, at time  $t$ , we create a copy  $\tilde{\mathcal{A}}_t$  of  $\mathcal{A}$  that is obtained by shuffling  $\mathcal{A}$  uniformly at random, and then output the arm  $\mathbf{BESA}(\tilde{\mathcal{A}}_t)$ .

### 3 Numerical Experiments

Our findings show that, surprisingly, **BESA** is a strong competitor of the state-of-the-art strategies from the bandit literature. Before providing one possible explanation for this striking performance, we now report these intriguing results more precisely.

In this section, the **BESA** algorithm is compared against well-known optimal algorithms such as **KL-UCB**, **KL-UCB+** ([11, 19, 9]), **Thompson sampling** ([23, 24]) with prior Beta(1,1), **KL-UCB-exp** ([9]) and **UCB-tuned** ([5]), on different scenarios with different set of arms. To avoid implementation bias, we use the open-source code available on-line at <http://mloss.org/software/view/415> for the implementation of these algorithms. Note that these algorithms do not need parameter tuning. In each one of the scenarios detailed below, the time horizon is set to  $T = 20,000$ , and



**Fig. 1.** Regret against time for the two-arm scenario, with  $\mu_\star = 0.9$  and  $\mu_a = 0.8$ .

**Table 1.** Performance measures for  $T = 20,000$  in the two-arm scenario, with  $\mu_\star = 0.9$  and  $\mu_a = 0.8$ . Complexity  $\alpha_{1/2}(M, 1) = O(0.9^M)$ .

	<b>BESA</b>	<b>KL-UCB</b>	<b>KL-UCB+</b>	<b>Thompson sampling</b>
Average regret at $T$	11.83	17.48	11.54	11.3
Beat <b>BESA</b>	—	1.82%	41.6%	58.28%
Average running time	2.86X	2.7X	3.12X	X

the scenario is run on 50,000 independent experiments. In each run, so as to have a fair comparison, the rewards of the arms are all drawn in advance, and all the algorithms are run on the same set of drawn rewards. In other words, on each of the 50,000 runs,  $\forall n \in \{1, \dots, 20000\}, a \in \mathcal{A}$  all the algorithms will observe the same reward on the  $n^{\text{th}}$  pull of arm  $a$ . This enables us to measure the percentage of runs on which one algorithm is better than a reference one, thus providing another measure of performance, besides the empirical average cumulative. We systematically report below in Table 1,  $\dots$ , 7 this percentage using **BESA** as a reference. In Figures 1,  $\dots$ , 7, the dark gray represents the plot quartiles, while the light gray represents the upper 5 percents quantile. Finally, in section 5, we introduce the so-called balance function  $\alpha_{1/2}$  (see definition 1) that acts as a complexity parameter. For clarity, we report the scaling of this function in most of the following scenarios as well.

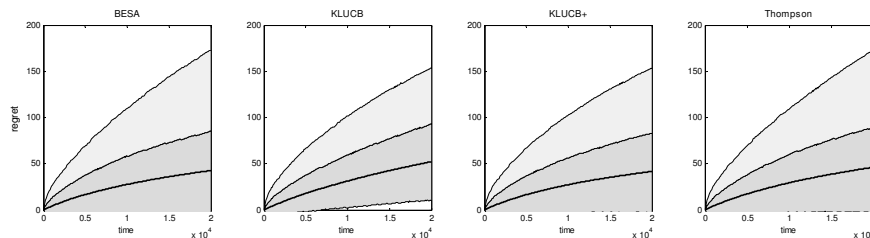
### 3.1 Scenario 1: Two Bernoulli Arms

In this scenario we consider the case of two Bernoulli arms  $\mathcal{A} = \{\star, a\}$ , with expectations  $\mu_\star = 0.9$  and  $\mu_a = 0.8$ , respectively. The empirical average cumulative regret of each algorithm is shown in the first row of Table 1, while the second row shows the percentage of the runs on which the algorithm gave a lower regret than **BESA**, and the third shows the average run time where X denotes the average run time of the fastest algorithm. In Figure 1 the average regret is shown as a function of time. The same scenario has been considered in [11]. On one hand, from figure 1 in [11] one can conclude that the average cumulative regret of **UCB-V** is larger than 50, while all the other algorithms but **KL-UCB** have average cumulative regret between 21 and 36. On the other hand, as reported in Table 1, the average cumulative regret of **BESA** is 11.38. Thus **BESA** outperforms all the algorithms considered in [11] such as e.g. **UCB-tuned**, **DMED**, **MOSS** on this scenario, including **KL-UCB**. Note that **KL-UCB+** does get a slightly lower expected regret, but does not beat **BESA** more than 50 per cent of the time. **Thompson sampling** here slightly outperforms **BESA**.

### 3.2 Scenario 2: Bernoulli with a Small $\Delta$

This scenario is similar to scenario 1 but with a smaller gap  $\Delta$ : We consider the case of two Bernoulli arms, with expectations  $\mu_1 = 0.81$  and  $\mu_2 = 0.8$  respectively. Similarly

to scenario 1 the average regret, the percentage of experiments on which **BESA** was beaten and the average run time are shown in Table 2, and the cumulative regret as a function of time is shown in Figure 2. Note that the average regret of **BESA** is close to that of **KL-UCB+** and smaller than that of **KL-UCB** and **Thompson sampling**. Interestingly enough, in addition the percentage of runs on which **BESA** is beaten by any of the state-of-the-art algorithms is smaller than 37%.



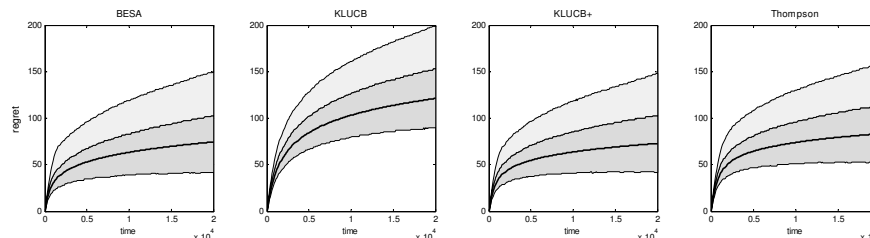
**Fig. 2.** Regret against time for the two-arm scenario, with  $\mu_* = 0.81$  and  $\mu_a = 0.8$ .

**Table 2.** Performance measures for  $T = 20,000$  for the two-arm scenario, with  $\mu_* = 0.81$  and  $\mu_a = 0.8$ . Complexity  $\alpha_{1/2}(M, 1) = O(0.9^M)$ .

	<b>BESA</b>	<b>KL-UCB</b>	<b>KL-UCB+</b>	<b>Thompson sampling</b>
Average regret at $T$	42.6	52.34	41.71	46.14
Beat <b>BESA</b>	--	25.61%	36.86%	35.2%
Average running time	4.56X	2.78X	3.47X	X

### 3.3 Scenario 3: Bernoulli with Low Means

In this scenario we consider the scenario used in [11], and inspired by a situation, frequent in applications like marketing or Internet advertising, where the mean reward of each arm is very low. More precisely we consider a harder case which has ten Bernoulli arms, the best arm has expectation 0.1, three arms have expectation 0.05, three arms expectation 0.02, and the rest with expectation 0.01. Table 3 summarizes the results of this experiment, and the regret as a function of time is shown in Figure 3. As can be seen from Table 3 the average regret of **BESA** is much smaller than **KL-UCB** and **Thompson sampling** regrets, and it is beaten by **KL-UCB** only in 1.57% of the runs and by **Thompson sampling** only in 3.09% of the runs. It is also beaten by **KL-UCB+** less than 36% of the time. As can be seen in figure 2 of [11] the regrets of all the algorithms but **DMED+** and **KL-UCB+**, which include e.g. **CP-UCB**, **DMED**, **UCB-Tuned**, are between 100 and 400. Thus we can conclude that **BESA**'s average is smaller than the average of those algorithms.



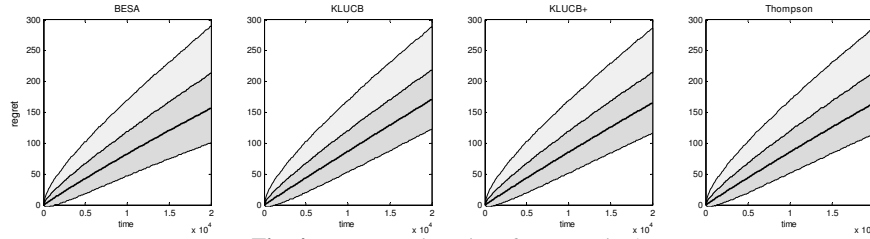
**Fig. 3.** Regret against time for scenario 3.

**Table 3.** Performance measures for  $T = 20,000$  for scenario 3.  $\alpha_{1/2}(M, 1) = O(0.5025^M)$ .

	<b>BESA</b>	<b>KL-UCB</b>	<b>KL-UCB+</b>	<b>Thompson sampling</b>
Average regret at $T$	74.41	121.21	72.84	83.36
Beat <b>BESA</b>	--	1.57%	35.41%	3.09%
Average running time	13.85X	2.83X	3.08X	X

### 3.4 Scenario 4: All Half but One

In this scenario we consider a case with ten Bernoulli arms, considered as being hard: The optimal arm has expectation 0.51 while all the others have expectation 0.5. The results of this experiment are shown in Table 4 and Figure 4. We note that **BESA** gets a smaller average regret than its competitors, and is not beaten by them more than 42% of the time. Thus **BESA** performs best in this hard setting.



**Fig. 4.** Regret against time for scenario 4.

**Table 4.** Performance measures for  $T = 20,000$  for scenario 4.  $\alpha_{1/2}(M, 1) = O(0.75^M)$ .

	<b>BESA</b>	<b>KL-UCB</b>	<b>KL-UCB+</b>	<b>Thompson sampling</b>
Average regret at $T$	156.7	170.82	165.28	165.08
Beat <b>BESA</b>	--	41.36%	41.57%	40.78%
Average running time	19.64X	2.78X	2.96X	X

### 3.5 Scenario 5: Truncated Exponential

In order to further demonstrate the flexibility of the **BESA** algorithm, we consider in this scenario the case of rewards coming from an exponential distribution. Five arms were considered with parameters  $\{\frac{1}{5}, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}, 1\}$ , truncated at 10 then divided by 10 (thus they are bounded in  $[0, 1]$ ). The results of this experiment are shown in Table 5 and Figure 5. Note that the regret of **KL-UCB-exp**, which is the version of **KL-UCB** specifically tuned for exponential families and achieving the state-of-the-art for this case, is lower than that of **BESA** only on 5.72% of the runs. Note that **BESA** need not know that the distributions are exponential, that is, we use exactly the same algorithm. Now, as can be seen in the figure of **BESA** the graph is not smooth: the reason is that **BESA** misses the optimal arm if the first reward that it gives is too low. In order to get a smoother behavior, we ran a slightly modified version of **BESA** to skip these corner cases: The modified algorithm is called **BESAT**, and simply pulls each arm ten times before starting running the regular **BESA**. As can be seen in the results this improved the regret dramatically. Now **KL-UCB-exp** beats **BESAT** only on 1.38% of the runs, and



similar numbers is achieved for **UCB-tuned**. In [11] a similar scenario is considered, with the difference that they didn't divided the reward by 10. It is actually easy to prove that both **BESA** and **BESAT** are actually invariant by rescaling, that is they pull the same arms in the same order wither we divide the reward by 10 or not. Thus, running the algorithms with the same runs without dividing the rewards by 10, the regret of **BESA** is 532.6 and the one of **BESAT** is 314.1, which is still better than the regrets reported in [11] at Figure 3 (they are above 600).

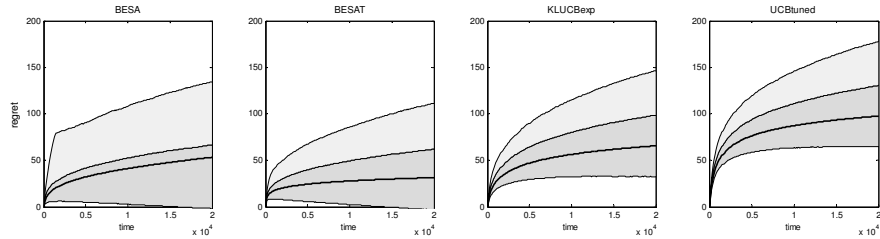


Fig. 5. Regret against time for scenario 5.

Table 5. Performance measures for  $T = 20,000$  for scenario 5.

	<b>BESA</b>	<b>BESAT</b>	<b>KL-UCB-exp</b>	<b>UCB-tuned</b>
Average regret at $T$	53.26	31.41	65.67	97.6
Beat <b>BESA</b>	--	40.59%	5.72%	4.33%
Beat <b>BESAT</b>	59.41%	--	1.38%	0.85%
Average running time	6.01X	7.09X	2.76X	X

### 3.6 Scenario 6: Truncated Poisson

In this scenario we consider the case of Poisson rewards, six Poisson arms are considered with parameters  $0.5 + \frac{i}{3}$ , where  $1 \leq i \leq 6$ , truncated at 10 then divided by 10. A similar scenario was considered in [9] where **KL-UCB-Poisson** is the leading algorithm. From Table 6, **BESA** and **BESAT** outperform **KL-UCB** and **KL-UCB-Poisson** on 95.95% of the runs for **BESA** and 97.99% for **BESAT**, with a much smaller average regret.

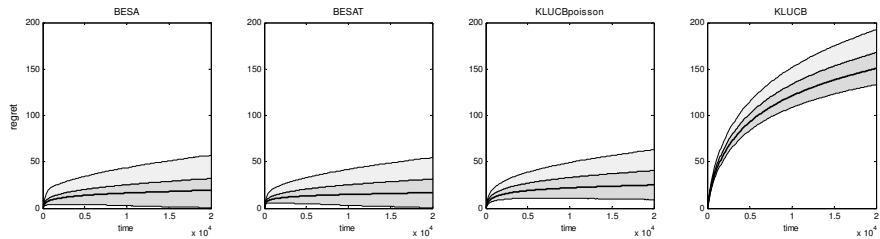


Fig. 6. Regret against time for scenario 6.

**Table 6.** Performance measures for  $T = 20,000$  for scenario 6.

	<b>BESA</b>	<b>BESAT</b>	<b>KL-UCB-Poisson</b>	<b>KL-UCB</b>
Average regret at $T$	19.37	16.72	25.05	150.56
Beat <b>BESA</b>	--	39.92%	4.05%	0.72%
Beat <b>BESAT</b>	59.51%	--	2.01%	0.17%
Average running time	3.53X	3.49X	1.15X	X

### 3.7 Scenario 7: Uniform Distributions

In this experiment, we consider a challenging setting where arms are uniformly distributed with  $X_a \sim \mathcal{U}([0.2, 0.4])$  and  $X_* \sim \mathcal{U}([0, 1])$ . Note that there is no natural **KL-UCB** nor **Thompson sampling** algorithm to deal with such family of distribution. In such a scenario, we note that  $\alpha_{1/2}(M, 1)$  does not decay exponentially to 0 with  $M$ , indicating that this is a difficult scenario for **BESA**. However, it holds that  $\alpha_{1/2}(1, n) = O(\beta^n)$  with  $\beta = 0.2$ . According to Theorem 1, we should initialize **BESA** by pulling  $n_0$  times each arm before applying **BESA** (that is, run **BESAT**), where for  $T = 20,000$ ,  $n_0 \simeq 6.15$ . We ran **BESAT** with different number of initialization pulls  $n_0 \in \{0, 3, 6, 7, 8, 9, 10\}$ , and we also ran **UCB KL-UCB** and **Thompson sampling** (with Beta(1,1) prior) on the same set of arms. The average regrets of each of the algorithms in addition to the percentage on which each non-**BESA** algorithm beats **BESA** with different  $n_0$  are provided in Table 7. The average regret of **BESA** improves with increasing  $n_0$  as expected, and as can be seen from the table, **BESA** with  $n_0 = 10$  gave a lower average regret than **UCB** and **KL-UCB** and a bit higher than **Thompson sampling** and it was beaten by the other algorithms on less than 0.8% of the runs.

**Table 7.** Performance measures for  $T = 20,000$  for scenario 7.

Alg.	<b>UCB</b>	<b>KL-UCB</b>	<b>Thompson sampling</b>
Average regret	21.23	20.72	13.18
Beat <b>BESA</b> $n_0 = 0$	24.26%	24.28%	24.7%
Beat <b>BESA</b> $n_0 = 3$	7.27%	7.3%	7.83%
Beat <b>BESA</b> $n_0 = 7$	1.56%	1.58%	1.76%
Beat <b>BESA</b> $n_0 = 10$	0.62%	0.63%	0.74%

<b>BESA</b>	$n_0 = 0$	$n_0 = 3$	$n_0 = 6$	$n_0 = 7$	$n_0 = 8$	$n_0 = 9$	$n_0 = 10$
Average regret	920.12	213.44	50.6	35.38	25.88	17.85	15.42

### 3.8 Summary of the Experimental Results

From the first four numerical experiments, we deduce that **BESA** is able to compete with the state-of-the-art bandit algorithms in the Bernoulli case. It becomes especially good when the gaps are *small*, or when the Bernoulli parameters are small, which are two main cases of practical interest (especially in web-advertising). Scenario 5, 6 and 7 highlight the flexibility of the algorithm: the same algorithm competes favorably against one of the best algorithm for exponential distributions as well as for Poisson distributions. Using a slight modification, we can beat them by an even larger margin.

## 4 Intuition and Properties

In this section, we provide an explanation for the striking performance of the **BESA** algorithm, and discuss further its advantages and drawbacks. In the next section, we use this intuition to derive a regret bound for the **BESA** algorithm.

### 4.1 Why Does It Work?

To give intuition why **BESA** works, let us focus on the two-arm bandit problem. The heuristic idea behind the algorithm is that a comparison between two empirical mean estimates built on a very different number of samples  $n_a$  and  $n_b$  is not really “fair”, and that it seems more natural to compare empirical means based on the same number of samples. Thus the algorithm we introduce is based on sub-sampling. In the rich sub-sampling literature, the works of [7] and [21], show that using sub-sampling without replacement ensures convergence guarantees in a strictly broader setting than sub-sampling with replacement (a.k.a bootstrap). This may provide some informal support for the soundness of the method. However we now provide a more direct justification for the striking performance of **BESA**.

On the theoretical side, one can justify the intuition by looking at the probability of repeatedly choosing a wrong action. If  $\mu_b > \mu_a$  and the number of plays of each arm satisfies  $n_a > n_b$ , the probability that **BESA** chooses a wrong action is approximately

$$\mathbb{P}\left[\widehat{\mu}\left(X_{1:n_a}^a(I(n_a; n_b))\right) > \widehat{\mu}\left(X_{1:n_b}^b\right)\right], \quad (3)$$

where  $I(n_a; n_b) \sim \text{Wr}(n_a; n_b)$ , and the probability of making  $M$  consecutive mistakes is essentially

$$\mathbb{P}\left[\forall m \in [M] \widehat{\mu}\left(X_{1:n_a^{(m)}}^a(I_m(n_a^{(m)}; n_b))\right) > \widehat{\mu}\left(X_{1:n_b}^b\right)\right], \quad (4)$$

where for all  $m \leq M$ ,  $I_m(n_a; n_b) \sim \text{Wr}(n_a; n_b)$ , and where we introduced for convenience the short-hand notation  $n_a^{(m)} = n_a + m - 1$ .

Now, for deterministic<sup>1</sup>  $n_a, n_b$ , (3) typically scales with  $\exp\left(-2n_b(\mu_b - \mu_a)^2\right)$ , by a standard Hoeffding inequality since  $n_b$  samples are involved. On the other hand, (4) can decrease at a much faster rate, intuitively of order  $\exp\left(-2n_b\tilde{M}(\mu_b - \mu_a)^2\right)$  where  $\tilde{M}$  is the number of non-overlapping sub-samples of size  $n_b$ . Indeed if  $I_{m'}(n_a^{(m')}; n_b) \cap I_m(n_a^{(m)}; n_b) = \emptyset$  for all  $m \neq m' \in \mathcal{M} \subset [M]$  where  $|\mathcal{M}| = \tilde{M}$ , then the corresponding empirical means are independent from each other, which leads to the intuitive improvement. Using sub-sampling, the later event is of high probability for a reasonable  $\tilde{M}$  provided that  $n_a/n_b$  is large enough. Note that in the case when we do not resort to sub-sampling, such a phenomenon will not happen, due to the strong dependency between the samples at two consecutive time steps. Thus  $\tilde{M}$  is essentially 1 which means that the probability of committing  $M$  successive mistakes, will stay big, of the order of  $\exp\left(-2n_b(\mu_b - \mu_a)^2\right)$ . Now in an ideal case, with only  $M = n_a/n_b$  subsets, we might get a mistake error scaling with  $\exp\left(-2n_a(\mu_b - \mu_a)^2\right)$ , even though  $b$  has only been pulled  $n_b$  times. As long as  $n_a/n_b \ll n_a - n_b$ , then we need less trials than the

<sup>1</sup> The exact argument needs to deal with the fact that  $n_a = N_t(a)$  and  $n_b = N_t(b)$  are both random stopping times.

procedure based only on confidence interval estimates before accurately discarding the wrong arm with the same probability. This intuitive idea is formalized and captured by Lemma 1, which we consider to be the key for the current analysis.

## 4.2 Properties of the Algorithm

We now highlight some of the main properties of **BESA**.

**Simplicity** We first note the simplicity of the **BESA** compared to previous methods, such as **KL-UCB** for instance that requires some fancy linear program in order to compute the upper confidence bound, or **Thompson sampling** that requires to be able to find an appropriate conjugate prior and implement the update of its parameters, or even the **UCB**-type strategies that generally require some free parameter to be adjusted. Here, **BESA** requires no parameter tuning, and no complicated prior/posterior relation is needed either. Thus the algorithm is directly applicable in a broad range of situation.

**Flexibility** A striking property of the algorithm is its flexibility to adapt to various situations. For instance, note that **BESA** does not need to know the support of the distributions, and is also invariant under rescaling. This is not the case of most bandit algorithms that explicit use the knowledge of the support  $[a, b]$  of distributions. We believe this can be a serious advantage in some situation. Moreover, both **Thompson sampling** and **KL-UCB** are dependent on a considered parametric set of distributions: in practice a different set of distribution leads to a different implementation. **BESA** does not need such parameters, keeps the same form in all situations, and more importantly still achieves excellent performance in a number of situations, as detailed in Section 3.

**Efficiency** Finally, one might wonder about the computational efficiency of **BESA** due to the use of a sub-sampling method, that is generally not memory less. We were a bit worried about this fact, and thus we implemented the algorithm in a naive way and reported the computational cost of the algorithm in each table for completeness. We conclude from these results that the computational cost of the algorithm is essentially not a problem. Moreover, note that, due to the i.i.d. nature of the data, one may use fancier but more efficient sub-sampler techniques. A naive implementation needs to save all the received rewards. In case the rewards take only finitely many values one can use instead a counter for each possible value. To avoid distracting the reader from the main message, we do not discuss possible tricks that could be used to improve further the numerical efficiency of the method.

## 5 Regret Upper-Bound

In this section, we provide a simple regret analysis of the **BESA** algorithm<sup>2</sup>. Formalizing further the heuristic intuition of Section 4, it is actually possible to derive a non-trivial regret bound for the **BESA** algorithm, given in Theorem 1. In order to characterize the difficulty of a bandit problem, we now define the following problem-dependent quantity

**Definition 1.** For integers  $M, n$  and  $\lambda \in [0, 1]$ , we define the balance function of the distributions  $(\nu_a, \nu_*)$  as

$$\alpha_\lambda(M, n) = \mathbb{E}_{Z \sim \nu_{*,n}} \left[ (1 - F_{\nu_{a,n}}(Z) + \lambda \nu_{a,n}(Z))^M \right],$$

<sup>2</sup> We study a slightly modified version, that break ties uniformly at random.

where  $\nu_{a,n}$  is the distribution of  $\sum_{i=1}^n X_i^a$  with  $X_i^a \stackrel{i.i.d}{\sim} \nu_a$  and  $F_\nu$  is the cdf of  $\nu$  (that is,  $F_\nu(x) = \mathbb{P}_{X \sim \nu}(X \leq x)$ ).

Let us provide some intuition on two examples. Note that  $\alpha_\lambda$  is not increasing both in  $M$  and  $n$ . First, let us consider two Bernoulli arms with  $X^a \sim \mathcal{B}(\mu_a)$  and  $X^* \sim \mathcal{B}(\mu_*)$ . We can compute easily that  $\alpha_\lambda(M, 1) = (\lambda\mu_a)^M \mu_* + (\lambda + \mu_a(1-\lambda))^M (1-\mu_*)$ . Now, since  $\mu_a < \mu_* \leq 1$ , we deduce that  $\alpha_{1/2}(M, 1) \xrightarrow{M \rightarrow \infty} 0$ , and more precisely that

$$\alpha_{1/2}(M, 1) = O\left(\left(\frac{\mu_a \vee (1-\mu_a)}{2}\right)^M\right).$$

Thus it converges exponentially fast to 0. Note, however that  $\alpha_1(M, 1) \xrightarrow{M \rightarrow \infty} 1 - \mu_*$ , which is non zero unless  $\mu_* = 1$ . Second, let us consider the case of two Uniform arms  $X^a \sim \mathcal{U}([0.2, 0.4])$  and  $X^* \sim \mathcal{U}([0, 1])$ . For all  $\lambda$  it holds that  $\alpha_\lambda(M, n) \xrightarrow{M \rightarrow \infty} 0.2^n$ . Thus there is no exponential decay with  $M$ . However, it holds that  $\alpha_\lambda(1, n) = O(0.2^n)$ .

We now prove the following

**Theorem 1.** *Let  $\mathcal{A} = \{*, a\}$  be a two-armed bandit with bounded rewards in  $[0, 1]$ , and  $\Delta = \mu^* - \mu_a$  be the mean gap. Let us moreover assume that there exists  $\alpha \in (0, 1)$  and  $c > 0$  such that  $\alpha_{1/2}(M, 1) \leq c\alpha^M$ . Then the regret of **BESA** at time  $T$  is controlled by*

$$\mathfrak{R}_T \leq \frac{11 \log(T)}{\Delta} + C_{\nu_a, \nu_*} + O(1),$$

where  $C_{\nu_a, \nu_*}$  depends on the parameters of the problem  $\alpha, c$  and  $\Delta$ , but not on  $T$ . Moreover if there exists some  $\beta \in (0, 1)$  and  $c > 0$  such that  $\alpha_{1/2}(1, n) \leq c\beta^n$ . Let us define

$$n_{0,T} = \left\lceil \frac{\ln(T) - \ln((1-\beta)C)}{\ln(1/\beta)} \right\rceil.$$

Then if **BESA** is initialized with  $n_{0,T}$  pulls of each arm, then its regret at time  $T$  is controlled by

$$\mathfrak{R}_T \leq \frac{11 \log(T)}{\Delta} + n_{0,T} + \tilde{C}_{\nu_a, \nu_*} + O(1).$$

where  $\tilde{C}_{\nu_a, \nu_*}$  depends on  $C$  and on the parameters  $\beta, c$  and  $\Delta$ , but not on  $T$ .

*Remark 1.* Up to Lemma 1 (see below) which is a purely probabilistic result, and is independent on the bandit setting, the proof is arguably simpler than the typical ones used for Thompson sampling. In particular, we do not need to resort to a fancy ‘‘Bernoulli-Beta’’ trick that is used in classical proofs of Thompson sampling and does not extend easily to general distributions (see for instance [17] that is entirely devoted to the extension to exponential families of dimension one).

*Remark 2.* Since one needs not use empirical confidence intervals in the analysis, but simply confidence intervals, one can hope to derive much tighter results in the future, using Kullback-Leibler-based Chernoff bounds, or event the sharpest Sanov bounds.

We provide the full proof of this result in the appendix. It mainly follows the proof of [16] that provides a sharp analysis of the **Thompson sampling** algorithm, with some simplifications: first, we consider only two arms, which enables us to skip a recurrence argument (but the same technique could be used to extend our analysis to the

case of  $K$ -arms); then, we only use mean-based arguments essentially for clarity of exposure. We believe it is more important at this point to provide a clear intuition about why the algorithm works than to provide a tight analysis based on Kullback-Leibler concentration results that are trickier to catch. Now for clarity, we summarize in the next Lemma what we believe is the key result for the regret analysis of **BESA**. This purely probabilistic result is specific to the properties of the sub-sampling procedure.

The sketch of proof is as follows: As usual, we express the regret in terms of the expected number of pulls of sub-optimal arms, that are further decomposed according to the event that the optimal arm has been pulled enough or not. Under the event that the optimal arm is pulled enough, we control easily the probability of mistake resorting to standard proof techniques based on concentration bounds. One difference with respect to standard bounds is that we use here a Serfling-Hoeffding ([22]) concentration inequality. This gives the first term of the regret. The next and difficult step is to show, as usual, that the optimal arm is indeed pulled enough with high probability. To that end, we borrow a proof technique considered in [16]: we introduce the random times  $\tau_j$  between the  $j^{\text{th}}$  and  $(j+1)^{\text{th}}$  pull of the optimal arm, and show that they cannot be too large for too many  $j$ , that is to show that the number of *consecutive* mistakes made by the algorithm must be small with high probability. This is one key of the regret analysis of **Thompson sampling** that enables to derive an optimal performance bound. The novelty that we introduce for the analysis of **BESA** is to relate this number of consecutive mistakes to the probability that many sub-samples of small size do not overlap, as explained in Section 4. The precise lemma that covers this part, and that eventually leads to our regret bound is the following:

**Lemma 1 (Maximal non-overlapping sub-samples).** *Let  $\mathcal{M} = \{p, \dots, q\} \subset \mathbb{N}$  be some interval. Let  $j, M \in \mathbb{N}$  be such that  $p \geq 2j$ , and  $M \leq |\mathcal{M}|$ . For all  $s \in \mathcal{M}$ , we introduce the random variable  $I_s(s-j; j) \sim \text{WR}(s-j; j)$ . Then the function defined by*

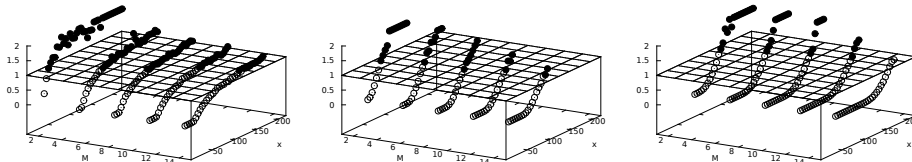
$$f_{\mathcal{M}}(M, j) = 1 - \mathbb{P} \left[ \exists s_1 < \dots < s_M \in \mathcal{M}, \right. \\ \left. \forall m \neq m' \in [M] : I_{s_m}(s_m - j; j) \cap I_{s_{m'}}(s_{m'} - j; j) = \emptyset \right]$$

*is decreasing with  $p$ . Moreover, for a sequence of intervals  $\mathcal{M}_t = \{p_t, \dots, q_t\}$ , such that  $\lim_{t \rightarrow \infty} \frac{q_t - p_t}{t} = C > 0$  and integers  $M_{t,j}$  such that  $M_{t,j} = O(\ln(q_t))$ , we have*

$$f_{\mathcal{M}_t}(M_{t,j}, j) = o(t^{-1}). \quad (5)$$

One way to show this lemma is by studying  $f_{\mathcal{M}}(M, j)$  formally, and trying to see how it behaves with the different parameters. This however turns out to be tedious and very technical. On the other hand, we can take advantage of the fact that this function is problem-independent and thus can be computed off-line. It can actually be simulated, and since it is decreasing with  $|\mathcal{M}|$ , it is enough to study its behavior for small  $|\mathcal{M}|$ . For our purpose in the analysis, we only need to look at small values of  $M = O(\log(T))$  (note that for a time horizon  $T = 20,000$ , then  $\log(T) < 10$ ). Similarly we use small value for  $n_0 \leq j \leq u_T = O(\log(T))$  as well. It is not difficult to simulate  $f_{\mathcal{M}}(M, j)$  for  $\mathcal{M}_x = [2j+1, x]$  with various values of  $M$  and  $x$ . We are interested in the ratio  $\alpha = -\log(f_{\mathcal{M}_x}(M, j))/\log(x)$ , and observe numerically that this ratio is increasing with  $x$  and quickly becomes larger than 1, that is  $f_{\mathcal{M}_x}(M, j) < x^{-1}$  for large enough

$x \geq C_M = o(M)$  that is slowly increasing with  $M$ . In the analysis of Theorem 1, we use  $x = O(t/\log(t))$ , and  $M = O(\log(T))$ , and thus as soon as  $t \geq C' \log(t)$ , which happens for  $t \geq C$  for some numerical constant  $C$ , then  $f_{\mathcal{M}_t}(M, j)$  starts decaying faster than  $t^{-1}$ , that is  $f_{\mathcal{M}_t}(\log(t)/j, j) = o(t^{-1})$ . In figure 7, we plot the function  $\alpha$  in terms of  $x$  and  $M$ , and for different values of  $j$ . Each point is the result obtained via 5,000 replications. We report especially in blue the regions when  $\alpha$  becomes larger than 1, which is the critical value to ensure that the lemma holds.



**Fig. 7.**  $\alpha$  as a function of  $x$  and  $M$  for  $j = 3$  (left),  $j = 8$  (middle) and  $j = 15$  (right). (Note that for  $\Delta = 0.3$  and  $T = 20,000$ , then  $\frac{1}{2\Delta} \log(T) \simeq 16.5$ .) The black circles indicate when  $\alpha_x > 1$  and the white circles when  $\alpha_x < 1$ .

## 6 Discussion and Conclusion

In this paper, we introduced a novel algorithm for the stochastic multi-armed bandit that is based on *sub-sampling*. We provided a careful experimental analysis of the **BESA** algorithm, by comparing it with the optimized versions of the state-of-the-art algorithms known in each situation. We demonstrated the advantage of **BESA** specifically in the case of Bernoulli distributions, including the case of small parameters and small gaps, as well as exponential and Poisson distributions. For completeness, we reported three measures of performance of the algorithm: plots of the cumulative reward, included quantiles, the percentage it is beaten by other standard algorithms and the numerical complexity with respect to the fastest method.

The algorithm has several striking properties: it is simple to implement, and is very flexible. It does not need to know a set of distributions in advance, unlike **Thompson sampling** or **KL-UCB** and does not even need to know the support, unlike **UCB** or **kl-UCB**. It is also invariant under rescaling of the rewards. This is thus a fully non-parametric algorithm, that competes favorably against standard algorithms.

We finally provide a regret analysis for **BESA**, which shows that the regret of the algorithm is logarithmic (we believe that the constants are not tight). More importantly, we provided a novel proof technique that we believe conveys the core intuition why the algorithm is working, and can lead to much tighter bounds in the future.

Now that we have introduced this algorithm and shown its flexibility, it seems natural to try to extend the **BESA** to other settings. One first direction of research is to consider the *contextual*-bandit problem, another one is to consider the *adversarial* multi-armed bandit setting.

**Acknowledgements** This work was supported by the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement 306638 (SUPREL) and the Technion.

## Bibliography

- [1] Rajeev Agrawal. Sample mean based index policies with  $O(\log n)$  regret for the multi-armed bandit problem. *Advances in Applied Probability*, 27(4):1054–1078, 1995.
- [2] Shipra Agrawal and Navin Goyal. Further optimal regret bounds for thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, volume 31 of *JMLR W&CP*, Scottsdale, AZ, US, 2013.
- [3] Jean-Yves Audibert and Sébastien Bubeck. Minimax policies for adversarial and stochastic bandits.
- [4] Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Exploration-exploitation trade-off using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410:1876–1902, 2009.
- [5] Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3:397–422, mar 2003.
- [6] Peter Auer and Ronald Ortner. UCB revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61(1-2):55–65, 2010.
- [7] Peter J. Bickel and Anat Sakov. On the choice of  $m$  in the  $m$  out of  $n$  bootstrap and confidence bounds for extrema. *Statistica Sinica*, 18:967–985, 2008.
- [8] Apostolos N. Burnetas and Michael N. Katehakis. Optimal adaptive policies for sequential allocation problems. *Adv. Appl. Math.*, 17(2):122–142, jun 1996.
- [9] Olivier Cappé, Aurélien Garivier, Odalric-Ambrym Maillard, Rémi Munos, and Gilles Stoltz. Kullback–leibler upper confidence bounds for optimal sequential allocation. *Ann. Statist.*, 41(3):1516–1541, 2013.
- [10] Fu Chang and Tze Leung Lai. Optimal stopping and dynamic allocation. *Advances in Applied Probability*, 19(4):829–853, 1987.
- [11] Aurélien Garivier and Olivier Cappé. The KL-UCB algorithm for bounded stochastic bandits and beyond. In *Proceedings of the 24th annual Conference On Learning Theory*, COLT ’11, 2011.
- [12] John C. Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, 41(2):148–177, 1979.
- [13] John C. Gittins and David M. Jones. A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika*, 66(3):561–565, 1979.
- [14] John C. Gittins, Richard Weber, and Kevin Glazebrook. *Multi-armed Bandit Allocation Indices*. Wiley, 1989.
- [15] Junya Honda and Akimichi Takemura. An asymptotically optimal bandit algorithm for bounded support models. pages 67–79.
- [16] Émilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: an asymptotically optimal finite-time analysis. In *Proceedings of the 23rd international conference on Algorithmic Learning Theory*, ALT’12, pages 199–213, Berlin, Heidelberg, 2012. Springer-Verlag.
- [17] Nathaniel Korda, Emilie Kaufmann, and Rémi Munos. Thompson sampling for 1-dimensional exponential family bandits. In Christopher J. C. Burges, Léon Bottou,



- Zoubin Ghahramani, and Kilian Q. Weinberger, editors, *NIPS*, number 26, pages 1448–1456, Lake Tahoe, Nevada, United States, 2013.
- [18] Tze L. Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
  - [19] Odalric-Ambrym Maillard, Rémi Munos, and Gilles Stoltz. Finite-time analysis of multi-armed bandits problems with kullback-leibler divergences. In *Proceedings of the 24th annual Conference On Learning Theory, COLT '11*, 2011.
  - [20] Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematics Society*, 58:527–535, 1952.
  - [21] Joseph P. Romano and Azeem M. Shaikh. On the uniform asymptotic validity of subsampling and the bootstrap. *The Annals of Statistics*, 40(6):2798–2822, 12 2012.
  - [22] Robert J. Serfling. Probability inequalities for the sum in sampling without replacement. *The Annals of Statistics*, 2(1):39–48, 1974.
  - [23] William R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25:285–294, 1933.
  - [24] William R. Thompson. On the theory of apportionment. *American Journal of Mathematics*, 57:450–456, 1935.

# Appendix beyond the page limit

## A Proof of Theorem 1

We now provide the proof of Theorem 1. It is divided in 6 steps, and is broadly inspired from the proof of [16] used to derive the state-of-the-art regret bounds for Thompson Sampling, with some simplifications due to the restriction to two arms.

STEP 1. Decomposition of the regret. For any  $u > 0$ , the expected regret  $\mathfrak{R}_T$  at time  $T$  can be decomposed as

$$\begin{aligned} \mathfrak{R}_T &= (\mu_\star - \mu_a)\mathbb{E}[N_T(a)] \\ &= \Delta \mathbb{E} \left[ \sum_{t=1}^T \mathbb{I}\{\tilde{\mu}_{t,a} > \tilde{\mu}_{t,\star}\} \right] \\ &\leq \Delta \sum_{t=1}^T \left( \mathbb{P} \left[ \tilde{\mu}_{t,a} > \tilde{\mu}_{t,\star} \cap N_t(\star) > u \right] + \mathbb{P} [N_t(\star) \leq u] \right). \end{aligned} \quad (6)$$

We now control separately each of the two events that appear in the right hand side of the inequality.

STEP 2. Probability of mistake. We combine here two results from concentration of measure. On the one hand, by application of a simple Serling-Hoeffding's inequality [22], it holds for all  $\delta \in (0, 1)$  that

$$\begin{aligned} \mathbb{P} \left[ \widehat{\mu}(X_{1:N_t(\star)}^\star(I_t^\star(N_t(\star), N_t(a)))) \leq \right. \\ \left. \widehat{\mu}(X_{1:N_t(\star)}^\star) - \sqrt{\frac{\rho_{\star,a} \log(1/\delta)}{2 \min\{N_t(a), N_t(\star)\}}} \middle| \mathcal{F}_{<t} \right] \leq \delta, \end{aligned} \quad (7)$$

where  $\rho_{\star,a} = 1 - \frac{N_t(a)-1}{N_t(\star)}$  if  $N_t(a) \leq N_t(\star)/2$ ,  $\rho_{\star,a} = 0$  if  $N_t(a) \geq N_t(\star)$ , and  $\rho_{\star,a} = (1 - \frac{N_t(a)}{N_t(\star)})(1 + \frac{1}{N_t(a)})$  else. On the other hand, we get by a simple Hoeffding bound that

$$\begin{aligned} \mathbb{P} \left[ \exists s \in \{\lceil u \rceil, \dots, t\} : \widehat{\mu}(X_{1:N_t(\star)}^\star) \leq \right. \\ \left. \mu_\star - \sqrt{\frac{\alpha \log(t)}{2s}} \right] \leq (t - \lceil u \rceil + 1)t^{-\alpha}, \end{aligned}$$

or by a more refined maximal Hoeffding's inequality combined with a peeling argument that for all  $\alpha > 1$ ,

$$\begin{aligned} \mathbb{P} \left[ \exists s \in \{\lceil u \rceil, \dots, t\} : \widehat{\mu}(X_{1:N_t(\star)}^\star) \leq \right. \\ \left. \mu_\star - \sqrt{\frac{\alpha \log(t)}{2s}} \right] \leq \left( \frac{\log(t - \lceil u \rceil + 1)}{\log((\alpha + 1)/2)} + 1 \right) t^{-\frac{2\alpha}{\alpha+1}}. \end{aligned}$$

Now, by definition of the **BESA** algorithm, it holds that  $A_t = a$  implies

$$\begin{aligned} & \tilde{\mu}_{t,a} - \widehat{\mu}(X_{1:N_t(a)}^a) + \widehat{\mu}(X_{1:N_t(a)}^*) - \mu_a + \mu_a \\ & \geq \tilde{\mu}_{t,\star} - \widehat{\mu}(X_{1:N_t(\star)}^*) + \widehat{\mu}(X_{1:N_t(\star)}^*) - \mu_\star + \mu_\star, \end{aligned}$$

that is, reorganizing the terms and using a union bound, we deduce that provided that  $N_t(\star) \geq u$ , then there exists an event of probability higher than  $1 - 2 \left( \frac{\log(t - \lceil u \rceil + 1)}{\log((\alpha + 1)/2)} + 2 \right) t^{-\frac{2\alpha}{\alpha + 1}}$  such that on this event

$$\begin{aligned} \Delta & \leq \left( \tilde{\mu}_{t,a} - \widehat{\mu}(X_{1:N_t(a)}^a) \right) + \left( \widehat{\mu}(X_{1:N_t(a)}^*) - \mu_a \right) \\ & + \left( \widehat{\mu}(X_{1:N_t(\star)}^*) - \tilde{\mu}_{t,\star} \right) + \left( \mu_\star - \widehat{\mu}(X_{1:N_t(\star)}^*) \right) \\ & \leq \sqrt{\frac{\rho_{a,\star} \alpha \log(t)}{(\alpha + 1) \min\{N_t(a), N_t(\star)\}}} + \sqrt{\frac{\alpha \log(t)}{2N_t(a)}} \\ & + \sqrt{\frac{\rho_{\star,a} \alpha \log(t)}{(\alpha + 1) \min\{N_t(a), N_t(\star)\}}} + \sqrt{\frac{\alpha \log(t)}{2N_t(\star)}}. \end{aligned} \quad (8)$$

Now, either  $N_t(a) < N_t(\star)$ , or  $N_t(a) \geq N_t(\star)$ . Let us first consider the case when  $N_t(a) < N_t(\star)$ . We then have  $\rho_{a,\star} = 0$ , and also

$$\rho_{\star,a} \leq 1 - \frac{N_t(a)}{N_t(\star)} + \frac{1}{N_t(\star)} \leq 2.$$

Thus, inequality (8) can be further simplified to

$$\Delta \leq \sqrt{\frac{2\alpha \log(t)}{N_t(a)}} \left( \sqrt{\frac{1}{\alpha + 1}} + 1 \right),$$

and thus the number of pulls of the sub-optimal arm must be controlled by

$$N_t(a) \leq \left( 1 + \sqrt{\frac{1}{\alpha + 1}} \right)^2 \frac{2\alpha \log(t)}{\Delta^2}. \quad (9)$$

In the case when  $N_t(a) \geq N_t(\star)$ , we deduce similarly that  $\rho_{\star,a} = 0$  and that

$$\Delta \leq \sqrt{\frac{2\alpha \log(t)}{N_t(\star)}} \left( \sqrt{\frac{1}{\alpha + 1}} + 1 \right),$$

from which we deduce that we must have

$$u \leq N_t(\star) \leq \left( 1 + \sqrt{\frac{1}{\alpha + 1}} \right)^2 \frac{2\alpha \log(t)}{\Delta^2}.$$

As a result, if we introduce  $u_t = \left( 1 + \sqrt{\frac{1}{\alpha + 1}} \right)^2 \frac{2\alpha \log(t)}{\Delta^2}$ , we deduce that provided that  $N_t(\star) > u_t$ , then with probability higher than  $1 - 2 \left( \frac{\log(t+1)}{\log((\alpha + 1)/2)} + 2 \right) t^{-\frac{2\alpha}{\alpha + 1}}$  we must

have  $N_t(a) < N_t(\star)$  when  $A_t = a$ , that is (9) must hold. This means that the number of pulls of the sub-optimal arms is controlled in this case. Thus, in order to control the regret, we only have to show that  $N_t(\star) > u_t$  indeed happens with sufficiently high probability for large enough  $t$  (such as  $t > 2u_t$  for instance).

STEP 3. In this step, we show that the optimal arm is pulled often with high probability. Let us introduce  $t_j$  to be the occurrence of the  $j^{\text{th}}$  play of arm  $\star$ , with  $t_0 = 0$ , and  $\tau_j = (t_{j+1} - 1) - t_j$  being the number of time steps during the  $j^{\text{th}}$  and the  $(j + 1)^{\text{th}}$  play of the optimal arm  $\star$ . Note that  $\tau_0 \leq 2$  since all arms are pulled once during the initialization step. Then, provided that  $t/u - 1 > 2$ , that is  $t > 3u$ , it holds that

$$\begin{aligned} \mathbb{P}[N_t(\star) \leq u] &= \mathbb{P}[N_t(a) \leq t - u] \\ &\leq \mathbb{P}[\exists j \in \{0, \dots, \lfloor u \rfloor\} : \tau_j \geq t/u - 1] \\ &= \mathbb{P}[\exists j \in \{1, \dots, \lfloor u \rfloor\} : \tau_j \geq t/u - 1] \\ &\leq \sum_{j=1}^{\lfloor u \rfloor} \mathbb{P}[\tau_j \geq t/u - 1]. \end{aligned} \quad (10)$$

We now focus on the event  $E_j = \{\tau_j \geq t/u - 1\}$ . We want to show that it is unlikely that  $\star$  is not chosen for a large amount of consecutive time steps, due to the high probability that two small sub-samples chosen from arm  $a$  are independent of each other. Note that  $a_s = a$  happens if either  $\tilde{\mu}_{s,a} > \tilde{\mu}_{s,\star}$  or  $\tilde{\mu}_{s,a} = \tilde{\mu}_{s,\star}$  and  $\xi_s = 1$ . Expliciting further the event  $\tau_j \geq t/u - 1$ , it holds that

$$\begin{aligned} \mathbb{P}[E_j] &\leq \mathbb{P}[\forall s \in \{t_j, \dots, t_j + \lfloor t/u \rfloor - 1\} : N_s(a) \geq j \\ &\cap \left\{ \hat{\mu}\left(X_{1:N_s(a)}^a(I_s(N_s(a); j))\right) > \hat{\mu}\left(X_{1:j}^\star\right) \cup \right. \\ &\left. \hat{\mu}\left(X_{1:N_s(a)}^a(I_s(N_s(a); j))\right) = \hat{\mu}\left(X_{1:j}^\star\right) \cap \xi_s = 1 \right\}] \\ &\quad + \mathbb{P}[\forall s \in \{t_j, \dots, t_j + \lfloor t/u \rfloor - 1\} : N_s(a) < j \\ &\cap \left\{ \hat{\mu}\left(X_{1:N_s(a)}^a\right) > \hat{\mu}\left(X_{1:j}^\star(I_s(j; N_s(a)))\right) \cup \right. \\ &\left. \hat{\mu}\left(X_{1:N_s(a)}^a\right) = \hat{\mu}\left(X_{1:j}^\star(I_s(j; N_s(a)))\right) \cap \xi_s = 1 \right\}]. \end{aligned}$$

We now show that the second term of the sum in right hand side of this inequality is 0 for large enough  $t$ . Indeed, since on the one hand  $N_{t_j + \lfloor t/u \rfloor - 1}(a) = t_j + \lfloor t/u \rfloor - 1 - j$ , and on the other hand  $t_j \geq j$ , we deduce that if  $t \geq u(1 + u)$ , then

$$t_j + \lfloor t/u \rfloor - 1 - j \geq \lfloor u \rfloor \geq j,$$

which shows the desired claim. Note that  $t \geq u_t(1 + u_t)$  happens for  $t \geq c_\Delta$  where  $c_\Delta$  is a constant depending on the sub-optimality gap  $\Delta$  only.

As a result, for  $t \geq u(1 + u)$ , we obtain for all  $1 \leq j \leq u$  the following control

$$\begin{aligned} \mathbb{P}[E_j] \leq & \mathbb{P}\left[\forall s \in \{t_j, \dots, t_j + \lfloor t/u \rfloor - 1\} : a_s = a \cap N_s(a) \geq j \right. \\ & \cap \left\{ \widehat{\mu}\left(X_{1:N_s(a)}^a(I_s(N_s(a); j))\right) > \widehat{\mu}\left(X_{1:j}^*\right) \cup \right. \\ & \left. \left. \widehat{\mu}\left(X_{1:N_s(a)}^a(I_s(N_s(a); j))\right) = \widehat{\mu}\left(X_{1:j}^*\right) \cap \xi_s = 1 \right\} \right]. \end{aligned}$$

Now, following a similar step as in [16], let us split the interval  $\mathcal{M}_{j,t} = [t_j, t_j + \lfloor t/u \rfloor - 1]$  in two parts

$$\begin{aligned} \mathcal{M}_{j,t}^{(1)} &= [t_j, t_j + \lfloor \frac{t/u - 1}{2} \rfloor] \\ \mathcal{M}_{j,t}^{(2)} &= [t_j + \lceil \frac{t/u - 1}{2} \rceil, t_j + \lfloor t/u \rfloor - 1]. \end{aligned}$$

When  $a$  is repeatedly chosen during all  $\mathcal{M}_{j,t}^{(1)}$ , then, for all  $s \in \mathcal{M}_{j,t}^{(2)}$ , we must have

$$\begin{aligned} N_s(a) &= s - j \geq t_j + \lceil \frac{t/u - 1}{2} \rceil - j \\ &\geq \lceil \frac{t/u - 1}{2} \rceil. \end{aligned}$$

For instance if  $t \geq u_t(1 + 2ku_t)$ , which happens for  $t \geq c_{k,\Delta}$  where  $c_{k,\Delta}$  is a constant independent on  $T$ , then for all  $s \in \mathcal{M}_{j,t}^{(2)}$  it holds that  $s - j \geq ku_t \geq kj$ . thus, it makes sense to introduce the deterministic quantity  $n_{t,kj} = \max\{\lceil \frac{t/u_t - 1}{2} \rceil, kj\}$ . Note that we have the following inequalities

$$\begin{aligned} & \mathbb{P}[E_j] \\ \leq & \mathbb{P}\left[\forall s \in \mathcal{M}_{j,t}^{(2)} : a_s = a \cap N_s(a) \geq n_{t,kj} \cap \right. \\ & \left\{ \widehat{\mu}\left(X_{1:N_s(a)}^a(I_s(N_s(a); j))\right) > \widehat{\mu}\left(X_{1:j}^*\right) \cup \right. \\ & \left. \left. \widehat{\mu}\left(X_{1:N_s(a)}^a(I_s(N_s(a); j))\right) = \widehat{\mu}\left(X_{1:j}^*\right) \cap \xi_s = 1 \right\} \right] \\ \leq & \mathbb{P}\left[\forall s \in \mathcal{M}_{j,t}^{(2)} : \widehat{\mu}\left(X_{1:s-j}^a(I_s(s-j; j))\right) > \widehat{\mu}\left(X_{1:j}^*\right) \right. \\ & \left. \cup \widehat{\mu}\left(X_{1:s-j}^a(I_s(s-j; j))\right) > \widehat{\mu}\left(X_{1:j}^*\right) \cap \xi_s = 1 \right], \end{aligned} \tag{11}$$

where  $|\mathcal{M}_{j,t}^{(2)}| \geq \overline{M}_t \stackrel{\text{def}}{=} \lfloor \frac{t/u_t - 1}{2} \rfloor$ .

STEP 4. So far, our analysis is not essentially different from what would be the regret analysis of the follow-the-leader algorithm, which is known to fail achieving reasonable regret bounds in a number of situations. Thus, the following is the key part

of the analysis, that proves **BESA** is a sound strategy. In order to control the event appearing in the last line of (11), we decompose it by considering the following event:

$$\mathcal{F}_{j,M} = \left\{ \exists s_1 < \dots < s_M \in \mathcal{M}_{j,t}^{(2)} : \forall m < m' \in [M], \right. \\ \left. I_{s_m}(s_m - j; j) \cap I_{s_{m'}}(s_{m'} - j; j) = \emptyset \right\}.$$

Using  $\mathcal{F}_{j,M}$  and the assumption of independence between the samples of arm  $a$ , we deduce the following bound, valid for any  $M \leq \overline{M}_t$

$$\begin{aligned} \mathbb{P}[E_j] &\leq \mathbb{P}[\mathcal{F}_{j,M}^c] \\ &+ \mathbb{P}[\forall s \in \mathcal{M}_{j,t}^{(2)} : \mathcal{F}_{j,M} \cap \\ &\left\{ \widehat{\mu}(X_{1:j}^{\prime a, (s)}) > \widehat{\mu}(X_{1:j}^*) \cup \widehat{\mu}(X_{1:j}^{\prime a, (s)}) = \widehat{\mu}(X_{1:j}^*) \cap \xi_s = 1 \right\}] \\ &\leq \mathbb{P}[\mathcal{F}_{j,M}^c] + \alpha_{1/2}(M, j), \end{aligned}$$

where  $X_{1:j}^{\prime a, (s)} = X_{1:s-j}^a(I_s(s-j; j))$ , with the sub-sampled index set  $I_s \sim \text{Wr}(s-j; j)$ , and where we introduced in the second line the following distribution-dependent quantity  $\alpha_{1/2}(M, j)$ , defined by

$$\alpha_{1/2}(M, j) = \mathbb{E}_{Z \sim \nu_{*,j}} \left[ \left( 1 - F_{\nu_{a,j}}(Z) + \frac{1}{2} \nu_{a,j}(Z) \right)^M \right],$$

where  $\nu_{a,j}$  is the distribution of the sum of  $j$  i.i.d variables following  $\nu_a$  and  $F_\nu(x) = \mathbb{P}_{X \sim \mu}(X \leq x)$  is the repartition function of  $\nu$ .

We now resort to Lemma 1 that enables to control the term  $\mathbb{P}[\mathcal{F}_{j,M}^c]$  in terms of  $n_{t,kj}$ ,  $j$ , and  $M$ . From Lemma 1, we deduce that

$$\begin{aligned} \mathbb{P}[\mathcal{F}_{j,M}^c | t_j] &= f_{\mathcal{M}_{j,t}^{(2)}}(M, j), \\ &= f_{\left[ t_j + \lceil \frac{t/u-1}{2} \rceil, t_j + \lfloor t/u \rfloor - 1 \right]}(M, j) \\ &\leq f_{\left[ j + \lceil \frac{t/u-1}{2} \rceil, j + \lfloor t/u \rfloor - 1 \right]}(M, j), \end{aligned}$$

since  $t_j \geq j$ . Thus, we deduce that

$$\mathbb{P}[\mathcal{F}_{j,M}^c | t_j] \leq f_{[p_t, q_t]}(M, j),$$

where  $p_t = j + \lceil \frac{t/u-1}{2} \rceil$  and  $q_t = O(t/u) = O(t/\log(t))$ .

We thus deduce the following bound

$$\mathbb{P}[E_j] \leq \alpha_{1/2}(M_{t,j}^*, j) + f_{[p_t, q_t]}(M_{t,j}^*, j). \quad (12)$$

STEP 5. We now combine the previous steps together. Note that  $t \geq u_t(u_t + 1)$  happens for  $t \geq c_\Delta$  where  $c_\Delta$  is a constant that depends only on  $\Delta$  (and not on  $T$ ). More precisely, combining (10) and (12), it holds that

$$\begin{aligned} & \sum_{t=1}^T \mathbb{P}[N_t(\star) \leq u_t] \leq c_\Delta \\ & + \sum_{t \geq c_\Delta} \sum_{j=1}^{\lfloor u_t \rfloor} \left( \alpha(M_{t,j}^*, j) + f_{[p_t, q_t]}(M_{t,j}^*, j) \right), \end{aligned} \quad (13)$$

where  $q_t - p_t = O(t/\log(t))$ . Under the assumption of the first part of the theorem, we have  $\alpha(M, 1) \leq c\alpha^M$  with  $\alpha < 1$ . Moreover, since  $\alpha(M, j) \leq \alpha(M, 1)$ , we can choose the value  $M_{t,j}^* = 1 \vee \frac{1}{\log(1/\alpha)} \left( (1+\varepsilon) \log(t) + \log \log(t) \right)$  for some small  $\varepsilon > 0$  and deduce by Lemma 1 that

$$\begin{aligned} \sum_{t=1}^T \mathbb{P}[N_t(\star) \leq u_t] & \leq c_\Delta + c_{\alpha, \Delta} \sum_{t=1}^T t^{-(1+\varepsilon)} + O(1) \\ & \leq c_\Delta + \frac{c_{\alpha, \Delta}}{\varepsilon} + O(1). \end{aligned}$$

where  $c_{\alpha, \Delta} = c \left( 1 + \sqrt{\frac{1}{\alpha+1}} \right)^2 \frac{2\alpha}{\Delta^2}$ .

Now, under the assumption that  $\alpha(1, n) \leq c\beta^n$  with  $\beta < 1$ , we use a slightly modified strategy that pulls each arm  $n_0$  times first, where

$$n_{0, T} \geq \frac{\ln(T) - \ln((1-\beta)C)}{\ln(1/\beta)}.$$

In that case, we have on the one hand that

$$\begin{aligned} \sum_{t \geq c_\Delta} \sum_{j \geq n_0}^{u_t} \alpha(M, j) & \leq c \frac{(1-\beta)C}{T} \frac{T - c_\Delta + 1}{1-\beta} \\ & \leq cC. \end{aligned}$$

We then control the second term by choosing  $M_{t,j}$  such that  $M_{t,j}j = O(\log(T))$ , and obtain

$$\sum_{t=1}^T \mathbb{P}[N_t(\star) \leq u_t] \leq c_\Delta + cC + O(1).$$

Thus, if we define  $n_{0, T} = 1$  for the first case, we can show that in the two cases,

$$\sum_{t=1}^T \mathbb{P}[N_t(\star) \leq u_t] \leq c_{\nu_a, \nu_\star} + n_{0, T} + O(1),$$

with  $c_{\nu_a, \nu_\star}$  a problem dependent bound, and a  $O(1)$  that hides a numerical, problem-independent constant.

STEP 6. Final regret bound. Plugging-in the result of step 2 in the decomposition (6), we deduce that the expected regret of **BESA** is upper-bounded at time  $T$  by

$$\begin{aligned}
\mathfrak{R}_T &\leq \Delta \sum_{t=1}^T \mathbb{P} \left[ \tilde{\mu}_{t,a} > \tilde{\mu}_{t,\star} \cap N_t(\star) > u_t \right] \\
&\quad + \Delta \sum_{t=1}^T \mathbb{P} [N_t(\star) \leq u_t] \\
&\leq \Delta \sum_{t=1}^T 2 \left( \frac{\log(t+1)}{\log((\alpha+1)/2)} + 2 \right) t^{-\frac{2\alpha}{\alpha+1}} \\
&\quad + \Delta \mathbb{E} \left[ \sum_{t=1}^T \mathbb{I}\{A_t = a \cap N_t(a) \leq u_t\} \right] \\
&\quad + \Delta \sum_{t=1}^T \mathbb{P} [N_t(\star) \leq u_t] .
\end{aligned}$$

Now, using (13) and reorganizing the terms, we obtain the regret bound

$$\begin{aligned}
\mathfrak{R}_T &\leq \Delta u_T + \Delta c_{\nu_a, \nu_\star} + n_{0,T} + O(1) \\
&\quad + 2\Delta \left( \frac{1}{\log((\alpha+1)/2)} + 2 \right) \sum_{t=1}^T \log(t+1) t^{-\frac{2\alpha}{\alpha+1}} ,
\end{aligned}$$

Let us simplify a little bit this bound. For  $\beta > 1$ , a simple sum-integral comparison lemma gives that

$$\begin{aligned}
\sum_{t=1}^T t^{-\beta} &\leq 1 + \int_{t=1}^T t^{-\beta} \\
&\leq 1 + \frac{1}{\beta-1} (1 - T^{1-\beta}) \\
&\leq \frac{\beta}{\beta-1} .
\end{aligned} \tag{14}$$

Thus, since  $\log(t+1) \leq t^{0.38}$ , we apply this result for  $\beta = \frac{2\alpha}{\alpha+1} - 0.38$ , provided that  $\alpha > 1.38/0.62$ . This leads to  $\frac{\beta}{\beta-1} = \frac{1.62\alpha - 0.38}{0.62\alpha - 1.38}$ . Now, on the other hand, we have

$$\sum_{j=0}^{\lfloor u_t \rfloor} \exp \left( -j\Delta^2/2 \right) \leq \frac{1}{1 - \exp(-\Delta^2/2)} \leq \frac{4}{\Delta^2} . \tag{15}$$

Thus, for  $\alpha \geq 2.23$ , from (14) and (15), the regret bound simplifies to

$$\begin{aligned}
\mathfrak{R}_T &\leq \left( 1 + \sqrt{\frac{1}{\alpha+1}} \right)^2 \frac{2\alpha \log(T)}{\Delta} + \frac{4}{\Delta} + n_{0,T} + \Delta c_{\nu_a, \nu_\star} \\
&\quad + 2\Delta \left( \frac{1}{\log((\alpha+1)/2)} + 2 \right) \frac{1.62\alpha - 0.38}{0.62\alpha - 1.38} + O(1) .
\end{aligned} \tag{16}$$



For a specific choice of  $\alpha$ , we get

$$\mathfrak{R}_T \leq \frac{11 \log(T)}{\Delta} + n_{0,T} + \frac{4}{\Delta} + (26 + c_{\nu_a, \nu_*})\Delta + O(1) \quad (17)$$

We conclude the proof by setting the constant  $C_{\nu_a, \nu_*} = \frac{4}{\Delta} + (26 + c_{\nu_a, \nu_*})\Delta$ .