# Sub-sampling for Efficient Non-Parametric Bandit Exploration 

Dorian Baudry<br>dorian.baudry@inria.fr

Emilie Kaufmann<br>emilie.kaufmann@univ-lille.fr

Odalric-Ambrym Maillard<br>odalric.maillard@inria.fr

Univ. Lille, CNRS, Inria, Centrale Lille, UMR 9198-CRIStAL, F-59000 Lille, France


#### Abstract

In this paper we propose the first multi-armed bandit algorithm based on resampling that achieves asymptotically optimal regret simultaneously for different families of arms (namely Bernoulli, Gaussian and Poisson distributions). Unlike Thompson Sampling which requires to specify a different prior to be optimal in each case, our proposal RB-SDA does not need any distribution-dependent tuning. RB-SDA belongs to the family of Sub-sampling Duelling Algorithms (SDA) which combines the sub-sampling idea first used by the BESA [1] and SSMC [2] algorithms with different sub-sampling schemes. In particular, RB-SDA uses Random Block sampling. We perform an experimental study assessing the flexibility and robustness of this promising novel approach for exploration in bandit models.


## 1 Introduction

A $K$-armed bandit problem is a sequential decision-making problem in which a learner sequentially samples from $K$ unknown distributions called arms. In each round the learner chooses an arm $A_{t} \in\{1, \ldots, K\}$ and obtains a random reward $X_{t}$ drawn from the distribution of the chosen arm, that has mean $\mu_{A_{t}}$. The learner should adjust her sequential sampling strategy $\mathcal{A}=\left(A_{t}\right)_{t \in \mathbb{N}}$ (or bandit algorithm) in order to maximize the expected sum of rewards obtained after $T$ selections. This is equivalent to minimizing the regret, defined as the difference between the expected total reward of an oracle strategy always selecting the arm with largest mean $\mu_{\star}$ and that of the algorithm:

$$
\mathcal{R}_{T}(\mathcal{A})=\mu_{\star} T-\mathbb{E}\left[\sum_{t=1}^{T} X_{t}\right]=\mathbb{E}\left[\sum_{t=1}^{T}\left(\mu_{\star}-\mu_{A_{t}}\right)\right]
$$

An algorithm with small regret needs to balance exploration (gain information about arms that have not been sampled a lot) and exploitation (select arms that look promising based on the available information). Many approaches have been proposed to solve this exploration-exploitation dilemma (see [3] for a survey), the most popular being Upper Confidence Bounds (UCB) algorithms [4, 5, 6] and Thompson Sampling (TS) [7, 8]. TS is a randomized Bayesian algorithm that selects arms according to their posterior probability of being optimal. These algorithms enjoy logarithmic regret under some assumptions on the arms, and some of them are even asymptotically optimal in that they attain the smallest possible asymptotic regret given by the lower bound of Lai \& Robbins [9], for some parametric families of distributions. For distributions that are continuously parameterized by their means, this lower bound states that under any uniformly efficient algorithm,

$$
\begin{equation*}
\liminf _{T \rightarrow \infty} \frac{\mathcal{R}_{T}(\mathcal{A})}{\log (T)} \geq \sum_{k: \mu_{k}<\mu_{\star}} \frac{\left(\mu_{\star}-\mu_{k}\right)}{\mathrm{kl}\left(\mu_{k}, \mu_{\star}\right)} \tag{1}
\end{equation*}
$$

where $\operatorname{kl}\left(\mu, \mu^{\prime}\right)$ is the Kullback-Leibler divergence between the distribution of mean $\mu$ and that of mean $\mu^{\prime}$ in the considered family of distributions. For arms that belong to a one-parameter exponential
family (e.g. Bernoulli, Gaussian, Poisson arms) kl-UCB using an appropriate divergence function [6] and Thompson Sampling using an appropriate prior distribution [10, 11, 12] are both asymptotically optimal in the sense that their regret matches that prescribed by the lower bound (1), for large values of $T$. Yet, a major drawback of theses algorithms is that their optimal tuning requires the knowledge of the families of distributions they operate on. In this paper, we overcome this issue and propose an algorithm that is simultaneously asymptotically optimal for several families of distributions.

In the past years, there has been a surge of interest in the design of non-parametric algorithms that directly use the empirical distribution of the data instead of trying to fit it in an already defined model, and are therefore good candidates to meet our goal. In [13], the authors propose the General Randomized Exploration (GRE) framework in which each arm $k$ is assigned an index $\hat{\mu}_{k, t}$ sampled from a distribution $p\left(\mathcal{H}_{k, t}\right)$ that depends on the history of past observed rewards for this arm $\mathcal{H}_{k, t}$, and the arm with largest index is selected. GRE includes Thompson Sampling (for which $p\left(\mathcal{H}_{k, t}\right)$ is the posterior distribution given a specified prior) but also allows for more general non-parametric re-sampling schemes. However, the authors of [13, 14] show that setting $p\left(\mathcal{H}_{k, t}\right)$ to be the nonparametric Bootstrap [15] leads to linear regret. They propose variants called GIRO and PHE which perturb the history by augmenting it with fake samples. History perturbation was already suggested by [16] and is also used by Reboot [17], with a slightly more complicated bootstrapping scheme. Finally, the recently proposed Non Parametric TS [18] does not use history perturbation but instead sets $\hat{\mu}_{k, t}$ as a weighted combination of all observations in $\mathcal{H}_{k, t}$ and the upper bound of the support, where the weights are chosen uniformly at random in the simplex of dimension $\left|\mathcal{H}_{k, t}\right|$.
Besides Reboot [17], which has been analyzed only for Gaussian distributions, all other algorithms have been analyzed for distributions with known bounded support, for which they are proved to have logarithmic regret. Among them, Non Parametric TS has strong optimality property as its regret is proved to match the lower bound of Burnetas and Katehakis [19] for (non-parametric) distribution that are bounded in $[0,1]$. In this paper, we propose the first re-sampling based algorithm that is asymptotically optimal for several classes of possibly un-bounded parametric distributions. We introduce a new family of algorithms called Sub-Sampling Duelling Algorithms, and provide a regret analysis for RB-SDA, an algorithm based on Random Block sub-sampling. In Theorem 3.1, we show that RB-SDA has logarithmic regret under some general conditions on the arms distributions. These conditions are in particular satisfied for Gaussian, Bernoulli and Poisson distribution, for which we further prove in Corollary 3.1.1 that RB-SDA is asymptotically optimal.
The general SDA framework that we introduce is inspired by two ideas first developed for the BESA algorithm by [1] and for the SSMC algorithm by [2]: 1) the arms pulled are chosen according to the outcome of pairwise comparison (duels) between arms, instead of choosing the maximum of some index computed for each arm as GRE algorithms do, and 2) the use of sub-sampling: the algorithm penalizes arms that have been pulled a lot by making them compete with the other arms with only a fraction of their history. More precisely, in a duel between two arms $A$ and $B$ selected $n_{A}$ and $n_{B}$ times respectively, with $n_{A}<n_{B}$, the empirical mean of arm $A$ is compared to the empirical mean of a sub-sample of size $n_{A}$ of the history of arm $B$. In BESA the duels are organized in a tournament and only the winner is sampled, while SSMC uses rounds of $K-1$ duels between an arm called leader and all other arms. Then the leader is pulled only if it wins all the duels, otherwise all the winning challengers are pulled. Second difference is that in BESA the sub-sample of the leader's history is obtained with Sampling Without Replacement, whereas SSMC selects this sub-sample as the block of consecutive observations with smallest empirical mean. Hence BESA uses randomization while SSMC does not. Finally, SSMC also uses some forced exploration (i.e. selects any arm drawn less than $\sqrt{\log r}$ times in round $r$ ). In SDA, we propose to combine the round structure for the duels used by SSMC with the use of a sub-sampling scheme assumed to be independent of the observations in the history (this generalizes the BESA duels), and we get rid of the use of forced exploration.
The rest of the paper is structured as follows. In Section 2 we introduce the SDA framework and present different instances that correspond to the choice of different sub-sampling algorithms, in particular RB-SDA. In Section 3 we present upper bounds on the regret of RB-SDA, showing in particular that the algorithm is asymptotically optimal for different exponential families. We sketch the proof of Theorem 3.1 in Section 4 , highlighting two important tools: First, a new concentration lemma for random sub-samples (Lemma 4.2). Second, an upper bound on the probability that the optimal arm is under-sampled, which decouples the properties of the sub-sampling algorithm used, and that of the arms' distributions (Lemma 4.3 . Finally, Section 5 presents the results of an empirical study comparing several instances of SDA to asymptotically optimal parametric algorithms and other
algorithms based on re-sampling or sub-sampling. These experiments reveal the robustness of the SDA approaches, which match the performance of Thompson Sampling, without exploiting the knowledge of the distribution.

## 2 Sub-sampling Duelling Algorithms

In this section, we introduce the notion of Sub-sampling Duelling Algorithm (SDA). We first introduce a few notation. For every integer $n$, we let $[n]=\{1, \ldots, n\}$. We denote by $\left(Y_{k, s}\right)_{s \in \mathbb{N}}$ the i.i.d. sequence of successive rewards from arm $k$, that are i.i.d. under a distribution $\nu_{k}$ with mean $\mu_{k}$. For every finite subset $\mathcal{S}$ of $\mathbb{N}$, we denote by $\hat{Y}_{k, \mathcal{S}}$ the empirical mean of the observations of arm $k$ indexed by $\mathcal{S}$ : if $|\mathcal{S}|>1, \hat{Y}_{k, \mathcal{S}}:=\frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} Y_{k, i}$. We also let $\hat{Y}_{k, n}$ as a shorthand notation for $\hat{Y}_{k,[n]}$.

A round-based algorithm Unlike index policies, a SDA algorithm relies on rounds, in which several arms can be played (at most once). In each round $r$ the learner selects a subset of arms $\mathcal{A}_{r}=\left\{k_{1}, \ldots, k_{i_{r}}\right\} \subseteq\{1, \ldots, K\}$, and receives the rewards $\mathcal{X}_{r}=\left\{Y_{k_{1}, N_{k_{1}}(r)}, \ldots, Y_{k_{i_{r}, N_{k_{i_{r}}}}(r)}\right\}$ associated to the chosen arms, where $N_{k}(r):=\sum_{s=1}^{r} \mathbb{1}\left(k \in \mathcal{A}_{s}\right)$ denotes the number of times arm $k$ was selected up to round $r$. Letting $\hat{r}_{T} \leq T$ be the (random) number of rounds used by algorithm $\mathcal{A}$ before the $T$-th arm selection, the regret of a round-based algorithm can be upper bounded as follows:

$$
\begin{align*}
\mathcal{R}_{T}(\mathcal{A}) & =\mathbb{E}\left[\sum_{t=1}^{T}\left(\mu_{\star}-\mu_{A_{t}}\right)\right] \leq \mathbb{E}\left[\sum_{s=1}^{\hat{r}_{T}} \sum_{k=1}^{K}\left(\mu_{\star}-\mu_{k}\right) \mathbb{1}\left(k \in \mathcal{A}_{s}\right)\right] \\
& \leq \mathbb{E}\left[\sum_{s=1}^{T} \sum_{k=1}^{K}\left(\mu_{\star}-\mu_{k}\right) \mathbb{1}\left(k \in \mathcal{A}_{s}\right)\right]=\sum_{k=1}^{K}\left(\mu_{\star}-\mu_{k}\right) \mathbb{E}\left[N_{k}(T)\right] \tag{2}
\end{align*}
$$

Hence upper bounding $\mathbb{E}\left[N_{k}(T)\right]$ for each sub-optimal arm provides a regret upper bound.
Sub-sampling Duelling Algorithms A SDA algorithm takes as input a sub-sampling algorithm $\mathrm{SP}(m, n, r)$ that depends on three parameters: two integers $m \geq n$ and a round $r$. A call to $\mathrm{SP}(m, n, r)$ at round $r$ produces a subset of $[m]$ that has size $n$, modeled as a random variable that is further assumed to be independent of the rewards generated from the arms, $\left(Y_{k, s}\right)_{k \in[K], s \in \mathbb{N}^{*}}$.

In the first round, a SDA algorithm selects $\mathcal{A}_{1}=[K]$ in order to initialize the history of all arms. For $r \geq 1$, at round $r+1$, a SDA algorithm based on a sampler SP, that we refer to as SP-SDA, first computes the leader, defined as the arm being selected the most in the first $r$ round: $\ell(r)=\operatorname{argmax}_{k} N_{k}(r)$. Ties are broken in favor of the arm with the largest mean, and if several arms share this mean then the previous leader is kept or one of these arms is chosen randomly. Then the set $\mathcal{A}_{r+1}$ is initialized to the empty set and $K-1$ duels are performed. For each "challenger" arm $k \neq \ell(r)$, a subset $\hat{\mathcal{S}}_{k}^{r}$ of $\left[N_{\ell(r)}(r)\right]$ of size $N_{k}(r)$ is obtained from $\operatorname{SP}\left(N_{\ell(r)}(r), N_{k}(r), r\right)$ and arm $k$ wins the duels if its empirical mean is larger than the empirical mean of the sub-sampled history of the leader. That is

$$
\hat{Y}_{k, N_{k}(r)}>\hat{Y}_{\ell(r), \hat{\mathcal{S}}_{k}^{r}} \Longrightarrow \mathcal{A}_{r+1}=\mathcal{A}_{r+1} \cup\{k\} .
$$

If the leader wins all the duels, that is if $\mathcal{A}_{r+1}$ is still empty after the $K-1$ duels, we set $\mathcal{A}_{r+1}=$ $\{\ell(r)\}$. Arms in $\mathcal{A}_{r+1}$ are then selected by the learner in a random order and are pulled if the total budget of pulls remains smaller than $T$. The pseudo-code of SP-SDA is given in Algorithm 1 .
To properly define the random variable $\hat{S}_{k}^{r}$ used in the algorithm, we introduce the following probabilistic modeling: for each round $r$, each arm $k$, we define a family $\left(S_{k}^{r}(m, n)\right)_{m \geq n}$ of independent random variables such that $S_{k}^{r}(m, n) \sim \operatorname{SP}(m, n, r)$. In words, $S_{k}^{r}(m, n)$ is the subset of the leader history used should arm $k$ be a challenger drawn $n$ times up to round $r$ dueling against a leader that has been drawn $m$ times. With this notation, for each arm $k \neq \ell(r)$ one has $\hat{S}_{k}^{r}=S_{k}^{r}\left(N_{\ell(r)}(r), N_{k}(r), r\right)$. We recall that in the SDA framework, it is crucial that those random variables are independent from the reward streams $\left(Y_{k, s}\right)$ of all arms $k$. We call such sub-sampling algorithms independent sampler.

Particular instances We now present a few sub-sampling algorithms that we believe are interesting to use within the SDA framework. Intuitively, these algorithms should ensure enough diversity in the output subsets when called in different rounds, so that the leader cannot always look good,

```
Algorithm 1 SP-SDA
Require: K arms, horizon T, Sampler SP
    \(t \leftarrow K, r \leftarrow 1, \forall k, N_{k} \leftarrow 1, \mathcal{H}_{k} \leftarrow\left\{Y_{k, 1}\right\}\) (Each arm is drawn once)
    while \(t<T\) do
        \(r \leftarrow r+1, \mathcal{A} \leftarrow\{ \}, \ell \leftarrow \operatorname{leader}(N, \mathcal{H}, \ell)\) (Initialize the round)
        for \(k \neq \ell \in 1, \ldots, K\) do
            Draw \(\hat{S}_{k}^{r} \sim \operatorname{SP}\left(N_{\ell}, N_{k}, r\right)\) (Choice of the sub-sample of \(\ell\) used for the duel with \(k\) )
            if \(\hat{Y}_{k, N_{k}}>\hat{Y}_{\ell, \hat{S}_{k}^{r}}\) then
                \(\mathcal{A} \leftarrow \mathcal{A} \cup\{k\}\) (Duel outcome)
            end if
        end for
        if \(|\mathcal{A}|=0\) then
            \(\mathcal{A} \leftarrow\{\ell\}\)
        end if
        if \(|\mathcal{A}|>T-t\) then
            \(\mathcal{A} \leftarrow \operatorname{choose}(\mathcal{A}, T-t)\) (Randomly selects a number of arm that does not exceed the budget)
        end if
        for \(a \in \mathcal{A}\) do
            Pull arm \(a\), observe reward \(Y_{a, N_{a}+1}\)
            \(t \leftarrow t+1, N_{a} \leftarrow N_{a}+1, \mathcal{H}_{a} \leftarrow \mathcal{H}_{a} \cup\left\{Y_{a, N_{a}}\right\}\) (Update step)
        end for
    end while
```

and challengers may win and be explored from time to time. The most intuitive candidates are random samplers like Sampling Without Replacement (WR) and Random Block Sampling (RB): the first one returns a subset of size $n$ selected uniformly at random in $[m$, while the second draws an element $n_{0}$ uniformly at random in $[m-n]$ and returns $\left\{n_{0}+1, \ldots, n_{0}+n\right\}$. But we also propose two deterministic sub-sampling: Last Block (LB) which returns $\{m-n+1, \ldots, m\}$, and Low Discrepancy Sampling (LDS) that is similar to RB with the first element $n_{0}$ of the block at a round $r$ defined as $\left\lceil u_{r}(m-n)\right\rceil$ with $u_{r}$ a predefined low discrepancy sequence [20] (Halton [21], Sobol [22]). We believe that these last two samplers may ensure enough diversity without the need for random sampling. These four variants of SDA will be compared in Section 5 in terms of empirical efficiency and numerical complexity. For RB-SDA, we provide a regret analysis in the next sections, highlighting what parts may or may not be extended to other sampling algorithms.

Links with existing algorithms The BESA algorithm [1] with $K=2$ coincides with WR-SDA. However beyond $K>2$, the authors of [1] rather suggest a tournament approach, without giving a regret analysis. WR-SDA can therefore be seen as an alternative generalization of BESA beyond 2 arms, which performs much better than the tournament, as can be seen in Section5. While the structure of SSDA is close to that of SSMC [2], SSMC is not a SP-SDA algorithm, as its sub-sampling algorithm heavily relies on the rewards, and is therefore not an independent sampler. Indeed, it outputs the set $\mathcal{S}=\left\{n_{0}+1, \ldots, n_{0}+n\right\}$ for which $\hat{Y}_{\ell(r), \mathcal{S}}$ is the smallest. The philosophy of SSMC is a bit different than that of SSDA: while the former tries to disadvantage the leader as much as possible, the latter only tries to make the leader use different parts of its history. Our experiments reveal that the SSMC approach seems to lead to a slightly larger regret, due to a bit more exploration in the beginning. Finally, we emphasize that alternative algorithms based on re-sampling (PHE, Reboot, Non-Parametric TS) are fundamentally different to SDA as they do not perform sub-sampling.

On the use of forced exploration $\operatorname{In}[2]$, SSMC additionally requires some forced exploration: each arm $k$ such that $N_{k}(r)$ is smaller than some value $f_{r}$ is added to $\mathcal{A}_{r+1}$. SSMC is proved to be asymptotically optimal for exponential families provided that $f_{r}=o(\log r)$ and $\log \log r=o\left(f_{r}\right)$. In the next section, we show that RB-SDA does not need forced exploration to be asymptotically optimal for Bernoulli, Gaussian and Poisson distributions. However, we show in Appendix H that adding forced exploration to RB-SDA is sufficient to prove its optimality for any exponential family.

## 3 Regret Upper Bounds for RB-SDA

In this section, we present upper bounds on the expected number of selections of each sub-optimal arm $k, \mathbb{E}\left[N_{k}(T)\right]$, for the RB-SDA algorithm. They directly yield an upper bound on the regret via (22). To ease the presentation, we assume that there is a unique optimal arm ${ }^{1}$, and denote it by 1 .
In Theorem 3.1, we first identify some conditions on the arms distribution under which RB-SDA has a regret that is provably logarithmic in $T$. In order to introduce these conditions, we recall the definition of the following balance function, first introduced by [1]. $\alpha_{k}(M, j)$ is equal to the probability that arm 1 loses a certain amount $M$ of successive duels against $M$ sub-samples from arm $k$ that have non-overlapping support, when arm 1 has been sampled $j$ times.
Definition. Letting $\nu_{k, j}$ denote the distribution of the sum of $j$ independent variables drawn from $\nu_{k}$, and $F_{\nu_{k, j}}$ its corresponding CDF, the balance function of arm $k$ is

$$
\alpha_{k}(M, j)=\mathbb{E}_{X \sim \nu_{1, j}}\left(\left(1-F_{\nu_{k, j}}(X)\right)^{M}\right) .
$$

Theorem 3.1 (Logarithmic Regret for RB-SDA). If the arms distributions $\nu_{1}, \ldots, \nu_{k}$ are such that

1. the empirical mean of each arm $k$ has exponential concentration given by a certain rate function $I_{k}(x)$ which is continuous and satisfies $I_{k}(x)=0$ if and only if $x=\mu_{k}$ :

$$
\forall x>\mu_{k}, \mathbb{P}\left(\hat{Y}_{k, n} \geq x\right) \leq e^{-n I_{k}(x)} \text { and } \forall x<\mu_{k}, \mathbb{P}\left(\hat{Y}_{k, n} \leq x\right) \leq e^{-n I_{k}(x)}
$$

2. the balance function of each sub-optimal arm $k$ satisfies

$$
\forall \beta \in(0,1), \quad \sum_{t=1}^{T} \sum_{j=1}^{\left\lfloor(\log t)^{2}\right\rfloor} \alpha_{k}\left(\left\lfloor\beta t /(\log t)^{2}\right\rfloor, j\right)=o(\log T) .
$$

Then, for all sub-optimal arm $k$, for all $\varepsilon>0$, under RB-SDA

$$
\mathbb{E}\left[N_{k}(T)\right] \leq \frac{1+\varepsilon}{I_{k}\left(\mu_{1}\right)} \log (T)+o(\log T)
$$

If the distributions belong to the same one-dimensional exponential family (see e.g. [6] for a presentation of some of their important properties), the Chernoff inequality tells us that the concentration condition 1. is satisfied with a rate function equal to $I_{k}(x)=\operatorname{kl}\left(x, \mu_{k}\right)$ where $\mathrm{kl}(x, y)$ is the Kullback-Leibler divergence between the distribution of mean $x$ and the distribution of mean $y$ in that exponential family. In Appendix $G$, we prove that Gaussian distribution with known variance, Bernoulli and Poisson distribution also satisfy the balance condition 2., which yields the following.
Corollary 3.1.1. Assume that the distribution of all arms belong to the family of Gaussian distributions with a known variance, Bernoulli or Poisson distributions. Then under RB-SDA for all $\varepsilon>0$, for all sub-optimal arm $k$,

$$
\mathbb{E}\left[N_{k}(T)\right] \leq \frac{1+\varepsilon}{\operatorname{kl}\left(\mu_{k}, \mu_{1}\right)} \log (T)+o_{\varepsilon, \mu}(\log (T))
$$

Corollary 3.1.1 permits to prove that $\lim \sup _{T \rightarrow} \frac{\mathcal{R}_{T}(\text { RB-SDA })}{\log (T)} \leq \sum_{k=2}^{K} \frac{\left(\mu_{1}-\mu_{k}\right)}{\mathrm{kl}\left(\mu_{k}, \mu_{1}\right)}$, which is matching the Lai \& Robbins lower bound (1) in each of these exponential families. In particular, RB-SDA is simultaneously asymptotically optimal for different examples of bounded (Bernoulli) and un-bounded (Poisson, Gaussian) distributions. In contrast, Non-Parametric TS is asymptotically optimal for any bounded distributions, but cannot be used for Gaussian or Poisson distributions. Note that the guarantees of Corollary 3.1.1] also hold for the SSMC algorithm [2], but we prove that RB-SDA can be asymptotically optimal without forced exploration for some distributions. Moreover, as will be seen in Section 5, algorithms based on randomized history-independent sub-sampling such as RB-SDA tend to perform better than deterministic algorithms such as SSMC.
Theorem 3.1 also shows that RB-SDA may have logarithmic regret for a wider range of distributions. For example, we conjecture that a truncated Gaussian distribution also satisfy the balance condition

[^0]2.. On the other hand, condition 2. does not hold for Exponential distribution, as discussed in Appendix G. 4 But we show in Appendix H. 2 that any distribution that belongs to a one-dimensional exponential family satisfies a slightly modified version of this condition, which permits to establish the asymptotic optimality of a variant of RB-SDA using forced exploration.

Finally, we note that it is possible to build on RB-SDA to propose a bandit algorithm that has logarithmic regret for any distribution that is bounded in $[0,1]$. To do so, we can use the binarization trick already proposed by [11] for Thompson Sampling, and run RB-SDA on top of a binarized history $\mathcal{H}_{k}^{\prime}$ for each arm $k$ in which a reward $Y_{k, s}$ is replaced by a binary pseudo-reward is $Y_{k, s}^{\prime}$ generated from a Bernoulli distribution with mean $Y_{k, s}$. The resulting algorithm inherits the regret guarantees of RB-SDA applied to Bernoulli distributions.
Characterizing the set of distributions for which the vanilla RB-SDA algorithm has logarithmic regret (without forced exploration or a binarization trick) is left as an interesting future work.

## 4 Sketch of Proof

In this section, we provide elements of proof for Theorem 3.1, postponing the proof of some lemmas to the appendix. The first step is given by the following lemma, which is proved in Appendix D.
Lemma 4.1. Under condition 1., for any SP-SSDA algorithm (using an independent sampler), for every $\varepsilon>0$, there exists a constant $C_{k}(\boldsymbol{\nu}, \varepsilon)$ with $\boldsymbol{\nu}=\left(\nu_{1}, \ldots, \nu_{k}\right)$ such that

$$
\mathbb{E}\left[N_{k}(T)\right] \leq \frac{1+\varepsilon}{I_{1}\left(\mu_{k}\right)} \log (T)+32 \sum_{r=1}^{T} \mathbb{P}\left(N_{1}(r) \leq(\log (r))^{2}\right)+C_{k}(\boldsymbol{\nu}, \varepsilon)
$$

The proof of this result follows essentially the same decomposition as the one proposed by [2] for the analysis of SSMC. However, it departs from this analysis in two significant ways. First, instead of using properties of forced exploration (that is absent in RB-SDA), we distinguish whether or not arm 1 has been selected a lot, which yields the middle term in the upper bound. Then, the argument relies on a new concentration result for sub-samples averages, that we state below. Lemma 4.2, proved in Appendix C, crucially exploits the fact that a SP-SDA algorithm is based on an independent sampler. Using condition 1. allows to further upper bound the right-hand side of the two inequalities in Lemma 4.2 by terms that decay exponentially and contribute to the constant $C_{k}(\boldsymbol{\nu}, \varepsilon)$.
Lemma 4.2 (concentration of a sub-sample). For all $(a, b)$ such that $\mu_{a}<\mu_{b}$, for all $\xi \in$ $\left(\mu_{a}, \mu_{b}\right)$ and $n_{0} \in \mathbb{N}$, under any instance of SP-SDA using an independent sampler, it holds that

$$
\begin{aligned}
& \sum_{s=1}^{r} \mathbb{P}\left(\hat{Y}_{a, N_{a}(s)} \geq \hat{Y}_{b, \hat{\mathcal{S}}_{b}^{s}\left(N_{b}(s), N_{a}(s)\right)}, N_{b}(s) \geq N_{a}(s), N_{a}(s) \geq n_{0}\right) \leq \sum_{j=n_{0}}^{r} \mathbb{P}\left(\hat{Y}_{a, j} \geq \xi\right)+r \sum_{j=n_{0}}^{r} \mathbb{P}\left(Y_{b, j} \leq \xi\right), \\
& \sum_{s=1}^{r} \mathbb{P}\left(\hat{Y}_{b, N_{b}(s)} \leq \hat{Y}_{a, \hat{\mathcal{S}}_{a}^{s}\left(N_{a}(s), N_{b}(s)\right)}, N_{a}(s) \geq N_{b}(s), N_{b}(s) \geq n_{0}\right) \leq \sum_{j=n_{0}}^{r} \mathbb{P}\left(\hat{Y}_{b, j} \leq \xi\right)+r \sum_{j=n_{0}}^{r} \mathbb{P}\left(Y_{a, j} \geq \xi\right)
\end{aligned}
$$

So far, we note that the analysis has not been specific to RB-SDA but applies to any instance of SDA. Then, we provide in Lemma 4.3 an upper bound on $\sum_{t=1}^{T} \mathbb{P}\left(N_{1}(t) \leq(\log t)^{2}\right)$ which is specific to RB-SDA. This sampler is randomized and independent of $r$, hence we use the notation $\mathrm{RB}(m, n)=\mathrm{RB}(m, n, r)$. The strength of this upper bound is that it decouples the properties of the sub-sampling algorithm and that of the arm distributions (through the balance function $\alpha_{k}$ ).
Lemma 4.3. Let $X_{m, H, j}$ be a random variable giving the number of non-overlapping sub-samples of size $j$ obtained in $m$ i.i.d. samples from $\operatorname{RB}(H, j)$ and define $c_{r}=\left\lfloor\frac{r /(\log r)^{2}-1}{2 K}\right\rfloor-1$. There exists $\gamma \in(0,1)$ and a constant $r_{K}$ such that with $\beta_{r, j}=\left\lfloor\gamma r / j(\log r)^{2}\right\rfloor$,

$$
\sum_{r=1}^{T} \mathbb{P}\left(N_{1}(r) \leq(\log r)^{2}\right) \leq r_{K}+\sum_{r=r_{K}}^{T} \sum_{j=1}^{\left\lfloor\log r^{2}\right\rfloor}\left[(K-1) \mathbb{P}\left(X_{c_{r}, c_{r}, j}<\beta_{r, j}\right)+\sum_{k=2}^{K} \alpha_{k}\left(\beta_{r, j}, j\right)\right]
$$

To prove Lemma 4.3 (see Appendix E], we extend the proof technique introduced by [1] for the analysis of BESA to handle more than 2 arms. The rationale is that if $N_{1}(r) \leq(\log r)^{2}$ then arm 1 is not the leader and has lost "many" duels, more precisely at least a number of successive duels
proportional to $r /(\log r)^{2}$. A fraction of these duels necessarily involves sub-samples of the leader history that have non-overlapping support. Exploiting the independence of these sub-samples brings in the balance function $\alpha_{k}$.

In order to conclude the proof, it remains to upper bound the right hand side of Lemma 4.3 Using condition 2. of balanced distributions the terms depending on $\alpha_{k}$ sum in $o(\log T)$ and negligibly contribute to the regret. Upper bounding the term featuring $X_{m, H, j}$ amounts to establishing the following diversity property of the random block sampler.
Definition (Diversity Property). Let $X_{m, H, j}$ be the random variable defined in Lemma 4.3 for a randomized sampler SP. SP satisfies the Diversity Property for a sequence $N_{r}$ of integers if

$$
\sum_{r=1}^{T} \sum_{j=1}^{(\log r)^{2}} \mathbb{P}\left(X_{N_{r}, N_{r}, j}<\gamma r /(\log r)^{2}\right)=o(\log T) .
$$

We prove in Appendix $F$ that the RB sampler satisfies the diversity property for the sequence $c_{r}$, which leads to $\sum_{t=1}^{T} \mathbb{P}\left(N_{1}(t) \leq(\log t)^{2}\right)=o(\log (T))$ and concludes the proof of Theorem 3.1 .
We believe that the WR sampler also satisfies the diversity property (as conjectured by [1]). While Lemma 4.3 should apply to WR-SDA as well, a different path has to be found for analyzing the LDS-SDA and LB-SDA algorithms, that are based on deterministic samplers and also perform well in practice. This is left for future work.

## 5 Experiments

In this section, we perform experiments on simulated data in order to illustrate the good performance of the four instances of SDA algorithms introduced in Section 2 for various distributions. The Python code used to perform these experiments is available on Github

Exponential families First, in order to illustrate Corollary 3.1.1, we investigate the performance of RB-SDA for both Bernoulli and Gaussian distributions (with known variance 1). Our first objective is to check that for a finite horizon the regret of RB-SDA is comparable with the regret of Thompson Sampling (with respectively a beta and improper uniform prior), which efficiently uses the knowledge of the distribution. Our second objective is to empirically compare different variants of SDA to other non-parametric approaches based on sub-sampling (BESA, SSMC) or on re-sampling. For Bernoulli and Gaussian distribution, Non-Parameteric TS coincides with Thompson Sampling, so we focus our study on algorithms based on history perturbation. We experiment with PHE [14] for Bernoulli bandits and ReBoot [17] for Gaussian bandits, as those two algorithms are guaranteed to have logarithmic regret in each of these settings. As advised by the authors, we use a parameter $a=1.1$ for PHE and $\sigma=1.5$ for ReBoot.

We ran experiments on 4 different Bernoulli bandit models: 1) $K=2, \mu=[0.8,0.9]$, 2) $K=2$, $\mu=[0.5,0.6]$, 3) $K=10, \mu_{1}=0.1, \mu_{2,3,4}=0.01, \mu_{5,6,7}=0.03, \mu_{8,9,10}=0.05$, 4) $K=8$ $\mu=[0.9,0.85, \ldots, 0.85]$ and 3 different bandits models with $\mathcal{N}\left(\mu_{k}, 1\right)$ arms with means: 1) $K=2$ $\mu=[0.5,0]$, 2) $K=4, \mu=[0.5,0,0,0], 3) K=4, \mu=[1.5,1,0.5,0]$. For each experiment, Table 1 and Table 2 report an estimate of the regret at time $T=20000$ based on 5000 independent runs (extended tables with standard deviations can be found in Appendix A.1). The best performing algorithms are highlighted in bold. In Figure 1 and Figure 2 we plot the regret of several algorithms as a function of time (in $\log$ scale) for $t \in[15000 ; 20000]$ for one Bernoulli and one Gaussian experiment respectively. We also add the Lai and Robbins lower bound $t \mapsto\left[\sum_{k}\left(\mu^{\star}-\mu_{k}\right) / \mathrm{kl}\left(\mu_{k}, \mu_{\star}\right)\right] \log (t)$.

Table 1: Regret at $T=20000$ for Bernoulli arms

| xp | Benchmark |  |  |  | SDA |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | TS | \|PHE | \|BESA | \|SSMC | \|RB | \|WR | \|LB | \|LDS |
| 1 | 11.2 | 125.9 | 11.7 | 12.3 | 11.5 | 11.6 | 12.2 | 11.4 |
| 2 | 22.9 | 24.0 | 22.1 | 24.3 | 22.0 | 21.5 | 24.0 | 21.8 |
| 3 | 94.2 | 248.1 | 88.1 | 100.1 | 89.0 | 86.9 | 100.7 | 89.2 |
| 4 | 108.1 | 216.5 | 147.5 | 119.9 | 105.1 | 106.9 | 119.6\| | 106.8 |

Figure 1: Regret as a function of time for Bernoulli experiment 3


Table 2: Regret at $T=20000$ for Gaussian arms

|  | Benchmark |  |  |  | SDA |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| xp | TS | $\mid$ ReBoot | BESA | SSMC | \|RB | \|WR | \|LB | \|LDS |
| 1 | 24.4 | 92.2 | 25.3 | 26.9 | 25.6 | 24.7 | 25.1 | 26.5 |
| 2 | 73.5 | 277.1 | 122.5 | 74.8 | 71.0 | 71.1 | 74.6 | 69.0 |
| 3 | 49.7 | 190.9 | 72.1 | 51.3 | 50.4 | 50.0 | 51.2 | 48.6 |

Figure 2: Regret as a function of time for Gaussian experiment 2


In all of these experiments, we notice that SDA algorithms are indeed strong competitors to Thompson Sampling (with appropriate prior) for both Bernoulli and Gaussian bandits. Figures 1 and 2 further show that RB-SDA is empirically matching the Lai and Robbins' lower bound on two instances, just like SSMC and Thompson Sampling, which can be seen from the parallel straight lines with the $x$ axis in $\log$ scale. The fact that the lower bound is above shows that it is really asymptotic and only captures the right first order term. The same observation was made for all experiments, but is not reported due to space limitation. Even if we only established the asymptotic optimality of RB-SDA, these results suggest that the other SDA algorithms considered in this paper may also be asymptotically optimal. Compared to SDA, re-sampling algorithms based on history perturbation seem to be much less robust. Indeed, in the Bernoulli case, PHE performs very well for experiment 2, but is significantly worse than Thompson Sampling on the three other instances. In the Gaussian case, ReBoot always performs significantly worse than other algorithms. This lack of robustness is also corroborated by additional experiments reported in Appendix ?? in which we average the performance of these algorithms over a large number of randomly chosen instances.

Turning our attention to algorithms based on sub-sampling, we first notice that WR-SDA seems to be a better generalization of BESA with 2 arms than the tournament approach currently proposed, as in experiments with $K>2$, WR-SDA often performs significantly better than BESA. Then we observe that SSMC and SDA algorithms have similar performance. Looking a bit closer, we see that the performance of SSMC is very close to that of LB-SDA, whereas SDA algorithms based on "randomized" (or pseudo-randomized for LDS-SDA) samplers tend to perform slightly better.

Truncated Gaussian Theorem 3.1 suggests that RB-SDA may attain logarithmic regret beyond exponential families. As an illustration, we present the results of experiments performed with Truncated Gaussian distributions (in which the distribution of arm $k$ is that of $Y_{k}=0 \vee\left(X_{k} \wedge 1\right)$ where $X_{k} \sim \mathcal{N}\left(\mu_{k}, \sigma^{2}\right)$ ). We report in Table 8 the regret at time $T=20000$ (estimated over 5000 runs) of various algorithms on four different problem instances: 1) $\mu=[0.5,0.6], \sigma=0.1$, 2) $\mu=[0,0.2], \sigma=0.3,3) \mu=[1.5,2], \sigma=14) \mu=[0.4,0.5,0.6,0.7], \sigma=1$. We include Non-Parametric TS which is known to be asymptotically optimal in this setting (while TS which uses a Beta prior and a binarization trick is not), PHE, and all algorithms based on sub-sampling. We again observe the good performance of SSMC and SDA algorithms across all experiments. They even outperform NP-TS in some experiments, which suggests SDA algorithms may be asymptotically optimal for a wider class of parametric distributions.

Table 3: Regret at $T=20000$ for Truncated Gaussian arms


Bayesian Experiments So far we tried our algorithms on specific instances of the distributions we considered. It is also interesting to check the robustness of the algorithms when the means of the arms are drawn at random according to some distribution. In this section we consider two examples:

Bernoulli bandits where the arms are drawn uniformly at random in $[0,1]$, and Gaussian distributions with the mean parameter of each arm itself drawn from a gaussian distribution $\mu_{k} \sim \mathcal{N}(0,1)$. In both cases we draw 10000 random problems with $K=10 \mathrm{arms}$ and run the algorithms for a time horizon $T=20000$. We experiment with TS, SSMC, RB-SDA and WR-SDA and also add the IMED algorithm ([23]) which is an asymptotically optimal algorithm that uses the knowledge of the distribution. We do not add LDS-SDA and LB-SDA as they are similar to RB-SDA and SSMC, respectively. In the Bernoulli case, we also run the PHE algorithm, which fails to compete with the other algorithms. This is not in contradiction with the results of [14] as in the Bayesian experiments of this paper, arms are drawn uniformly in $[0.25,0.75]$ instead of $[0,1]$. Actually, we noticed that PHE with parameter $a=1.1$ has some difficulties when several arms are close to 1 .

Table 4: Average Regret on 10000 random experiments with Bernoulli Arms

| T | TS | IMED | PHE | SSMC | RB | WR |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 100 | 13.8 | 15.1 | 16.7 | 16.5 | 14.8 | 14.3 |
| 1000 | 27.8 | 31.9 | 39.5 | 34.2 | 31.8 | 30.9 |
| 10000 | 45.8 | 51.2 | 72.3 | 55.0 | 51.1 | 50.6 |
| 20000 | 52.2 | 57.6 | 85.6 | 61.9 | 57.7 | 57.3 |

Table 5: Average Regret on 10000 random experiments with Gaussian Arms

| T | TS | IMED | WR | RB | SSMC |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 100 | 41.2 | 45.1 | 38.3 | 38.1 | 40.6 |
| 1000 | 76.4 | 82.1 | 72.7 | 70.4 | 76.2 |
| 10000 | 118.5 | 124.0 | 115.8 | 111.8 | 120.1 |
| 20000 | 132.6 | 138.1 | 130.2 | 125.7 | 135.1 |

Results reported in Tables 4 and 5 show that RB-SDA and WR-SDA are strong competitors to TS and IMED for both Bernoulli and Gaussian bandits. Recall that these algorithm operate without the need for a specific tuning for each distribution, unlike TS and IMED. Moreover, observe that in the Bernoulli case, TS further uses the same prior as that from which the means are drawn.

Computational aspects To choose a sub-sampling based algorithm, numerical consideration can be taken into account. First, compared to Thompson Sampling, all sub-sampling based algorithm require to store the history of the observation. But then, the cost of sub-sampling varies across algorithms: in the general case RB-SDA is more efficient than WR-SDA as the latter requires to draw a random subset while the former only needs to draw the random integer starting the block. However, for distributions with finite supports WR-SDA can be made as efficient as TS using multivariate geometric distributions, just like PHE does. If one does not want to use randomization then LDS-SDA could be preferred to RB-SDA as it uses a deterministic sequence. Finally, LB-SDA has the smallest computational cost in the SDA family and while its performance is very close to that of SSMC, it can avoid the cost of scanning all the sub-sample means in this algorithm. The computational cost of these two algorithms is difficult to evaluate precisely. Indeed, they can be made very efficient when the leader does not change, but each change of leader is costly, in particular for SSMC. The expected number of such changes is proved to be finite, but for experiments with a finite time horizon the resulting constant can be big. Finally, Non-Parametric TS has a good performance for Truncated Gaussian, but the cost of drawing a random probability vector over a large history is very high.

More experiments In Appendix Awe enhance this empirical study: we show some limitations of SDA for exponential distributions and propose a fix using forced exploration as in SSMC.

## 6 Conclusion

We introduced the SDA framework for exploration in bandits models. We proved that one particular instance, RB-SDA, combines both optimal theoretical guarantees and good empirical performance for several distributions, possibly with unbounded support. Moreover, SDA can be associated with other samplers that seem to achieve similar performance, with their own specificity in terms of computation time. The empirical study presented in the paper also shows the robustness of sub-sampling approach over other types of re-sampling algorithms. This new approach to exploration may be generalized in many directions, for example to contextual bandits or reinforcement learning, where UCB and Thompson Sampling are still the dominant approaches. It is also particularly promising to develop new algorithm for non-stationary bandit, as such algorithms already store the full history of rewards.

## Broader Impact

This work is advertising a new way to do non-parametric exploration in bandit models, that enjoy good empirical performance and strong theoretical guarantees. First, bandit problems are at the heart of numerous applications to online content recommendation, hence the good performance of SDA algorithms may inspire new algorithms for more realistic models used for these applications, such as contextual bandits. Then, exploration is a central question in the broader field of reinforcement learning, hence new ideas for bandits may lead to new ideas for reinforcement learning.

## Acknowledgments and Disclosure of Funding

The PhD of Dorian Baudry is funded by a CNRS80 grant. The authors acknowledge the funding of the French National Research Agency under projects BADASS (ANR-16-CE40-0002) and BOLD (ANR-19-CE23-0026-04).

Experiments presented in this paper were carried out using the Grid' 5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (seehttps://www.grid5000.fr).

## References

[1] Akram Baransi, Odalric-Ambrym Maillard, and Shie Mannor. Sub-sampling for multi-armed bandits. In Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2014. Proceedings, Part I, 2014.
[2] Hock Peng Chan. The multi-armed bandit problem: An efficient nonparametric solution. The Annals of Statistics, 48(1), Feb 2020.
[3] Tor Lattimore and Csaba Szepesvari. Bandit Algorithms. Cambridge University Press, 2019.
[4] R. Agrawal. Sample mean based index policies with $O(\log n)$ regret for the multi-armed bandit problem. Advances in Applied Probability, 27(4), 1995.
[5] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. Machine learning, 47(2-3), 2002.
[6] Olivier Cappé, Aurélien Garivier, Odalric-Ambrym Maillard, Rémi Munos, and Gilles Stoltz. Kullback-leibler upper confidence bounds for optimal sequential allocation. The Annals of Statistics, 41(3), Jun 2013.
[7] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. Biometrika, 25(3-4), 121933.
[8] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In COLT 2012-The 25th Annual Conference on Learning Theory, 2012.
[9] T.L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. Advances in Applied Mathematics, 6(1), 1985.
[10] Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In Algorithmic Learning Theory - 23rd International Conference, ALT 2012, 2012.
[11] Shipra Agrawal and Navin Goyal. Further optimal regret bounds for thompson sampling. In Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2013, 2013.
[12] N. Korda, E. Kaufmann, and R. Munos. Thompson Sampling for 1-dimensional Exponential family bandits. In Advances in Neural Information Processing Systems (NIPS), 2013.
[13] Branislav Kveton, Csaba Szepesvari, Zheng Wen, Mohammad Ghavamzadeh, and Tor Lattimore. Garbage in, reward out: Bootstrapping exploration in multi-armed bandits. In ICML, 2019.
[14] Branislav Kveton, Csaba Szepesvári, Mohammad Ghavamzadeh, and Craig Boutilier. Perturbedhistory exploration in stochastic multi-armed bandits. In Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, 2019.
[15] Bradley Efron and Robert J Tibshirani. An introduction to the bootstrap. CRC press, 1994.
[16] Ian Osband and Benjamin Van Roy. Bootstrapped thompson sampling and deep exploration. CoRR, abs/1507.00300, 2015.
[17] Chi-Hua Wang, Yang Yu, Botao Hao, and Guang Cheng. Residual bootstrap exploration for bandit algorithms. CoRR, abs/2002.08436, 2020.
[18] Charles Riou and Junya Honda. Bandit algorithms based on thompson sampling for bounded reward distributions. In Proceedings of the 31st International Conference on Algorithmic Learning Theory, 2020.
[19] A.N Burnetas and M. Katehakis. Optimal adaptive policies for sequential allocation problems. Advances in Applied Mathematics, 17(2), 1996.
[20] Michael Drmota and Robert Tichy. Sequences, discrepancies and applications, volume 1651 of Lecture Notes in Mathematics. Springer Verlag, Deutschland, 1 edition, 1997.
[21] J. H. Halton. Algorithm 247: Radical-inverse quasi-random point sequence. Commun. ACM, 7(12), December 1964.
[22] I.M Sobol'. On the distribution of points in a cube and the approximate evaluation of integrals. USSR Computational Mathematics and Mathematical Physics, 7(4), 1967.
[23] Junya Honda and Akimichi Takemura. Non-asymptotic analysis of a new bandit algorithm for semi-bounded rewards. Journal of Machine Learning Research, 16(113), 2015.
[24] Aurélien Garivier and Olivier Cappé. The KL-UCB algorithm for bounded stochastic bandits and beyond. In 24th Annual Conference on Learning Theory (COLT), 2011.
[25] Alex Mendelson, Maria Zuluaga, Brian Hutton, and Sébastien Ourselin. What is the distribution of the number of unique original items in a bootstrap sample? CoRR, abs/1602.05822, 2016.
[26] B.C. Rennie and A.J. Dobson. On stirling numbers of the second kind. Journal of Combinatorial Theory, 7(2), 1969.
[27] Anthony W. Ledford and Jonathan A. Tawn. Modelling dependence within joint tail regions. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 59(2), May 1997.

## A Complement of Experiments

## A. 1 Additional Figures for Bernoulli, Gaussian and Truncated Gaussian arms

We enhance the tables of Section 5 with the standard deviation (reported in parenthesis) of the regret at time $T=20000$ on the 5000 trajectories.

Table 6: Regret and at $T=20000$ for Bernoulli arms, with standard deviation

|  | Benchmark |  |  |  | SSDA |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| xp | \|TS | \|PHE | \|BESA | SSMC | RB | \|WR | \|LB | \|LDS |
| 1 | $\left\lvert\, \begin{aligned} & \mathbf{1 1 . 2} \\ & (10 .) \end{aligned}\right.$ | $\left\|\begin{array}{l} 25.9 \\ (87.9) \end{array}\right\|$ | $\begin{aligned} & 11.7 \\ & (12.1) \end{aligned}$ | $\left\lvert\, \begin{aligned} & \mathbf{1 2 . 3} \\ & (7.3) \end{aligned}\right.$ | $\begin{aligned} & \mathbf{1 1 . 5} \\ & (10.1) \end{aligned}$ | $\left\lvert\, \begin{aligned} & \mathbf{1 1 . 6} \\ & (10.2) \end{aligned}\right.$ | 12.2 <br> $(7.4)$ | $\left\lvert\, \begin{aligned} & 11.4 \\ & (9.0)\end{aligned}\right.$ |
| 2 | $\left\|\begin{array}{l} \mathbf{2 2 . 9} \\ (29.2) \end{array}\right\|$ | $\begin{aligned} & \mathbf{2 4 . 0} \\ & (22.0) \end{aligned}$ | $\left\lvert\, \begin{aligned} & \mathbf{2 2 . 1} \\ & (25.2) \end{aligned}\right.$ | $\begin{aligned} & \mathbf{2 4 . 3} \\ & (38.2) \end{aligned}$ | $\begin{aligned} & \mathbf{2 2 . 0} \\ & (34.5) \end{aligned}$ | $\left\lvert\, \begin{aligned} & 21.5 \\ & (17.3) \end{aligned}\right.$ | $\left\|\begin{array}{l} \mathbf{2 4 . 0} \\ (24.6) \end{array}\right\|$ | $\left\lvert\, \begin{aligned} & 21.8 \\ & (24.5) \end{aligned}\right.$ |
| 3 | $\left\|\begin{array}{l} 94.2 \\ (15.8) \end{array}\right\|$ | $\left\|\begin{array}{l} 248.1 \\ (25.5) \end{array}\right\|$ | $\begin{aligned} & \mathbf{8 8 . 1} \\ & (89.2) \end{aligned}$ | $\left\lvert\, \begin{aligned} & 100.1 \\ & (20.0) \end{aligned}\right.$ | $\left\|\begin{array}{l} \mathbf{8 9 . 0} \\ (19.8) \end{array}\right\|$ | $\left\lvert\, \begin{aligned} & \mathbf{8 6 . 9} \\ & (21.7) \end{aligned}\right.$ | $\left\|\begin{array}{l} 100.7 \\ (21.3) \end{array}\right\|$ | $\begin{aligned} & \mathbf{8 9 . 2} \\ & (21.8) \end{aligned}$ |
| 4 | $\left\|\begin{array}{l}\text { 108.1 } \\ (45.1)\end{array}\right\|$ | $\left\lvert\, \begin{aligned} & 216.5 \\ & (89.8)\end{aligned}\right.$ | $\left\lvert\, \begin{aligned} & 147.5 \\ & (209.8)\end{aligned}\right.$ | 119.9 <br> $(40.8)$ | $\left\lvert\, \begin{aligned} & 105.1 \\ & (41.1)\end{aligned}\right.$ | $\left\|\begin{array}{l}\text { 106.9 } \\ (42.1)\end{array}\right\|$ | $\left\|\begin{array}{l}119.6 \\ (42.7)\end{array}\right\|$ | 106.8 <br> (47.7) |

Table 7: Regret and at $T=20000$ for Gaussian arms, with standard deviation

|  |  | Benchmark |  |  | SDA |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | , | ReBoot | BESA | \|S | RB \| | \|WR | \|LB | \|LDS |
| 1 | $\left\lvert\, \begin{aligned} & 24.4 \\ & (17.1) \end{aligned}\right.$ | $\left\lvert\, \begin{aligned} & 92.2 \\ & (23.4) \end{aligned}\right.$ | $\left\lvert\, \begin{aligned} & \mathbf{2 5 . 3} \\ & (27.1) \end{aligned}\right.$ | $\left\lvert\, \begin{aligned} & \mathbf{2 6 . 9} \\ & (52.8) \end{aligned}\right.$ | $\left\lvert\, \begin{aligned} & \mathbf{2 5 . 6} \\ & (62.8) \end{aligned}\right.$ | $\left\lvert\, \begin{aligned} & \mathbf{2 4 . 7} \\ & (20.6) \end{aligned}\right.$ | $\left\lvert\, \begin{aligned} & \mathbf{2 5 . 1} \\ & (17.9) \end{aligned}\right.$ | $\left\lvert\, \begin{aligned} & \mathbf{2 6 . 5} \\ & (140.2) \end{aligned}\right.$ |
| 2 | $\left\lvert\, \begin{aligned} & \mathbf{7 3 . 5} \\ & (107.8) \end{aligned}\right.$ | $\left\lvert\, \begin{aligned} & 277.1 \\ & (41.3) \end{aligned}\right.$ | $\left\|\begin{array}{l} 122.5 \\ (585.5) \end{array}\right\|$ | $\begin{aligned} & 74.8 \\ & (34.7) \end{aligned}$ | $\left\lvert\, \begin{aligned} & \mathbf{7 1 . 0} \\ & (152.2) \end{aligned}\right.$ | $\left\lvert\, \begin{aligned} & 71.1 \\ & (50.2) \end{aligned}\right.$ | $\begin{aligned} & 74.6 \\ & (35.1) \end{aligned}$ | $\begin{aligned} & \mathbf{6 9 . 0} \\ & (50.4) \end{aligned}$ |
| 3 | $\left\lvert\, \begin{aligned} & 49.7 \\ & (26.9) \end{aligned}\right.$ | $\begin{aligned} & 190.9 \\ & (29.6) \end{aligned}$ | $\left\|\begin{array}{l} 72.1 \\ (410.3) \end{array}\right\|$ | $\left\lvert\, \begin{aligned} & \mathbf{5 1 . 3} \\ & (23.7) \end{aligned}\right.$ | $\left\lvert\, \begin{aligned} & \mathbf{5 0 . 4} \\ & (156.5) \end{aligned}\right.$ | $\begin{array}{\|l\|} \mathbf{5 0 . 0} \\ (33.3) \end{array}$ | $\left\lvert\, \begin{aligned} & \mathbf{5 1 . 2} \\ & (22.4) \end{aligned}\right.$ | $\begin{aligned} & 48.6 \\ & (41.6) \end{aligned}$ |

Table 8: Regret at $T=20000$ for Truncated Gaussian arms

|  | Benchmark |  |  |  |  | SDA |  |  | LDS |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| xp | TS | NP-TS | PHE | BESA | SSMC | RB | WR | LB |  |
| 1 | $\left\|\begin{array}{l} 21.9 \\ (20.4) \end{array}\right\|$ | $\left\lvert\, \begin{aligned} & 4.2 \\ & (0.6) \end{aligned}\right.$ | $\left\lvert\, \begin{aligned} & 22.3 \\ & (2.6) \end{aligned}\right.$ | $\begin{aligned} & \mathbf{1 . 4} \\ & (1.7) \end{aligned}$ | $\begin{array}{\|l\|} \hline 1.5 \\ (0.7) \end{array}$ | $\begin{aligned} & \mathbf{1 . 4} \\ & (1.1) \end{aligned}$ | $\begin{aligned} & \mathbf{1 . 4} \\ & (0.8) \end{aligned}$ | $\left\lvert\, \begin{aligned} & \mathbf{1 . 5} \\ & (0.7) \end{aligned}\right.$ | $\begin{aligned} & \mathbf{1 . 4} \\ & (0.8) \end{aligned}$ |
| 2 | $\begin{array}{\|l} 13.3 \\ (7) \end{array}$ | $\left\lvert\, \begin{aligned} & 8 \\ & (1.8) \end{aligned}\right.$ | $\begin{aligned} & 19.5 \\ & (3.8) \end{aligned}$ | $\begin{array}{\|l\|} \hline 4.6 \\ (3.3) \end{array}$ | $\left\lvert\, \begin{aligned} & 4.7 \\ & (2.3) \end{aligned}\right.$ | $\begin{array}{\|l\|} \hline 4.4 \\ (4.6) \end{array}$ | $\begin{array}{\|l\|} \hline 4.5 \\ (3.1) \end{array}$ | $\left\lvert\, \begin{aligned} & \mathbf{4 . 6} \\ & (2.4) \end{aligned}\right.$ | $\begin{array}{\|l\|} \hline 4.3 \\ (2.9) \end{array}$ |
| 3 | $\left\|\begin{array}{l} 9.7 \\ (10.1) \end{array}\right\|$ | $\left\lvert\, \begin{aligned} & 7.8 \\ & (4.5) \end{aligned}\right.$ | $\left\|\begin{array}{l} 48.5 \\ (217.8) \end{array}\right\|$ | $\begin{array}{\|l\|} \hline 7.8 \\ (9.4) \end{array}$ | (5) | $\left\lvert\, \begin{aligned} & 7.1 \\ & (10) \end{aligned}\right.$ | $\begin{array}{\|l\|} \hline 7.7 \\ (13.4) \end{array}$ | $\left\lvert\, \begin{aligned} & 8.2 \\ & (27.5) \end{aligned}\right.$ | $\begin{aligned} & 7.1 \\ & (5.8) \end{aligned}$ |
| 4 | $\begin{aligned} & 86.6 \\ & (57.8) \end{aligned}$ | $\left\lvert\, \begin{aligned} & \mathbf{7 0} \\ & (39.4) \end{aligned}\right.$ | $\begin{aligned} & 86 \\ & (53.7) \end{aligned}$ | $\left\|\begin{array}{l} 76.5 \\ (113.9) \end{array}\right\|$ | $\begin{aligned} & \mathbf{6 9 . 5} \\ & (40.9) \end{aligned}$ | $\left\|\begin{array}{l} \mathbf{6 4 . 9} \\ (60.5) \end{array}\right\|$ | $\begin{aligned} & \mathbf{6 4 . 8} \\ & (43.9) \end{aligned}$ | $\left\lvert\, \begin{aligned} & \mathbf{6 8 . 7} \\ & (39.1) \end{aligned}\right.$ | $\begin{aligned} & \mathbf{6 3 . 2} \\ & (51.1) \end{aligned}$ |

For Bernoulli arms and Truncated Gaussian, the standard deviations of SDA are very similar to that of Thompson Sampling, while the trajectories of PHE and BESA have much more variance in experiment 1 and 4 , and on experiments 3 and 4 respectively. For Gaussian arms we remark the low variability of ReBoot, but at the cost of a non-competitive regret. SDA are less homogeneous in this
case: some algorithms have large variance for some instances (LDS-SDA on experiment 1, RB-SDA on experiments 2 and 3). Note that TS also has a high variability in experiment 2.
We believe that this is due to the nature of the Gaussian distribution, and in particular to its balance function: in Appendix G we prove that $\alpha_{k}(M, j)$ does satisfy Assumption 2. of Theorem 3.1 . however the upper bound derived for $\alpha_{k}(M, j)$ is much larger than the one for Bernoulli distribution, which justifies that "bad runs" in which a good arm looses many duels are more likely to happen in that case, and can explain the larger variance. If one wants to reduce the variance of the regret of SDA we recommend the use of some asymptotically negligible forced exploration, as presented for exponential distribution in Appendix A.2, and for which we prove that the algorithm remains asymptotically optimal in Appendix H

Finally, as in Section 55, we plot the regret of several algorithms as a function of time (in log scale) for $t \in[10000,20000]$, this time for the Truncated Gaussian distributions. These plots illustrate the fact that some SDA algorithms may achieve asymptotic optimality for this distribution too, even if it does not belong to a one-parameter exponential family. Indeed, the rate of the regret of all SDA seem too match both the rate of the regret of Non-Parametric TS, which is optimal for this family, and the Burnetas and Katehakis lower bound whose expression is $\left(\sum_{k \neq k^{*}} \frac{\mathbb{E}_{X \sim \nu_{k^{*}}}[X]-\mathbb{E}_{X \sim \nu_{k}}[X]}{\operatorname{KL}\left(\nu_{k}, \nu_{k^{*}}\right)}\right) \log (T)$ in this particular case, with $\operatorname{KL}\left(\nu_{k}, \nu_{k^{*}}\right)=$ $p_{0, k} \log \left(\frac{p_{0, k}}{p_{0, *}}\right)+\left(1-p_{1, k}\right) \log \left(\frac{1-p_{1, k}}{1-p_{1, *}}\right)+\int_{0}^{1} f_{k}(x) \log \left(\frac{f_{k}(x)}{f_{*}(x)}\right) d x . p_{x, k}$ is the value of the CDF of the underlying Gaussian random variable associated with $\nu_{k}$ in $x$, and $f_{k}(x)$ the density of this variable in $x$.

Figure 3: SDA vs NP-TS on TG expe 2


Figure 4: SDA vs NP-TS on TG expe 4


## A. 2 Experiments with Exponential Arms

In Appendix G, we prove that exponential distributions are not balanced (i.e. do not satisfy Assumption 2. of Theorem 3.1), so our theoretical results on the regret of RB-SDA do not apply. However, it is still interesting to test our algorithms for these distributions in order to see if it still achieves a good performance. We performed 6 experiments, with the following mean parameters: 1) $\mu=[1.5,1], 2$ ) $\mu=[0.2,0.1]$, 3) $\mu=[11,10]$, 4) $\mu=[4,3,2,1]$, 5) $\mu=[0.4,0.3,0.2,0.1], 6) ~ \mu=[5,4,4,4]$. It is interesting to remark that the standard deviation of an exponential distribution is equal to its mean, so with similar gaps problems are harder when the means are high.

Table 9: Average Regret with Exponential Arms (with std)

| xp | TS | IMED | BESA | SSMC | RB | WR | LB | LDS |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 1 | 48.2 | $\mathbf{4 0 . 0}$ | 45.7 | $\mathbf{4 1 . 9}$ | 44.8 | 45.4 | 46.6 | 45.5 |
|  | $(191.8)$ | $(78.4)$ | $(114.1)$ | $(84.2)$ | $(121.4)$ | $(134.4)$ | $(176.8)$ | $(109.7)$ |
| 2 | 3.8 | $\mathbf{3 . 4}$ | 4.2 | $\mathbf{3 . 6}$ | 4.1 | 3.9 | 3.9 | 5.4 |
|  | $(9.9)$ | $(3.6)$ | $(25.1)$ | $(41.9)$ | $(14.3)$ | $(13.4)$ | $(8.7)$ | $(49.5)$ |
| 3 | 832.8 | $\mathbf{7 7 9 . 9}$ | 820.5 | 856.9 | 848.4 | $\mathbf{7 7 8 . 4}$ | 846.7 | 877.7 |
|  | $(1065.1)$ | $(896.9)$ | $(1304.6)$ | $(1111.0)$ | $(1533.3)$ | $(1118.7)$ | $(1150.1)$ | $(1708.7)$ |
| 4 | 258.3 | $\mathbf{2 3 4 . 6}$ | 525.4 | $\mathbf{2 5 1 . 3}$ | 272.6 | 262.1 | 263.8 | 258.4 |
|  | $(519.6)$ | $(126.6)$ | $(2115.1)$ | $(328.3)$ | $(692.2)$ | $(524.4)$ | $(477.9)$ | $(599.0)$ |
| 5 | $\mathbf{2 5 . 6}$ | $\mathbf{2 4 . 0}$ | 55.7 | $\mathbf{2 5 . 6}$ | $\mathbf{2 5 . 5}$ | $\mathbf{2 5 . 0}$ | 26.5 | $\mathbf{2 4 . 7}$ |
|  | $(51.2)$ | $(33.6)$ | $(219.9)$ | $(23.6)$ | $(46.7)$ | $(24.0)$ | $(36.8)$ | $(37.6)$ |
| 6 | $\mathbf{6 1 8 . 7}$ | $\mathbf{6 0 3 . 6}$ | 1184.2 | $\mathbf{6 2 7 . 9}$ | $\mathbf{5 9 5 . 7}$ | $\mathbf{6 1 6 . 0}$ | 652.6 | $\mathbf{6 0 5 . 9}$ |
|  | $(672.3)$ | $(576.8)$ | $(3096.4)$ | $(755.6)$ | $(790.7)$ | $(780.2)$ | $(685.3)$ | $(871.4)$ |

First, we notice that the performance of the SDA in terms of the average regret are reasonable, although less impressive than with the other distributions we tested. IMED is almost always the best algorithm in these experiments, and SSMC performs pretty well on many examples (which is not surprising as SSMC is proved to be asymptotically optimal for exponential distributions). We remark that there is much more variability in the results of RB-SDA, WR-SDA and LDS-SDA than before, where they performed quite similarly. For instance, we notice that on example 3, LDS-SDA and RB-SDA are much worse than WR-SDA. A look at the quantile table for this experiment, which displays the empirical quantiles of $R_{T}$ estimated over 5000 runs, shows that this is due to a small number of "bad" trajectories for these algorithms:

Table 10: Quantile Table for Experiment 3 with Exponential Arms

| \% of runs | TS | IMED | SSMC | RB | WR | LB | LDS |
| :---: | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| $20 \%$ | 319.8 | 336.0 | 335.0 | 261.0 | 290.0 | 326.0 | 261.8 |
| $50 \%$ | 626.0 | 650.0 | 661.0 | 532.0 | 568.5 | 642.0 | 536.0 |
| $80 \%$ | 1122.0 | 1080.0 | 1142.0 | 1006.0 | 1019.0 | 1143.2 | 1020.2 |
| $95 \%$ | 1924.1 | 1704.0 | 1846.0 | $\mathbf{2 1 9 9 . 0}$ | 1817.2 | 1869.1 | $\mathbf{2 1 3 4 . 1}$ |
| $99 \%$ | 4209.4 | 2632.9 | 3536.8 | $\mathbf{6 8 1 3 . 1}$ | 4146.0 | 3762.3 | $\mathbf{7 3 9 6 . 7}$ |

We see that up to the $80 \%$ quantile, RB-SDA and LDS-SDA are even significantly better than IMED. This is very different when we look at the $95 \%$ and $99 \%$ quantiles, which are much greater for our 2 algorithms (even 2.5 times greater for the $99 \%$ quantile).

We believe that this very high variability prevents RB-SDA to have a logarithmic regret for exponential arms. Still, the regret is not as bad as being linear, as using the fact that the balance function $\alpha_{k}(M, j)$ is of order $\exp (-j C) / M$ permits to prove that $\sum_{r=1}^{T} \mathbb{P}\left(N_{1}(r)<\log ^{2}(r)\right)=\mathcal{O}\left(\log ^{2}(T)\right)$ (which requires to choose a different $\beta_{r, j}$ in Lemma 4.3). But we also found a solution to obtain (asymptotically optimal) linear regret, which consists in adding an asymptotically negligible amount of forced exploration as the SSMC algorithm does. This exploration in $o(\log T)$ avoids trajectories where the optimal arm has a very bad first observation and is not drawn for a very long time. In Appendix H , we prove the asymptotic optimality of RB-SDA with forced exploration $f_{r}=\sqrt{\log r}$ for any one-dimensional exponential family. In practice, adding this amount of forced exploration to SDA algorithms leads to the following results:

Table 11: Average Regret with Exponential Arms: SDA with forced exploration

| xp | RB | WR | LB | LDS |
| :--- | :--- | :--- | :--- | :--- |
| 1 | 44.9 | $\mathbf{4 2 . 5}$ | $\mathbf{4 2 . 4}$ | 45.0 |
|  | $(167.3)$ | $(107.4)$ | $(60.5)$ | $(176.0)$ |
| 2 | $\mathbf{3 . 6}$ | $\mathbf{3 . 4}$ | 4.0 | $\mathbf{3 . 6}$ |
|  | $(9.2)$ | $(2.2)$ | $(27.9)$ | $(11.2)$ |
| 3 | 837.5 | $\mathbf{7 8 8 . 5}$ | 827.7 | 832.3 |
|  | $(1466.1)$ | $(1222.1)$ | $(1055.3)$ | $(1514.6)$ |
| 4 | $\mathbf{2 4 4 . 8}$ | $\mathbf{2 3 8 . 9}$ | 251.7 | $\mathbf{2 4 6 . 0}$ |
|  | $(403.3)$ | $(250.8)$ | $(248.5)$ | $(323.4)$ |
| 5 | $\mathbf{2 3 . 6}$ | $\mathbf{2 5 . 1}$ | $\mathbf{2 5 . 4}$ | $\mathbf{2 4 . 9}$ |
|  | $(33.4)$ | $(41.0)$ | $(23.4)$ | $(42.2)$ |
| 6 | $\mathbf{5 7 8 . 9}$ | $\mathbf{5 9 5 . 1}$ | 631.2 | $\mathbf{5 7 7 . 8}$ |
|  | $(651.9)$ | $(561.3)$ | $(446.4)$ | $(652.7)$ |

Hence, adding forced exploration results in a noticeable improvement for SDA algorithms, with RB-SDA, WR-SDA and LDS-SDA becoming competitive with IMED (or even slightly better) on most examples. Observe that LB-SDA has again comparable performance with SSMC with this new feature. This is not surprising as we implemented the SSMC algorithm with the same amount of forced exploration $f_{r}=\sqrt{\log r}$.

## B Notation for the Proof

General notations:

- $K$ number of arms
- $\nu_{k}$ distribution of the arm $k$, with mean $\mu_{k}$
- we assume that $\mu_{1}=\max _{k \in[K]} \mu_{k}$ so we call the (unique) optimal arm "arm 1"
- $I_{k}(x)$ some rate function of the arm $k$, evaluated in $x$. For 1-parameter exponential families this function will always be the KL-divergence between $\nu_{k}$ and the distribution from the same family with mean $x$.
- $N_{k}(r)$ number of pull of arm $k$ up to (and including) round $r$.
- $Y_{k, i}$ reward obtained at the i-th pull of arm $k$.
- $\hat{Y}_{k, i}$ mean of the i-th first reward of arm $k, \hat{Y}_{k, \mathcal{S}}$ mean of the rewards of $k$ on a subset of indices $\mathcal{S} \subset\left[N_{k}(r)\right]: \hat{Y}_{k, \mathcal{S}}=\frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} Y_{k, s}$. If $|\mathcal{S}|=i$, then $Y_{k, i}$ and $Y_{k, \mathcal{S}}$ have the same distribution.
- $\ell(r)$ leader at round $r+1, \ell(r)=\operatorname{argmax}_{k \in[K]} N_{k}(r)$.
- $\operatorname{SP}(m, n, r)$ sub-sampling algorithm, or Sampler, which returns a sequence of $n$ unique elements out of $[m]$.
- $\left(S_{k}^{r}(m, n)\right)_{m \geq n}$ a family of independent random variables such that $S_{k}^{r}(m, n) \sim$ $\mathrm{SP}(m, n, r)$.
- $\mathcal{A}_{r}$ set of arms pulled at a round $r$.
- $\mathcal{R}_{r}$ regret at the end of round $r$.

Notations for the regret analysis, part relying on concentration:

- $\mathcal{G}_{k}^{r}=\cup_{s=1}^{r-1}\{\ell(s)=1\} \cap\left\{k \in \mathcal{A}_{s+1}\right\} \cap\left\{N_{k}(s) \geq(1+\varepsilon) \xi_{k} \log r\right\}$
- $\mathcal{H}_{k}^{r}=\cup_{s=1}^{r-1}\{\ell(s)=1\} \cap\left\{k \in \mathcal{A}_{s+1}\right\} \cap\left\{N_{k}(s) \geq J_{k} \log r\right\}$
- $\mathcal{Z}^{r}=\{\ell(r) \neq 1\}$, the leader used for the duels in round $r+1$ is sub-optimal
- $\mathcal{D}^{r}=\{\exists u \in[\lfloor r / 4\rfloor, r]$ such that $\ell(u-1)=1\}$, the leader has been optimal at least once between $\lfloor r / 4\rfloor$ and $r$
- $\mathcal{B}^{u}=\left\{\ell(u)=1, k \in \mathcal{A}_{u+1}, N_{k}(u)=N_{1}(u)-1\right.$ for some arm $\left.k\right\}$, the optimal arm is leader in $u$ but loses its duel again some arm $k$, that have been pulled enough to possibly take over the leadership at next round
- $\mathcal{C}^{u}=\left\{\exists k \neq 1, N_{k}(u) \geq N_{1}(u), \hat{Y}_{k, S_{1}^{u}\left(N_{k}(u), N_{1}(u)\right)} \geq \hat{Y}_{1, N_{1}(u)}\right\}$, the optimal arm is not the leader and has lost its duel against the sub-optimal leader.
- $\mathcal{L}^{r}=\sum_{u=\lfloor r / 4\rfloor}^{r} \mathbb{1}_{\mathcal{C}^{u}}$

Notations for the regret analysis, control of the number of pulls of the optimal arm:

- $r_{j}$ round of the j -th play of the optimal arm
- $\tau_{j}=r_{j+1}-r_{j}$
- $\mathcal{E}_{j}^{r}:=\left\{\tau_{j} \geq r / \log r^{2}-1\right\}$
- $\mathcal{M}_{j, r}^{1}=\left[r_{j}+1, r_{j}+\left\lfloor\frac{r / \log r^{2}-1}{2}\right\rfloor\right]$
- $\mathcal{M}_{j, r}^{2}=\left[t_{j}+\left\lceil\frac{r / \log r^{2}-1}{2}\right\rceil, r_{j}+\left\lfloor r / \log r^{2}\right\rfloor-1\right]$
- $\mathcal{I}_{j, r}^{k}=\left\{s \in \mathcal{M}_{j, r}^{2}: \ell(s-1)=k\right\}$
- $\mathcal{W}_{s, j}^{k}=\left\{\left\{\hat{Y}_{1, j}<\hat{Y}_{k, S_{1}^{s}\left(N_{k}(s), j\right)}\right\}, N_{k}(s) \geq c_{r, K}, N_{1}(s)=j\right\}$
- $\mathcal{F}_{j, M}^{k, r}=\left\{\exists i_{1}, \ldots, i_{M} \in I_{j, r}^{k}: \forall m<m^{\prime} \in[M], S_{1}^{i_{m}}\left(N_{k}\left(i_{m}\right), j\right) \cap S_{1}^{i_{m^{\prime}}}\left(N_{k}\left(i_{m^{\prime}}\right), j\right)=\emptyset\right\}$
- CDF: Cumulative Distribution Function, PDF: Probability Density Function and PMF: Probability Mass Function.


## C Concentration Result: Proof of Lemma 4.2

We first recall the probabilistic model introduced in Section 2 for each round $r$, each arm $k$, we define a family $\left(S_{k}^{r}(n, m)\right)_{n>m}$ of independent random variables such that $S_{k}^{r}(n, m) \sim \mathrm{SP}(n, m, r)$. Those random variables are also independent from the reward streams $\left(Y_{k, s}\right)_{s \geq 0}$ of all arms $k$.
$S_{k}^{r}(n, m)$ is the subset of the leader history that is used should arm $k$ be a challenger drawn $m$ times up to round $r$ duelling against a leader that has been drawn $n$ times. With this notation, letting $\ell(r)$ be the leader after $r$ rounds, at round $r+1$, for all $k \neq \ell(r)$,

$$
\left(k \in \mathcal{A}_{r+1}\right) \Leftrightarrow\left(\hat{Y}_{k, N_{k}(r)}>\hat{Y}_{\ell(r), S_{k}^{r}\left(N_{\ell(r)}(r), N_{k}(r)\right)}\right) .
$$

Let $k$ be an arm such that $\mu_{k}<\mu_{1}$. We denote by $\left[n_{1}, n_{k}\right]$ the set of subset of $\left\{1, \ldots, n_{1}\right\}$ of size $n_{k}$. We define an event

$$
\mathcal{Q}_{k}^{s}=\left\{N_{k}(s) \geq n_{0}, \ell(s)=1, \hat{Y}_{k, N_{k}(s)}>\hat{Y}_{\ell(s), S_{k}^{s}\left(N_{1}(s), N_{k}(s)\right)}\right\}
$$

Noting that $\left\{\hat{Y}_{k, N_{k}(s)}>\hat{Y}_{\ell(s), S_{k}^{s}\left(N_{1}(s), N_{k}(s)\right)}\right\} \subset\left\{\hat{Y}_{k, N_{k}(s)} \geq \xi\right\} \cup\left\{\hat{Y}_{\ell(s), S_{k}^{s}\left(N_{1}(s), N_{k}(s)\right)} \leq \xi\right\}$ for all $\xi \in \mathbb{R}$, we can write $\mathcal{Q}_{k}^{s} \subset \mathcal{Q}_{k}^{s, 1} \cup \mathcal{Q}_{k}^{s, 2}$ where

$$
\begin{aligned}
\mathcal{Q}_{k}^{s, 1} & =\left\{N_{k}(s) \geq n_{0}, \ell(s)=1, \hat{Y}_{k, N_{k}(s)}>\xi\right\} \\
\text { and } \mathcal{Q}_{k}^{s, 2} & =\left\{N_{k}(s) \geq n_{0}, \ell(s)=1, \hat{Y}_{\ell(s), S_{k}^{s}\left(N_{1}(s), N_{k}(s)\right)} \leq \xi\right\}
\end{aligned}
$$

This yields $\sum_{s=1}^{r} \mathbb{P}\left(\mathcal{Q}_{k}^{s}\right) \leq \sum_{s=1}^{r} \mathbb{P}\left(\mathcal{Q}_{k}^{r, 1}\right)+\sum_{s=1}^{r} \mathbb{P}\left(\mathcal{Q}_{k}^{r, 2}\right)$, which will later provide the two terms in the bound of the lemma. The first one does not involve sub-sampling and can be upper bounded as:

$$
\begin{aligned}
& \sum_{s=1}^{r} \mathbb{P}\left(\mathcal{Q}_{k}^{s, 1}\right) \leq \mathbb{E} \sum_{s=1}^{r} \mathbb{1}\left(N_{k}(s) \geq n_{0}\right) \mathbb{1}\left(N_{1}(s)>N_{k}(s)\right) \mathbb{1}\left(\hat{Y}_{k, N_{k}(s)} \geq \xi\right) \mathbb{1}\left(k \in \mathcal{A}_{s+1}\right) \\
& \leq \mathbb{E} \sum_{s=n_{0}}^{r} \sum_{n_{k}=n_{0}}^{r} \mathbb{1}\left(N_{k}(s)=n_{k}, k \in \mathcal{A}_{s+1}\right) \mathbb{1}\left(\hat{Y}_{k, n_{k}} \geq \xi\right) \\
& \leq \mathbb{E} \sum_{n_{k}=n_{0}}^{r} \mathbb{1}\left(\hat{Y}_{k, n_{k}} \geq \xi\right) \underbrace{\sum_{s=n_{0}}^{r} \mathbb{1}\left(N_{k}(s)=n_{k}, k \in \mathcal{A}_{s+1}\right)}_{\leq 1} \\
& \leq \sum_{n_{k}=n_{0}}^{r} \mathbb{P}\left(\hat{Y}_{k, n_{k}} \geq \xi\right)
\end{aligned}
$$

where in the last inequality we use that the event $\left(N_{k}(s)=n\right) \cap\left(k \in \mathcal{A}_{s+1}\right)$ can happen at most once for $s \in\left\{n_{0}, \ldots, r\right\}$ (a similar trick was used for example in the analysis of kl-UCB [24]).
Upper bounding the second term $B_{r}=\sum_{s=1}^{r} \mathbb{P}\left(\mathcal{Q}_{k}^{r, 2}\right)$ is more intricate as it involves both $N_{k}(s)$ and $N_{1}(s)$. With a similar method we get:

$$
\begin{aligned}
& B_{r} \leq \mathbb{E} \sum_{s=n_{0}}^{r} \sum_{n_{k}=n_{0}}^{r} \sum_{n_{1}=n_{k}}^{r} \sum_{\mathcal{S} \in\left[n_{1}, n_{k}\right]} \mathbb{1}\left(N_{k}(s)=n_{k}, k \in \mathcal{A}_{s+1}\right) \mathbb{1}\left(N_{1}(s)=n_{1}\right) \mathbb{1}\left(S_{k}^{S}\left(n_{1}, n_{k}\right)=\mathcal{S}\right) \mathbb{1}\left(\hat{Y}_{\ell, \mathcal{S}} \leq \xi\right) \\
& \leq \mathbb{E} \sum_{n_{k}=n_{0}}^{r} \sum_{n_{1}=n_{k}}^{r} \sum_{\mathcal{S} \in\left[n_{1}, n_{k}\right]} \mathbb{1}\left(\hat{Y}_{\ell, \mathcal{S}} \leq \xi\right) \sum_{s=n_{0}}^{r} \mathbb{1}\left(N_{k}(s)=n_{k}, k \in \mathcal{A}_{s+1}\right) \mathbb{1}\left(S_{k}^{s}\left(n_{1}, n_{k}\right)=\mathcal{S}\right) \\
& =\mathbb{E} \sum_{n_{k}=n_{0}}^{r} \sum_{n_{1}=n_{k}}^{r} \sum_{\mathcal{S} \in\left[n_{1}, n_{k}\right]} \mathbb{1}\left(\hat{Y}_{\ell, \mathcal{S}} \leq \xi\right) \sum_{s=n_{0}}^{r} \mathbb{E}\left[\mathbb{1}\left(N_{k}(s)=n_{k}, k \in \mathcal{A}_{s+1}\right) \mathbb{1}\left(S_{k}^{s}\left(n_{1}, n_{k}\right)=\mathcal{S}\right) \mid \mathcal{F}\right],
\end{aligned}
$$

where $\mathcal{F}$ is the filtration generated by the reward streams. $N_{k}(s)$ may have a complicated distribution with respect to this filtration but this is not a problem here. Indeed, $S_{k}^{s}\left(n_{1}, n_{k}\right)$ is by design independent of this filtration, and one can write

$$
\begin{aligned}
& B_{r} \leq \mathbb{E} \sum_{n_{k}=n_{0}}^{r} \sum_{n_{1}=n_{k}}^{r} \sum_{\mathcal{S} \in\left[n_{1}, n_{k}\right]} \mathbb{1}\left(\hat{Y}_{1, \mathcal{S}} \leq \xi\right) \sum_{s=n_{0}}^{r} \mathbb{P}\left(S_{k}^{s}\left(n_{1}, n_{k}\right)=\mathcal{S}\right) \mathbb{E}\left[\mathbb{1}\left(N_{k}(s)=n_{k}, k \in \mathcal{A}_{s+1}\right) \mid \mathcal{F}\right] \\
& =\mathbb{E} \sum_{n_{k}=n_{0}}^{r} \sum_{n_{1}=n_{k}}^{r} \sum_{\mathcal{S} \in\left[n_{1}, n_{k}\right]} \mathbb{1}\left(\hat{Y}_{1, \mathcal{S}} \leq \xi\right) \sum_{s=n_{0}}^{r} \mathbb{P}\left(S_{k}^{s}\left(n_{1}, n_{k}\right)=\mathcal{S}\right) \mathbb{1}\left(N_{k}(s)=n_{k}, k \in \mathcal{A}_{s+1}\right) \\
& =\sum_{n_{k}=n_{0}}^{r} \sum_{n_{1}=n_{k}}^{r} \sum_{\mathcal{S} \in\left[n_{1}, n_{k}\right]} \mathbb{P}\left(\hat{Y}_{1, \mathcal{S}} \leq \xi\right) \sum_{s=n_{0}}^{r} \mathbb{P}\left(S_{k}^{s}\left(n_{1}, n_{k}\right)=\mathcal{S}\right) \mathbb{E}\left(\mathbb{1}\left(N_{k}(s)=n_{k}, k \in \mathcal{A}_{s+1}\right)\right) \\
& =\sum_{n_{k}=n_{0}}^{r} \sum_{n_{1}=n_{k}}^{r} \mathbb{P}\left(\hat{Y}_{1, n_{1}}^{r} \leq \xi\right) \sum_{s=n_{0}}^{r}\left(\sum_{\mathcal{S} \in\left[n_{1}, n_{k}\right]} \mathbb{P}\left(S_{k}^{s}\left(n_{1}, n_{k}\right)=\mathcal{S}\right)\right) \mathbb{E}\left(\mathbb{1}\left(N_{k}(s)=n_{k}, k \in \mathcal{A}_{s+1}\right)\right) \\
& =\sum_{n_{k}=n_{0}}^{r} \sum_{n_{1}=n_{k}}^{r} \mathbb{P}\left(\hat{Y}_{1, n_{1}}^{r} \leq \xi\right) \mathbb{E} \sum_{s=n_{0}}^{r} \mathbb{1}\left(N_{k}(s)=n_{k}, k \in \mathcal{A}_{s+1}\right) \\
& \leq \sum_{n_{k}=n_{0}}^{r} \sum_{n_{1}=n_{k}}^{r} \mathbb{P}\left(\hat{Y}_{1, n_{1}} \leq \xi\right) \\
& \leq r \sum_{n_{1}=n_{k}}^{r} \mathbb{P}\left(\hat{Y}_{1, n_{1}}^{r} \leq \xi\right) .
\end{aligned}
$$

Here we have used the independence of the $S_{k}^{s}(m, n)$ from the reward streams and the fact that for every subset $\mathcal{S}$ of size $n_{1}, \hat{Y}_{k, n_{1}}$ and $\hat{Y}_{k, \mathcal{S}}$ have the same distribution. We can conclude as follows, proving the lemma:

$$
\begin{aligned}
\sum_{s=1}^{r} \mathbb{P}\left(\mathcal{Q}_{k}^{r}\right) & \leq \sum_{s=1}^{r} \mathbb{P}\left(\mathcal{Q}_{k}^{r, 1}\right)+\sum_{s=1}^{r} \mathbb{P}\left(\mathcal{Q}_{k}^{r, 2}\right) \\
& \leq \sum_{n_{k}=n_{0}}^{r} \mathbb{P}\left(\hat{Y}_{k, n_{k}} \geq \xi\right)+r \sum_{n_{1}=n_{0}}^{r} \mathbb{P}\left(\hat{Y}_{1, n_{1}} \leq \xi\right)
\end{aligned}
$$

## D Regret Decomposition: Proof of Lemma 4.1

We recall that we assume that arm 1 is the only optimal arm: $\mu_{1}=\max _{k \in[K]} \mu_{k}$. The proof in this section follows the path of the proof in [2] for SSMC, but hinges on the new concentration result of Lemma 4.2 Moreover, some parts need to be adapted to handle the properties of an independent sampler instead of the duelling rule used in SSMC. As in [2], we introduce the following events:

- $\mathcal{G}_{k}^{T}=\cup_{r=1}^{T-1}\{\ell(r)=1\} \cap\left\{k \in \mathcal{A}_{r+1}\right\} \cap\left\{N_{k}(r) \geq(1+\varepsilon) \xi_{k} \log T\right\}$
- $\mathcal{H}_{k}^{T}=\cup_{r=1}^{T-1}\{\ell(r)=1\} \cap\left\{k \in \mathcal{A}_{r+1}\right\} \cap\left\{N_{k}(r) \geq J_{k} \log T\right\}$
- $\mathcal{Z}^{r}=\{\ell(r) \neq 1\}$, the leader used at round $r+1$ is sub-optimal.

These events directly provide an upper bound of the number of pulls of a sub-optimal arm $k$ :

$$
\begin{align*}
\mathbb{E}\left[N_{k}(T)\right] & =\mathbb{E}\left[N_{k}(T) \mathbb{1}_{\mathcal{H}_{k}^{T}}\right]+\mathbb{E}\left[N_{k}(T) \mathbb{1}_{\mathcal{G}_{k}^{T}} \mathbb{1}_{\overline{\mathcal{H}}_{k}^{T}}\right]+\mathbb{E}\left[N_{k}(T) \mathbb{1}_{\overline{\mathcal{G}}_{k}^{T}}\right] \\
& \leq T \mathbb{P}\left(\mathcal{H}_{k}^{T}\right)+\left(1+J_{k} \log T\right) \mathbb{P}\left(\mathcal{G}_{k}^{T}\right)+1+(1+\varepsilon) \xi_{k} \log T+2 \sum_{r=1}^{T-1} \mathbb{P}\left(\mathcal{Z}^{r}\right) \tag{3}
\end{align*}
$$

Indeed, due to the definition of each event we have:

$$
\begin{aligned}
N_{k}(T) \mathbb{1}_{\overline{\mathcal{G}}_{k}^{T}} & \leq 1+\sum_{r=1}^{T-1} \mathbb{1}_{\left(k \in \mathcal{A}_{r+1}\right)} \mathbb{1}_{(\ell(r) \neq 1) \cup\left(N_{k}(r)<(1+\varepsilon) \xi_{k} \log (T)\right)} \\
& \leq 1+\sum_{r=1}^{T-1} \mathbb{1}_{\left(k \in \mathcal{A}_{r+1}\right)} \mathbb{1}_{\left(N_{k}(r)<(1+\varepsilon) \xi_{k} \log (T)\right)}+\sum_{r=1}^{T-1} \mathbb{1}_{(\ell(r) \neq 1)} \\
& \leq 1+(1+\varepsilon) \xi_{k} \log T+\sum_{r=1}^{T-1} \mathbb{1}_{\mathcal{Z}^{r}}
\end{aligned}
$$

and similarly

$$
\begin{aligned}
N_{k}(T) \mathbb{1}_{\mathcal{G}_{k}^{T}} \mathbb{1}_{\overline{\mathcal{H}}_{k}^{T}} & \leq\left(1+J_{k} \log T+\sum_{r=1}^{T-1} \mathbb{1}_{\mathcal{Z}^{r}}\right) \mathbb{1}_{\mathcal{G}_{k}^{T}} \\
& \leq\left(1+J_{k} \log T\right) \mathbb{1}_{\mathcal{G}_{k}^{T}}+\sum_{r=1}^{T-1} \mathbb{1}_{\mathcal{Z}^{r}}
\end{aligned}
$$

Choosing $\xi_{k}=1 / I_{1}\left(\mu_{k}\right)$ the bound in (3) exhibits the term in $\frac{1+\varepsilon}{I_{1}\left(\mu_{k}\right)} \log T$ in Lemma 4.1 To obtain the result, it remains to upper bound

$$
T \mathbb{P}\left(\mathcal{H}_{k}^{T}\right)+\left(1+J_{k} \log T\right) \mathbb{P}\left(\mathcal{G}_{k}^{T}\right)+2 \sum_{r=1}^{T-1} \mathbb{P}\left(\mathcal{Z}^{r}\right)
$$

for an appropriate choice of $J_{k}$. To do so, we shall first use the concentration inequality Lemma 4.2 to upper bound the terms involving $\mathcal{G}_{k}^{T}$ and $\mathcal{H}_{k}^{T}$ by problem-dependent constants (Appendix D.1), and then we carefully handle the terms in $\mathcal{Z}^{r}$ (Appendix D.2.

## D. 1 Upper Bounds on $\mathbb{P}\left(\mathcal{G}_{k}^{T}\right)$ and $\mathbb{P}\left(\mathcal{H}_{k}^{T}\right)$

We first fix some real numbers $J_{k}$ and $\omega, \omega_{k}$ in $\left(\mu_{k}, \mu_{1}\right)$ to be specified later. We also fix $\xi_{k}=$ $1 / I_{1}\left(\mu_{k}\right)$. Starting with $\mathcal{G}_{k}^{T}$, we apply the second statement in Lemma 4.2 for arm $k$ and arm 1 with $n_{0}=(1+\varepsilon) \xi_{k} \log r$ and $\xi=\omega_{k}$ :

$$
\begin{aligned}
\mathbb{P}\left(\mathcal{G}_{k}^{T}\right) & \leq \sum_{r=1}^{T-1} \mathbb{P}\left(\ell(r)=1, k \in \mathcal{A}_{r+1}, N_{k}(r) \geq(1+\varepsilon) \xi_{k} \log T\right) \\
& \leq \sum_{r=1}^{T-1} \mathbb{P}\left(N_{1}(r) \geq N_{k}(r), \hat{Y}_{k, N_{k}(r)}>\hat{Y}_{1, S_{1}^{r}\left(N_{1}(r), N_{k}(r)\right)}, N_{k}(r) \geq(1+\varepsilon) \xi_{k} \log T\right) \\
& \leq \frac{T}{1-e^{-I_{1}\left(\omega_{k}\right)}} e^{-(1+\varepsilon) \xi_{k} I_{1}\left(\omega_{k}\right) \log T}+\frac{1}{1-e^{-I_{k}\left(\omega_{k}\right)}} e^{-(1+\varepsilon) \xi_{k} I_{k}\left(\omega_{k}\right) \log T}
\end{aligned}
$$

Similarly, we obtain

$$
\mathbb{P}\left(\mathcal{H}_{k}^{T}\right) \leq \frac{T}{1-e^{-I_{1}(\omega)}} e^{-J_{k} I_{1}(\omega) \log T}+\frac{1}{1-e^{-I_{k}(\omega)}} e^{-J_{k} I_{k}(\omega) \log T}
$$

Our objective is to bound $T \mathbb{P}\left(\mathcal{H}_{k}^{T}\right)$ and $\log (T) \mathbb{P}\left(\mathcal{G}_{k}^{T}\right)$ by constants. This is achieved for instance if $T \mathbb{P}\left(\mathcal{H}_{k}^{T}\right) \underset{T \rightarrow+\infty}{\longrightarrow} 0$ and $\log (T) \mathbb{P}\left(\mathcal{G}_{k}^{T}\right) \underset{T \rightarrow+\infty}{\longrightarrow} 0$. The following conditions are sufficient to ensure these properties:

- $(1+\varepsilon) \xi_{k} I_{1}\left(\omega_{k}\right)>1$
- $(1+\varepsilon) \xi_{k} I_{k}\left(\omega_{k}\right)>0$
- $J_{k} I_{1}(\omega)>2$
- $J_{k} I_{k}(\omega)>1$

These conditions are met with the following values:

- $\omega=\frac{1}{2}\left(\mu_{1}+\max _{k \neq 1} \mu_{k}\right)$
- $J_{k}>\max \left(\frac{1}{I_{k}(\omega)}, \frac{2}{I_{1}(\omega)}\right)$
- $\mu_{k}<\omega_{k}<\mu_{1}$ chosen such that $(1+\varepsilon) I_{1}\left(\omega_{k}\right)>I_{k}\left(\omega_{k}\right)$. We are sure that such value exists if we choose $\omega_{k}$ close enough to $\mu_{k}$, thanks to the continuity of the rate functions and the fact that $I_{k}\left(\mu_{k}\right)=0$ and $I_{1}\left(\mu_{k}\right) \neq 0$ (assumed in Assumption 1. of Theorem 3.1).

Choosing these values, both the terms in $\mathcal{G}_{k}^{T}$ and $\mathcal{H}_{k}^{T}$ in (3) are part of the constant $C_{k}(\boldsymbol{\nu}, \varepsilon)$ in Lemma 4.1 We can now focus on upper bounding $\sum_{r=1}^{T-1} \mathbb{P}\left(\mathcal{Z}^{r}\right)$, which is more challenging.
D. 2 Upper Bound on $\sum_{r=1}^{T-1} \mathbb{P}\left(\mathcal{Z}^{r}\right)$

The first steps of this part of the proof are again similar to [2]. The definition of the leader as the arm with the largest history gives the following property, that will be very useful for the analysis:

$$
\ell(r)=k \Rightarrow N_{k}(r) \geq\left\lfloor\frac{r}{K}\right\rfloor-1
$$

So if an arm $k$ is the leader at a given round it has been drawn a linear amount of time at this round. Intuitively, this will provide very interesting concentration guarantees for the leader after a reasonable amount of rounds, that we are going to use in this section. For every $r \geq 8$, we define $a_{r}=\left\lfloor\frac{r}{4}\right\rfloor$ and use the decomposition

$$
\begin{equation*}
\mathbb{P}\left(\mathcal{Z}^{r}\right)=\mathbb{P}\left(\mathcal{Z}^{r} \cap \mathcal{D}^{r}\right)+\mathbb{P}\left(\mathcal{Z}^{r} \cap \overline{\mathcal{D}}^{r}\right) \tag{4}
\end{equation*}
$$

where $\mathcal{D}^{r}$ is the event that the optimal has been leader at least once in $\left[a_{r}, r\right]$ :

$$
\mathcal{D}^{r}=\left\{\exists u \in\left[a_{r}, r\right] \text { such that } \ell(u)=1\right\} .
$$

We now explain how to upper bound the sum of the two terms in the left hand side of (4).

## D.2.1 Controlling $\mathbb{P}\left(\mathcal{Z}^{r} \cap \mathcal{D}^{r}\right)$ : arm 1 has been leader between $\lfloor r / 4\rfloor$ and $r$

We introduce a new event

$$
\mathcal{B}^{u}=\left\{\ell(u)=1, k \in \mathcal{A}_{u+1}, N_{k}(u)=N_{1}(u)-1 \text { for some arm } k\right\}
$$

If $\mathcal{D}^{r}$ happens, then the event $\mathcal{Z}^{r}$ can be true only if the leadership has been taken over by a sub-optimal arm at some round between $a_{r}$ and $r$, that is

$$
\mathcal{Z}^{r} \cap \mathcal{D}^{r} \subset \cup_{u=a_{r}}^{r}\left\{\overline{\mathcal{Z}}_{u}, \mathcal{Z}_{u+1}\right\} \subset \cup_{u=a_{r}}^{r} \mathcal{B}^{u}
$$

We now upper bound $\sum_{r=8}^{T-1} \sum_{u=a_{r}}^{r} \mathbb{P}\left(\mathcal{B}^{u}\right)$. We use the notation $b_{r}=\left\lfloor a_{r} / K\right\rfloor$, where we recall $a_{r}=\lfloor r / 4\rfloor$. Then we write $\left.\mathcal{B}^{u}=\cup_{k=2}^{K} \mathcal{B}_{k}^{u}:=\left\{\ell(u)=1, k \in \mathcal{A}_{u+1}, N_{k}(u)=N_{1}(u)-1\right\}\right\}$, which fixes a specific suboptimal arm. For any $w_{k}$ in $\left(\mu_{k}, \mu_{1}\right)$, one can write

$$
\begin{align*}
& \sum_{r=8}^{T-1} \sum_{u=a_{r}}^{r} \mathbb{P}\left(\mathcal{B}_{k}^{u}\right)=\mathbb{E} \sum_{r=8}^{T-1} \sum_{u=a_{r}}^{r} \mathbb{1}(\ell(u)=1) \mathbb{1}\left(k \in \mathcal{A}_{u+1}\right) \mathbb{1}\left(N_{1}(u)=N_{k}(u)+1\right) \\
& \leq \mathbb{E} \sum_{r=8}^{T-1} \sum_{u=a_{r}}^{r} \mathbb{1}\left(N_{1}(u) \geq b_{r}\right) \mathbb{1}\left(\bar{Y}_{k, N_{k}(u)} \geq \bar{Y}_{1, S_{k}^{u}\left(N_{1}(u), N_{k}(u)\right)}\right) \mathbb{1}\left(N_{1}(u)=N_{k}(u)+1\right) \mathbb{1}\left(k \in \mathcal{A}_{u+1}\right) \\
& \leq \mathbb{E} \sum_{r=8}^{T-1} \sum_{u=a_{r}}^{r} \mathbb{1}\left(N_{1}(u) \geq b_{r}\right) \mathbb{1}\left(\bar{Y}_{k, N_{k}(u)}<w_{k}\right) \mathbb{1}\left(N_{1}(u)=N_{k}(u)+1\right) \mathbb{1}\left(k \in \mathcal{A}_{u+1}\right)  \tag{5}\\
& +\mathbb{E} \sum_{r=8}^{T-1} \sum_{u=a_{r}}^{r} \mathbb{1}\left(N_{1}(u) \geq b_{r}\right) \mathbb{1}\left(\bar{Y}_{1, S_{k}^{u}\left(N_{1}(u), N_{k}(u)\right)}>w_{k}\right) \mathbb{1}\left(N_{1}(u)=N_{k}(u)+1\right) \mathbb{1}\left(k \in \mathcal{A}_{u+1}\right) \tag{6}
\end{align*}
$$

We now separately upper bound each of these two terms. First,

$$
\begin{aligned}
\text { (5) } & \leq \mathbb{E} \sum_{r=8}^{T-1} \sum_{u=a_{r}}^{r} \sum_{n_{k}=b_{r}-1}^{r} \mathbb{1}\left(N_{k}(u)=n_{k}\right) \mathbb{1}\left(k \in \mathcal{A}_{u+1}\right) \mathbb{1}\left(\bar{Y}_{k, n_{k}}<w_{k}\right) \\
& \leq \mathbb{E} \sum_{r=8}^{T-1} \sum_{n_{k}=b_{r}-1}^{r} \mathbb{1}\left(\bar{Y}_{k, n_{k}}<w_{k}\right) \underbrace{}_{\underbrace{\sum_{n}}_{u=a_{r}} \mathbb{1}\left(N_{k}(u)=n_{k}\right) \mathbb{1}\left(k \in \mathcal{A}_{u+1}\right)} \\
& \leq \sum_{r=8}^{T-1} \sum_{n_{k}=b_{r}-1}^{r} \mathbb{P}\left(\bar{Y}_{k, n_{k}}<w_{k}\right) \\
& \leq \sum_{r=8}^{T-1} \sum_{n_{k}=b_{r}-1}^{r} \exp \left(-n_{k} I_{k}\left(w_{k}\right)\right) \\
& \leq \frac{e^{(2+1 / K) I_{k}\left(\omega_{k}\right)}}{\left(1-e^{-I_{k}\left(\omega_{k}\right)}\right)\left(1-e^{-I_{k}\left(\omega_{k}\right) / 4 K}\right)}
\end{aligned}
$$

Then, letting $[m, n]$ denote the set of subset of $[n]$ of size $m$,

$$
\begin{aligned}
\text { (6) } & \leq \mathbb{E} \sum_{r=8}^{T-1} \sum_{u=a_{r}}^{r} \sum_{n_{k}=b_{r}-1}^{r} \mathbb{1}\left(\bar{Y}_{1, S_{k}^{u}\left(n_{k}+1, n_{k}\right)}>w_{k}\right) \mathbb{1}\left(N_{k}(u)=n_{k}\right) \mathbb{1}\left(k \in \mathcal{A}_{u+1}\right) \\
& \leq \mathbb{E} \sum_{r=8}^{T-1} \sum_{u=a_{r}}^{r} \sum_{n_{k}=b_{r}-1}^{r} \sum_{\mathcal{S} \in\left[n_{k}, n_{k}+1\right]} \mathbb{1}\left(\bar{Y}_{1, \mathcal{S}}>w_{k}\right) \mathbb{1}\left(S_{k}^{u}\left(n_{k}+1, n_{k}\right)=\mathcal{S}\right) \mathbb{1}\left(N_{k}(u)=n_{k}\right) \mathbb{1}\left(k \in \mathcal{A}_{u+1}\right) \\
& \leq \mathbb{E} \sum_{r=8}^{T-1} \sum_{u=a_{r}}^{r} \sum_{n_{k}=b_{r}-1}^{r} \sum_{\mathcal{S} \in\left[n_{k}, n_{k}+1\right]} \mathbb{1}\left(\bar{Y}_{1, \mathcal{S}}>w_{k}\right) \mathbb{1}\left(N_{k}(u)=n_{k}\right) \mathbb{1}\left(k \in \mathcal{A}_{u+1}\right) \\
& \leq \mathbb{E} \sum_{r=8}^{T-1} \sum_{n_{k}=b_{r}-1}^{r} \sum_{\mathcal{S} \in\left[n_{k}, n_{k}+1\right]} \mathbb{1}\left(\bar{Y}_{1, \mathcal{S}}>w_{k}\right) \sum_{u=a_{r}}^{r} \mathbb{1}\left(N_{k}(u)=n_{k}\right) \mathbb{1}\left(k \in \mathcal{A}_{u+1}\right) \\
& \leq \sum_{r=8}^{T-1} \sum_{n_{k}=b_{r}-1}^{r} \sum_{\mathcal{S} \in\left[n_{k}, n_{k}+1\right]} \mathbb{P}\left(\bar{Y}_{1, \mathcal{S}}>w_{k}\right) \\
& \leq \sum_{r=8}^{T-1} \sum_{n_{k}=b_{r}-1}^{r} n_{\left(n_{k}+1\right) \mathbb{P}\left(\bar{Y}_{1, n_{k}}>w_{k}\right)}^{r} \\
& \leq \sum_{r=8}^{T-1}(r+1) \sum_{n_{k}=b_{r}-1}^{r} \exp \left(-n_{k} I_{1}\left(w_{k}\right)\right) \\
& \leq \frac{e^{(2+1 / K) I_{1}\left(\omega_{k}\right)}}{\left(1-e^{-I_{1}\left(\omega_{k}\right)}\right)\left(1-e^{-I_{1}\left(\omega_{k}\right) / 4 K}\right)^{2}}
\end{aligned}
$$

Here we have used that there are $n_{k}+1$ subsets in $\left[n_{k}, n_{k}+1\right]$ and that $\mathbb{P}\left(\bar{Y}_{1, \mathcal{S}}>w_{k}\right)=\mathbb{P}\left(\bar{Y}_{1, n_{k}}>\right.$ $w_{k}$ ) for all such subsets. Choosing $\omega_{k}$ such that $I_{1}\left(\omega_{k}\right)=I_{k}\left(\omega_{k}\right)$ (which is possible given the continuity assumptions on the two rate functions), we obtain

$$
\begin{equation*}
\sum_{r=8}^{T-1} \mathbb{P}\left(\mathcal{Z}^{r} \cap \mathcal{D}^{r}\right) \leq \sum_{r=8}^{T-1} \sum_{u=a_{r}}^{r} \mathbb{P}\left(\mathcal{B}^{u}\right) \leq \sum_{k=2}^{K} \frac{2 e^{(2+1 / K) I_{1}\left(\omega_{k}\right)}}{\left(1-e^{-I_{1}\left(\omega_{k}\right)}\right)\left(1-e^{-I_{1}\left(\omega_{k}\right) / 4 K}\right)^{2}} \tag{7}
\end{equation*}
$$

## D.2.2 Controlling $\mathbb{P}\left(\mathcal{Z}^{r} \cap \mathcal{D}^{r}\right)$ : arm 1 has not been leader between $\lfloor r / 4\rfloor$ and $r$

The idea in this part is to leverage the fact that if the optimal arm is not leader between $\lfloor s / 4\rfloor$ and $s$, then it has necessarily lost a lot of duels against the current leader at each round. We then use the
fact that when the leader has been drawn "enough", concentration prevents this situation with large probability. We introduce

$$
\mathcal{L}^{r}=\sum_{u=s_{0}}^{r} \mathbb{1}_{\mathcal{C}^{u}}
$$

for the event $\mathcal{C}^{u}=\left\{\exists a \neq 1, N_{a}(u) \geq N_{1}(u), \hat{Y}_{a, S_{1}^{u}\left(N_{a}(u), N_{1}(u)\right)} \geq \hat{Y}_{1, N_{1}(u)}\right\}$. One can prove the following inequality:

$$
\mathbb{P}\left(\mathcal{Z}^{r} \cap \overline{\mathcal{D}}^{r}\right) \leq \mathbb{P}\left(\mathcal{L}^{r} \geq r / 4\right)
$$

Proof. Under $\overline{\mathcal{D}}^{r}$ arm 1 is a challenger for every round $u \in\left[a_{r}, r\right]$. Then, each time $\mathcal{C}^{u}$ is not true arm 1 wins its duel against the current leader and is pulled. Hence, if $\left\{\mathcal{L}_{r}<r / 4\right\}$ then we necessarily have $\left\{N_{1}(r)>r / 2\right\}$ and arm 1 is leader in round $r$. Hence, $\left\{\mathcal{Z}^{r} \cap \overline{\mathcal{D}}^{r}\right\} \cap\left\{\mathcal{L}_{r}<r / 4\right\}=\emptyset$, which justifies the inequality.

Now, as in [2] we use the Markov inequality to get:

$$
\mathbb{P}\left(\mathcal{L}^{r} \geq r / 4\right) \leq \frac{\mathbb{E}\left(\mathcal{L}^{r}\right)}{r / 4}=\frac{4}{r} \sum_{u=\lfloor r / 4\rfloor}^{r} \mathbb{P}\left(\mathcal{C}^{u}\right)
$$

By further decomposing the probability of $\mathbb{P}\left(\mathcal{C}^{u}\right)$ in two parts depending on the value of the number of selections of arm 1, we obtain the upper bound

$$
\mathbb{P}\left(\mathcal{Z}^{r} \cap \overline{\mathcal{D}}^{r}\right) \leq \frac{4}{r} \sum_{u=\lfloor r / 4\rfloor}^{r} \mathbb{P}\left(N_{1}(u) \leq(\log u)^{2}\right)+\frac{4}{r} \underbrace{\sum_{u=\lfloor r / 4\rfloor}^{r} \mathbb{P}\left(\mathcal{C}^{u}, N_{1}(u) \geq(\log u)^{2}\right)}_{B_{r}}
$$

We now upper bound the quantity $B_{r}$ defined above by using Lemma4.2 For each $a$, for any $\omega_{a}$ such that $\omega_{a} \in\left(\mu_{a}, \mu_{1}\right)$, one can write

$$
\begin{aligned}
B_{r} & \leq \sum_{u=\lfloor r / 4\rfloor}^{r} \mathbb{P}\left(\mathcal{C}^{u}, N_{1}(u) \geq(\log \lfloor r / 4\rfloor)^{2}\right) \\
& \leq \sum_{a=2}^{K} \sum_{u=\lfloor r / 4\rfloor}^{r} \mathbb{P}\left(Y_{a, S_{1}^{u}\left(N_{a}(u), N_{1}(u)\right)}^{r} \hat{Y}_{1, N_{1}(u)}, N_{1}(u) \geq \log (\lfloor r / 4\rfloor)^{2}, N_{a}(u)>N_{1}(u)\right) \\
& \leq \sum_{a=2}^{K}\left(\frac{1}{1-e^{-I_{1}\left(\omega_{a}\right)}} e^{-(\log \lfloor r / 4\rfloor)^{2} I_{1}\left(\omega_{a}\right)}+\frac{r}{1-e^{-I_{k}\left(\omega_{a}\right)}} e^{-(\log \lfloor r / 4\rfloor)^{2} I_{a}\left(\omega_{a}\right)}\right)
\end{aligned}
$$

Choosing each $\omega_{a}$ such that $I_{1}\left(\omega_{a}\right)=I_{a}\left(\omega_{a}\right)$, we obtain:

$$
\begin{aligned}
\frac{4}{r} \sum_{r=8}^{T} B_{r} & \leq \sum_{r=8}^{T} \sum_{a=2}^{K} \frac{4(r+1)}{r\left(1-e^{-I_{1}\left(\omega_{a}\right)}\right)} e^{-(\log \lfloor r / 4\rfloor)^{2} I_{a}\left(\omega_{a}\right)} \\
& \leq \sum_{a=2}^{K} \sum_{r=8}^{T} \frac{6}{1-e^{-I_{1}\left(\omega_{a}\right)}} e^{-(\log \lfloor r / 4\rfloor)^{2} I_{a}\left(\omega_{a}\right)}
\end{aligned}
$$

and for each $a$ the series in $r$ is convergent as for any constant $C, C \log (r) \leq(\log \lfloor r / 4\rfloor)^{2}$ for $r$ large enough. Hence, there exists some constant $D(\boldsymbol{\nu})$ where $\boldsymbol{\nu}=\left(\nu_{1}, \ldots, \nu_{K}\right)$ such that $\frac{4}{r} \sum_{r=8}^{T} B_{r} \leq$ $D(\boldsymbol{\nu})$. It follows that

$$
\sum_{r=8}^{T} \mathbb{P}\left(\mathcal{Z}^{r} \cap \overline{\mathcal{D}}^{r}\right) \leq \sum_{r=8}^{T} \frac{4}{r} \sum_{u=\lfloor r / 4\rfloor}^{r} \mathbb{P}\left(N_{1}(u) \leq(\log u)^{2}\right)+D(\boldsymbol{\nu})
$$

We now transform the double sum in the right-hand side into a simple sum by counting the number of times each term appears in the double sum:

$$
\sum_{r=8}^{T} \frac{4}{r} \sum_{u=\lfloor r / 4\rfloor}^{r} \mathbb{P}\left(N_{1}(u) \leq(\log u)^{2}\right)=\sum_{r=8}^{T}\left(\sum_{t=1}^{r} \frac{4}{t} \mathbb{1}(t \in[r, 4 r])\right) \mathbb{P}\left(N_{1}(r) \leq(\log r)^{2}\right)
$$

Noting that $\sum_{t=1}^{r} \frac{4}{t} \mathbb{1}(t \in[s, 4 s]) \leq(4 s-s+1) \times \frac{4}{s} \leq 16$, we finally obtain:

$$
\begin{equation*}
\sum_{r=8}^{T} \mathbb{P}\left(\mathcal{Z}^{r} \cap \overline{\mathcal{D}}^{r}\right) \leq 16 \sum_{r=1}^{T} \mathbb{P}\left(N_{1}(r) \leq(\log (r))^{2}\right)+D(\boldsymbol{\nu}) \tag{8}
\end{equation*}
$$

Combining (7) and (8) yields

$$
\sum_{r=1}^{T} \mathbb{P}\left(\mathcal{Z}^{r}\right) \leq 16 \sum_{r=1}^{T} \mathbb{P}\left(N_{1}(r) \leq(\log (r))^{2}\right)+D_{k}^{\prime}(\boldsymbol{\nu})
$$

for some constant $D_{k}^{\prime}(\boldsymbol{\nu})$ that depends on $k$ and $\boldsymbol{\nu}$, which contributes to the final constant $C_{k}(\boldsymbol{\nu}, \varepsilon)$ in Lemma 4.1. Plugging this inequality in Equation (3) concludes the proof of Lemma 4.1

## E Probability that the Optimal Arm is not Drawn Enough: Proof of Lemma 4.3

We start with a decomposition that follows the steps of [1] for BESA with 2 arms that we generalize for $K$ arms.
We first denote by $r_{j}$ the round of the $j^{\text {th }}$ play of arm 1 with $r_{0}=0$ and let $\tau_{j}=r_{j+1}-r_{j}$. We notice that $\tau_{0} \leq K$ as all arms are initialized once. Then:

$$
\begin{aligned}
\mathbb{P}\left(N_{1}(r) \leq(\log r)^{2}\right) & \leq \mathbb{P}\left(\exists j \in\left\{1, \ldots, \log r^{2}\right\}: \tau_{j} \geq r /(\log r)^{2}-1\right) \\
& \leq \sum_{j=1}^{(\log r)^{2}} \mathbb{P}\left(\tau_{j} \geq r /(\log r)^{2}-1\right)
\end{aligned}
$$

Proof. If we assume that $\forall j \tau_{j} \leq r /(\log r)^{2}-1$ then $t_{\log r^{2}}=\sum_{j=0}^{\log r^{2}} \tau_{j}<r$, which yields $N_{\ell}(r)>\log r^{2}+1$.

We now fix $j \leq(\log r)^{2}$ and upper bound the probability of the event

$$
\mathcal{E}_{j}:=\left\{\tau_{j} \geq r / \log r^{2}-1\right\}
$$

On this event arm 1 lost at least $r / \log r^{2}$ consecutive duels between $r_{j}+1$ and $r_{j+1}$ (either as a challenger of as the leader) which yields

$$
\begin{aligned}
\mathbb{P}\left(\mathcal{E}_{j}\right) \leq \mathbb{P}\left(\forall s \in\left\{r_{j}+1, \ldots, r_{j}+\left\lfloor r / \log r^{2}-1\right\rfloor\right\}:\right. & \left\{\hat{Y}_{1, j} \leq \hat{Y}_{\ell(s), S_{1}^{s}\left(N_{\ell(s)}(s), j\right)}, N_{1}(s)=j, N_{\ell(s)}(s) \geq j\right\} \\
& \left.\cup\left\{\ell(s)=1, N_{1}(s)=j\right\}\right)
\end{aligned}
$$

The important change compared to the proof of [1] is that with $K>2,1$ ) we don't know the identity of the leader and 2) the leader is not necessarily pulled if it wins its duel against 1.
Now we notice that when $r$ is large, the time range considered in $\mathcal{E}_{j}$ is large. By looking at the second half of this time range only, we can ensure that the leader has been drawn a large number of times. More precisely, introducing the two intervals

$$
\begin{aligned}
& \mathcal{M}_{j, r}^{1}=\left[r_{j}+1, r_{j}+\left\lfloor\frac{r / \log r^{2}-1}{2}\right\rfloor\right] \\
& \mathcal{M}_{j, r}^{2}=\left[t_{j}+\left\lceil\frac{r / \log r^{2}-1}{2}\right\rceil, t_{j}+\left\lfloor r / \log r^{2}\right\rfloor-1\right]
\end{aligned}
$$

it holds that
$\mathbb{P}\left(\mathcal{E}_{j}\right) \leq \mathbb{P}\left(\forall s \in \mathcal{M}_{j, r}^{2}:\left\{\hat{Y}_{1, j} \leq \hat{Y}_{\ell(s), S_{1}^{s}\left(N_{\ell(s)}(s), j\right)}, N_{1}(s)=j, N_{\ell(s)}(s) \geq j\right\} \cup\left\{\ell(s)=1, N_{1}(s)=j\right\}\right)$.
But we know that on $\mathcal{M}_{j, r}^{2}$ the leader must has been selected at least $\frac{1}{K}\left(j+\left\lceil\frac{r / \log r^{2}-1}{2}\right\rceil\right)$ times. Let $r_{K}$ be the first integer such that $\log ^{2}(r)<\frac{1}{K-1}\left\lceil\frac{r / \log r^{2}-1}{2}\right\rceil$, for every $r \geq r_{K}$, as $j \leq \log ^{2}(r)$,
the leader has been selected strictly more than $j$ times, which prevents arm 1 from being the leader for any round in $\mathcal{M}_{j, r}^{2}$. Hence, for $r \geq r_{K}$, for all $j \leq \log ^{2}(r)$,

$$
\mathbb{P}\left(\mathcal{E}_{j}\right) \leq \mathbb{P}\left(\forall s \in \mathcal{M}_{j, r}^{2}:\left\{\hat{Y}_{1, j} \leq \hat{Y}_{\ell(s), S_{1}^{s}\left(N_{\ell(s)}(s), j\right)}, N_{1}(s)=j, N_{\ell(s)}(s) \geq j\right\}\right)
$$

To remove the problem of the identity of the leader we would like to find a way to fix our attention on one arm. To this extent, we notice that during an interval of length $\left|\mathcal{M}_{j, r}^{2}\right|$, if there are only $K-1$ candidates for the leader then one of them must have been leader at least $m_{r}:=\left|\mathcal{M}_{j, r}^{2}\right| /(K-1)-1$ times during this range. We also know that at any round in $\mathcal{M}_{j, r}^{2}$, the leader satisfies $N_{\ell(s)}(s) \geq$ $\left(t_{j}+\left\lfloor\frac{r / \log r^{2}-1}{2}\right\rfloor\right) / K-1 \geq\left(\left\lfloor\frac{r / \log r^{2}-1}{2}\right\rfloor\right) / K-1=\frac{\left|\mathcal{M}_{j, r}^{1}\right|}{K}-1:=c_{r}$. Observe that $m_{r}>c_{r}$. Finally, we introduce the notation

$$
I_{j, r}^{k}=\left\{s \in \mathcal{M}_{j, r}^{2}: \ell(s)=k\right\}
$$

for the set of rounds in $\mathcal{M}_{j, r}^{2}$ in which a particular arm $k$ is leader. From the above discussion, we know that there exists an arm $k$ such that $\left|I_{j, r}^{k}\right| \geq m_{r}$.
To ease the notation, we introduce the event

$$
\mathcal{W}_{s, j}^{k}=\left\{\left\{\hat{Y}_{1, j}<\hat{Y}_{k, S_{1}^{s}\left(N_{k}(s), j\right)}\right\}, N_{k}(s) \geq c_{r}, N_{1}(s)=j\right\}
$$

and write

$$
\begin{aligned}
\mathbb{P}\left(\mathcal{E}_{j}\right) & \left.\leq \mathbb{P}\left(\bigcap_{s \in \mathcal{M}_{j, r}^{2}} \bigcup_{k=2}^{K}\left\{\ell(s)=k, 1 \notin \mathcal{A}_{s}\right)\right\}\right) \\
& \leq \mathbb{P}\left(\bigcap_{k=2}^{K} \bigcap_{s \in I_{j, r}^{k}} \mathcal{W}_{s, j}^{k}\right) \\
& \leq \mathbb{P}\left(\bigcup_{k=2}^{K}\left\{\left|I_{j, r}^{k}\right|>m_{r}, \bigcap_{s \in I_{j, r}^{k}} \mathcal{W}_{s, j}^{k}\right\}\right) \\
& \leq \sum_{k=2}^{K} \mathbb{P}\left(\left|I_{j, r}^{k}\right|>m_{r}, \bigcap_{s \in I_{j, r}^{k}} \mathcal{W}_{s, j}^{k}\right)
\end{aligned}
$$

Finally, we define for any integer $M$ the event that we can find $M$ pairwise non-overlapping subsamples in the set of the sub-samples of arm $k$ drawn in rounds $s \in I_{j, r}^{k}$ :

$$
\mathcal{F}_{j, M}^{k, r}=\left\{\exists i_{1}, \ldots, i_{M} \in I_{j, r}^{k}: \forall m<m^{\prime} \in[M], S_{1}^{i_{m}}\left(N_{k}\left(i_{m}\right), j\right) \cap S_{1}^{i_{m^{\prime}}}\left(N_{k}\left(i_{m^{\prime}}\right), j\right)=\emptyset\right\}
$$

Introducing $H_{j, r}^{k}=\min _{s \in I_{j, r}^{k}} N_{k}(s)$, the minimal size of the history of arm $k$ during rounds in $I_{j, r}^{k}$ (which is known to be larger than $c_{r}$ as $k$ is leader in these rounds), one has

$$
\begin{align*}
& \mathbb{P}\left(\mathcal{E}_{j}\right) \leq \sum_{k=2}^{K} \mathbb{P}\left(\left|I_{j, r}^{k}\right|>m_{r}, \cap_{s \in I_{j, r}^{k}} \mathcal{W}_{s, j} \cap\left\{\mathcal{F}_{j, M}^{k, r} \cup \overline{\mathcal{F}}_{j, M}^{k, r}\right\}\right) \\
& \leq \sum_{k=2}^{K} \mathbb{P}\left(\left|I_{j, r}^{k}\right| \geq m_{r}, H_{j, r}^{k} \geq c_{r}, \overline{\mathcal{F}}_{j, M}^{k, r}\right)+\sum_{k=2}^{K} \mathbb{P}\left(\left|I_{j, r}^{k}\right|>m_{r}, \cap_{s \in I_{j, r}^{k}} \mathcal{W}_{s, j} \cap \mathcal{F}_{j, M}^{k, r}\right) \tag{9}
\end{align*}
$$

Upper bound on the first term in (9) The probability $\mathbb{P}\left(\left|I_{j, r}^{k}\right| \geq m_{r}, H_{j, r}^{k} \geq c_{r}, \overline{\mathcal{F}}_{j, M}^{k, r}\right)$ can be upper bounded by
$\mathbb{P}\left(\#\left\{\right.\right.$ pairwise non-overlapping subsets in $\left.\left.\left(S_{1}^{s}\left(N_{k}(s), j\right)\right)_{s \in I_{j, r}^{k}}\right\}<M \mid\left\{\left|I_{j, r}^{k}\right|>m_{r}, H_{j, r}^{k} \geq c_{r}\right\}\right)$.
This probability can be related to some intrinsic properties of the sampler $\operatorname{SP}(H, j)$. To formalize this, we introduce the following definition.

Definition. For every integers $N, H, j$ such that $H>j, X_{N, H, j}$ is a random variable which counts the maximum number of non-overlapping subsets among $N$ i.i.d. samples from $\operatorname{SP}(H, j)$.

Letting $H_{1}, \ldots, H_{m_{r}}$ be integers that are all larger than $c_{r}$, and letting $S_{1}, \ldots, S_{m_{r}}$ be independent subsets such that $S_{i} \sim \mathrm{SP}\left(H_{i}, j\right)$, the above probability is upper bounded by

$$
\left.\mathbb{P}\left(\# \text { pairwise non-overlapping subsets in }\left(S_{i}\right)_{i=1}^{m_{r}}\right\}<M\right)
$$

which is itself upper bounded by $\mathbb{P}\left(X_{m_{r}, c_{r}, j}<M\right)$.
This last inequality is quite intuitive: if one draws subsets of size $j$ from histories that may be larger than $c_{r}$, there is more "room" for non-overlapping subsets than if we always draw them from the same history of size $c_{r}$. For Random Block sampling, where the drawn subset is fully determined by the random position of its first element, to formalize this intuition it is sufficient to prove that if $X_{i}, Y_{i}$ are two sequences of random variables such that $X_{i}$ is uniform in $\left[H_{i}-j\right]$ and $Y_{i}$ is uniform in $[H-j]$, where $H_{i} \geq H$, the random variable that counts the maximal number of elements in the sequence $\left(Y_{i}\right)$ whose pairwise distance are larger than $j$ is stochastically dominated by that the same random variable but for the sequence $\left(X_{i}\right)$. We performed numerical experiments that confirm that this last condition holds.

Upper bound on the second term in (9) On the event $\left(\left|I_{j, r}^{k}\right|>m_{r}, \cap_{s \in I_{j, r}^{k}} \mathcal{W}_{s, j} \cap \mathcal{F}_{j, M}^{k, r}\right)$, one can define $\tilde{\imath}_{1}, \ldots, \tilde{\imath}_{M}$ the first $M$ rounds in $I_{j, r}^{k}$ for which the subsets $\tilde{S}_{m}:=S^{\tilde{\imath}_{m}}\left(N_{k}\left(\tilde{\imath}_{m}\right), j\right)$ are pairwise non-overlapping and we get

$$
\mathbb{P}\left(\left|I_{j, r}^{k}\right|>m_{r}, \cap_{s \in I_{j, r}^{k}} \mathcal{W}_{s, j} \cap \mathcal{F}_{j, M}^{k, r}\right) \leq \mathbb{P}\left(\forall m \in[M], \hat{Y}_{1, j} \leq \hat{Y}_{k, \tilde{S}_{m}}\right)
$$

By definition the subsets $\tilde{S}_{m}$ are pairwise non-overlapping, hence the sub-samples $\hat{Y}_{k, \tilde{S}_{m}}$ are independent. We prove that this probability can be in fact upper bound by the balance function we defined in section 3
Indeed, introducing $X \sim \nu_{1, j}$ and an independent i.i.d. sequence $Z_{i} \sim \nu_{k, j}$, one can write

$$
\begin{aligned}
\mathbb{P}\left(\left|I_{j, r}^{k}\right|>m_{r}, \cap_{s \in I_{j, r}^{k}} \mathcal{W}_{s, j} \cap \mathcal{F}_{j, M}^{k, r}\right) & \leq \mathbb{P}\left(X<\min _{i \in[M]} Z_{i}\right) \\
& =\mathbb{E}_{X \sim \nu_{1, j}}\left[\prod_{Z \sim \nu_{k, j}^{\otimes, j}}^{M} \mathbb{1}_{X \leq Z_{i}}\right] \\
& =\mathbb{E}_{X \sim \nu_{1, j}}\left[\mathbb{E}_{Z \sim \nu_{k, j}^{\otimes j}}\left[\prod_{i} \mathbb{1}_{X \leq Z_{i}} \mid X\right]\right] \\
& =\mathbb{E}_{X \sim \nu_{1, j}}\left[\left(1-F_{k, j}(X)\right)^{M}\right] \\
& =\alpha_{k}(M, j)
\end{aligned}
$$

Conclusion Putting things together, we have proved that

$$
\mathbb{P}\left(\mathcal{E}_{j}\right) \leq(K-1) \mathbb{P}\left(X_{m_{r}, c_{r}, j}<M\right)+\sum_{k=2}^{K} \alpha_{k}(M, j)
$$

where $X_{N, H, j}$ and $\alpha_{k}(M, j)$ are introduced in Definition Eand respectively. If we replace $M$ by the sequence $\beta_{r, j}$ we have

$$
\begin{aligned}
\sum_{r=1}^{T} \mathbb{P}\left(N_{1}(r) \leq \log r^{2}\right) & \leq r_{K}+\sum_{r=r_{K}}^{T} \sum_{j=1}^{\log r^{2}}\left[(K-1) \mathbb{P}\left(X_{m_{r}, c_{r}, j}<\beta_{r, j}\right)+\sum_{k=2}^{K} \alpha_{k}\left(\beta_{r, j}, j\right)\right] \\
& \leq r_{K}+\sum_{r=r_{K}}^{T} \sum_{j=1}^{\log r^{2}}\left[(K-1) \mathbb{P}\left(X_{c_{r}, c_{r}, j}<\beta_{r, j}\right)+\sum_{k=2}^{K} \alpha_{k}\left(\beta_{r, j}, j\right)\right]
\end{aligned}
$$

as $c_{r} \leq m_{r}$, which proves Lemma 4.3
This definition allows to analyze separately the properties of the sub-sampling algorithms and the properties of the distribution family for randomized samplers.

## F Proof that RB-SDA Satisfies the Diversity Property

We recall that $X_{m, H, j}$ denotes the maximal number of pairwise non-overlapping subsets obtained in $m$ i.i.d. samples from $\operatorname{RB}(H, j)$. In this section we aim at upper bounding the probability of

$$
\mathbb{P}\left(X_{m, H, j} \leq \gamma r /(\log r)^{2}\right)
$$

for some values of $m, H, j$, that will be fixed later. This probability depends on several parameters, with straightforward effects:

- The probability decreases with the length of the history size $H$.
- The probability increases with the size $j$ of each sub-sample.
- The probability decreases with the total number of sub-samples we draw $m$. Intuitively if $m$ is large enough every sample of size $j$ in the history will be drawn.

First step with $j=1$ : in this case the distribution of the $m$ subsets of size 1 is actually the distribution of sampling with replacement in $H$. The question of the number of different items drawn with sampling without replacement has been studied in [25], from which we use the following result:
Result 1: for any $k \in[H], \mathbb{P}\left(X_{m, H, 1}=k\right)=\frac{H!}{(H-k)!\times H^{m}} \times S_{k, m}$, where $S_{k, m}$ is the Stirling number of the second kind for $k, m$.

We use this result with further assumptions that are specific to our problem and will ease the computation: $H=m=O\left(r /(\log r)^{2}\right)$. To ease the notation we continue to use $H$, and write $\gamma t /(\log t)^{2}=\alpha H$ for some $\alpha \in(0,1)$.
We first look at $\mathbb{P}\left(X_{H, H, 1}=\alpha H\right)$. According to [26] the following inequality holds

$$
S_{k, H} \leq \frac{1}{2}\binom{H}{k} k^{H-k}
$$

This allows to upper bound the expression in result 1 :

$$
\mathbb{P}\left(X_{H, H, 1}=k\right) \leq \frac{1}{2}\left(\frac{k}{H}\right)^{H-k}\binom{H}{k}
$$

We now want to bound $\binom{H}{k}$. As $k$ is small compared with $H$, it is natural to use

$$
\binom{H}{k} \leq \frac{H^{k}}{k!}
$$

We then bound $1 / k$ ! by its Stirling approximation and add a multiplicative constant $c$ along the way:

$$
\binom{H}{k} \leq c \frac{H^{k}}{\sqrt{2 \pi k} \times k^{k}} e^{k}
$$

Refactoring provides

$$
\mathbb{P}\left(X_{H, H, 1}=k\right) \leq \frac{c}{2}\left(\frac{k}{H}\right)^{H-2 k} \frac{e^{k}}{\sqrt{2 \pi k}}
$$

Then we notice that if $k \leq H-2 k$ we get:

$$
\mathbb{P}\left(X_{H, H, 1}=k\right) \leq \frac{c}{2 \sqrt{2 \pi k}}\left(\frac{k e}{H}\right)^{H-2 k}
$$

Now we can replace $k$ by $\alpha H$ (assume it's an integer for the simplicity of notations), such that 1 ) $\alpha \leq \frac{1}{3} \Rightarrow H(1-3 \alpha)>0$ and $\alpha e<1$.

If $k \leq \alpha H, \alpha e \leq 1$ and $(1-2 \alpha)>0$ then: $\left(\frac{k e}{H}\right)^{H-2 k} \leq(\alpha e)^{H-2 k}$. We have:

$$
\begin{align*}
\mathbb{P}\left(X_{H, H, 1} \leq \alpha H\right) & \leq \frac{c}{2 \sqrt{2 \pi}} \sum_{k=0}^{\lfloor\alpha H\rfloor}(\alpha e)^{H-2 k} \\
& \leq \frac{c}{2 \sqrt{2 \pi}} \sum_{k=0}^{\lfloor\alpha H\rfloor}(\alpha e)^{H-2(\lfloor\alpha H\rfloor-k))}  \tag{10}\\
& \leq \frac{c}{2 \sqrt{2 \pi}}(\alpha e)^{H-2\lfloor\alpha H\rfloor} \frac{1}{1-(\alpha e)^{2}} \\
& \leq \frac{c}{2 \sqrt{2 \pi}} \frac{1}{1-(\alpha e)^{2}} \exp (-(1-2 \alpha) H \log (1 /(\alpha e)))
\end{align*}
$$

From $X_{H, H, 1}$ to $X_{H, H, j}$ This result is enough to get general properties for Random Block Sampling. Indeed, as the process of RBS consists in only drawing the first element of the block used in the duel we can see that the previous bound also applies to the number of unique starting points. With this property, the Random Block Sampler satisfies for all $x>0$ :

$$
\mathbb{P}\left(X_{m, H, j} \leq\left\lfloor\frac{x}{j}\right\rfloor\right) \leq \mathbb{P}\left(X_{m, H, 1} \leq x\right)
$$

Proof. Assume that the Random Block Sampler provides $x$ blocks with different starting points. Let's further assume that $x$ is an integer and try to identify the sequence of starting times $t_{i}=\left(t_{1}, \ldots, t_{x}\right)$ that minimizes the number of mutually non-overlapping samples: the value of $t_{1}$ is not important due to the symmetry of the problem. Then if we want to reduce the possibilities to get non-overlapping sample we want to choose a value for $t_{2}$ that 1 ) makes the blocks $\left[t_{1}, t_{1}+j\right]$ and $\left[t_{2}, t_{2}+k\right]$ nonoverlapping and 2) makes things easier to continue this process for $t_{3}, \ldots, t_{m}$. It seems intuitive to choose either the block starting at $t_{1}+1$ or at $t_{1}-1$ as we cover the minimum amount of space with the constraint that $t_{2} \neq t_{1}$. If we repeat this choice until $m$ blocks are chosen and reorder the blocks properly, we get a sequence of starting points $\left[t_{1}, t_{1}+1, \ldots t_{1}+m\right]$ that are all different and minimize the total amount of space covered by the block. Even in this setup, we can find exactly $\left\lfloor\frac{m}{j}\right\rfloor$ mutually non-overlapping blocks as for instance all $\left[t_{1}+k j, t_{1}+(k+1) j-1\right],\left[t_{1}+k^{\prime} j, t_{1}+\left(k^{\prime}+1\right) j-1\right]$ blocks are non-overlapping for $k \neq k^{\prime}$ and $\left(k, k^{\prime}\right) \in\left[0,\left\lfloor\frac{m}{j}\right\rfloor-1\right]$.

We can finally prove the following for Random Block sampling.
Lemma F. 1 (Diversity Property for Random Block Sampling). If we choose a constant $\gamma \leq 1 / 3 \times \frac{1}{2 K}$ then Random Block Sampling satisfies the diversity property.

Proof. For $\gamma \leq\left\lfloor 1 / 3 \times \frac{1}{2 K}\right\rfloor$, there exists $\alpha>0$ such that:

$$
\begin{aligned}
\mathbb{P}\left(X_{c_{r}, c_{r}, j} \leq \gamma / j\left(r /(\log r)^{2}\right)\right) & \leq \mathbb{P}\left(X_{c_{r}, c_{r}, j} \leq \alpha / j c_{r}\right) \\
& \leq \mathbb{P}\left(X_{c_{r}, c_{r}, 1} \leq \alpha c_{r}\right) \\
& =o\left(r^{-2}\right)
\end{aligned}
$$

The last line comes from the expression obtained in Equation 10), and allows to conclude that $\sum_{r=1}^{T} \sum_{j=1}^{(\log r)^{2}} \mathbb{P}\left(X_{c_{r}, c_{r}, j} \leq \alpha / j(r /(\log r))^{2}\right)=o(\log T)$

## G Analysis of the Balance Function for Some Distributions

For the simplicity of the notation we write the balance function $\alpha(M, j)$ for any distribution and any instance of these distributions. The family of distributions and the notation for their parameter will always mentioned at the beginning of the corresponding sub-section.
In the next parts we use the notation $G(x)=1-F(x)$ where $F$ denotes the CDF of some distribution. For some arm distribution $\nu_{i}$ the distribution of the sum of $j$ independent observations drawn from $\nu_{i}$ is denoted by $\nu_{i, j}$. With this notation, for two arms 1 and 2 we write:

$$
\alpha(M, j)=\mathbb{E}_{Z \sim \nu_{1, j}}\left(G_{2, j}(Z)^{M}\right)
$$

## G. 1 The Bernoulli Distribution is Balanced

We prove the following lemma, which bears strong similarity with an upper bound given by [10] for a similar quantity in their analysis of Thompson Sampling.

Lemma G. 1 (Bound on $\alpha(M, j)$ ). For two Binomial distributions $\nu_{1} \sim \mathcal{B}\left(j, \mu_{1}\right)$ and $\nu_{2} \sim \mathcal{B}\left(j, \mu_{2}\right)$ such that $\mu_{1}>\mu_{2}$ and for any integer $M>1: \exists \lambda>1$ such as

$$
\mathbb{E}_{X \sim \nu_{1, j}}\left(\left(1-F_{j, \mu_{2}}(X)\right)^{M}\right) \leq C_{\lambda_{0}, \lambda} \frac{1}{M^{\lambda}} e^{-j d_{\lambda, \mu_{1}, \mu_{2}}}+\left(\frac{1}{2}\right)^{M}
$$

Where $C_{\lambda, \mu_{1}, \mu_{2}}>0$, and $F_{j, \mu_{2}}$ is the CDF of a Binomial $\mathcal{B}\left(j, \mu_{2}\right)$.

Proof. We use the same notation as before: $G_{2}(k)=1-F_{2}(k)$ and $f_{1}, f_{2}$ as the PMF of $\nu_{1}, \nu_{2}$.
We first use a common property of Binomial distributions, $\forall k>\left\lceil j \mu_{2}\right\rceil: G(k) \leq \frac{1}{2}$. So we can directly write:

$$
\mathbb{E}_{X \sim \nu_{1}}\left(\left(1-F_{j, \mu_{2}}(X)\right)^{M}\right) \leq\left(\frac{1}{2}\right)^{M}+\underbrace{\sum_{k=0}^{\left\lfloor j \mu_{2}\right\rfloor} f_{1}(k) G_{2}(k)^{M}}_{(A)}
$$

Using convexity we get: $G(k)^{M} \leq \exp \left(-M F_{2}(k)\right)$, hence

$$
(A) \leq \sum_{k=0}^{\left\lfloor j \mu_{2}\right\rfloor} f_{1}(k) \exp \left(-M F_{2}(k)\right)
$$

Then we use that for $\lambda>1, \forall x>0: x^{\lambda} e^{-x} \leq\left(\frac{\lambda}{e}\right)^{\lambda}=C_{\lambda}$, so:

$$
(A) \leq \frac{C_{\lambda}}{M^{\lambda}} \sum_{k=0}^{\left\lceil j \mu_{2}\right\rceil} \frac{f_{1}(k)}{F_{2}(k)^{\lambda}} \leq \frac{C_{\lambda}}{M^{\lambda}} \sum_{k=0}^{\left\lceil j \mu_{2}\right\rceil} \frac{f_{1}(k)}{f_{2}(k)^{\lambda}}
$$

As in [10], we compute:

$$
\begin{aligned}
\frac{f_{1}(k)}{f_{2}(k)^{\lambda}} & \leq \frac{\mu_{1}^{k}\left(1-\mu_{1}\right)^{j-k}}{\mu_{2}^{\lambda k}\left(1-\mu_{2}\right)^{\lambda\left(j_{k}\right)}} \\
& \leq\left(\frac{1-\mu_{1}}{\left(1-\mu_{2}\right)^{\lambda}}\right)^{j}\left(\frac{\mu_{1}\left(1-\mu_{2}\right)^{\lambda}}{\mu_{2}^{\lambda}\left(1-\mu_{1}\right)}\right)^{k} \\
& =\left(\frac{1-\mu_{1}}{\left(1-\mu_{2}\right)^{\lambda}}\right)^{j} R_{\lambda}\left(\mu_{1}, \mu_{2}\right)^{k}
\end{aligned}
$$

with $R_{\lambda}\left(\mu_{1}, \mu_{2}\right)=\frac{\mu_{1}\left(1-\mu_{2}\right)^{\lambda}}{\mu_{2}^{\lambda}\left(1-\mu_{1}\right)}$. We then notice that we can choose $\lambda>1$ such that $R_{\lambda}\left(\mu_{1}, \mu_{2}\right)>1$. It is true for any $\lambda>1$ if $\mu_{2} \leq 0.5$, and for $1<\lambda<\log \frac{\mu_{1}}{1-\mu_{1}} / \log \frac{\mu_{2}}{1-\mu_{2}}$ if $\mu_{2}>0.5$.

Plugging that expression into the sum gives:

$$
\begin{aligned}
(A) & \leq \frac{C_{\lambda}}{M^{\lambda}} \sum_{k=0}^{\left\lceil j \mu_{2}\right\rceil} \frac{f_{1}(k)}{f_{2}(k)^{\lambda}} \\
& \leq \frac{C_{\lambda}}{M^{\lambda}}\left(\frac{1-\mu_{1}}{\left(1-\mu_{2}\right)^{\lambda}}\right)^{j} \sum_{k=0}^{\left\lceil j \mu_{2}\right\rceil} R_{\lambda}\left(\mu_{1}, \mu_{2}\right)^{k} \\
& =\frac{C_{\lambda}}{M^{\lambda}}\left(\frac{1-\mu_{1}}{\left(1-\mu_{2}\right)^{\lambda}}\right)^{j} \frac{R_{\lambda}\left(\mu_{1}, \mu_{2}\right)^{\left\lfloor j \mu_{2}\right\rfloor+1}-1}{R_{\lambda}\left(\mu_{1}, \mu_{2}\right)-1} \\
& \leq \frac{C_{\lambda}}{M^{\lambda}}\left(\frac{1-\mu_{1}}{\left(1-\mu_{2}\right)^{\lambda}}\right)^{j} \frac{R_{\lambda}\left(\mu_{1}, \mu_{2}\right)}{R_{\lambda}\left(\mu_{1}, \mu_{2}\right)-1} R_{\lambda}\left(\mu_{1}, \mu_{2}\right)^{j \mu_{2}} \\
& =\frac{C_{\lambda}}{M^{\lambda}} \frac{R_{\lambda}\left(\mu_{1}, \mu_{2}\right)}{R_{\lambda}\left(\mu_{1}, \mu_{2}\right)-1}\left(\frac{1-\mu_{1}}{\left(1-\mu_{2}\right)^{\lambda}}\right)^{j\left(1-\mu_{2}\right)}\left(\frac{\mu_{1}}{\mu_{2}^{\lambda}}\right)^{j \mu_{2}} \\
& =\frac{C_{\lambda}}{M^{\lambda}} \frac{R_{\lambda}\left(\mu_{1}, \mu_{2}\right)}{R_{\lambda}\left(\mu_{1}, \mu_{2}\right)-1} e^{-j d_{\lambda}\left(\mu_{2}, \mu_{1}\right)} \\
& =\frac{C_{\lambda, \mu_{1}, \mu_{2}}^{M^{\lambda}}}{M^{-j d_{\lambda}\left(\mu_{2}, \mu_{1}\right)}}
\end{aligned}
$$

where $d_{\lambda}\left(\mu_{2}, \mu_{1}\right)=\lambda\left(\mu_{2} \log \mu_{2}+\left(1-\mu_{2}\right) \log \left(1-\mu_{2}\right)\right)-\left(\mu_{2} \log \mu_{1}+\left(1-\mu_{2}\right) \log \left(1-\mu_{1}\right)\right)=$ $\mathrm{KL}\left(\mu_{2}, \mu_{1}\right)-(\lambda-1) \mathrm{H}\left(\mu_{2}\right), \mathrm{KL}\left(\mu_{2}, \mu_{1}\right)$ denotes the KL-divergence between $\nu_{2}$ and $\nu_{1}$, and $\mathrm{H}\left(\mu_{2}\right)=$ $\mathbb{E}_{X \sim \nu_{2, j}}\left(\log f_{2}(X)\right)$. We need to choose $\lambda$ as:

$$
\lambda<1+\frac{\mathrm{KL}\left(\mu_{2}, \mu_{1}\right)}{\mathrm{H}\left(\mu_{2}\right)}=\lambda_{0}\left(\mu_{1}, \mu_{2}\right)
$$

Note that those quantities correspond to the Bernoulli distributions, the $j$ is not involved here. In [10], the authors explain that this condition is more restrictive than the previous one so we can state that $\forall \lambda<\lambda_{0}\left(\mu_{1}, \mu_{2}\right)$ :

$$
\mathbb{E}_{X \sim \nu_{1, j}}\left(\left(1-F_{j, \mu_{2}}(X)\right)^{M}\right) \leq\left(\frac{1}{2}\right)^{M}+\frac{C_{\lambda, \mu_{1}, \mu_{2}}}{M^{\lambda}} e^{-j d_{\lambda}\left(\mu_{2}, \mu_{1}\right)}
$$

This is enough to prove that the Bernoulli distribution is balanced by replacing $M$ by $\left\lfloor\beta t /(\log t)^{2}\right\rfloor$ in the expression in Lemma G. 1 and summing on $t$ and $j$. The power term is in $o\left(t\left(\log t^{2}\right)\right)$, while the other term is the term of a convergent geometric series in $j$ multiplied by a term in $o(1 / t)$ in $t$, which is enough to get the result.

## G. 2 The Poisson Distribution is Balanced

We can actually use the same sketch of proof as for Bernoulli distributions, using that for 2 Poisson random variables:

$$
\frac{p_{1, j}(k)}{p_{2, j}(k)^{\lambda}}=e^{-j\left(\theta_{1}-\lambda \theta_{2}\right)}\left(\frac{k!}{n^{k}}\right)^{\lambda}\left(\frac{\theta_{1}}{\theta_{2}^{\lambda}}\right)^{k} \leq e^{-j\left(\theta_{1}-\lambda \theta_{2}\right)}\left(\frac{\theta_{1}}{\theta_{2}^{\lambda}}\right)^{k}
$$

So:

$$
\begin{aligned}
\sum_{k=0}^{d_{0, j}} \frac{p_{1, j}(k)}{p_{2, j}(k)^{\lambda}} & \leq e^{-j\left(\theta_{1}-\lambda \theta_{2}\right)} \sum_{k=0}^{d_{0, j}}\left(\frac{\theta_{1}}{\theta_{2}^{\lambda}}\right)^{k} \\
& \leq \frac{\theta_{2}^{\lambda}}{\left|\theta_{1}-\theta_{2}^{\lambda}\right|} e^{-j\left(\theta_{1}-\lambda \theta_{2}\right)} \times \max \left\{1,\left(\frac{\theta_{1}}{\theta_{2}^{\lambda}}\right)^{d_{0, j}}\right\},
\end{aligned}
$$

where $d_{0, j}=\theta_{2} j-1$. Now we remark that if we choose $\lambda \in\left(1, \theta_{1} / \theta_{2}\right)$ we have 2 possibilities: 1$)$ we can choose $\lambda$ such that the second term equals one, hence we can bound the whole term by a constant without further conditions, or 2$) \forall \lambda \in\left(0, \theta_{1}, \theta_{2}\right):\left(\frac{\theta_{1}}{\theta_{2}^{\lambda}}\right)>1$. Let us focus on the second case, we study the term $e^{-j\left(\left(\theta_{1}-\lambda \theta 2\right)-\theta_{2}\left(\log \theta_{1}-\lambda \log \theta_{2}\right)\right)}$. As for Bernoulli distributions, we identify the KL-divergence between $\nu_{2}$ and $\nu_{1}$ and write:

$$
\left(\theta_{1}-\lambda \theta 2\right)-\theta_{2}\left(\log \theta_{1}-\lambda \log \theta_{2}\right)=\operatorname{KL}\left(\nu_{2}, \nu_{1}\right)-(\lambda-1) \theta_{2}\left(1-\log \theta_{2}\right)
$$

So if $\log \theta_{2}>1$ we can choose any $\lambda>1$. In the other case we have to restrict our choice of $\lambda$ to get:

$$
\lambda<1+\frac{\operatorname{KL}\left(\nu_{2}, \nu_{1}\right)}{\theta_{2}\left(1-\log \left(\theta_{2}\right)\right)}
$$

So with an appropriate choice for $\lambda$ Poisson distributions are balanced with the same argument that makes Bernoulli distributions balanced.

## G. 3 The Gaussian Distribution is Balanced

For the Gaussian distribution we leverage the fact that both the PDF and CDF of any Gaussian distribution can be expressed with the PDF and CDF of the standard normal distribution. With such decomposition, we can express $\alpha(M, j)$ as a function of these $\mathrm{CDF} / \mathrm{PDF}$ and use some properties of the normal distribution.

We use the notations $f$ and $G$ for the PDF and CDF of the $\mathcal{N}(0,1)$ distribution, $\Delta$ for the gap between the two arms, and compute the expectation:

$$
\begin{aligned}
\alpha(M, j) & =\int_{-\infty}^{+\infty} f_{1, j}(x) G_{2, j}(x)^{M} d x \\
& \leq \int_{-\infty}^{z} f_{1, j}(x) G_{2, j}(x)^{M} d x+G_{2, j}(z)^{M}, \forall z \in \mathbb{R} \\
& \leq \int_{-\infty}^{z} f\left(\frac{x-\mu_{1, j}}{\sqrt{j}}\right) G\left(\frac{x-\mu_{2, j}}{\sqrt{j}}\right)^{M} d x+G_{2, j}(z)^{M} \\
& \leq \int_{-\infty}^{\frac{z-\mu_{2, j}}{\sqrt{j}}} f(y-\sqrt{j} \Delta) G(y)^{M} d y+G_{2, j}(z)^{M}
\end{aligned}
$$

At this step we use two things: 1) the normal distribution satisfies $f(x-a)=e^{-a^{2}+2 a x} f(x)$ for all $a, x$, and 2) $h: x \rightarrow(M+1) f(x) G(x)^{M}$ is a probability distribution of CDF $x \rightarrow 1-G(x)^{M+1}$. We continue the computation with:

$$
\begin{aligned}
\alpha(M, j) & \leq \frac{e^{-j \Delta^{2}}}{M+1} \int_{-\infty}^{\frac{z-\mu_{2} j}{\sqrt{j}}} e^{\sqrt{j} \Delta y} h(y) d y+G_{2, j}(z)^{M} \\
& \leq \frac{e^{-j \Delta^{2}}}{M+1} e^{\sqrt{j} \Delta \frac{z-\mu_{2} j}{\sqrt{j}}}\left(1-G\left(\frac{z-\mu_{2} j}{\sqrt{j}}\right)^{M+1}\right)+G\left(\frac{z-\mu_{2} j}{\sqrt{j}}\right)^{M} \\
& \leq \frac{e^{-j \Delta^{2}}}{M+1} e^{\sqrt{j} \Delta \frac{z-\mu_{2} j}{\sqrt{j}}}+G\left(\frac{z-\mu_{2} j}{\sqrt{j}}\right)^{M}
\end{aligned}
$$

As the inequality is true for all $z \in \mathbb{R}$, it holds that,

$$
\forall y \in \mathbb{R}, \quad \alpha(M, j) \leq \frac{e^{-j \Delta^{2}}}{M+1} e^{\sqrt{j} \Delta y}+G(y)^{M}
$$

Now let $y_{M}$ be such as $G\left(y_{M}\right)=1-\frac{1}{\sqrt{M}}$. This value ensures that the second term satisfies $G\left(y_{M}\right)^{M} \leq e^{-\sqrt{M}}=o\left(M^{-2}\right)$. Observe that $y_{M}=F^{-1}\left(\frac{1}{\sqrt{M+1}}\right)$. Using the following equivalent of the quantile function of the normal distribution when the quantile is small (see for instance [27]):

$$
F^{-1}(p)=-\sqrt{\log \frac{1}{p^{2}}-\log \log \frac{1}{p}+\log 2 \pi}+o_{p \rightarrow 0}(1)
$$

there exists a constant $C \in \mathbb{R}$ such that $y_{M} \leq-C \sqrt{\log M-\log \log M+\log 4 \pi}$. This yields

$$
\alpha(M, j) \leq \frac{e^{-j \Delta^{2}}}{M+1} e^{-C \sqrt{j} \Delta \sqrt{\log M-\log \log M+\log 4 \pi}}+e^{-\sqrt{M}}
$$

Noting that for all $k \in \mathbb{N}^{*}$,

$$
k \log \log M=o(C \sqrt{j} \Delta \sqrt{\log M-\log \log M+\log 4 \pi})
$$

we get that $\forall k \in \mathbb{N}^{*}$ :

$$
\alpha(M, j)=o\left(\frac{e^{-j \Delta^{2}}}{(M+1)(\log M)^{k}}\right)
$$

This is sufficient to prove that the Gaussian distribution is balanced. Indeed, as for the Bernoulli distribution this term sums as a convergent geometric series in $j$, and with $M=O\left(t /(\log t)^{2}\right)$ we can make the sum in $t$ a convergent Bertrand Series.

## G. 4 The Exponential Distribution is Not Balanced

For $j=1$, a direct calculation yields

$$
\alpha(M, 1)=\frac{1}{1+\left(\frac{\mu_{1}}{\mu_{2}}\right) M}
$$

Using this, we now prove that the series in Assumption 2. of Theorem 3.1 is in $\Omega(\log (T))$, hence the balance condition is not satisfied. As all the $\alpha_{k}(M, j)$ are positive, one can write

$$
\begin{aligned}
\sum_{t=1}^{T} \sum_{j=1}^{\left\lfloor(\log t)^{2}\right\rfloor} \alpha_{k}\left(\left\lfloor\beta t /(\log t)^{2}\right\rfloor, j\right) & \geq \sum_{t=1}^{T} \alpha_{k}\left(\left\lfloor\beta t /(\log t)^{2}\right\rfloor, 1\right) \\
& =\sum_{t=2}^{T} \frac{1}{1+\left(\frac{\mu_{1}}{\mu_{k}}\right)\left\lfloor\beta t /(\log t)^{2}\right\rfloor} \\
& \geq \sum_{t=2}^{T} \frac{1}{1+\left(\frac{\mu_{1}}{\mu_{k}}\right) \beta t /(\log t)^{2}} \\
& \geq C \sum_{t=2}^{T} \frac{1}{t}=\mathcal{O}(\log (T)),
\end{aligned}
$$

where $C$ is some small enough constant that depend on $\mu_{1}, \mu_{k}$ and $\beta$.

## H Sketch of Proof with Forced Exploration

In this section, we explain how the proof of Theorem 3.1 is modified when we add forced exploration with $f_{r}=\sqrt{\log r}$, that is when in every round $r+1$ we add to $\mathcal{A}_{r+1}$ every arm $k$ such that $N_{k}(r) \leq f_{r}$. It is easy to verify that the proof of Lemma 4.1 remains unchanged, as it is inspired by the analysis of SSMC which also uses forced exploration. We now explain how forced exploration modifies the proof of Lemma 4.3 and how we upper bound the resulting new terms for any exponential family.

## H. 1 Handling Forced Exploration in Lemma 4.3

The idea is to use the same proof sketch as without forced exploration. We note $f(r)=\sqrt{\log r}$ the forced exploration rate and $f^{-1}(r)=\exp \left(r^{2}\right)$ its inverse function.
Let us consider the round $a_{r}=f^{-1}(f(r)-1)$. At this round, the value of exploration function is $f(r)-1=\sqrt{\log r}-1$, which means that the number of pulls of arm 1 is at least $\lfloor\sqrt{\log r}-1\rfloor \geq$ $\sqrt{\log r}-2$.

Now we look at the length of the interval $r-a_{r}$ :

$$
\begin{aligned}
r-a_{r} & =r-f^{-1}(f(r)-1) \\
& =r-\exp \left((\sqrt{\log r}-1)^{2}\right) \\
& =r-\exp (\log r-2 \sqrt{\log r}+1) \\
& =r(1-\exp (-2 \sqrt{\log r}+1)) \\
& \sim r \text { when } r \rightarrow+\infty
\end{aligned}
$$

As $r-a_{r}$ is equivalent to $r$ when $r$ is large, for any constant $\gamma>0$ there exists a round $r_{\gamma}$ such that for $r>r_{\gamma}: r-a_{r}>\gamma r$. This means that after the round $a_{r}$ arm 1 faces a linear amount of duels, and has an history of at least $j=\lfloor\sqrt{\log r}-1\rfloor$ samples. Introducing $b_{r}$ the random variable giving the first time when $N_{1}\left(b_{r}\right)=\lfloor\sqrt{\log r}-1\rfloor$, we necessarily have $b_{r} \leq a_{r}$. We now use that $N_{1}(r) \leq(\log r)^{2} \Rightarrow \sum_{j=1}^{\left\lfloor(\log r)^{2}\right\rfloor-1} \tau_{j} \leq r$, which further implies

$$
b_{r}+\sum_{j=\sqrt{\log r}-1}^{\left\lfloor(\log r)^{2}\right\rfloor-1} \tau_{j} \leq r \Rightarrow \sum_{j=\sqrt{\log r}-1}^{\left\lfloor(\log r)^{2}\right\rfloor-1} \tau_{j} \leq r-b_{r}
$$

We can then use the same proof as in Appendix E

$$
\begin{aligned}
\mathbb{P}\left(N_{1}(r) \leq(\log r)^{2}\right) & \left.\leq \mathbb{P}\left(\exists j \in\left\{\lfloor\sqrt{\log r}-1\rfloor, \ldots,\left\lfloor(\log r)^{2}\right\rfloor-1\right\}: \tau_{j} \geq \frac{r-b_{r}}{(\log r)^{2}-\lfloor\sqrt{\log r}-1\rfloor}\right\}\right) \\
& \left.\leq \mathbb{P}\left(\exists j \in\left\{\lfloor\sqrt{\log r}-1\rfloor, \ldots,\left\lfloor(\log r)^{2}\right\rfloor-1\right\}: \tau_{j} \geq \frac{r-a_{r}}{(\log r)^{2}-\lfloor\sqrt{\log r}-1\rfloor}\right\}\right) \\
& \left.\leq \mathbb{P}\left(\exists j \in\left\{\lfloor\sqrt{\log r}-1\rfloor, \ldots,\left\lfloor(\log r)^{2}\right\rfloor-1\right\}: \tau_{j} \geq \frac{r-a_{r}}{(\log r)^{2}}\right\}\right) \\
& \left.\leq \mathbb{P}\left(\exists j \in\left\{\lfloor\sqrt{\log r}-1\rfloor, \ldots,\left\lfloor(\log r)^{2}\right\rfloor-1\right\}: \tau_{j} \geq \frac{\gamma r}{(\log r)^{2}}\right\}\right)
\end{aligned}
$$

The constant $\gamma$ does not change the sketch of proof, and we finally have:

$$
\begin{equation*}
\left.\sum_{r=1}^{T} \mathbb{P}\left(N_{1}(r)\right) \leq(\log r)^{2}\right) \leq r_{K}^{\prime}+\sum_{r=r_{K}^{\prime}}^{T} \sum_{j=\left\lfloor f_{r}\right\rfloor-1}^{(\log r)^{2}}\left[(K-1) \mathbb{P}\left(X_{c_{r}, c_{r}, j}<M_{r, j}\right)+\sum_{k=2}^{K} \alpha_{k}\left(M_{r, j}, j\right)\right] \tag{11}
\end{equation*}
$$

for any sequence $M_{r, j}$, and a new constant $r_{K}^{\prime}$. Observe that the sum in $j$ does not start in 1 as it does in the statement of Lemma 4.3 in the absence of forced exploration. This justifies the introduction of the generalized balance condition in Appendix H. 2

## H. 2 Exponential Families Satisfy a Generalized Balanced Condition

To conclude the proof as in Theorem 3.1, as Random Block Sampling satisfies the Diversity Property, from (11) (with the choice $M_{r, j}=\left\lfloor\beta r /(\log r)^{2}\right\rfloor$ ) it is sufficient to prove that one-dimensional exponential families satisfy the following generalized balance condition.
Definition (generalized balance condition). If SDA is defined with a forced exploration rate $f_{r}$ then the generalized balance condition for the rate $f_{r}$ is:

$$
\forall \beta \in(0,1), \quad \sum_{r=1}^{T} \sum_{j=\mathbf{f}_{\mathbf{r}}}^{\left\lfloor(\log r)^{2}\right\rfloor} \alpha_{k}\left(\left\lfloor\beta r /(\log r)^{2}\right\rfloor, j\right)=o(\log T) .
$$

The following lemma proves that this holds in particular for the choice $f_{r}=\sqrt{\log (r)}$, which permits to prove that RB-SDA with this forced exploration sequence is asymptotically optimal for distributions that belong to any one-dimensional exponential family.
Lemma H. 1 (Generalized balance condition on exponential families). If the exploration rate $f_{r}$ satisfies $\frac{f_{r}}{\log \log r} \rightarrow+\infty$ then any exponential family of distributions with one parameter satisfies the generalized balance condition.

Proof. A distribution that belong to a one-dimensional exponential family has a density $f_{\theta}(y)=$ $f(x, 0) e^{\eta(\theta) y-\psi(\theta)}$ for some natural parameter $\theta \in \mathbb{R}$.
We observe that for any $y_{1}, \ldots, y_{j} \in \mathbb{R}^{j}$, if $\sum_{i=1}^{j} y_{i} \leq \mu_{k}$ :

$$
\prod_{u=1}^{j} f_{\theta_{1}}\left(y_{u}\right)=\prod_{u=1}^{j} e^{\left(\eta\left(\theta_{1}\right)-\eta\left(\theta_{k}\right)\right) y_{u}-\left(\psi\left(\theta_{1}\right)-\psi\left(\theta_{k}\right)\right)} f_{\theta_{k}}\left(y_{u}\right) \leq e^{-j I_{1}\left(\mu_{k}\right)} \prod_{u=1}^{j} f_{\theta_{k}}\left(y_{u}\right)
$$

This inequality ensures that for all $x, u \in \mathbb{R}$, if $F_{k, j}^{-1}(u) \leq \mu_{k}$ :

$$
F_{1, j}(x) \leq e^{-j I_{1}\left(\mu_{k}\right)} F_{k, j}(x) \Rightarrow F_{1, j}\left(F_{k, j}^{-1}(u)\right) \leq e^{-j I_{1}\left(\mu_{k}\right)} u
$$

So for exponential families a strictly positive gap between two distributions leads to an exponential decrease of the ratio of the CDF of the sum. If we use the fact that for all $u \in \mathbb{R}$ :

$$
\begin{aligned}
\alpha_{k}(M, j) & =\int_{-\infty}^{+\infty} f_{1, j}(x) G_{2, j}(x)^{M} d \mathbb{P}(x) \\
& \leq \int_{-\infty}^{u} f_{1, j}(x) G_{2, j}(x)^{M}+\int_{u}^{+\infty} f_{1, j}(x) G_{2, j}(x)^{M} d \mathbb{P}(x) \\
& \leq F_{1, j}(u)+G_{2, j}(u)^{M}
\end{aligned}
$$

Then $\forall \beta \in(0,1)$ and for all sequence $u_{r}$ :

$$
\begin{aligned}
\sum_{r=1}^{T} \sum_{j=\mathbf{f}_{\mathbf{r}}}^{\left\lfloor(\log r)^{2}\right\rfloor} \alpha_{k}\left(\left\lfloor\beta r /(\log r)^{2}\right\rfloor, j\right) & \leq \sum_{r=1}^{T} \sum_{j=\mathbf{f}_{\mathbf{r}}}^{\left\lfloor(\log r)^{2}\right\rfloor}\left(1-u_{r}\right)^{\left\lfloor\beta r /(\log r)^{2}\right\rfloor}+e^{-j I_{1}\left(\mu_{k}\right)} u_{r} \\
& \leq \sum_{r=1}^{T} \log (r)^{2}\left(1-u_{r}\right)^{\left\lfloor\beta r /(\log r)^{2}\right\rfloor}+\sum_{r=1}^{T} \frac{e^{-f_{r} I_{1}\left(\mu_{k}\right)}}{1-e^{-I_{1}\left(\mu_{k}\right)}} u_{r}
\end{aligned}
$$

We now choose $u_{r}$ of the form $u_{r}=\frac{(\log r)^{k}}{r}$. Indeed, for the first term we get:

$$
\begin{aligned}
\log (r)^{2}\left(1-u_{r}\right)^{\left\lfloor\beta r /(\log r)^{2}\right\rfloor} & \leq \log (r)^{2} \exp \left(-\left\lfloor\beta r /(\log r)^{2}\right\rfloor \frac{(\log r)^{k}}{r}\right) \\
& \leq \log (r)^{2} \exp \left(-\left(\beta r /(\log r)^{2}-1\right) \frac{(\log r)^{k}}{r}\right) \\
& \leq \zeta_{k}(\log r)^{2} \exp \left(-\beta(\log r)^{k-2}\right) \\
& =o\left(r^{-1}\right) \text { for } k>2
\end{aligned}
$$

where $\zeta_{k}$ is an upper bound for $\exp \left(\frac{(\log r)^{k}}{r}\right)$. From now on we work with $k=3$. For the second term we have to study $u_{r} e^{-I_{1}\left(\mu_{k}\right) f_{r}}$ :

$$
\begin{aligned}
u_{r} e^{-I_{1}\left(\mu_{k}\right) f_{r}} & =\exp \left(\log u_{r}-I_{1}\left(\mu_{k}\right) f_{r}\right) \\
& =\exp \left(3 \log \log r-\log r-I_{1}\left(\mu_{k}\right) f_{r}\right)
\end{aligned}
$$

We see that $u_{r} e^{-I_{1}\left(\mu_{k}\right) f_{r}}=o\left(r^{-1}\right)$ if $3 \log \log r-I_{1}\left(\mu_{k}\right) f_{r} \rightarrow 0$. This condition is achieved if $f_{r} / \log \log r \rightarrow+\infty$, hence an exploration rate satisfying this condition ensures the generalized balance condition for any exponential family of distributions with one parameter for this rate.

We point out the fact that this forced exploration is not necessary in SDA, as we proved that some distributions (Bernoulli, Gaussian, Poisson) directly satisfy the balance condition defined in Assumption 2. of Theorem 3.1. We leave for future research an in-depth analysis of the properties of different families of distribution that could exhibit general conditions for the use of forced exploration (or not) in the SDA family.


[^0]:    ${ }^{1}$ as can be seen in the analysis of SSMC [2], treating the general case only requires some additional notation.

