



HAL
open science

A Finite-Time Analysis of Multi-armed Bandits Problems with Kullback-Leibler Divergences

Odalric-Ambrym Maillard, Rémi Munos, Gilles Stoltz

► **To cite this version:**

Odalric-Ambrym Maillard, Rémi Munos, Gilles Stoltz. A Finite-Time Analysis of Multi-armed Bandits Problems with Kullback-Leibler Divergences. 24th Annual Conference on Learning Theory : COLT'11, Jul 2011, Budapest, Hungary. pp.18. inria-00574987v2

HAL Id: inria-00574987

<https://hal.inria.fr/inria-00574987v2>

Submitted on 27 May 2011

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

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

A Finite-Time Analysis of Multi-armed Bandits Problems with Kullback-Leibler Divergences

Odalric-Ambrym Maillard
INRIA Lille Nord-Europe
France
odalric.maillard@inria.fr

Rémi Munos
INRIA Lille Nord-Europe
France
remi.munos@inria.fr

Gilles Stoltz
Ecole normale supérieure*, Paris
& HEC Paris
France
gilles.stoltz@ens.fr

Abstract

We consider a Kullback-Leibler-based algorithm for the stochastic multi-armed bandit problem in the case of distributions with finite supports (not necessarily known beforehand), whose asymptotic regret matches the lower bound of Burnetas and Katehakis (1996). Our contribution is to provide a finite-time analysis of this algorithm; we get bounds whose main terms are smaller than the ones of previously known algorithms with finite-time analyses (like UCB-type algorithms).

1 Introduction

The *stochastic* multi-armed bandit problem, introduced by Robbins (1952), formalizes the problem of decision-making under uncertainty, and illustrates the fundamental tradeoff that appears between *exploration*, i.e., making decisions in order to improve the knowledge of the environment, and *exploitation*, i.e., maximizing the payoff.

Setting. In this paper, we consider a multi-armed bandit problem with *finitely* many arms indexed by \mathcal{A} , for which each arm $a \in \mathcal{A}$ is associated with an unknown and fixed probability distribution ν_a over $[0, 1]$. The game is *sequential* and goes as follows: at each round $t \geq 1$, the player first picks an arm $A_t \in \mathcal{A}$ and then receives a stochastic payoff Y_t drawn at random according to ν_{A_t} . He only gets to see the payoff Y_t .

For each arm $a \in \mathcal{A}$, we denote by μ_a the expectation of its associated distribution ν_a and we let a^* be any optimal arm, i.e., $a^* \in \operatorname{argmax}_{a \in \mathcal{A}} \mu_a$.

We write μ^* as a short-hand notation for the largest expectation μ_{a^*} and denote the gap of the expected payoff μ_a of an arm $a \in \mathcal{A}$ to μ^* as $\Delta_a = \mu^* - \mu_a$. In addition, the number of times each arm $a \in \mathcal{A}$ is pulled between the rounds 1 and T is referred to as $N_T(a)$,

$$N_T(a) \stackrel{\text{def}}{=} \sum_{t=1}^T \mathbb{I}_{\{A_t=a\}}.$$

The quality of a strategy will be evaluated through the standard notion of *expected regret*, which we recall now. The expected regret (or simply regret) at round $T \geq 1$ is defined as

$$R_T \stackrel{\text{def}}{=} \mathbb{E} \left[T\mu^* - \sum_{t=1}^T Y_t \right] = \mathbb{E} \left[T\mu^* - \sum_{t=1}^T \mu_{A_t} \right] = \sum_{a \in \mathcal{A}} \Delta_a \mathbb{E} [N_T(a)], \quad (1)$$

where we used the tower rule for the first equality. Note that the expectation is with respect to the random draws of the Y_t according to the ν_{A_t} and also to the possible auxiliary randomizations that the decision-making strategy is resorting to.

The regret measures the cumulative loss resulting from pulling sub-optimal arms, and thus quantifies the amount of exploration required by an algorithm in order to find a best arm, since, as (1) indicates, the regret scales with the expected number of pulls of sub-optimal arms. Since the formulation of the problem by Robbins (1952) the regret has been a popular criterion for assessing the quality of a strategy.

*CNRS – Ecole normale supérieure, Paris – INRIA, within the project-team CLASSIC

Known lower bounds. Lai and Robbins (1985) showed that for some (one-dimensional) parametric classes of distributions, any consistent strategy (i.e., any strategy not pulling sub-optimal arms more than in a polynomial number of rounds) will despite all asymptotically pull in expectation any sub-optimal arm a at least

$$\mathbb{E}[N_T(a)] \geq \left(\frac{1}{\mathcal{K}(\nu_a, \nu^*)} + o(1) \right) \log(T)$$

times, where $\mathcal{K}(\nu_a, \nu^*)$ is the Kullback-Leibler (KL) divergence between ν_a and ν^* ; it measures how close distributions ν_a and ν^* are from a theoretical information perspective.

Later, Burnetas and Katehakis (1996) extended this result to some classes of multi-dimensional parametric distributions and proved the following generic lower bound: for a given family \mathcal{P} of possible distributions over the arms,

$$\mathbb{E}[N_T(a)] \geq \left(\frac{1}{\mathcal{K}_{\text{inf}}(\nu_a, \mu^*)} + o(1) \right) \log(T), \quad \text{where } \mathcal{K}_{\text{inf}}(\nu_a, \mu^*) \stackrel{\text{def}}{=} \inf_{\nu \in \mathcal{P}: E(\nu) > \mu^*} \mathcal{K}(\nu_a, \nu),$$

with the notation $E(\nu)$ for the expectation of a distribution ν . The intuition behind this improvement is to be related to the goal that we want to achieve in bandit problems; it is not detecting whether a distribution is optimal or not (for this goal, the relevant quantity would be $\mathcal{K}(\nu_a, \nu^*)$), but rather achieving the optimal rate of reward μ^* (i.e., one needs to measure how close ν_a is to any distribution $\nu \in \mathcal{P}$ whose expectation is at least μ^*).

Known upper bounds. Lai and Robbins (1985) provided an algorithm based on the KL divergence, which has been extended by Burnetas and Katehakis (1996) to an algorithm based on \mathcal{K}_{inf} ; it is asymptotically optimal since the number of pulls of any sub-optimal arm a satisfies

$$\mathbb{E}[N_T(a)] \leq \left(\frac{1}{\mathcal{K}_{\text{inf}}(\nu_a, \mu^*)} + o(1) \right) \log(T).$$

This result holds for finite-dimensional parametric distributions under some assumptions, e.g., the distributions having a finite and known support or belonging to a set of Gaussian distributions with known variance. Recently Honda and Takemura (2010a) extended this asymptotic result to the case of distributions \mathcal{P} with support in $[0, 1]$ and such that $\mu^* < 1$; the key ingredient in this case is that $\mathcal{K}_{\text{inf}}(\nu_a, \mu^*)$ is equal to $\mathcal{K}_{\text{min}}(\nu_a, \mu^*) \stackrel{\text{def}}{=} \inf_{\nu \in \mathcal{P}: E(\nu) \geq \mu^*} \mathcal{K}(\nu_a, \nu)$.

Motivation. All the results mentioned above provide asymptotic bounds only. However, any algorithm is only used for a finite number of rounds and it is thus essential to provide a finite-time analysis of its performance. Auer et al. (2002) initiated this work by providing an algorithm (UCB1) based on a Chernoff-Hoeffding bound; it pulls any sub-optimal arm, till any time T , at most $(8/\Delta_a^2) \log T + 1 + \pi^2/3$ times, in expectation. Although this yields a logarithmic regret, the multiplicative constant depends on the gap $\Delta_a^2 = (\mu^* - \mu_a)^2$ but not on $\mathcal{K}_{\text{inf}}(\nu_a, \mu^*)$, which can be seen to be larger than $\Delta_a^2/2$ by Pinsker's inequality; that is, this non-asymptotic bound does not have the right dependence in the distributions. (How much is gained of course depends on the specific families of distributions at hand.) Audibert et al. (2009) provided an algorithm (UCB-V) that takes into account the empirical variance of the arms and exhibited a strategy such that $\mathbb{E}[N_T(a)] \leq 10(\sigma_a^2/\Delta_a^2 + 2/\Delta_a) \log T$ for any time T (where σ_a^2 is the variance of arm a); it improves over UCB1 in case of arms with small variance. Other variants include the MOSS algorithm by Audibert and Bubeck (2010) and Improved UCB by Auer and Ortner (2010).

However, all these algorithms only rely on one moment (for UCB1) or two moments (for UCB-V) of the empirical distributions of the obtained rewards; they do not fully exploit the empirical distributions. As a consequence, the resulting bounds are expressed in terms of the means μ_a and variances σ_a^2 of the sub-optimal arms and not in terms of the quantity $\mathcal{K}_{\text{inf}}(\nu_a, \mu^*)$ appearing in the lower bounds. The numerical experiments reported in Filippi (2010) confirm that these algorithms are less efficient than those based on \mathcal{K}_{inf} .

Our contribution. In this paper we analyze a \mathcal{K}_{inf} -based algorithm inspired by the ones studied in Lai and Robbins (1985), Burnetas and Katehakis (1996), Filippi (2010); it indeed takes into account the full empirical distribution of the observed rewards. The analysis is performed (with explicit bounds) in the case of Bernoulli distributions over the arms. Less explicit but finite-time bounds are obtained in the case of finitely supported distributions (whose supports do not need to be known in advance). Finally, we pave the way for handling the case of general finite-dimensional parametric distributions. These results improve on the ones by Burnetas and Katehakis (1996), Honda and Takemura (2010a) since finite-time bounds (implying their asymptotic results)

are obtained; and on Auer et al. (2002), Audibert et al. (2009) as the dependency of the main term scales with $\mathcal{K}_{\text{inf}}(\nu_a, \mu^*)$. The proposed \mathcal{K}_{inf} -based algorithm is also more natural and more appealing than the one presented in Honda and Takemura (2010a).

Recent related works. Since our initial submission of the present paper, we got aware of two papers that tackle problems similar to ours. First, a revised version of Honda and Takemura (2010b, personal communication) obtains finite-time regret bounds (with prohibitively large constants) for a *randomized* (less natural) strategy in the case of distributions with finite supports (also not known in advance). Second, another paper at this conference (Garivier and Cappé, 2011) also deals with the \mathcal{K} -strategy which we study in Theorem 3; they however do not obtain second-order terms in closed forms as we do and later extend their strategy to exponential families of distributions (while we extend our strategy to the case of distributions with finite supports). On the other hand, they show how the \mathcal{K} -strategy can be extended in a straightforward manner to guarantee bounds with respect to the family of all bounded distributions on a known interval; these bounds are suboptimal but improve on the ones of UCB-type algorithms.

2 Definitions and tools

Let \mathcal{X} be a Polish space; in the next sections, we will consider $\mathcal{X} = \{0, 1\}$ or $\mathcal{X} = [0, 1]$. We denote by $\mathcal{P}(\mathcal{X})$ the set of probability distributions over \mathcal{X} and equip $\mathcal{P}(\mathcal{X})$ with the distance d induced by the norm $\|\cdot\|$ defined by $\|\nu\| = \sup_{f \in \mathcal{L}} |\int_{\mathcal{X}} f d\nu|$, where \mathcal{L} is the set of Lipschitz functions over \mathcal{X} , taking values in $[-1, 1]$ and with Lipschitz constant smaller than 1.

Kullback-Leibler divergence: For two elements $\nu, \kappa \in \mathcal{P}(\mathcal{X})$, we write $\nu \ll \kappa$ when ν is absolutely continuous with respect to κ and denote in this case by $d\nu/d\kappa$ the density of ν with respect to κ . We recall that the Kullback-Leibler divergence between ν and κ is defined as

$$\mathcal{K}(\nu, \kappa) = \int_{[0,1]} \frac{d\nu}{d\kappa} \log \frac{d\nu}{d\kappa} d\kappa \quad \text{if } \nu \ll \kappa; \quad \text{and } \mathcal{K}(\nu, \kappa) = +\infty \quad \text{otherwise.} \quad (2)$$

Empirical distribution: We consider a sequence X_1, X_2, \dots of random variables taking values in \mathcal{X} , independent and identically distributed according to a distribution ν . For all integers $t \geq 1$, we denote the empirical distribution corresponding to the first t elements of the sequence by

$$\hat{\nu}_t = \frac{1}{t} \sum_{s=1}^t \delta_{X_s}.$$

Non-asymptotic Sanov's Lemma: The following lemma follows from a straightforward adaptation of Dinwoodie (1992, Theorem 2.1 and comments on page 372). Details of the proof are provided in the appendix.

Lemma 1 *Let \mathcal{C} be an open convex subset of $\mathcal{P}(\mathcal{X})$ such that $\Lambda(\mathcal{C}) = \inf_{\kappa \in \mathcal{C}} \mathcal{K}(\kappa, \nu) < \infty$.*

Then, for all $t \geq 1$, one has $\mathbb{P}_{\nu}\{\hat{\nu}_t \in \bar{\mathcal{C}}\} \leq e^{-t\Lambda(\bar{\mathcal{C}})}$ where $\bar{\mathcal{C}}$ is the closure of \mathcal{C} .

This lemma should be thought of as a deviation inequality. The empirical distribution converges (in distribution) to ν . Now, if (and only if) ν is not in the closure of \mathcal{C} , then $\Lambda(\mathcal{C}) > 0$ and the lemma indicates how unlikely it is that $\hat{\nu}_t$ is in this set $\bar{\mathcal{C}}$ not containing the limit ν . The probability of interest decreases at a geometric rate, which depends on $\Lambda(\mathcal{C})$.

3 Finite-time analysis for Bernoulli distributions

In this section, we start with the case of Bernoulli distributions. Although this case is a special case of the general results of Section 4, we provide here a complete and self-contained analysis of this case, where, in addition, we are able to provide closed forms for all the terms in the regret bound. Note however that the resulting bound is slightly worse than what could be derived from the general case (for which more sophisticated tools are used). This result is mainly provided as a warm-up.

3.1 Reminder of some useful results for Bernoulli distributions

We denote by \mathcal{B} the subset of $\mathcal{P}([0, 1])$ formed by the Bernoulli distributions; it corresponds to $\mathcal{B} = \mathcal{P}(\{0, 1\})$. A generic element of \mathcal{B} will be denoted by $\beta(p)$, where $p \in [0, 1]$ is the probability mass put on 1. We consider a sequence X_1, X_2, \dots of independent and identically distributed random

Parameters: A non-decreasing function $f : \mathbb{N} \rightarrow \mathbb{R}$

Initialization: Pull each arm of \mathcal{A} once

For rounds $t + 1$, where $t \geq |\mathcal{A}|$,

- compute for each arm $a \in \mathcal{A}$ the quantity

$$B_{a,t}^+ = \max \left\{ q \in [0, 1] : N_t(a) \mathcal{K}(\beta(\hat{\mu}_{a,N_t(a)}), \beta(q)) \leq f(t) \right\},$$

where
$$\hat{\mu}_{a,N_t(a)} = \frac{1}{N_t(a)} \sum_{s \leq t: A_s = a} Y_s;$$

- in case of a tie, pick an arm with largest value of $\hat{\mu}_{a,N_t(a)}$;
 - pull any arm $A_{t+1} \in \operatorname{argmax}_{a \in \mathcal{A}} B_{a,t}^+$.
-

Figure 1: The \mathcal{K} -strategy.

variables, with common distribution $\beta(p)$; for the sake of clarity we will index, in this subsection only, all probabilities and expectations with p .

For all integers $t \geq 1$, we denote by $\hat{p}_t = \frac{1}{t} \sum_{s=1}^t X_s$ the empirical average of the first t elements of the sequence.

The lemma below follows from an adaptation of Garivier and Leonardi (2010, Proposition 2). The details of the adaptation (and simplification) can be found in the appendix.

Lemma 2 *For all $p \in [0, 1]$, all $\varepsilon > 1$, and all $t \geq 1$,*

$$\mathbb{P}_p \left(\bigcup_{s=1}^t \left\{ s \mathcal{K}(\beta(\hat{p}_s), \beta(p)) \geq \varepsilon \right\} \right) \leq 2e \lceil \varepsilon \log t \rceil e^{-\varepsilon}.$$

In particular, for all random variables N_t taking values in $\{1, \dots, t\}$,

$$\mathbb{P}_p \left\{ N_t \mathcal{K}(\beta(\hat{p}_{N_t}), \beta(p)) \geq \varepsilon \right\} \leq 2e \lceil \varepsilon \log t \rceil e^{-\varepsilon}.$$

Another immediate fact about Bernoulli distributions is that for all $p \in (0, 1)$, the mappings $\mathcal{K}_{.,p} : q \in (0, 1) \mapsto \mathcal{K}(\beta(p), \beta(q))$ and $\mathcal{K}_{p,.} : q \in [0, 1] \mapsto \mathcal{K}(\beta(q), \beta(p))$ are continuous and take finite values. In particular, we have, for instance, that for all $\varepsilon > 0$ and $p \in (0, 1)$, the set

$$\left\{ q \in [0, 1] : \mathcal{K}(\beta(p), \beta(q)) \leq \varepsilon \right\}$$

is a closed interval containing p . This property still holds when $p \in \{0, 1\}$, as in this case, the interval is reduced to $\{p\}$.

3.2 Strategy and analysis

We consider the so-called \mathcal{K} -strategy of Figure 1, which was already considered in the literature, see Burnetas and Katehakis (1996), Filippi (2010). The numerical computation of the quantities $B_{a,t}^+$ is straightforward (by convexity of \mathcal{K} in its second argument, by using iterative methods) and is detailed therein.

Before proceeding, we denote by $\sigma_a^2 = \mu_a(1 - \mu_a)$ the variance of each arm $a \in \mathcal{A}$ (and take the short-hand notation $\sigma^{*,2}$ for the variance of an optimal arm).

Theorem 3 *When $\mu^* \in (0, 1)$, for all non-decreasing functions $f : \mathbb{N} \rightarrow \mathbb{R}_+$ such that $f(1) \geq 1$, the expected regret R_T of the strategy of Figure 1 is upper bounded by the infimum, as the $(c_a)_{a \in \mathcal{A}}$ describe $(0, +\infty)$, of the quantities*

$$\sum_{a \in \mathcal{A}} \Delta_a \left(\frac{(1 + c_a) f(T)}{\mathcal{K}(\beta(\mu_a), \beta(\mu^*))} + 4e \sum_{t=|\mathcal{A}|}^{T-1} \lceil f(t) \log t \rceil e^{-f(t)} + \frac{(1 + c_a)^2}{8 c_a^2 \Delta_a^2 \min\{\sigma_a^4, \sigma^{*,4}\}} \mathbb{I}_{\{\mu_a \in (0,1)\}} + 3 \right).$$

For $\mu^ = 0$, its regret is null. For $\mu^* = 1$, it satisfies $R_T \leq 2(|\mathcal{A}| - 1)$.*

A possible choice for the function f is $f(t) = \log((et) \log^3(et))$, which is non decreasing, satisfies $f(1) \geq 1$, and is such that the second term in the sum above is bounded (by a basic result about so-called Bertrand's series). Now, as the constants c_a in the bound are parameters of the analysis (and not of the strategy), they can be optimized. For instance, with the choice of $f(t)$ mentioned above, taking each c_a proportional to $(\log T)^{-1/3}$ (up to a multiplicative constant that depends on the distributions ν_a) entails the regret bound

$$\sum_{a \in \mathcal{A}} \Delta_a \frac{\log T}{\mathcal{K}(\beta(\mu_a), \beta(\mu^*))} + \varepsilon_T,$$

where it is easy to give an explicit and closed-form expression of ε_T ; in this conference version, we only indicate that ε_T is of order of $(\log T)^{2/3}$ but we do not know whether the order of magnitude of this second-order term is optimal.

Proof: We first deal with the case where $\mu^* \notin \{0, 1\}$ and introduce an additional notation. In view of the remark at the end of Section 3.1, for all arms a and rounds t , we let $B_{a,t}^-$ be the element in $[0, 1]$ such that

$$\left\{ q \in [0, 1] : N_t(a) \mathcal{K}(\beta(\widehat{\mu}_{a, N_t(a)}), \beta(q)) \leq f(t) \right\} = [B_{a,t}^-, B_{a,t}^+]. \quad (3)$$

As (1) indicates, it suffices to bound $N_T(a)$ for all suboptimal arms a , i.e., for all arms such that $\mu_a < \mu^*$. We will assume in addition that $\mu_a > 0$ (and we also have $\mu_a \leq \mu^* < 1$); the case where $\mu_a = 0$ will be handled separately.

Step 1: A decomposition of the events of interest. For $t \geq |\mathcal{A}|$, when $A_{t+1} = a$, we have in particular, by definition of the strategy, that $B_{a,t}^+ \geq B_{a^*,t}^+$. On the event

$$\{A_{t+1} = a\} \cap \left\{ \mu^* \in [B_{a^*,t}^-, B_{a^*,t}^+] \right\} \cap \left\{ \mu_a \in [B_{a,t}^-, B_{a,t}^+] \right\},$$

we therefore have, on the one hand, $\mu^* \leq B_{a^*,t}^+ \leq B_{a,t}^+$ and on the other hand, $B_{a,t}^- \leq \mu_a \leq \mu^*$, that is, the considered event is included in $\left\{ \mu^* \in [B_{a,t}^-, B_{a,t}^+] \right\}$. We thus proved that

$$\{A_{t+1} = a\} \subseteq \left\{ \mu^* \notin [B_{a^*,t}^-, B_{a^*,t}^+] \right\} \cup \left\{ \mu_a \notin [B_{a,t}^-, B_{a,t}^+] \right\} \cup \left\{ \mu^* \in [B_{a,t}^-, B_{a,t}^+] \right\}.$$

Going back to the definition (3), we get in particular the inclusion

$$\begin{aligned} \{A_{t+1} = a\} \subseteq & \left\{ N_t(a^*) \mathcal{K}(\beta(\widehat{\mu}_{a^*, N_t(a^*)}), \beta(\mu^*)) > f(t) \right\} \\ & \cup \left\{ N_t(a) \mathcal{K}(\beta(\widehat{\mu}_{a, N_t(a)}), \beta(\mu_a)) > f(t) \right\} \\ & \cup \left(\left\{ N_t(a) \mathcal{K}(\beta(\widehat{\mu}_{a, N_t(a)}), \beta(\mu^*)) \leq f(t) \right\} \cap \{A_{t+1} = a\} \right). \end{aligned}$$

Step 2: Bounding the probabilities of two elements of the decomposition. We consider the filtration (\mathcal{F}_t) , where for all $t \geq 1$, the σ -algebra \mathcal{F}_t is generated by $A_1, Y_1, \dots, A_t, Y_t$. In particular, A_{t+1} and thus all $N_{t+1}(a)$ are \mathcal{F}_t -measurable. We denote by $\tau_{a,1}$ the deterministic round at which a was pulled for the first time and by $\tau_{a,2}, \tau_{a,3}, \dots$ the rounds $t \geq |\mathcal{A}| + 1$ at which a was then played; since for all $k \geq 2$,

$$\tau_{a,k} = \min\{t \geq |\mathcal{A}| + 1 : N_t(a) = k\},$$

we see that $\{\tau_{a,k} = t\}$ is \mathcal{F}_{t-1} -measurable. Therefore, for each $k \geq 1$, the random variable $\tau_{a,k}$ is a (predictable) stopping time. Hence, by a well-known fact in probability theory (see, e.g., Chow and Teicher 1988, Section 5.3), the random variables $\widetilde{X}_{a,k} = Y_{\tau_{a,k}}$, where $k = 1, 2, \dots$ are independent and identically distributed according to ν_a . Since on $\{N_t(a) = k\}$, we have the rewriting

$$\widehat{\mu}_{a, N_t(a)} = \widetilde{\mu}_{a,k} \quad \text{where} \quad \widetilde{\mu}_{a,k} = \frac{1}{k} \sum_{j=1}^k \widetilde{X}_{a,j},$$

and since for $t \geq |\mathcal{A}| + 1$, one has $N_t(a) \geq 1$ with probability 1, we can apply the second statement in Lemma 2 and get, for all $t \geq |\mathcal{A}| + 1$,

$$\mathbb{P}\left\{N_t(a) \mathcal{K}\left(\beta(\widehat{\mu}_{a,N_t(a)}), \beta(\mu_a)\right) > f(t)\right\} \leq 2e \lceil f(t) \log t \rceil e^{-f(t)}.$$

A similar argument shows that for all $t \geq |\mathcal{A}| + 1$,

$$\mathbb{P}\left\{N_t(a^*) \mathcal{K}\left(\beta(\widehat{\mu}_{a^*,N_t(a^*)}), \beta(\mu^*)\right) > f(t)\right\} \leq 2e \lceil f(t) \log t \rceil e^{-f(t)}.$$

Step 3: Rewriting the remaining terms. We therefore proved that

$$\mathbb{E}[N_T(a)] \leq 1 + 4e \sum_{t=|\mathcal{A}|}^{T-1} \lceil f(t) \log t \rceil e^{-f(t)} + \sum_{t=|\mathcal{A}|}^{T-1} \mathbb{P}\left(\left\{N_t(a) \mathcal{K}\left(\beta(\widehat{\mu}_{a,N_t(a)}), \beta(\mu^*)\right) \leq f(t)\right\} \cap \{A_{t+1} = a\}\right)$$

and deal now with the last sum. Since f is non decreasing, it is bounded by

$$\sum_{t=|\mathcal{A}|}^{T-1} \mathbb{P}\left(K_t \cap \{A_{t+1} = a\}\right) \quad \text{where} \quad K_t = \left\{N_t(a) \mathcal{K}\left(\beta(\widehat{\mu}_{a,N_t(a)}), \beta(\mu^*)\right) \leq f(T)\right\}.$$

$$\text{Now,} \quad \sum_{t=|\mathcal{A}|}^{T-1} \mathbb{P}\left(K_t \cap \{A_{t+1} = a\}\right) = \mathbb{E}\left[\sum_{t=|\mathcal{A}|}^{T-1} \mathbb{I}_{\{A_{t+1}=a\}} \mathbb{I}_{K_t}\right] = \mathbb{E}\left[\sum_{k \geq 2} \mathbb{I}_{\{\tau_{a,k} \leq T\}} \mathbb{I}_{K_{\tau_{a,k}-1}}\right].$$

We note that, since $N_{\tau_{a,k}-1}(a) = k - 1$, we have that

$$K_{\tau_{a,k}-1} = \left\{(k-1) \mathcal{K}\left(\beta(\widetilde{\mu}_{a,k-1}), \beta(\mu^*)\right) \leq f(T)\right\}.$$

All in all, since $\tau_{a,k} \leq T$ implies $k \leq T - |\mathcal{A}| + 1$ (as each arm is played at least once during the first $|\mathcal{A}|$ rounds), we have

$$\mathbb{E}\left[\sum_{k \geq 2} \mathbb{I}_{\{\tau_{a,k} \leq T\}} \mathbb{I}_{K_{\tau_{a,k}-1}}\right] \leq \mathbb{E}\left[\sum_{k=2}^{T-|\mathcal{A}|+1} \mathbb{I}_{K_{\tau_{a,k}-1}}\right] = \sum_{k=2}^{T-|\mathcal{A}|+1} \mathbb{P}\left\{(k-1) \mathcal{K}\left(\beta(\widetilde{\mu}_{a,k-1}), \beta(\mu^*)\right) \leq f(T)\right\}. \quad (4)$$

Step 4: Bounding the probabilities of the latter sum via Sanov's lemma. For each $\gamma > 0$, we define the convex open set $\mathcal{C}_\gamma = \left\{\beta(q) \in \mathcal{B} : \mathcal{K}(\beta(q), \beta(\mu^*)) < \gamma\right\}$, which is a non-empty set (since $\mu^* < 1$); by continuity of the mapping $\mathcal{K}(\cdot, \mu^*)$ defined after the statement of Lemma 2 when $\mu^* \in (0, 1)$, its closure equals $\overline{\mathcal{C}}_\gamma = \left\{\beta(q) \in \mathcal{B} : \mathcal{K}(\beta(q), \beta(\mu^*)) \leq \gamma\right\}$.

In addition, since $\mu_a \in (0, 1)$, we have that $\mathcal{K}(\beta(q), \beta(\mu_a)) < \infty$ for all $q \in [0, 1]$. In particular, for all $\gamma > 0$, the condition $\Lambda(\mathcal{C}_\gamma) < \infty$ of Lemma 1 is satisfied. Denoting this value by

$$\theta_a(\gamma) = \inf\left\{\mathcal{K}(\beta(q), \beta(\mu_a)) : \beta(q) \in \mathcal{B} \text{ such that } \mathcal{K}(\beta(q), \beta(\mu^*)) \leq \gamma\right\},$$

we get by the indicated lemma that for all $k \geq 1$,

$$\mathbb{P}\left\{\mathcal{K}\left(\beta(\widetilde{\mu}_{a,k}), \beta(\mu^*)\right) \leq \gamma\right\} = \mathbb{P}\left\{\beta(\widetilde{\mu}_{a,k}) \in \overline{\mathcal{C}}_\gamma\right\} \leq e^{-k \theta_a(\gamma)}.$$

Now, since (an open neighborhood of) $\beta(\mu_a)$ is not included in $\overline{\mathcal{C}}_\gamma$ as soon as $0 < \gamma < \mathcal{K}(\beta(\mu_a), \beta(\mu^*))$, we have that $\theta_a(\gamma) > 0$ for such values of γ . To apply the obtained inequality to the last sum in (4),

we fix a constant $c_a > 0$ and denote by k_0 the following upper integer part, $k_0 = \left\lceil \frac{(1+c_a)f(T)}{\mathcal{K}(\beta(\mu_a), \beta(\mu^*))} \right\rceil$,

so that $f(T)/k \leq \mathcal{K}(\beta(\mu_a), \beta(\mu^*)) / (1+c_a) < \mathcal{K}(\beta(\mu_a), \beta(\mu^*))$ for $k \geq k_0$, hence,

$$\begin{aligned} \sum_{k=2}^{T-|\mathcal{A}|+1} \mathbb{P}\left\{(k-1) \mathcal{K}\left(\beta(\widetilde{\mu}_{a,k-1}), \beta(\mu^*)\right) \leq f(T)\right\} &\leq \sum_{k=1}^T \mathbb{P}\left\{\mathcal{K}\left(\beta(\widetilde{\mu}_{a,k}), \beta(\mu^*)\right) \leq \frac{f(T)}{k}\right\} \\ &\leq k_0 - 1 + \sum_{k=k_0}^T \exp\left(-k \theta_a(f(T)/k)\right). \end{aligned}$$

Since θ_a is a non-increasing function,

$$\begin{aligned} \sum_{k=k_0}^T \exp\left(-k \theta_a(f(T)/k)\right) &\leq \sum_{k=k_0}^T \exp\left(-k \theta_a(\mathcal{K}(\beta(\mu_a), \beta(\mu^*)) / (1 + c_a))\right) \\ &\leq \Gamma_a(c_a) \exp\left(-k_0 \theta_a(\mathcal{K}(\beta(\mu_a), \beta(\mu^*)) / (1 + c_a))\right) \leq \Gamma_a(c_a), \end{aligned}$$

where $\Gamma_a(c_a) = \left[1 - \exp\left(-\theta_a(\mathcal{K}(\beta(\mu_a), \beta(\mu^*)) / (1 + c_a))\right)\right]^{-1}$.

Putting all pieces together, we thus proved so far that

$$\mathbb{E}[N_T(a)] \leq 1 + \frac{(1 + c_a) f(T)}{\mathcal{K}(\beta(\mu_a), \beta(\mu^*))} + 4e \sum_{t=|A|}^{T-1} [f(t) \log t] e^{-f(t)} + \Gamma_a(c_a)$$

and it only remains to deal with $\Gamma_a(c_a)$.

Step 5: Getting an upper bound in closed form for $\Gamma_a(c_a)$. We will make repeated uses of Pinsker's inequality: for $p, q \in [0, 1]$, one has $\mathcal{K}(\beta(p), \beta(q)) \geq 2(p - q)^2$.

In what follows, we use the short-hand notation $\Theta_a = \theta_a(\mathcal{K}(\beta(\mu_a), \beta(\mu^*)) / (1 + c_a))$ and therefore need to upper bound $1/(1 - e^{-\Theta_a})$. Since for all $u \geq 0$, one has $e^{-u} \leq 1 - u + u^2/2$, we get $\Gamma_a(c_a) \leq \frac{1}{\Theta_a(1 - \Theta_a/2)} \leq \frac{2}{\Theta_a}$ for $\Theta_a \leq 1$, and $\Gamma_a(c_a) \leq \frac{1}{1 - e^{-1}} \leq 2$ for $\Theta_a \geq 1$. It thus only remains to lower bound Θ_a in the case when it is smaller than 1.

By the continuity properties of the Kullback-Leibler divergence, the infimum in the definition of θ_a is always achieved; we therefore let $\tilde{\mu}$ be an element in $[0, 1]$ such that

$$\Theta_a = \mathcal{K}(\beta(\tilde{\mu}), \beta(\mu_a)) \quad \text{and} \quad \mathcal{K}(\beta(\tilde{\mu}), \beta(\mu^*)) = \frac{\mathcal{K}(\beta(\mu_a), \beta(\mu^*))}{1 + c_a};$$

it is easy to see that we have the ordering $\mu_a < \tilde{\mu} < \mu^*$. By Pinsker's inequality, $\Theta_a \geq 2(\tilde{\mu} - \mu_a)^2$ and we now lower bound the latter quantity. We use the short-hand notation $f(p) = \mathcal{K}(\beta(p), \beta(\mu^*))$ and note that the thus defined mapping f is convex and differentiable on $(0, 1)$; its derivative equals $f'(p) = \log((1 - \mu^*)/(\mu^*)) + \log(p/(1 - p))$ for all $p \in (0, 1)$ and is therefore non positive for $p \leq \mu^*$. By the indicated convexity of f , using a sub-gradient inequality, we get $f(\tilde{\mu}) - f(\mu_a) \geq f'(\mu_a)(\tilde{\mu} - \mu_a)$, which entails, since $f'(\mu_a) < 0$,

$$\tilde{\mu} - \mu_a \geq \frac{f(\tilde{\mu}) - f(\mu_a)}{f'(\mu_a)} = \frac{c_a}{1 + c_a} \frac{f(\mu_a)}{-f'(\mu_a)}, \quad (5)$$

where the equality follows from the fact that by definition of μ , we have $f(\tilde{\mu}) = f(\mu_a)/(1 + c_a)$. Now, since f' is differentiable as well on $(0, 1)$ and takes the value 0 at μ^* , a Taylor's equality entails that there exists a $\xi \in (\mu_a, \mu^*)$ such that

$$-f'(\mu_a) = f'(\mu^*) - f'(\mu_a) = f''(\xi)(\mu^* - \mu_a) \quad \text{where} \quad f''(\xi) = 1/\xi + 1/(1 - \xi) = 1/(\xi(1 - \xi)).$$

Therefore, by convexity of $\tau \mapsto \tau(1 - \tau)$, we get that

$$\frac{1}{-f'(\mu_a)} \geq \frac{\min\{\mu_a(1 - \mu_a), \mu^*(1 - \mu^*)\}}{\mu^* - \mu_a}.$$

Substituting this into (5) and using again Pinsker's inequality to lower bound $f(\mu_a)$, we have proved

$$\tilde{\mu} - \mu_a \geq 2 \frac{c_a}{1 + c_a} (\mu^* - \mu_a) \min\{\mu_a(1 - \mu_a), \mu^*(1 - \mu^*)\}.$$

Putting all pieces together, we thus proved that

$$\Gamma_a(c_a) \leq 2 \max \left\{ \frac{(1 + c_a)^2}{8 c_a^2 (\mu^* - \mu_a)^2 \left(\min\{\mu_a(1 - \mu_a), \mu^*(1 - \mu^*)\} \right)^2}, 1 \right\};$$

bounding the maximum of the two quantities by their sum concludes the main part of the proof.

Step 6: For $\mu^* \in \{0, 1\}$ and/or $\mu_a = 0$. When $\mu^* = 1$, then $\widehat{\mu}_{a^*, N_t(a^*)} = 1$ for all $t \geq |\mathcal{A}| + 1$, so that $B_{a^*, t}^+ = 1$ for all $t \geq |\mathcal{A}| + 1$. Thus, the arm a is played after round $t \geq |\mathcal{A}| + 1$ only if $B_{a, t}^+ = 1$ and $\widehat{\mu}_{a, N_t(a)} = 1$ (in view of the tie-breaking rule of the considered strategy). But this means that a is played as long as it gets payoffs equal to 1 and is stopped being played when it receives the payoff 0 for the first time. Hence, in this case, we have that the sum of payoffs equals at least $T - 2(|\mathcal{A}| - 1)$ and the regret $R_T = \mathbb{E}[T\mu^* - (Y_1 + \dots + Y_t)]$ is therefore bounded by $2(|\mathcal{A}| - 1)$.

When $\mu^* = 0$, a Dirac mass over 0 is associated with all arms and the regret of all strategies is equal to 0.

We consider now the case $\mu^* \in (0, 1)$ and $\mu_a = 0$, for which the first three steps go through; only in the upper bound of step 4 we used the fact that $\mu_a > 0$. But in this case, we have a deterministic bound on (4). Indeed, since $\mathcal{K}(\beta(0), \beta(\mu^*)) = -\log \mu^*$, we have $k \mathcal{K}(\beta(0), \beta(\mu^*)) \leq f(T)$ if and only if

$$k \leq \frac{f(T)}{-\log \mu^*} = \frac{f(T)}{\mathcal{K}(\beta(\mu_a), \beta(\mu^*))},$$

which improves on the general bound exhibited in step 4. ■

Remark 1 Note that Step 5 in the proof is specifically designed to provide an upper bound on $\Gamma_a(c_a)$ in the case of Bernoulli distributions. In the general case, getting such an explicit bound seems more involved.

4 A finite-time analysis in the case of distributions with finite support

Before stating and proving our main result, Theorem 9, we introduce the quantity \mathcal{K}_{inf} and list some of its properties.

4.1 Some useful properties of \mathcal{K}_{inf} and its level sets

We now introduce the key quantity in order to generalize the previous algorithm to handle the case of distributions with finite support. To that end, we introduce $\mathcal{P}_F([0, 1])$, the subset of $\mathcal{P}([0, 1])$ that consists of distributions with finite support.

Definition 4 For all distributions $\nu \in \mathcal{P}_F([0, 1])$ and $\mu \in [0, 1]$, we define

$$\mathcal{K}_{\text{inf}}(\nu, \mu) = \inf \left\{ \mathcal{K}(\nu, \nu') : \nu' \in \mathcal{P}_F([0, 1]) \text{ s.t. } E(\nu') > \mu \right\},$$

where $E(\nu') = \int_{[0, 1]} x d\nu'(x)$ denotes the expectation of the distribution ν' .

We now remind some useful properties of \mathcal{K}_{inf} . Honda and Takemura (2010b, Lemma 6) can be reformulated in our context as follows.

Lemma 5 For all $\nu \in \mathcal{P}_F([0, 1])$, the mapping $\mathcal{K}_{\text{inf}}(\nu, \cdot)$ is continuous and non decreasing in its argument $\mu \in [0, 1]$. Moreover, the mapping $\mathcal{K}_{\text{inf}}(\cdot, \mu)$ is lower semi-continuous on $\mathcal{P}_F([0, 1])$ for all $\mu \in [0, 1]$.

The next two lemmas bound the variation of \mathcal{K}_{inf} , respectively in its first and second arguments. (For clarity, we denote the expectations with respect to ν by \mathbb{E}_ν .) Their proofs are both deferred to the appendix. We denote by $\|\cdot\|_1$ the ℓ^1 -norm on $\mathcal{P}([0, 1])$ and recall that the ℓ^1 -norm of $\nu - \nu'$ corresponds to twice the distance in variation between ν and ν' .

Lemma 6 For all $\mu \in (0, 1)$ and for all $\nu, \nu' \in \mathcal{P}_F([0, 1])$, the following holds true.

- In the case when $\mathbb{E}_\nu[(1 - \mu)/(1 - X)] > 1$, then $\mathcal{K}_{\text{inf}}(\nu, \mu) - \mathcal{K}_{\text{inf}}(\nu', \mu) \leq M_{\nu, \mu} \|\nu - \nu'\|_1$, for some constant $M_{\nu, \mu} > 0$.
- In the case when $\mathbb{E}_\nu[(1 - \mu)/(1 - X)] \leq 1$, the fact that $\mathcal{K}_{\text{inf}}(\nu, \mu) - \mathcal{K}_{\text{inf}}(\nu', \mu) \geq \alpha \mathcal{K}_{\text{inf}}(\nu, \mu)$ for some $\alpha \in (0, 1)$ entails that

$$\|\nu - \nu'\|_1 \geq \frac{1 - \mu}{(2/\alpha)((2/\alpha) - 1)}.$$

Lemma 7 We have that for any $\nu \in \mathcal{P}_F([0, 1])$, provided that $\mu \geq \mu - \varepsilon > E(\nu)$, the following inequalities hold true:

$$\varepsilon/(1 - \mu) \geq \mathcal{K}_{\text{inf}}(\nu, \mu) - \mathcal{K}_{\text{inf}}(\nu, \mu - \varepsilon) \geq 2\varepsilon^2$$

Moreover, the first inequality is also valid when $E(\nu) \geq \mu > \mu - \varepsilon$ or $\mu > E(\nu) \geq \mu - \varepsilon$.

Parameters: A non-decreasing function $f : \mathbb{N} \rightarrow \mathbb{R}$

Initialization: Pull each arm of \mathcal{A} once

For rounds $t + 1$, where $t \geq |\mathcal{A}|$,

- compute for each arm $a \in \mathcal{A}$ the quantity

$$B_{a,t}^+ = \max \left\{ q \in [0, 1] : N_t(a) \mathcal{K}_{\text{inf}}(\widehat{\nu}_{a, N_t(a)}, q) \leq f(t) \right\},$$

where
$$\widehat{\nu}_{a, N_t(a)} = \frac{1}{N_t(a)} \sum_{s \leq t: A_s = a} \delta_{Y_s};$$

- in case of a tie, pick an arm with largest value of $\widehat{\mu}_{a, N_t(a)}$;

- pull any arm $A_{t+1} \in \operatorname{argmax}_{a \in \mathcal{A}} B_{a,t}^+$.
-

Figure 2: The strategy \mathcal{K}_{inf} .

Level sets of \mathcal{K}_{inf} : For each $\gamma > 0$ and $\mu \in (0, 1)$, we consider the set

$$\begin{aligned} \mathcal{C}_{\mu, \gamma} &= \left\{ \nu' \in \mathcal{P}_F([0, 1]) : \mathcal{K}_{\text{inf}}(\nu', \mu) < \gamma \right\} \\ &= \left\{ \nu' \in \mathcal{P}_F([0, 1]) : \exists \nu'_\mu \in \mathcal{P}_F([0, 1]) \text{ s.t. } E(\nu'_\mu) > \mu \text{ and } \mathcal{K}(\nu', \nu'_\mu) < \gamma \right\}. \end{aligned}$$

We detail a property in the following lemma, whose proof is also deferred to the appendix.

Lemma 8 *For all $\gamma > 0$ and $\mu \in (0, 1)$, the set $\mathcal{C}_{\mu, \gamma}$ is a non-empty open convex set. Moreover,*

$$\overline{\mathcal{C}}_{\mu, \gamma} \supseteq \left\{ \nu' \in \mathcal{P}_F([0, 1]) : \mathcal{K}_{\text{inf}}(\nu', \mu) \leq \gamma \right\}.$$

4.2 The \mathcal{K}_{inf} -strategy and a general performance guarantee

For each arm $a \in \mathcal{A}$ and round t with $N_t(a) > 0$, we denote by $\widehat{\nu}_{a, N_t(a)}$ the empirical distribution of the payoffs obtained till round t when picking arm a , that is,

$$\widehat{\nu}_{a, N_t(a)} = \frac{1}{N_t(a)} \sum_{s \leq t: A_s = a} \delta_{Y_s},$$

where for all $x \in [0, 1]$, we denote by δ_x the Dirac mass on x . We define the corresponding empirical averages as

$$\widehat{\mu}_{a^*, N_t(a^*)} = E(\widehat{\nu}_{a^*, N_t(a^*)}) = \frac{1}{N_t(a^*)} \sum_{s \leq t: A_s = a^*} Y_s.$$

We then consider the \mathcal{K}_{inf} -strategy defined in Figure 2. Note that the use of maxima in the definitions of the $B_{a,t}^+$ is justified by Lemma 5.

As explained in Honda and Takemura (2010b), the computation of the quantities \mathcal{K}_{inf} can be done efficiently in this case, i.e., when we consider only distributions with finite supports. This is because in the computation of \mathcal{K}_{inf} , it is sufficient to consider only distributions with the same support as the empirical distributions (up to one point). Note that the knowledge of the support of the distributions associated with the arms is not required.

Theorem 9 *Assume that ν^* is finitely supported, with expectation $\mu^* \in (0, 1)$ and with support denoted by \mathcal{S}^* . Let $a \in \mathcal{A}$ be a suboptimal arm such that $\mu_a > 0$ and ν_a is finitely supported. Then, for all $c_a > 0$ and all*

$$0 < \varepsilon < \min \left\{ \Delta_a, \frac{c_a/2}{1 + c_a} (1 - \mu^*) \mathcal{K}_{\text{inf}}(\nu_a, \mu^*) \right\},$$

the expected number of times the \mathcal{K}_{inf} -strategy, run with $f(t) = \log t$, pulls arm a satisfies

$$\mathbb{E}[N_T(a)] \leq 1 + \frac{(1 + c_a) \log T}{\mathcal{K}_{\text{inf}}(\nu_a, \mu^*)} + \frac{1}{1 - e^{-\Theta_a(c_a, \varepsilon)}} + \frac{1}{\varepsilon^2} \log \left(\frac{1}{1 - \mu^* + \varepsilon} \right) \sum_{k=1}^T (k+1)^{|\mathcal{S}^*|} e^{-k\varepsilon^2} + \frac{1}{(\Delta_a - \varepsilon)^2},$$

where

$$\Theta_a(c_a, \varepsilon) = \theta_a \left(\frac{\log T}{k_0} + \frac{\varepsilon}{1 - \mu^*} \right) \quad \text{with} \quad k_0 = \left\lceil \frac{(1 + c_a) \log T}{\mathcal{K}_{\inf}(\nu_a, \mu^*)} \right\rceil.$$

and for all $\gamma > 0$,

$$\theta_a(\gamma) = \inf \left\{ \mathcal{K}(\nu', \nu_a) : \nu' \text{ s.t. } \mathcal{K}_{\inf}(\nu', \mu^*) < \gamma \right\}.$$

As a corollary, we get (by taking some common value for all c_a) that for all $c > 0$,

$$\bar{R}_T \leq \sum_{a \in \mathcal{A}} \Delta_a \frac{(1 + c) \log T}{\mathcal{K}_{\inf}(\nu_a, \mu^*)} + h(c),$$

where $h(c) < \infty$ is a function of c (and of the distributions associated with the arms), which is however independent of T . As a consequence, we recover the asymptotic results of Burnetas and Katehakis (1996), Honda and Takemura (2010a), i.e., the guarantee that

$$\limsup_{T \rightarrow \infty} \frac{\bar{R}_T}{\log T} \leq \sum_{a \in \mathcal{A}} \frac{\Delta_a}{\mathcal{K}_{\inf}(\nu_a, \mu^*)}.$$

Of course, a sharper optimization can be performed by carefully choosing the constants c_a , that are parameters of the analysis; similarly to the comments after the statement of Theorem 3, we would then get a dominant term with a constant factor 1 instead of $1 + c$ as above, plus an additional second-order term. Details are left to a journal version of this paper.

Proof: By arguments similar to the ones used in the first step of the proof of Theorem 3, we have

$$\{A_{t+1} = a\} \subseteq \left\{ \mu^* - \varepsilon < \hat{\mu}_{a, N_t(a)} \right\} \cup \left\{ \mu^* - \varepsilon > B_{a^*, t}^+ \right\} \cup \left\{ \mu^* - \varepsilon \in [\hat{\mu}_{a, N_t(a)}, B_{a, t}^+] \right\};$$

indeed, on the event $\{A_{t+1} = a\} \cap \left\{ \mu^* - \varepsilon \geq \hat{\mu}_{a, N_t(a)} \right\} \cap \left\{ \mu^* - \varepsilon \leq B_{a^*, t}^+ \right\}$,

we have, $\hat{\mu}_{a, N_t(a)} \leq \mu^* - \varepsilon \leq B_{a^*, t}^+ \leq B_{a, t}^+$ (where the last inequality is by definition of the strategy). Before proceeding, we note that

$$\left\{ \mu^* - \varepsilon \in [\hat{\mu}_{a, N_t(a)}, B_{a, t}^+] \right\} \subseteq \left\{ N_t(a) \mathcal{K}_{\inf}(\hat{\nu}_{a, N_t(a)}, \mu^* - \varepsilon) \leq f(t) \right\},$$

since \mathcal{K}_{\inf} is a non-decreasing function in its second argument and $\mathcal{K}_{\inf}(\nu, E(\nu)) = 0$ for all distributions ν . Therefore,

$$\begin{aligned} \mathbb{E}[N_T(a)] &\leq 1 + \sum_{t=|A|}^{T-1} \mathbb{P} \left\{ \mu^* - \varepsilon < \hat{\mu}_{a, N_t(a)} \text{ and } A_{t+1} = a \right\} + \sum_{t=|A|}^{T-1} \mathbb{P} \left\{ \mu^* - \varepsilon > B_{a^*, t}^+ \right\} \\ &\quad + \sum_{t=|A|}^{T-1} \mathbb{P} \left\{ N_t(a) \mathcal{K}_{\inf}(\hat{\nu}_{a, N_t(a)}, \mu^* - \varepsilon) \leq f(t) \text{ and } A_{t+1} = a \right\}; \end{aligned}$$

now, the two sums with the events “and $A_{t+1} = a$ ” can be rewritten by using the stopping times $\tau_{a, k}$ introduced in the proof of Theorem 3; more precisely, by mimicking the transformations performed in its step 3, we get the simpler bound

$$\begin{aligned} \mathbb{E}[N_T(a)] &\leq 1 + \sum_{k=2}^{T-|A|+1} \mathbb{P} \left\{ \mu^* - \varepsilon < \tilde{\mu}_{a, k-1} \right\} + \sum_{t=|A|}^{T-1} \mathbb{P} \left\{ \mu^* - \varepsilon > B_{a^*, t}^+ \right\} \\ &\quad + \sum_{k=2}^{T-|A|+1} \mathbb{P} \left\{ (k-1) \mathcal{K}_{\inf}(\tilde{\nu}_{a, k-1}, \mu^* - \varepsilon) \leq f(t) \right\}, \quad (6) \end{aligned}$$

where the $\tilde{\nu}_{a, s}$ and $\tilde{\mu}_{a, s}$ are respectively the empirical distributions and empirical expectations computed on the first s elements of the sequence of the random variables $\tilde{X}_{a, j} = Y_{\tau_{a, j}}$, which are i.i.d. according to ν_a .

Step 1: The first sum in (6) is bounded by resorting to Hoeffding's inequality, whose application is legitimate since $\mu^* - \mu_a - \varepsilon > 0$;

$$\begin{aligned} \sum_{k=2}^{T-|\mathcal{A}|+1} \mathbb{P}\left\{\mu^* - \varepsilon < \tilde{\mu}_{a,k-1}\right\} &= \sum_{k=1}^{T-|\mathcal{A}|} \mathbb{P}\left\{\mu^* - \mu_a - \varepsilon < \tilde{\mu}_{a,k} - \mu_a\right\} \\ &\leq \sum_{k=1}^{T-|\mathcal{A}|} e^{-2k(\mu^* - \mu_a - \varepsilon)^2} \leq \frac{1}{1 - e^{-2(\mu^* - \mu_a - \varepsilon)^2}} \leq \frac{1}{(\mu^* - \mu_a - \varepsilon)^2}, \end{aligned}$$

where we used for the last inequality the general upper bounds provided at the beginning of step 5 in the proof of Theorem 3.

Step 2: The second sum in (6) is bounded by first using the definition of $B_{a^*,t}^+$, then, decomposing the event depending on the values taken by $N_t(a^*)$; and finally using the fact that on $\{N_t(a^*) = k\}$, we have the rewriting $\hat{\nu}_{a^*,N_t(a^*)} = \tilde{\nu}_{a^*,k}$ and $\hat{\mu}_{a^*,N_t(a^*)} = \tilde{\mu}_{a^*,k}$; more precisely,

$$\begin{aligned} \sum_{t=|\mathcal{A}|}^{T-1} \mathbb{P}\left\{\mu^* - \varepsilon > B_{a^*,t}^+\right\} &\leq \sum_{t=|\mathcal{A}|}^{T-1} \mathbb{P}\left\{N_t(a^*) \mathcal{K}_{\text{inf}}(\hat{\nu}_{a^*,N_t(a^*)}, \mu^* - \varepsilon) > f(t)\right\} \\ &= \sum_{t=|\mathcal{A}|}^{T-1} \sum_{k=1}^t \mathbb{P}\left\{N_t(a^*) = k \text{ and } k \mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^* - \varepsilon) > f(t)\right\} \\ &\leq \sum_{k=1}^T \sum_{t=|\mathcal{A}|}^{T-1} \mathbb{P}\left\{k \mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^* - \varepsilon) > f(t)\right\}. \end{aligned}$$

Since $f = \log$ is increasing, we can rewrite the bound, using a Fubini-Tonelli argument, as

$$\begin{aligned} \sum_{t=|\mathcal{A}|}^{T-1} \mathbb{P}\left\{\mu^* - \varepsilon > B_{a^*,t}^+\right\} &\leq \sum_{k=1}^T \sum_{t=|\mathcal{A}|}^{T-1} \mathbb{P}\left\{f^{-1}\left(k \mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^* - \varepsilon)\right) > t\right\} \\ &\leq \sum_{k=1}^T \mathbb{E}\left[f^{-1}\left(k \mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^* - \varepsilon)\right) \mathbb{I}_{\{\mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^* - \varepsilon) > 0\}}\right]. \end{aligned}$$

Now, Honda and Takemura (2010a, Lemma 13) indicates that, since $\mu^* - \varepsilon \in [0, 1)$,

$$\sup_{\nu \in \mathcal{P}_F([0,1])} \mathcal{K}_{\text{inf}}(\nu, \mu^* - \varepsilon) \leq \log(1/(1 - \mu^* + \varepsilon)) \stackrel{\text{def}}{=} K_{\max};$$

we define $Q = K_{\max}/\varepsilon^2$ and introduce the following sets $(V_q)_{1 \leq q \leq Q}$:

$$V_q = \left\{\nu \in \mathcal{P}_F([0,1]) : (q-1)\varepsilon^2 < \mathcal{K}_{\text{inf}}(\nu, \mu^* - \varepsilon) \leq q\varepsilon^2\right\}.$$

A peeling argument (and by using that $f^{-1} = \exp$ is increasing as well) entails, for all $k \geq 1$,

$$\begin{aligned} &\mathbb{E}\left[f^{-1}\left(k \mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^* - \varepsilon)\right) \mathbb{I}_{\{\mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^* - \varepsilon) > 0\}}\right] \tag{7} \\ &= \sum_{q=1}^Q \mathbb{E}\left[f^{-1}\left(k \mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^* - \varepsilon)\right) \mathbb{I}_{\{\tilde{\nu}_{a^*,k} \in V_q\}}\right] \\ &\leq \sum_{q=1}^Q \mathbb{P}\{\tilde{\nu}_{a^*,k} \in V_q\} f^{-1}(kq\varepsilon^2) \leq \sum_{q=1}^Q \mathbb{P}\left\{\mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^* - \varepsilon) > (q-1)\varepsilon^2\right\} f^{-1}(kq\varepsilon^2), \tag{8} \end{aligned}$$

where we used the definition of V_q to obtain each of the two inequalities. Now, by Lemma 7, when $E(\tilde{\nu}_{a^*,k}) < \mu^* - \varepsilon$, which is satisfied whenever $\mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^* - \varepsilon) > 0$, we have

$$\mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^* - \varepsilon) \leq \mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^*) - 2\varepsilon^2 \leq \mathcal{K}(\tilde{\nu}_{a^*,k}, \nu^*) - 2\varepsilon^2,$$

where the last inequality is by mere definition of \mathcal{K}_{inf} . Therefore,

$$\mathbb{P}\left\{\mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^* - \varepsilon) > (q-1)\varepsilon^2\right\} \leq \mathbb{P}\left\{\mathcal{K}(\tilde{\nu}_{a^*,k}, \nu^*) > (q+1)\varepsilon^2\right\}.$$

We note that for all $k \geq 1$, $\mathbb{P}\left\{\mathcal{K}(\tilde{\nu}_{a^*,k}, \nu^*) > (q+1)\varepsilon^2\right\} \leq (k+1)^{|\mathcal{S}^*|} e^{-k(q+1)\varepsilon^2}$, where we recall that \mathcal{S}^* denotes the finite support of ν^* and where we applied Corollary 12 of the appendix. Now, (8) then yields, via the choice $f = \log$ and thus $f^{-1} = \exp$, that

$$\mathbb{E}\left[f^{-1}\left(k \mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^* - \varepsilon)\right) \mathbb{I}\left\{\mathcal{K}_{\text{inf}}(\tilde{\nu}_{a^*,k}, \mu^* - \varepsilon) > 0\right\}\right] \leq \underbrace{\sum_{q=1}^Q (k+1)^{|\mathcal{S}^*|} e^{-k(q+1)\varepsilon^2} e^{kq\varepsilon^2}}_{=Q(k+1)^{|\mathcal{S}^*|} e^{-k\varepsilon^2}}.$$

Substituting the value of Q , we therefore have proved that

$$\sum_{t=|\mathcal{A}|}^{T-1} \mathbb{P}\left\{\mu^* - \varepsilon > B_{a^*,t}^+\right\} \leq \frac{1}{\varepsilon^2} \log\left(\frac{1}{1 - \mu^* + \varepsilon}\right) \sum_{k=1}^T (k+1)^{|\mathcal{S}^*|} e^{-k\varepsilon^2}.$$

Step 3: The third sum in (6) is first upper bounded by Lemma 7, which states that

$$\mathcal{K}_{\text{inf}}(\tilde{\nu}_{a,k-1}, \mu^*) - \varepsilon/(1 - \mu^*) \leq \mathcal{K}_{\text{inf}}(\tilde{\nu}_{a,k-1}, \mu^* - \varepsilon),$$

for all $k \geq 1$, and by using $f(t) \leq f(T)$; this gives

$$\begin{aligned} \sum_{k=1}^{T-|\mathcal{A}|} \mathbb{P}\left\{k \mathcal{K}_{\text{inf}}(\tilde{\nu}_{a,k}, \mu^* - \varepsilon) \leq f(t)\right\} \\ \leq \sum_{k=1}^{T-|\mathcal{A}|} \mathbb{P}\left\{k \mathcal{K}_{\text{inf}}(\tilde{\nu}_{a,k}, \mu^*) \leq f(T) + \frac{k\varepsilon}{1 - \mu^*}\right\} = \sum_{k=1}^{T-|\mathcal{A}|} \mathbb{P}\left\{\tilde{\nu}_{a,k} \in \bar{\mathcal{C}}_{\mu^*, \gamma_k}\right\}, \end{aligned}$$

where $\gamma_k = f(T)/k + \varepsilon/(1 - \mu^*)$ and where the set $\bar{\mathcal{C}}_{\mu^*, \gamma_k}$ was defined in Section 4.1. For all $\gamma > 0$, we then introduce

$$\theta_a(\gamma) = \inf\left\{\mathcal{K}(\nu', \nu_a) : \nu' \in \mathcal{C}_{\mu^*, \gamma}\right\} = \inf\left\{\mathcal{K}(\nu', \nu_a) : \nu' \in \bar{\mathcal{C}}_{\mu^*, \gamma}\right\},$$

(where the second equality follows from the lower semi-continuity of \mathcal{K}) and aim at bounding $\mathbb{P}\left\{\tilde{\nu}_{a,k} \in \bar{\mathcal{C}}_{\mu^*, \gamma}\right\}$.

As shown in Section 4.1, the set $\mathcal{C}_{\mu^*, \gamma}$ is a non-empty open convex set. If we prove that $\theta_a(\gamma)$ is finite for all $\gamma > 0$, then all the conditions will be required to apply Lemma 1 and get the upper bound

$$\sum_{k=1}^{T-|\mathcal{A}|} \mathbb{P}\left\{\tilde{\nu}_{a,k} \in \bar{\mathcal{C}}_{\mu^*, \gamma_k}\right\} \leq \sum_{k=1}^{T-|\mathcal{A}|} e^{-k\theta_a(\gamma_k)}.$$

To that end, we use the fact that ν_a is finitely supported. Now, either the probability of interest is null and we are done; or, it is not null, which implies that there exists a possible value of $\tilde{\nu}_{a,k}$ that is in $\bar{\mathcal{C}}_{\mu^*, \gamma}$; since this value is a distribution with a support included in the one of ν_a , it is absolutely continuous with respect to ν_a and hence, the Kullback-Leibler divergence between this value and ν_a is finite; in particular, $\theta_a(\gamma)$ is finite.

Finally, we bound the $\theta_a(\gamma_k)$ for values of k larger than $k_0 = \left\lceil \frac{(1+c_a)f(T)}{\mathcal{K}_{\text{inf}}(\nu_a, \mu^*)} \right\rceil$;

we have that for all $k \geq k_0$, in view of the bound put on ε ,

$$\gamma_k \leq \gamma_{k_0} = \frac{f(T)}{k_0} + \frac{\varepsilon}{1 - \mu^*} < \frac{\mathcal{K}_{\text{inf}}(\nu_a, \mu^*)}{1 + c_a} + \frac{c_a/2}{1 + c_a} \mathcal{K}_{\text{inf}}(\nu_a, \mu^*) = \frac{1 + c_a/2}{1 + c_a} \mathcal{K}_{\text{inf}}(\nu_a, \mu^*). \quad (9)$$

Since θ_a is non increasing, we have

$$\sum_{k=1}^{T-|\mathcal{A}|} e^{-k\theta_a(\gamma_k)} \leq k_0 - 1 + \sum_{k=k_0}^{T-|\mathcal{A}|} e^{-k\theta_a(\gamma_{k_0})} \leq k_0 - 1 + \frac{1}{1 - e^{-\Theta_a(c_a, \varepsilon)}},$$

provided that the quantity $\Theta_a(c_a, \varepsilon) = \theta_a(\gamma_{k_0})$ is positive, which we prove now.

Indeed for all $\nu' \in \mathcal{C}_{\mu^*, \gamma_{k_0}}$, we have by definition and by (9) that

$$\mathcal{K}_{\text{inf}}(\nu', \mu^*) - \mathcal{K}_{\text{inf}}(\nu_a, \mu^*) < \gamma_{k_0} - \mathcal{K}_{\text{inf}}(\nu_a, \mu^*) < -((c_a/2)/(1 + c_a))\mathcal{K}_{\text{inf}}(\nu_a, \mu^*).$$

Now, in the case where $\mathbb{E}_{\nu_a}[(1-\mu^*)/(1-X)] > 1$, we have, first by application of Pinsker's inequality and then by Lemma 6, that

$$\mathcal{K}(\nu', \nu_a) \geq \frac{\|\nu' - \nu_a\|_1^2}{2} \geq \frac{1}{2M_{\nu_a, \mu^*}^2} (\mathcal{K}_{\text{inf}}(\nu_a, \mu^*) - \mathcal{K}_{\text{inf}}(\nu', \mu^*))^2 > \frac{c_a^2 (\mathcal{K}_{\text{inf}}(\nu_a, \mu^*))^2}{8(1+c_a)^2 M_{\nu_a, \mu^*}^2};$$

since, again by Pinsker's inequality, $\mathcal{K}_{\text{inf}}(\nu_a, \mu^*) \geq (\mu_a - \mu^*)^2/2 > 0$, we have exhibited a lower bound independent of ν' in this case. In the case where $\mathbb{E}_{\nu_a}[(1-\mu^*)/(1-X)] \leq 1$, we apply the second part of Lemma 6, with $\alpha_a = (c_a/2)/(1+c_a)$, and get

$$\mathcal{K}(\nu', \nu_a) \geq \frac{\|\nu' - \nu_a\|_1^2}{2} \geq \frac{1}{2} \left(\frac{1-\mu^*}{(2/\alpha_a)((2/\alpha_a)-1)} \right)^2 > 0.$$

Thus, in both cases we found a positive lower bound independent of ν' , so that the infimum over $\nu' \in \mathcal{C}_{\mu^*, \gamma_{k_0}}$ of the quantities $\mathcal{K}_{\text{inf}}(\nu', \mu^*)$, which precisely equals $\theta_a(\gamma_{k_0})$, is also positive. This concludes the proof. \blacksquare

Conclusion. We provided a finite-time analysis of the (asymptotically optimal) \mathcal{K}_{inf} -strategy in the case of finitely supported distributions. One could think that the extension to the case of general distributions is straightforward. However this extension appears somewhat difficult (at least when using the current definition of \mathcal{K}_{inf}) for the following reasons: (1) Step 2 in the proof uses the method of types, that would require some extension of Sanov's non-asymptotic Theorem to this case. (2) Step 3 requires to have both $\theta_a(\gamma) < \infty$ for all $\gamma > 0$ and $\theta_a(\gamma) > 0$ for $\gamma < \mathcal{K}_{\text{inf}}(\nu_a, \mu^*)$, which does not seem to be always the case for general distributions. Exploring other directions for such extensions is left for future work; for instance, histogram-based approximations of general distributions could be considered.

Acknowledgements. The authors wish to thank *Peter Auer* and *Daniil Ryabko* for insightful discussions. They acknowledge support from the French National Research Agency (ANR) under grant EX-PLO/RA ("Exploration-exploitation for efficient resource allocation") and by the PASCAL2 Network of Excellence under EC grant no. 506778.

References

- J.-Y. Audibert, R. Munos, and C. Szepesvari. Exploration-exploitation trade-off using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410:1876–1902, 2009.
- J.Y. Audibert and S. Bubeck. Regret bounds and minimax policies under partial monitoring. *Journal of Machine Learning Research*, 11:2635–2686, 2010.
- P. Auer and R. Ortner. UCB revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61(1-2):55–65, 2010.
- P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002.
- A.N. Burnetas and M.N. Katehakis. Optimal adaptive policies for sequential allocation problems. *Advances in Applied Mathematics*, 17(2):122–142, 1996.
- Y. Chow and H. Teicher. *Probability Theory*. Springer, 1988.
- I.H. Dinwoodie. Mesures dominantes et théorème de Sanov. *Annales de l’Institut Henri Poincaré – Probabilités et Statistiques*, 28(3):365–373, 1992.
- S. Filippi. *Stratégies optimistes en apprentissage par renforcement*. PhD thesis, Télécom ParisTech, 2010.
- A. Garivier and O. Cappé. The KL-UCB algorithm for bounded stochastic bandits and beyond. In *Proceedings of COLT*, 2011.
- A. Garivier and F. Leonardi. Context tree selection: A unifying view. arXiv:1011.2424, 2010.
- J. Honda and A. Takemura. An asymptotically optimal bandit algorithm for bounded support models. In *Proceedings of COLT*, pages 67–79, 2010a.
- J. Honda and A. Takemura. An asymptotically optimal policy for finite support models in the multiarmed bandit problem. arXiv:0905.2776, 2010b.
- T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6:4–22, 1985.
- H. Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematics Society*, 58:527–535, 1952.

A Appendix beyond the COLT page limit

A conference version of this paper was published in the *Proceedings of the Twenty-Fourth Annual Conference on Learning Theory (COLT'11)*; this appendix details some material which was alluded at in this conference version but could not be published therein because of the page limit.

A.1 Proof of Lemma 2

We only provide it for the convenience of the readers since it is similar to the one presented in Garivier and Leonardi (2010, Proposition 2) or in Garivier and Cappé (2011); it was however somewhat simplified by noting that the proof technique used leads to a maximal inequality, as stated in Lemma 2, and not only to an inequality for a self-normalized average, as stated in the original reference.

Proof: The result is straightforward in the cases $p = 0$ or $p = 1$, since then, $\hat{p}_s = p$ almost surely; in the rest of the proof, we therefore only consider the case where $p \in (0, 1)$.

It suffices to show the first bound stated in the lemma, since the second one follows by a decomposition of the probability space according to the values of N_t . Actually, we will show

$$\mathbb{P}_p \left(\bigcup_{s=1}^t \left\{ s \mathcal{K}(\beta(\hat{p}_s), \beta(p)) \geq \varepsilon \text{ and } \hat{p}_s > p \right\} \right) \leq e \lceil \varepsilon \log t \rceil e^{-\varepsilon},$$

and the desired result will follow by symmetry and a union bound.

Step 1: A martingale. For all $\lambda > 0$, we consider the log-Laplace transform

$$\psi_p(\lambda) = \log \mathbb{E}_p [e^{\lambda X_1}] = \log((1-p) + p e^\lambda),$$

with which we define the martingale

$$W_s(\lambda) = \exp(\lambda(X_1 + \dots + X_s) - s \psi_p(\lambda)).$$

Step 2: A peeling argument. We introduce $t_0 = 1$ and $t_k = \lfloor \gamma^k \rfloor$, for some $\gamma > 1$ that will be defined by the analysis. We also denote by $K = \lceil (\log t) / (\log \gamma) \rceil$ an upper bound on the number of elements in the peeling.

We also note that by continuity of the Kullback-Leibler divergence in the case of Bernoulli distributions, for all $\varepsilon > 0$, there exists a unique element $p_\varepsilon \in (p, 1)$ such that $\mathcal{K}(\beta(p_\varepsilon), \beta(p)) = \varepsilon$; this element satisfies that

$$\mathcal{K}(\beta(q), \beta(p)) \geq \varepsilon \text{ and } q \geq p \quad \text{entails} \quad q \geq p_\varepsilon.$$

Denoting by $\varepsilon_k = \varepsilon / t_k$, a union bound using the described peeling then yields

$$\begin{aligned} & \mathbb{P}_p \left(\bigcup_{s=1}^t \left\{ s \mathcal{K}(\beta(\hat{p}_s), \beta(p)) \geq \varepsilon \text{ and } \hat{p}_s > p \right\} \right) \\ & \leq \sum_{k=1}^K \mathbb{P}_p \left(\bigcup_{s=t_{k-1}}^{t_k} \left\{ s \mathcal{K}(\beta(\hat{p}_s), \beta(p)) \geq \varepsilon \text{ and } \hat{p}_s > p \right\} \right) \\ & \leq \sum_{k=1}^K \mathbb{P}_p \left(\bigcup_{s=t_{k-1}}^{t_k} \left\{ \mathcal{K}(\beta(\hat{p}_s), \beta(p)) \geq \frac{\varepsilon}{t_k} \text{ and } \hat{p}_s > p \right\} \right) \\ & = \sum_{k=1}^K \mathbb{P}_p \left(\bigcup_{s=t_{k-1}}^{t_k} \left\{ \hat{p}_s \geq p_{\varepsilon_k} \right\} \right) = \sum_{k=1}^K \mathbb{P}_p \left(\bigcup_{s=t_{k-1}}^{t_k} \left\{ X_1 + \dots + X_s - s p_{\varepsilon_k} \geq 0 \right\} \right) \end{aligned}$$

Now, the variational formula for Kullback-Leibler divergences shows that for all k , there exists a λ_k such that

$$\varepsilon_k = \mathcal{K}(\beta(p_{\varepsilon_k}), \beta(p)) = \lambda_k p_{\varepsilon_k} - \psi_p(\lambda_k);$$

actually, a straightforward calculation shows that $\lambda_k = \log(p_{\varepsilon_k}(1-p) - \log(p(1-p_{\varepsilon_k}))) > 0$ is a

suitable value. Thus,

$$\begin{aligned}
& \sum_{k=1}^K \mathbb{P}_p \left(\bigcup_{s=t_{k-1}}^{t_k} \{X_1 + \dots + X_s - s p_{\varepsilon_k} \geq 0\} \right) \\
&= \sum_{k=1}^K \mathbb{P}_p \left(\bigcup_{s=t_{k-1}}^{t_k} \{ \exp(\lambda_k(X_1 + \dots + X_s) - \lambda_k s p_{\varepsilon_k}) \geq 1 \} \right) \\
&= \sum_{k=1}^K \mathbb{P}_p \left(\bigcup_{s=t_{k-1}}^{t_k} \{ \exp(\lambda_k(X_1 + \dots + X_s) - s \psi_p(\lambda_k)) \geq e^{s \varepsilon_k} \} \right) \\
&\leq \sum_{k=1}^K \mathbb{P}_p \left(\bigcup_{s=t_{k-1}}^{t_k} \{ W_s(\lambda_k) \geq e^{t_{k-1} \varepsilon_k} \} \right) \\
&\leq \sum_{k=1}^K e^{-t_{k-1} \varepsilon_k} = K e^{-\varepsilon/\gamma},
\end{aligned}$$

where in the last step, we resorted to Doob's maximal inequality.

Step 3: Choosing γ . The obtained bound equals, by substituting the value of K and by choosing $\gamma = \varepsilon/(\varepsilon - 1)$,

$$K e^{-\varepsilon/\gamma} = \lceil (\log t)/(\log \gamma) \rceil e^{-\varepsilon+1} = \left\lceil \frac{\log t}{\log(\varepsilon/(\varepsilon - 1))} \right\rceil e^{-\varepsilon+1};$$

the proof is concluded by noting that $\varepsilon > 1 \mapsto \log(\varepsilon/(\varepsilon - 1)) - 1/\varepsilon$ is decreasing (its derivative is negative), with limit 0 at $+\infty$. \blacksquare

A.2 Details of the adaptation leading to Lemma 1

The exact statement of Dinwoodie (1992, Theorem 2.1 and comments on page 372) is the following.

Lemma 10 (Non-asymptotic Sanov's lemma) *Let \mathcal{C} be an open convex subset of $\mathcal{P}(\mathcal{X})$ such that*

$$\Lambda(\mathcal{C}) = \inf_{\kappa \in \mathcal{C}} \mathcal{K}(\kappa, \nu) < \infty.$$

Then, for all $t \geq 1$,

$$\mathbb{P}_\nu \{ \hat{\nu}_t \in \mathcal{C} \} \leq e^{-t\Lambda(\mathcal{C})}.$$

We show how it entails Lemma 1. Let \mathcal{C} be an open convex subset of $\mathcal{P}(\mathcal{X})$ and let $\bar{\mathcal{C}}$ be its closure. We denote by

$$\mathcal{C}_\delta = \{ \nu \in \mathcal{C} : d(\nu, \mathcal{C}) < \delta \}$$

the δ -open neighborhood of \mathcal{C} , we have $\bar{\mathcal{C}} \subseteq \mathcal{C}_\delta$ for all $\delta > 0$. Therefore, by the lemma above, since $\Lambda(\mathcal{C}_\delta) \leq \Lambda(\mathcal{C}) < \infty$,

$$\mathbb{P}_\nu \{ \hat{\nu}_t \in \bar{\mathcal{C}} \} \leq \mathbb{P}_\nu \{ \hat{\nu}_t \in \mathcal{C}_\delta \} \leq e^{-t\Lambda(\mathcal{C}_\delta)}.$$

We pick for each integer $n \geq 1$ an element κ_n such that $\Lambda(\mathcal{C}_{1/n}) = \mathcal{K}(\kappa_n, \nu) - 1/n$; by Dinwoodie (1992, proof of Proposition 1.1), the sequence of the κ_n admits a converging subsequence $\kappa_{\varphi(n)}$, whose limit point κ_∞ belongs to $\bar{\mathcal{C}}$ and which satisfies

$$\mathcal{K}(\kappa_\infty, \nu) \leq \liminf_{n \rightarrow \infty} \mathcal{K}(\kappa_n, \nu) = \liminf_{\delta \rightarrow 0} \Lambda(\mathcal{C}_\delta).$$

Therefore, by taking limits in the above inequality, we have proved the desired inequality,

$$\mathbb{P}_\nu \{ \hat{\nu}_t \in \bar{\mathcal{C}} \} \leq e^{-t\mathcal{K}(\kappa_\infty, \nu)} \leq e^{-t\Lambda(\bar{\mathcal{C}})}.$$

A.3 Useful properties of \mathcal{K}_{inf} and its level sets

Proof of Lemma 6: We resort to the formulation of \mathcal{K}_{inf} in terms of a convex optimization problem as introduced in Honda and Takemura (2010b); more precisely, it is shown therein that

$$\mathcal{K}_{\text{inf}}(\nu, \mu) = \max \left\{ \mathbb{E}_{\nu} \left[\log(1 + \lambda(\mu - X)) \right] : \lambda \in [0, 1/(1 - \mu)] \right\} \quad (10)$$

(where X denotes a random variable distributed according to ν), as well as the following alternative. The optimal value λ_{ν} of the parameter λ indexing the set is equal to $1/(1 - \mu)$ if and only if $\mathbb{E}_{\nu}[(1 - \mu)/(1 - X)] \leq 1$, and lies in $[0, 1/(1 - \mu))$ if $\mathbb{E}_{\nu}[(1 - \mu)/(1 - X)] > 1$.

For all $\lambda \in [0, 1/(1 - \mu)]$, we now introduce the function

$$\phi_{\lambda} : x \in [0, 1] \mapsto \log(1 + \lambda(\mu - x)),$$

which is always continuous on $[0, 1]$; we note also that it is continuous and finite at $x = 1$ when $\lambda < 1/(1 - \mu)$. In the latter case, ϕ_{λ} is bounded; since it is decreasing, it is easy to get a uniform bound: for all x ,

$$|\phi_{\lambda}(x)| \leq |\phi(0)| + |\phi(1)| = \log \frac{1 + \lambda\mu}{1 + \lambda(\mu - 1)} \stackrel{\text{def}}{=} M_{\lambda}.$$

It then follows that for all $\lambda \in [0, 1/(1 - \mu))$,

$$\mathbb{E}_{\nu}[\phi_{\lambda}(X)] - \mathbb{E}_{\nu'}[\phi_{\lambda}(X)] \leq M_{\lambda} \|\nu - \nu'\|_1. \quad (11)$$

In the case when $\lambda_{\nu} < 1/(1 - \mu)$, we have from the variational formulation (10) that

$$\mathcal{K}_{\text{inf}}(\nu, \mu) - \mathcal{K}_{\text{inf}}(\nu', \mu) \leq \mathbb{E}_{\nu}[\phi_{\lambda_{\nu}}(X)] - \mathbb{E}_{\nu'}[\phi_{\lambda_{\nu}}(X)] \leq M_{\lambda_{\nu}} \|\nu - \nu'\|_1.$$

Thus, the constant $M_{\nu, \mu}$ in the statement of the lemma corresponds to our quantity $M_{\lambda_{\nu}}$ in this case.

We now consider the case where $\lambda_{\nu} = 1/(1 - \mu)$. By (11) and variational formulation (10), we have that for all $\lambda \in [0, 1/(1 - \mu))$,

$$\begin{aligned} \mathcal{K}_{\text{inf}}(\nu, \mu) - \mathcal{K}_{\text{inf}}(\nu', \mu) &\leq \mathcal{K}_{\text{inf}}(\nu, \mu) - \mathbb{E}_{\nu'}[\phi_{\lambda}(X)] \\ &= \left(\mathcal{K}_{\text{inf}}(\nu, \mu) - \mathbb{E}_{\nu}[\phi_{\lambda}(X)] \right) + \left(\mathbb{E}_{\nu}[\phi_{\lambda}(X)] - \mathbb{E}_{\nu'}[\phi_{\lambda}(X)] \right). \end{aligned}$$

The second difference is bounded according to (11); the first difference is bounded by concavity of $\lambda < 1/(1 - \mu) \mapsto \phi_{\lambda}(x)$, for all x :

$$\begin{aligned} \mathbb{E}_{\nu}[\phi_{\lambda}(X)] &\geq (1 - \lambda(1 - \mu)) \mathbb{E}_{\nu}[\phi_0(X)] + \lambda(1 - \mu) \mathbb{E}_{\nu}[\phi_{1/(1 - \mu)}(X)] \\ &= \lambda(1 - \mu) \mathbb{E}_{\nu}[\phi_{1/(1 - \mu)}(X)] = \lambda(1 - \mu) \mathcal{K}_{\text{inf}}(\nu, \mu), \end{aligned}$$

since ϕ_0 is the null function and $\lambda_{\nu} = 1/(1 - \mu)$. Putting all pieces together, we have proved that for all $\lambda \in [0, 1/(1 - \mu))$,

$$\mathcal{K}_{\text{inf}}(\nu, \mu) - \mathcal{K}_{\text{inf}}(\nu', \mu) \leq (1 - \lambda(1 - \mu)) \mathcal{K}_{\text{inf}}(\nu, \mu) + M_{\lambda} \|\nu - \nu'\|_1. \quad (12)$$

We recall that by assumption, $\mathcal{K}_{\text{inf}}(\nu, \mu) - \mathcal{K}_{\text{inf}}(\nu', \mu) \geq \alpha \mathcal{K}_{\text{inf}}(\nu, \mu)$ with $\alpha \in (0, 1)$, so that the choice $\lambda = (1 - \alpha/2)/(1 - \mu)$, which indeed lies in $(0, 1/(1 - \mu))$, is such that

$$M_{\lambda} = \log \left(1 + \frac{\lambda}{1 + \lambda(\mu - 1)} \right) = \log \left(1 + \frac{\lambda}{\alpha/2} \right) \leq \frac{2\lambda}{\alpha},$$

so that (12) entails

$$\alpha \mathcal{K}_{\text{inf}}(\nu, \mu) \leq \frac{\alpha}{2} \mathcal{K}_{\text{inf}}(\nu, \mu) + \frac{2\lambda}{\alpha} \|\nu - \nu'\|_1,$$

and finally

$$\|\nu - \nu'\|_1 \geq \frac{\alpha^2}{4\lambda} = \frac{\alpha^2(1 - \mu)}{1 - \alpha/2} = \frac{1 - \mu}{(2/\alpha)((2/\alpha) - 1)};$$

which concludes the proof. ■

Proof of Lemma 7: In Honda and Takemura (2010b) it is shown that in this case, $\mathcal{K}_{\text{inf}}(\nu, \mu)$ is differentiable in $\mu \in (E(\nu), 1)$ with

$$\frac{1}{1-\mu} \geq \frac{\partial}{\partial \mu} \mathcal{K}_{\text{inf}}(\nu, \mu) \geq \frac{\mu - E(\nu)}{\mu(1-\mu)}. \quad (13)$$

We apply this result to the rewriting

$$\mathcal{K}_{\text{inf}}(\nu, \mu) - \mathcal{K}_{\text{inf}}(\nu, \mu - \varepsilon) = \int_{\mu-\varepsilon}^{\mu} \frac{\partial}{\partial \mu} \mathcal{K}_{\text{inf}}(\nu, u) \, du,$$

which already gives one part of the bound. For the lower bound, we note that by assumption $-E(\nu) > -(\mu - \varepsilon)$ and that $u(1-u) \leq 1/4$ (since we consider distributions with support included in $[0, 1]$); so that, for all u ,

$$\frac{u - E(\nu)}{u(1-u)} \geq 4(u - (\mu - \varepsilon)).$$

Integrating the bound concludes the main part of the proof.

Now, to see that the first inequality in the statement is always valid, we need to consider the case when $E(\nu) \geq \mu$, for which the statement is trivial since then $\mathcal{K}_{\text{inf}}(\nu, \mu) = 0$, and the case when $\mu > E(\nu) \geq \mu - \varepsilon$. But in the latter case, it is shown in Honda and Takemura (2010b, Lemma 6, case 2) that

$$\mathcal{K}_{\text{inf}}(\nu, \mu) \leq \frac{\mu - E(\nu)}{1-\mu},$$

which concludes the proof. ■

Proof of Lemma 8: First, $\mathcal{C}_\mu(\gamma)$ is non empty as it always contains δ_μ , the Dirac mass on μ .

The fact that $\mathcal{C}_\mu(\gamma)$ is convex follows from the convexity of \mathcal{K} in the pair of probability distributions that it takes as an argument. Indeed, for all $\alpha \in [0, 1]$, $\nu', \nu'' \in \mathcal{C}_\mu(\gamma)$, denoting by ν'_μ, ν''_μ some distributions such that the defining conditions in $\mathcal{C}_\mu(\gamma)$ are satisfied, we have that

$$E(\alpha \nu'_\mu + (1-\alpha) \nu''_\mu) > \mu$$

and

$$\mathcal{K}(\alpha \nu' + (1-\alpha) \nu'', \alpha \nu'_\mu + (1-\alpha) \nu''_\mu) \leq \alpha \mathcal{K}(\nu', \nu'_\mu) + (1-\alpha) \mathcal{K}(\nu'', \nu''_\mu) < \gamma.$$

We prove that $\mathcal{C}_\mu(\gamma)$ is an open set. With each $\nu' \in \mathcal{C}_\mu(\gamma)$, we associate a distribution ν'_μ satisfying the defining constraints in $\mathcal{C}_\mu(\gamma)$; by choosing

$$\alpha = \frac{1 - \mu/E(\nu'_\mu)}{2} \in (0, 1/2),$$

we have that the open set formed by the

$$(1-\alpha) \nu' + \alpha \nu'', \quad \nu'' \in \text{B}(\nu', 1)$$

is contained in $\mathcal{C}_{\mu, \gamma}$, where $\text{B}(\nu', 1)$ denotes the ball with center ν' and radius 1 in the norm $\|\cdot\|$ over $\mathcal{P}(\mathcal{X})$. Indeed, we have on the one hand,

$$E((1-\alpha) \nu' + \alpha \nu'') \geq (1-\alpha) E(\nu') \geq \left(1 - \frac{1 - \mu/E(\nu'_\mu)}{2}\right) E(\nu') = \frac{E(\nu'_\mu) + \mu}{2} > \mu,$$

and on the other hand, by convexity of the Kullback-Leibler divergence,

$$\mathcal{K}((1-\alpha) \nu' + \alpha \nu'', (1-\alpha) \nu'_\mu + \alpha \nu''_\mu) \leq (1-\alpha) \mathcal{K}(\nu', \nu'_\mu) < (1-\alpha) \gamma.$$

To prove the desired inclusion, we first note that in the case of $\mathcal{P}_F([0, 1])$, Honda and Takemura (2010b) show that one has the rewriting

$$\mathcal{K}_{\text{inf}}(\nu, \mu) = \min \left\{ \mathcal{K}(\nu, \nu') : \nu' \in \mathcal{P}_F([0, 1]) \text{ s.t. } E(\nu') \geq \mu \right\};$$

in particular, the infimum is achieved with this new formulation. Hence,

$$\mathcal{C}_{\mu, \gamma} = \left\{ \nu' \in \mathcal{P}_F([0, 1]) : \exists \nu'_\mu \in \mathcal{P}_F([0, 1]) \text{ s.t. } E(\nu'_\mu) \geq \mu \text{ and } \mathcal{K}(\nu', \nu'_\mu) < \gamma \right\}.$$

Also, an element of the set of interest is therefore a $\nu' \in \mathcal{P}_F([0, 1])$ such that $\mathcal{K}_{\text{inf}}(\nu', \mu) \leq \gamma$, that is, such that there exists $\nu'_\mu \in \mathcal{P}([0, 1])$ with $E(\nu'_\mu) \geq \mu$ and $\mathcal{K}(\nu', \nu'_\mu) \leq \gamma$. Now, the distributions

$$\nu'_n = \left(1 - \frac{1}{n}\right) \nu' + \frac{1}{n} \delta_1, \quad \text{thanks to the} \quad \nu'_{\mu, n} = \left(1 - \frac{1}{n}\right) \nu'_\mu + \frac{1}{n} \delta_1,$$

all belong to \mathcal{C}_γ , as, similarly to the above argument,

$$E(\nu'_n) \geq \mu + \frac{1 - \mu}{n} > \mu \quad \text{and} \quad \mathcal{K}(\nu'_n, \nu'_{\mu, n}) \leq \left(1 - \frac{1}{n}\right) \mathcal{K}(\nu', \nu'_\mu) < \gamma.$$

In addition, we have by construction that the ν'_n converge to ν' , hence, $\nu' \in \overline{\mathcal{C}_\gamma}$. ■

A.4 The method of types

Let X_1, X_2, \dots be a sequence of random variables that are i.i.d. according to a distribution denoted by ν . In this subsection, we will index all probabilities and expectations by ν .

For all $k \geq 1$, we denote by \mathcal{E}_k the set of possible values (the so-called types) of the empirical distribution

$$\widehat{\nu}_k = \sum_{j=1}^k \delta_{X_j}.$$

If ν has a finite support denoted by \mathcal{S} , then the cardinality $|\mathcal{E}_k|$ of \mathcal{E}_k is bounded by $(k+1)^{|\mathcal{S}|}$.

Lemma 11 *In the case where ν has a finite support, for all $k \geq 1$ and $\kappa \in \mathcal{E}_k$,*

$$\mathbb{P}_\nu\{\widehat{\nu}_k = \kappa\} \leq e^{-k \mathcal{K}(\kappa, \nu)}.$$

Corollary 12 *In the case where ν has a finite support, for all $k \geq 1$, all $\gamma > 0$,*

$$\begin{aligned} \mathbb{P}\left\{\mathcal{K}(\widehat{\nu}_k, \nu) > \gamma\right\} &= \sum_{\kappa \in \mathcal{E}_k} \mathbb{I}_{\{\mathcal{K}(\kappa, \nu) > \gamma\}} \mathbb{P}_\nu\{\widehat{\nu}_k = \kappa\} \\ &\leq \sum_{\kappa \in \mathcal{E}_k} \mathbb{I}_{\{\mathcal{K}(\kappa, \nu) > \gamma\}} e^{-k \mathcal{K}(\kappa, \nu)} \leq |\mathcal{E}_k| e^{-k\gamma} \leq (k+1)^{|\mathcal{S}|} e^{-k\gamma}. \end{aligned}$$