



HAL
open science

Polynomial Cost of Adaptation for X -Armed Bandits

Hédi Hadiji

► **To cite this version:**

Hédi Hadiji. Polynomial Cost of Adaptation for X -Armed Bandits. NeurIPS 2019 - Thirty-third Conference on Neural Information Processing Systems, Dec 2019, Vancouver, Canada. hal-02138492v2

HAL Id: hal-02138492

<https://hal.science/hal-02138492v2>

Submitted on 7 Dec 2019

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.

Polynomial Cost of Adaptation for \mathcal{X} -Armed Bandits

Hédi Hadiji

Laboratoire de Mathématiques d'Orsay
Université Paris-Sud, Orsay, France
hedi.hadiji@math.u-psud.fr

Abstract

In the context of stochastic continuum-armed bandits, we present an algorithm that adapts to the unknown smoothness of the objective function. We exhibit and compute a *polynomial cost of adaptation* to the Hölder regularity for regret minimization. To do this, we first reconsider the recent lower bound of Locatelli and Carpentier [21], and define and characterize admissible rate functions. Our new algorithm matches any of these minimal rate functions. We provide a finite-time analysis and a thorough discussion about asymptotic optimality.

1 Introduction

Multi-armed bandits are a well-known sequential learning problem. When the number of available decisions is large, some assumptions on the environment have to be made. In a vast line of work (see the literature discussion in Section 1.1), these assumptions show up as nonparametric regularity conditions on the mean-payoff function. If this function is Hölder continuous with constant L and exponent α , and if the values of L and α are given to the player, then natural strategies can ensure that the regret is upper bounded by

$$L^{1/(2\alpha+1)}T^{(\alpha+1)/(2\alpha+1)}. \quad (1)$$

Of course, assuming that the player knows α and L is often not realistic. Thus the need for *adaptive* methods, that are agnostic with respect to the true regularity of the mean-payoff function. Unfortunately, Locatelli and Carpentier [21] recently showed that full adaptation is impossible, and that no algorithm can enjoy the same minimax guarantees as when the regularity is given to the player. We persevere and address the question:

What can the player achieve when the true regularity is completely unknown?

A polynomial cost of adaptation In statistics, minimax adaptation for nonparametric function estimation is a deep and active research domain. In many contexts, sharp adaptation is possible; often, an additional logarithmic factor in the error has to be paid when the regularity is unknown: this is known as the *cost of adaptation*. See e.g., Lepskii [20], Birgé and Massart [4], Massart [22] for adaptive methods, and Cai [9] for a detailed survey of the topic. Under some more exotic assumptions —see e.g., Example 3 of Cai and Low [10] — adapting is significantly harder: there may be a *polynomial cost of adaptation*.

In this paper, we show that in the sequential setting of multi-armed bandits, the necessary exploration forces a similar phenomenon, and we exhibit this polynomial cost of adaptation. To do so, we revisit the lower bounds of Locatelli and Carpentier [21], and design a new algorithm that matches these lower bounds.

As a representative example of our results, our algorithm can achieve, without the knowledge of α and L , an unimprovable (up to logarithmic factors) regret bound of order

$$L^{1/(1+\alpha)}T^{(\alpha+2)/(2\alpha+2)}. \quad (2)$$

1.1 Related work

Continuum-armed bandits Continuum-armed bandits, with nonparametric regularity assumptions, were introduced by Agrawal [1]. Kleinberg [16] established the minimax rates in the Hölder setting and introduced the CAB1 algorithm. Auer et al. [3] studied the problem with additional regularity assumptions under which the minimax rates are improved. Via different roads, Bubeck et al. [7] and Kleinberg et al. [17] explored further generalizations of these types of regularity, namely the zooming dimension and the near-optimality dimension. Bull [8] exhibited an algorithm that essentially adapts to some cases when the near-optimality dimension is zero.

In all these articles, the mean-payoff function needs to satisfy simultaneously two sets of regularity conditions. The first type is a usual Hölder condition, which ensures that the function does not vary too much around (one of) its maxima. The second type is a “margin condition” that lower bounds the number of very suboptimal arms; in the literature these are defined in many technically different ways. Adapting to the margin conditions is often possible when the Hölder regularity is known. However, all these algorithms require some prior knowledge about the Hölder regularity.

In this paper, we focus on the problem of adapting to Hölder regularity. Accordingly, we call *adaptive* the algorithms that assume no knowledge of the Hölder exponent nor of the Lipschitz constant.

Adaptation for cumulative regret Bubeck et al. [5] introduced the problem of adaptation, and adapted to the Lipschitz constant under extra requirements. An important step was made in Locatelli and Carpentier [21], where it is shown that adaptation at the classical minimax rates is impossible. In the same paper, the authors exhibited some conditions under which full adaptation is achievable, e.g., with knowledge of the value of the maximum, or when the near-optimality dimension is zero.

Other settings For simple regret, the objections against adaptation do not hold, as the objective does not penalize exploration. Adaptation up to polylog factors is done with various (meta-)algorithms. Locatelli and Carpentier [21] sketch out an aggregation approach inspired by Lepski’s method, while Valko et al. [24], Grill et al. [14], Shang et al. [23] describe cross-validation methods thanks to which they adapt to the near-optimality dimension with unknown smoothness. As it turns out, this last approach yields clean results with our smoothness assumptions; we write the details in Appendix E.

Recently, Krishnamurthy et al. [18] studied continuum-armed contextual bandits and use a sophisticated aggregation scheme to derive an algorithm that adapts to the Lipschitz constant when $L \geq 1$.

1.2 Contributions and outline

In this paper, we fully compute the cost of adaptation for bandits with Hölder regularity. In Section 2 we discuss the adaptive (and nonadaptive) lower bounds. We take an asymptotic stance in order to precisely define the objective of adaptation. Doing so, we uncover a family of noncomparable lower bounds for adaptive algorithms (Theorem 1), and define the corresponding notion of optimality: admissibility.

Section 3 contains our main contribution: an admissible adaptive algorithm. We first recall the CAB1 algorithm, which is nonadaptive minimax, and use it as a building block for our new algorithm (Subsection 3.1). This algorithm works in a regime-based fashion. Between successive regimes of doubling lengths, we reset the algorithm and use a new discretization with fewer arms. In order to carry information between the different stages, we use CAB1 in a clever way: besides partitioning the arm space, we add summaries of previous regimes by allowing the algorithm to play according to the empirical distributions of past plays. This is formally described in Subsection 3.2.

A salient difference with all previous approaches is that we zoom out by using fewer and fewer arms. To our knowledge, this is unique, as all other algorithms for bandits zoom in in a way that crucially depends on the regularity parameters. Another important feature of our analysis is that we adapt both to the Hölder exponent α and to the Lipschitz constant L . On a technical level, this is thanks to the fact that we do not explicitly choose a grid of regularity parameters, which means that we implicitly handle all values (L, α) simultaneously.

We first give a regret bound in the known horizon case (Subsection 3.2), then we provide an anytime version and we show that they match the lower bounds of adaptation (Subsection 3.3). Finally Section 4 provides the proof of our main regret bound.

2 Setup, preliminary discussion

2.1 Notation and known results

Let us reintroduce briefly the standard bandit terminology. We consider the arm space $\mathcal{X} = [0, 1]$. The environment sets a reward function $f : \mathcal{X} \rightarrow [0, 1]$. At each time step t , the player chooses an arm $X_t \in \mathcal{X}$, and the environment then displays a reward Y_t such that $\mathbb{E}[Y_t | X_t] = f(X_t)$, independently from the past. We assume that the variables $Y_t - f(X_t)$ are $(1/4)$ -subgaussian conditionnally on X_t ; this is satisfied if the payoffs are bounded in $[0, 1]$ by Hoeffding's lemma.

The objective of the player is to find a strategy that minimizes her *expected cumulative (pseudo-)regret*. If $M(f)$ denotes the maximum value of f , the regret at time T is defined as

$$\bar{R}_T = TM(f) - \mathbb{E} \left[\sum_{t=1}^T Y_t \right] = TM(f) - \mathbb{E} \left[\sum_{t=1}^T f(X_t) \right]. \quad (3)$$

In this paper, we assume that the function f satisfies a Hölder assumption around one of its maxima:

Definition 1. For $\alpha > 0$ and $L > 0$, we denote by $\mathcal{H}(L, \alpha)$ the set of functions that satisfy

$$\exists x^* \in [0, 1] \text{ s.t. } f(x^*) = M(f) \text{ and } \forall x \in [0, 1] \quad |f(x^*) - f(x)| \leq L|x^* - x|^\alpha. \quad (4)$$

We are interested in minimax rates of regret when the mean-payoff function f belongs to these Hölder-type classes, i.e., the quantity $\inf_{\text{algorithms}} \sup_{f \in \mathcal{H}(L, \alpha)} \bar{R}_T$.

MOSS Throughout this paper, we exploit discretization arguments and use a minimax optimal algorithm for finite-armed bandits: MOSS, from Audibert and Bubeck [2]. When run for T rounds on a K -armed bandit problem with $(1/4)$ -subgaussian noise, and when $T \geq K$, its regret is upper-bounded by $18\sqrt{KT}$ (the improved constant is from Garivier et al. [12]).

Non-adaptive minimax rates When the regularity is given to the player, for any α, L and T :

$$0.001 L^{1/(2\alpha+1)} T^{(\alpha+1)/(2\alpha+1)} \leq \inf_{\text{algorithms}} \sup_{f \in \mathcal{H}(L, \alpha)} \bar{R}_T \leq 28 L^{1/(2\alpha+1)} T^{(\alpha+1)/(2\alpha+1)}. \quad (5)$$

This is well-known since Kleinberg [16]. For completeness, we recall how to derive the upper bound in Section 3.1, and the lower bound in Section 2.2.

2.2 Lower bounds: adaptation at usual rates is not possible

Locatelli and Carpentier [21] prove a version of the following theorem; see our reshuffled and slightly improved proof in Appendix F.

Theorem (Variation on Th.3 from [21]). *Let $B > 0$ be a positive number. Let $\alpha, \gamma > 0$ and $L, \ell > 0$ be regularity parameters that satisfy $\alpha \leq \gamma$ and $L \geq \ell$.*

Assume moreover that $2^{-3} 12^\alpha B^{-1} \leq L \leq \ell^{1+\alpha} T^{\alpha/2} 2^{(1+\alpha)(8-2\gamma)}$. If an algorithm is such that $\sup_{f \in \mathcal{H}(\ell, \gamma)} \bar{R}_T \leq B$, then the regret of this algorithm is lower bounded on $\mathcal{H}(L, \alpha)$:

$$\sup_{f \in \mathcal{H}(L, \alpha)} \bar{R}_T \geq 2^{-10} T L^{1/(\alpha+1)} B^{-\alpha/(\alpha+1)}. \quad (6)$$

Remark (Bibliographical note). *Locatelli and Carpentier [21] consider a more general setting where additional margin conditions are exploited. In our setting, we slightly improve their result by dealing with the dependence on the Lipschitz constant, and by removing a requirement on B .*

In a different context, Krishnamurthy et al. [18] show a variation of this bound where the Lipschitz constant is considered, but only in the case where $\alpha = \gamma = 1$, for $\ell = 1$ and $L \geq 1$.

As explained in Locatelli and Carpentier [21] this forbids adaptation at the usual minimax rates over two regularity classes; we recall how in the paragraph that follows Theorem 1. However this is not the end of the story, as one naturally wonders what is the best the player can do.

To further investigate this question, we discuss it asymptotically by considering the rates at which the minimax regret goes to infinity, therefore focusing on the dependence on T . Our main results are completely nonasymptotic, yet we feel the asymptotic analysis of optimality is clearer.

Definition 2. Let $\theta : [0, 1] \rightarrow [0, 1]$ denote a nonincreasing function. We say an algorithm achieves adaptive rates θ if

$$\forall \varepsilon > 0, \forall \alpha, L > 0, \quad \limsup_{T \rightarrow \infty} \frac{\sup_{f \in \mathcal{H}(L, \alpha)} \bar{R}_T}{T^{\theta(\alpha) + \varepsilon}} < +\infty.$$

We include the ε in the definition in order to neglect the potential logarithmic factors.

As rate functions are not always comparable for pointwise order, the good notion of optimality is the standard statistical notion of *admissibility* (akin to ‘‘Pareto optimality’’ for game-theorists).

Definition 3. A rate function is said to be *admissible* if it is achieved by some algorithm, and if no other algorithm achieves strictly smaller rates for pointwise order. An algorithm is *admissible* if it achieves an admissible rate function.

We recall that a function θ' is strictly smaller than θ for pointwise order if $\theta'(\alpha) \leq \theta(\alpha)$ for all α and $\theta'(\alpha_0) < \theta(\alpha_0)$ for at least one value of α_0 .

It turns out we can fully characterize the admissible rate functions by inspecting the lower bounds (6).

Theorem 1. The admissible rate functions are exactly the family

$$\theta_m : \alpha \mapsto \max \left(m, 1 - m \frac{\alpha}{\alpha + 1} \right), \quad m \in [1/2, 1]. \quad (7)$$

This theorem contains two assertions. The lower bound side states that no smaller rate function may be achieved by any algorithm. This side is derived from an asymptotic rewording of lower bound (6), see Proposition 1 stated below (proofs are in Appendix A). The second statement is that the θ_m ’s are indeed achieved by an algorithm, which is the subject of Section 3.2.

Figure 1 illustrates how these admissible rates compare to each other, and to the usual minimax rates.

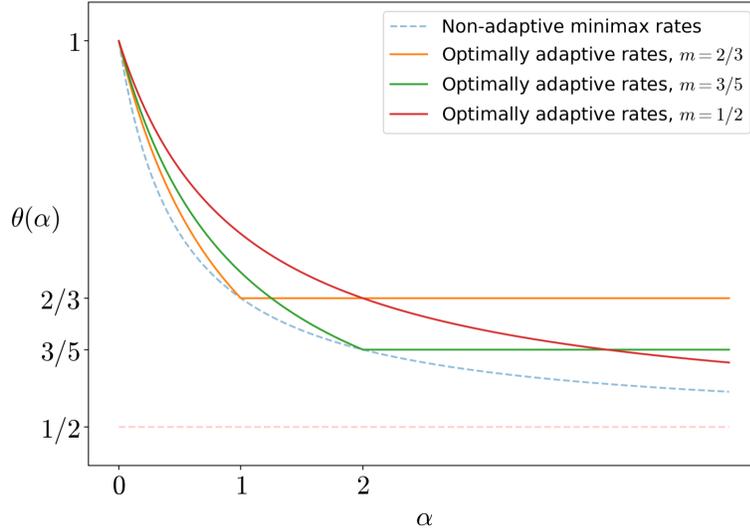


Figure 1: The lower bounds on adaptive rates: plots of the admissible rate functions $\alpha \mapsto \theta_m(\alpha)$. If an algorithm has regret of order $\mathcal{O}(T^{\theta(\alpha)})$, then θ is everywhere above one of these curves.

In particular, we see that reaching the nonadaptive minimax rates for multiple values of α is impossible. Moreover, at $m = (\gamma + 1)/(2\gamma + 1)$, we have $\theta_m(\gamma) = (\gamma + 1)/(2\gamma + 1)$, which is the usual minimax rate (1) when γ is known. This yields an alternative parameterization of the family θ_m : one may choose to parameterize the functions either by their value at infinity $m \in [1/2, 1]$, or by the only point $\gamma \in [0, +\infty]$ at which they coincide with the usual minimax rates function (1).

Proposition 1. Assume an algorithm achieves adaptive rates θ . Then θ satisfies the functional inequation

$$\forall \gamma > 0, \quad \forall \alpha \leq \gamma, \quad \theta(\alpha) \geq 1 - \theta(\gamma) \frac{\alpha}{\alpha + 1}. \quad (8)$$

2.3 Yet can we adapt in some way?

We have described in (7) the minimal rate functions that are compatible with the lower bounds of adaptation: no algorithm can enjoy uniformly better rates. Of course, at this point, the next natural question is whether any of these adaptive rate functions may indeed be reached by an algorithm.

All previous algorithms for continuum-armed bandits require the regularity as an input in some way (see the literature discussion in Section 1.1). Such algorithms are flawed: if the true regularity is underestimated then we only recover the guarantees that correspond to the smaller regularity, which is often far worse than the lower bounds of Theorem 1. More dramatically, if the true regularity is overestimated, then, a priori, no guarantees hold at all.

We prove that all these rate functions may be achieved by a new algorithm. More precisely, if the player wishes to reach one of the lower bounds θ_m , she may select a value of the input accordingly and match the chosen θ_m . This is our main contribution and is described in the next section.

3 An admissible adaptive algorithm and its analysis

We discuss in Section 3.1 how the well-known CAB1 algorithm can be generalized for our purpose. In Section 3.2 we describe our algorithm and the main upper bound on its regret. Section 3.3 is devoted to the anytime version of the algorithm and to a discussion on optimality.

3.1 An abstract version of CAB1 as a building block towards adaptation

We describe a generalization of the CAB1 algorithm from Kleinberg [16], where we include arbitrary measures in the discretization. Although this extension is straightforward, we detail it as we will use this algorithm repeatedly further in this paper. In the original CAB1, the space of arms is discretized into a partition of K subsets, and an algorithm for finite-armed bandits plays on the K midpoints of the sets. Auer et al. [3] replace the midpoints by a random point uniformly chosen in the subset.

We introduce a generic version of this algorithm we call CAB1.1. We consider K arbitrary probability distributions over \mathcal{X} , which we denote by $(\pi_i)_{1 \leq i \leq K}$. Denote also by $\pi(f)$ the expectation of $f(X)$ when $X \sim \pi$. At each time step, the decision maker chooses one distribution, π_{I_t} , and plays an arm picked according to that distribution. By the tower rule, she receives a reward such that

$$\mathbb{E}[Y_t | I_t] = \mathbb{E}[f(X_t) | I_t] = \pi_{I_t}(f).$$

As the player uses a finite-arm algorithm \mathcal{A} to select I_t , the regret she suffers can be decomposed as the sum of two terms (denoting by \tilde{R}_T the expected regret of the finite-armed algorithm):

$$\bar{R}_T = T(M(f) - \max_{i=1, \dots, K} \pi_i(f)) + \tilde{R}_T((\pi_i(f))_{1 \leq i \leq K}; \mathcal{A}). \quad (9)$$

This identity is central to the construction of our algorithm. Using terminology from Auer et al. [3], the first term measures an *approximation error* of the maximum of f , and the other the actual *cost of learning* in the approximate problem. Parameters are chosen to balance these two sources of error.

Algorithm 1 CAB1.1 (Continuum-Armed Bandit, adapted from Kleinberg [16])

- 1: **Input:** T the time horizon, K probability measures over \mathcal{X} denoted by π_1, \dots, π_K , discrete K -armed bandit algorithm \mathcal{A}
 - 2: **for** $t = 1, 2, \dots, T$ **do**
 - 3: Define I_t the arm in $\{1, \dots, K\}$ recommended by \mathcal{A}
 - 4: Play $X_t \in \mathcal{X}$ drawn according to π_{I_t} , and receive Y_t such that $\mathbb{E}[Y_t | X_t] = f(X_t)$
 - 5: Give Y_t as input to \mathcal{A} corresponding to I_t
 - 6: **end for**
-

The canonical example is that for which the space of arms is cut into a partition. Denote by $\mathbf{Disc}(K)$ the family of the uniform measures over the intervals $[(i-1)/K, i/K]$ for $1 \leq i \leq K$. We state this results (and prove it in Appendix A.1) to recall the non-adaptive minimax bound (1).

Proposition 2. *Let $\alpha > 0$ and $L > 1/\sqrt{T}$ be regularity parameters, and define the number of discrete arms $K^* = \min(\lceil L^{2/(2\alpha+1)} T^{1/(2\alpha+1)} \rceil, T)$. Algorithm CAB1.1 run with the uniform discretization $\mathbf{Disc}(K^*)$ and $\mathcal{A} = \text{MOSS}$ enjoys the bound*

$$\sup_{f \in \mathcal{H}(L, \alpha)} \bar{R}_T \leq 28 L^{1/(2\alpha+1)} T^{(\alpha+1)/(2\alpha+1)}.$$

3.2 Memorize past plays, Discretize the arm space, and Zoom Out: the MeDZO algorithm

To achieve adaptation, we combine two tricks: *going from fine to coarser discretizations* while *keeping a summary of past plays in memory*.

Our algorithm works in successive regimes. At each time regime i , we reset the algorithm and start over a new regime of length double the previous one ($\Delta T_i = 2^{p+i}$), and with fewer discrete arms ($K_i = 2^{p+2-i}$). While doing this, we keep in memory the previous plays: in addition to the uniform distributions over the subsets of partitions, we include the empirical measures $\hat{\nu}_j$ of the actions played in the past regimes, for $j < i$.

Algorithm 2 MeDZO (Memorize, Discretize, Zoom Out)

- 1: **Input:** parameter B , time horizon T
 - 2: **Set:** $p = \lceil \log_2 B \rceil$, $K_i = 2^{p+2-i}$ and $\Delta T_i = 2^{p+i}$
 - 3: **for** $i = 1, \dots, p$ **do**
 - 4: For ΔT_i rounds, run algorithm CAB1.1 with the uniform discretization in K_i pieces *and* the empirical measures of the previous plays $\hat{\nu}_j$ for $j < i$; use MOSS as the discrete algorithm.^a
 - 5: **Set:** $\hat{\nu}_i$ the empirical measure of the plays during regime i .
 - 6: **end for**
-

^aNo $\hat{\nu}$ is used for $i = 0$

Appendix C provides a figure illustrating the behavior of the algorithm.

Our construction is based on the following remark. Consider the approximation error suffered during regime i . Denoting the by Π_i the set of measures given to the player during regime i , that is, the uniform measures over the regular K_i -partition and the empirical measures of arms played during the regimes $j < i$, the approximation error is bounded as follows:

$$\Delta T_i \left(M(f) - \mathbb{E} \left[\max_{\pi \in \Pi_i} \pi(f) \right] \right) \leq \Delta T_i (M(f) - \mathbb{E}[\hat{\nu}_j(f)]) = \frac{\Delta T_i}{\Delta T_j} \sum_{t \in \text{Regime } j} (M(f) - \mathbb{E}[f(X_t)]) \quad (10)$$

and this bound is proportional to the regret suffered during regime j . This means that even though we zoom out by using fewer arms, we can make sure that the average approximation error in regime i is less than the regret previously suffered. Moreover, the first discretizations are fine enough to ensure a small regret in the first regimes, thanks to the Hölder property. This argument is formalized in the proof (Lemma 1), and shows that MeDZO maintains a balance between approximation and cost of learning that yields optimal regret.

A surprising fact here is that we go from finer to coarser discretizations during the different phases. Thus, paradoxically, *the algorithm zooms out as time passes*. Note also that although this regime-based approach is reminiscent of the doubling trick, there is an essential difference in that information is carried between the regimes via the distribution of the previous plays.

We first state our central result, a generic bound that holds for any input parameter B . We discuss the optimality of these adaptive bounds in the next subsection.

Theorem 2. *Algorithm 2 run with the knowledge of T and input $B \geq \sqrt{T}$ enjoys the following guarantee: for all $\alpha > 0$ and $L > 0$,*

$$\sup_{f \in \mathcal{H}(L, \alpha)} \bar{R}_T \leq 412 (\log_2 B)^{3/2} \max(B, TL^{1/(\alpha+1)} B^{-\alpha/(\alpha+1)}). \quad (11)$$

We provide some illustrative numerical experiments in Appendix D, comparing the results of MeDZO with other non-adaptive algorithms.

3.3 Discussion: anytime version and admissibility

Anytime version via the doubling trick The dependence of Algorithm 2 on the parameter B makes it horizon-dependent. We use the doubling trick to build an anytime version of the algorithm. At each new doubling-trick regime, we input a value of B that depends on the length of the k -th regime. If it is of length $T^{(k)}$, one typically thinks of $B_k = (T^{(k)})^m$ for some exponent m . In that case, we get the following bound —see the proof and description of the algorithm in Appendix B.

Corollary 1 (Doubling trick version). *Choose $m \in [1/2, 1]$. The doubling-trick version of MeDZO, run with m as sole input (and without the knowledge of T) ensures that for all regularity parameters $\alpha > 0$ and $L > 0$ and for $T \geq 1$*

$$\sup_{f \in \mathcal{H}(L, \alpha)} \bar{R}_T \leq 4000 (\log_2 T^m)^{3/2} \max(T^m, TL^{1/(\alpha+1)}(T^m)^{-\alpha/(\alpha+1)}) = \mathcal{O}((\log T)^{3/2} T^{\theta_m(\alpha)}).$$

Admissibility of Algorithm 2 The next result is a direct consequence of Corollary 1. This echoes the discussion following Theorem 1, and shows that for any input parameter m , the anytime version of MeDZO cannot be improved uniformly for all α .

Corollary 2. *For any $m \in [1/2, 1]$, the doubling trick version of MeDZO (see App. B) with input m achieves rate function θ_m , and is therefore admissible.*

3.4 About the remaining parameter: the $B = \sqrt{T}$ case

Tuning the value of B amounts to selecting one of the minimal curves in Figure 1. Therefore this parameter is a feature of the algorithm, as it allows the player to choose between the possible optimal behaviors. The tuning of this parameter is an unavoidable choice for the player to make.

The next example illustrates well the performance of MeDZO, as it is easily comparable to the usual minimax bounds. Looking at Figure 1, this choice corresponds to $m = 1/2$, i.e., the only choice of parameter that reaches the usual minimax rates as $\alpha \rightarrow \infty$. In other words, if the players wishes to ensure that her regret on very regular functions is of order \sqrt{T} , then she has to pay a polynomial cost of adaptation for not knowing α and that price is exactly the ratio between (1) and (2).

Corollary 3. *Set a horizon T and run Algorithm 2 with $B = \sqrt{T}$. Then for $\alpha > 0$ and $L > 1/\sqrt{T}$,*

$$\sup_{f \in \mathcal{H}(L, \alpha)} \bar{R}_T \leq 146 (\log_2 T)^{3/2} L^{1/(\alpha+1)} T^{(\alpha+2)/(2\alpha+2)}. \quad (12)$$

This is straightforward from Theorem 2, since the inequality $B = \sqrt{T} \leq TL^{1/(\alpha+1)}\sqrt{T}^{-\alpha/(\alpha+1)}$ holds whenever $L \geq 1/\sqrt{T}$. An anytime version of this result can be obtained from Corollary 1.

4 Proof of Theorem 2

Full proof of Theorem 2. Let $\mathcal{F}_t = \sigma(I_1, X_1, Y_1, \dots, I_t, X_t, Y_t)$ be the σ -algebra corresponding to the information available at the end of round t . Define also the transition times $T_i = \sum_{j=1}^i \Delta T_j$ with the convention $T_0 = 0$. Let us first verify that T is smaller than the total length of the regimes. By definition of p , we have $B \leq 2^p < 2B$. Thus $T_p = 2^{p+1}(2^p - 1) \geq 2B(B - 1) > B^2 > T$, and the algorithm is indeed well-defined up to time T .

Consider the regret suffered during the i -th regime $\bar{R}_{T_{i-1}, T_i} := \Delta T_i M(f) - \sum_{t=T_{i-1}+1}^{T_i} \mathbb{E}[f(X_t)]$. We bound this quantity thanks to the decomposition (9), by first conditioning on the past up to time T_{i-1} . Since there are $K_i + i$ discrete actions, the regret bound on MOSS ensures that

$$\mathbb{E} \left[\sum_{t=T_{i-1}+1}^{T_i} (M(f) - f(X_t)) \mid \mathcal{F}_{T_{i-1}} \right] \leq \Delta T_i (M(f) - M_i^*) + 18\sqrt{(K_i + i)\Delta T_i} \quad (13)$$

where $M_i^* = \max\{\pi_j^{(i)}(f) \mid \pi_j^{(i)} \in \mathbf{Disc}(K_i)\} \cup \{\hat{v}_\ell(f) \mid \ell = 0, \dots, i-1\}$. Notice that this bound holds even though M_i^* is a random variable, as the algorithm is completely reset, and the measures $(\hat{v}_j)_{j < i}$ are fixed at time $T_{i-1} + 1$ (i.e., they are $\mathcal{F}_{T_{i-1}}$ -measurable). Integrating once more, we obtain

$$\bar{R}_{T_{i-1}, T_i} \leq \Delta T_i (M(f) - \mathbb{E}[M_i^*]) + 18\sqrt{(K_i + i)\Delta T_i}. \quad (14)$$

Bounding the cost of learning. By definition of K_i and ΔT_i , we have $K_i \Delta T_i = 2^{2p+2} \leq 16B^2$. Therefore, since p and K_i are integers greater than 1, using $a + b - 1 \leq ab$ for positive integers,

$$\sqrt{(K_i + i)\Delta T_i} \leq \sqrt{(K_i + p - 1)\Delta T_i} \leq \sqrt{pK_i \Delta T_i} \leq 4\sqrt{p}B. \quad (15)$$

Bounding the approximation error. The key ingredient for this part is the following fact, that synthesizes the benefits of our construction as hinted in (10) and the surrounding discussion.

Lemma 1. *The total approximation error of MeDZO in regime i is controlled by the Hölder bound on the grid of mesh size $1/K_i$, and by the regret suffered during the previous regimes,*

$$\Delta T_i (M(f) - \mathbb{E}[M_i^*]) \leq \Delta T_i \min \left(L \frac{1}{K_i^\alpha}, \min_{j < i} \left(\frac{\bar{R}_{T_{j-1}, T_j}}{\Delta T_j} \right) \right) \quad (16)$$

Proof. This derives easily from the construction of the algorithm, i.e., from the definition of M_i^* . Considering an interval in the regular K_i -partition that contains a maximum of f , by the Hölder property, $M(f) - M_i^* \leq L/K_i^\alpha$. For the second minimum, as described in Eq. (10), for $j < i$,

$$M(f) - M_i^* \leq M(f) - \hat{\nu}_j(f) = \frac{1}{\Delta T_j} \sum_{t=T_{j-1}+1}^{T_j} (M(f) - f(X_t)).$$

Taking an expectation, \bar{R}_{T_{j-1}, T_j} appears, and we conclude by taking the minimum over j . \square

Remember that since $K_i \Delta T_i = 2^{2p+2}$, we have $L \Delta T_i / K_i^\alpha = L 2^{2p+2} / K_i^{1+\alpha}$. Therefore, the first bound on the approximation error in (16) increases with i , as K_i decreases with i . Denote by i_0 the last time regime i for which

$$L \frac{\Delta T_{i_0}}{K_{i_0}^\alpha} \leq B. \quad (17)$$

If this is never satisfied, i.e., not even for $i = 1$, then $L 2^{2p+1} / 2^{\alpha(p+1)} > B$ which yields, using $B \leq 2^p \leq 2B$, that $4LB \geq 2^{\alpha+1} B^\alpha B$ and then $L > B^\alpha / 2$. In that case, $L^{1/(\alpha+1)} B^{-\alpha/(\alpha+1)} \geq 1$ and the total regret bound (11) is true as it is weaker than the trivial bound $R_T \leq T$.

Hence we may assume that $i_0 \geq 1$ is well defined. By comparing i to i_0 , we now show the inequality

$$\sum_{i=1}^p \Delta T_i (M(f) - \mathbb{E}[M_i^*]) \leq \sum_{i=1}^{i_0} B + \sum_{i=i_0+1}^p 2(1 + 72\sqrt{p}) \Delta T_i L^{1/(\alpha+1)} B^{-\alpha/(\alpha+1)}. \quad (18)$$

For all $i \leq i_0$ the approximation error is smaller than the first argument of the minimum in (16), and this term is smaller than B . Therefore $\Delta T_i (M(f) - \mathbb{E}[M_i^*]) \leq B$. In particular, this together with (14) and (15) implies that the total regret suffered during regime i_0 is $\bar{R}_{T_{i_0-1}, T_{i_0}} \leq (1 + 72\sqrt{p})B$.

For the later time regimes $i_0 < i \leq p$, we use the fact that preceding empirical measures were kept as discrete actions, and in particular the one of the i_0 -th regime: (16) instantiated with $j = i_0$ yields

$$\Delta T_i (M(f) - \mathbb{E}[M_i^*]) \leq \Delta T_i \frac{\bar{R}_{T_{i_0-1}, T_{i_0}}}{\Delta T_{i_0}} \leq (1 + 72\sqrt{p}) \Delta T_i \frac{B}{\Delta T_{i_0}}. \quad (19)$$

Solving equations $L \Delta T_{i_0} / K_{i_0}^\alpha \approx B \approx 4\sqrt{\Delta T_{i_0} K_{i_0}}$, we get $B / \Delta T_{i_0} \leq 2 L^{1/(\alpha+1)} B^{-\alpha/(\alpha+1)}$, (details are in Appendix A.4). Therefore for $i_0 < i \leq p$, using (19),

$$\Delta T_i (M(f) - \mathbb{E}[M_i^*]) \leq 2(1 + 72\sqrt{p}) \Delta T_i L^{1/(\alpha+1)} B^{-\alpha/(\alpha+1)},$$

and we obtain (18) by summing over i .

Conclusion We conclude with some crude boundings. First, as $i_0 \leq p$ and the sum of the ΔT_i 's is smaller than T , the total approximation error is less than $pB + 2(1 + 72\sqrt{p})TL^{1/(\alpha+1)}B^{-\alpha/(\alpha+1)}$. Let us include the cost of learning, which is smaller than $72p\sqrt{p}B$ and conclude, using $a + b \leq \max(a, b)$

$$\begin{aligned} \bar{R}_T &\leq 2(1 + 72\sqrt{p})TL^{1/(\alpha+1)}B^{-\alpha/(\alpha+1)} + pB + 72p^{3/2}B \\ &= 2(1 + 72\sqrt{p})TL^{1/(\alpha+1)}B^{-\alpha/(\alpha+1)} + p(1 + 72\sqrt{p})B \\ &\leq \left(2(1 + 72\sqrt{p}) + p(1 + 72\sqrt{p}) \right) \max(B, TL^{1/(\alpha+1)}B^{-\alpha/(\alpha+1)}) \end{aligned} \quad (20)$$

from which the desired bound follows, using $1 \leq p$, and $p \leq 2 \log_2 B$ and $4(1 + 72\sqrt{2}) \leq 412$. \square

5 Further considerations

Local regularity assumption Theorem 2 holds under a relaxed smoothness assumption, namely that the function satisfies the Hölder condition only in a small cell containing the maximum. By looking carefully at the proof, we observe that the condition is only required up to the i_0 -th epoch (defined in (17)), at which the size of the cells in the discretization is of order $1/K_{i_0} \approx (LB)^{-1/(1+\alpha)}$. Therefore we only need condition (4) to be satisfied for points x in an interval of size $(LB)^{-1/(1+\alpha)}$ around the maximum.

Higher dimension Our results can be generalized to functions $[0, 1]^d \rightarrow [0, 1]$ that are $\|\cdot\|_\infty$ -Hölder. For MeDZO to be well-defined, take $K_i = 2^{d(p+2-i)}$ and $\Delta T_i = 2^{d(p+i)}$, with $p \approx (\log B)/d$. The bounds are similar to their one-dimensional counterparts, up to replacing α by α/d in the exponents, but the constants are deteriorated by a factor that is exponential in d . The bound in Theorem 2 changes to $\max(B, L^{d/(\alpha+d)} T B^{-\alpha/(\alpha+d)})$.

Acknowledgements

We would like to thank Gilles Stoltz and Pascal Massart for their valuable comments and suggestions.

References

- [1] Rajeev Agrawal. The Continuum-Armed Bandit Problem. *SIAM Journal on Control and Optimization*, 33(6):1926–1951, 1995. ISSN 0363-0129, 1095-7138. doi: 10.1137/S0363012992237273. URL <http://epubs.siam.org/doi/10.1137/S0363012992237273>.
- [2] Jean-Yves Audibert and Sébastien Bubeck. Minimax policies for adversarial and stochastic bandits. In *COLT*, pages 217–226.
- [3] Peter Auer, Ronald Ortner, and Csaba Szepesvári. Improved rates for the stochastic continuum-armed bandit problem. In *International Conference on Computational Learning Theory*, pages 454–468. Springer.
- [4] Lucien Birgé and Pascal Massart. Estimation of Integral Functionals of a Density. *The Annals of Statistics*, 23(1):11–29, 1995. URL <https://doi.org/10.1214/aos/1176324452>.
- [5] Sébastien Bubeck, Gilles Stoltz, and Jia Yuan Yu. Lipschitz bandits without the Lipschitz constant. In *International Conference on Algorithmic Learning Theory*, pages 144–158. Springer.
- [6] Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure exploration in finitely-armed and continuous-armed bandits. *Theoretical Computer Science*, 412(19):1832–1852, 2011. ISSN 03043975. doi: 10.1016/j.tcs.2010.12.059. URL <https://linkinghub.elsevier.com/retrieve/pii/S030439751000767X>.
- [7] Sébastien Bubeck, Rémi Munos, Gilles Stoltz, and Csaba Szepesvári. X-armed bandits. *Journal of Machine Learning Research*, 12:1655–1695, 2011.
- [8] Adam D. Bull. Adaptive-treed bandits. *Bernoulli*, 21(4):2289–2307, 2015. doi: 10.3150/14-BEJ644. URL <https://doi.org/10.3150/14-BEJ644>.
- [9] T. Tony Cai. Minimax and Adaptive Inference in Nonparametric Function Estimation. *Statistical Science*, 27(1):31–50, 2012. doi: 10.1214/11-STS355. URL <https://doi.org/10.1214/11-STS355>.
- [10] T. Tony Cai and Mark G. Low. On adaptive estimation of linear functionals. 33(5):2311–2343. doi: 10.1214/009053605000000633. URL <https://doi.org/10.1214/009053605000000633>.
- [11] Pierre-Arnaud Coquelin and Rémi Munos. Bandit algorithms for tree search. In *Uncertainty in Artificial Intelligence*, 2007.
- [12] Aurélien Garivier, Hédi Hadiji, Pierre Menard, and Gilles Stoltz. KL-UCB-switch: Optimal regret bounds for stochastic bandits from both a distribution-dependent and a distribution-free viewpoints. URL <http://arxiv.org/abs/1805.05071>.
- [13] Aurélien Garivier, Pierre Ménard, and Gilles Stoltz. Explore First, Exploit Next: The True Shape of Regret in Bandit Problems. *Mathematics of Operations Research*, 2018. ISSN 0364-765X. doi: 10.1287/moor.2017.0928. URL <https://pubsonline.informs.org/doi/10.1287/moor.2017.0928>.

- [14] Jean-Bastien Grill, Michal Valko, and Rémi Munos. Black-box optimization of noisy functions with unknown smoothness. In *Advances in Neural Information Processing Systems*, pages 667–675.
- [15] Olav Kallenberg. *Foundations of modern probability*. Springer Science & Business Media, 2006.
- [16] Robert Kleinberg. Nearly Tight Bounds for the Continuum-armed Bandit Problem. In *Proceedings of the 17th International Conference on Neural Information Processing Systems*, pages 697–704. MIT Press.
- [17] Robert Kleinberg, Aleksandrs Slivkins, and Eli Upfal. Bandits and Experts in Metric Spaces. abs/1312.1277. URL <http://arxiv.org/abs/1312.1277>.
- [18] Akshay Krishnamurthy, John Langford, Aleksandrs Slivkins, and Chicheng Zhang. Contextual Bandits with Continuous Actions: Smoothing, Zooming, and Adapting. URL <http://arxiv.org/abs/1902.01520>.
- [19] Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*, volume Bubeck, Sébastien and Munos, Rémi and Stoltz, Gilles and Szepesvári, Csaba. Cambridge University Press (preprint), 2019.
- [20] O. Lepskii. On a Problem of Adaptive Estimation in Gaussian White Noise. *Theory of Probability & Its Applications*, 35(3):454–466, 1991. ISSN 0040-585X. doi: 10.1137/1135065. URL <https://epubs.siam.org/doi/10.1137/1135065>.
- [21] Andrea Locatelli and Alexandra Carpentier. Adaptivity to Smoothness in X-armed bandits. In *Conference on Learning Theory*, pages 1463–1492.
- [22] Pascal Massart. *Concentration Inequalities and Model Selection*. Ecole d’Eté de Probabilités de Saint-Flour XXXIII - 2003. Springer Berlin Heidelberg. ISBN 978-3-540-48503-2. URL <https://doi.org/10.1007/978-3-540-48503-2>.
- [23] Xuedong Shang, Emilie Kaufmann, and Michal Valko. General parallel optimization without a metric. In *Algorithmic Learning Theory*, volume 98. URL <https://hal.inria.fr/hal-02047225>.
- [24] Michal Valko, Alexandra Carpentier, and Rémi Munos. Stochastic Simultaneous Optimistic Optimization. In *International Conference on Machine Learning*.

Supplementary Material for "Polynomial Cost of Adaptation for \mathcal{X} -Armed Bandits"

A Omitted proofs

A.1 Proposition 2: Regret bound for non-adaptive CAB1.1

This proof is a straightforward application of the Hölder bound and of the bound of MOSS, together with the approximation/cost of learning decomposition of the regret. Some extra care is needed to handle the boundary cases.

Proof of Proposition 2. Choose $f \in \mathcal{H}(L, \alpha)$. Let us denote by i^* an integer such that there exists an optimal arm x^* in the interval $[(i^* - 1)/K^*, i^*/K^*]$. By the Hölder assumption

$$\frac{1}{K^*} \int_{(i^*-1)/K^*}^{i^*/K^*} (f(x^*) - f(x)) dx \leq L \left(\frac{1}{K^*} \right)^\alpha,$$

and this upper bounds the approximation error of the discretization. Moreover, since $T \geq K^*$, the cost of learning is smaller than $18\sqrt{K^*T}$. Thus by (9)

$$\bar{R}_T \leq TL \left(\frac{1}{K^*} \right)^\alpha + 18\sqrt{K^*T}.$$

K^* was chosen to minimize this quantity. We distinguish cases depending on the value of K^* .

If $1 < K^* < T$, then $L^{2/(2\alpha+1)}T^{1/(2\alpha+1)} \leq K^* \leq 2L^{2/(2\alpha+1)}T^{1/(2\alpha+1)}$ (the bound $\lceil x \rceil \leq 2x$, which is valid when $x \geq 1$, is more practical to handle the multiplicative constants), we deduce the upper bound:

$$(1 + 18\sqrt{2}) L^{1/(2\alpha+1)} T^{(\alpha+1)/(2\alpha+1)}.$$

Since we assumed that $L > 1/\sqrt{T}$, we have always $K^* > 1$. Therefore the last case to consider is if $K^* = T$. Then $L^{2/(2\alpha+1)}T^{1/(2\alpha+1)} \geq T/2$ and thus $L \geq 2^{-(2\alpha+1)/2}T^\alpha$. In this case $L^{1/(2\alpha+1)}T^{(\alpha+1)/(2\alpha+1)} \geq (\sqrt{2}/2)T$ and the claimed bound is met since in that case, we have by a trivial bound $\bar{R}_T \leq T \leq \sqrt{2}L^{1/(2\alpha+1)}T^{(\alpha+1)/(2\alpha+1)}$. \square

A.2 Proposition 1: Lower bound on the adaptive rates

Proof of Proposition 1. Choose α, γ such that $\alpha \leq \gamma$, and $\varepsilon > 0$. Set $L > 0$. There exist constants c_1 and c_2 (depending on L, α, γ and ε) such that for T large enough,

$$\sup_{f \in \mathcal{H}(L, \alpha)} \bar{R}_T \leq c_1 T^{\theta(\alpha) + \varepsilon} \quad \text{and} \quad \sup_{f \in \mathcal{H}(L, \gamma)} \bar{R}_T \leq c_2 T^{\theta(\gamma) + \varepsilon}.$$

Moreover, for T large enough, the assumptions for lower bound (6) hold. Hence applying the lower bound with $B = c_2 T^{\theta(\gamma) + \varepsilon}$, for some constant c_3 :

$$c_1 T^{\theta(\alpha) + \varepsilon} \geq 0.0001 T (c_2 T^{\theta(\gamma) + \varepsilon})^{-\alpha/(\alpha+1)} \geq c_3 T^{1 - \theta(\gamma)\alpha/(\alpha+1) - \varepsilon\alpha/(\alpha+1)}$$

Since the above inequality holds for any T sufficiently large, this implies that for all $\varepsilon > 0$

$$\theta(\alpha) + \varepsilon \geq 1 - \theta(\gamma) \frac{\alpha}{\alpha+1} - \varepsilon \frac{\alpha}{\alpha+1},$$

which yields the desired result as $\varepsilon \rightarrow 0$. \square

A.3 Theorem 1: Admissible rate functions

We prove here that all the admissible rate functions belong to the family (θ_m) , by relying on Proposition 1. The proof is done through a careful inspection of the functional inequation defining the lower bound.

Proof of Theorem 1. First of all, by Corollary 2, the appropriately tuned MeDZO may achieve all the θ'_m s. Thus we are left to prove the lower bound side, i.e., that all the admissible rate functions belong to the family θ_m .

The best way to see this is to first notice that for θ nonincreasing and positive, the inequation in Proposition 1 is equivalent to

$$\forall \alpha > 0, \quad \theta(\alpha) \geq 1 - \theta(\infty) \frac{\alpha}{\alpha + 1}. \quad (21)$$

Notice that taking $\gamma = +\infty$ is always valid in what follows, as θ is assumed to be nonincreasing and lower bounded by $1/2$. Now if θ satisfies (8), then it satisfies (21) by taking $\gamma = +\infty$. For the converse, consider $\alpha \leq \gamma$, then $\theta(\gamma) \geq \theta(\infty)$, thus $1 - \theta(\infty)\alpha/(\alpha + 1) \geq 1 - \theta(\gamma)\alpha/(\alpha + 1)$.

Now consider an admissible θ . Since θ is achieved by some algorithm, by Proposition 1 and the remark above, it satisfies Eq. (21). As θ is nonincreasing, and by Eq. (21), we have $\theta(\alpha) \geq \theta(\infty)$ and $\theta(\alpha) \geq 1 - \theta(\infty)\alpha/(\alpha + 1)$. In other words, $\theta \geq \theta_{m_\theta}$, where $m_\theta = \theta(\infty) \in [1/2, 1]$. By the admissibility of θ , this implies that $\theta = \theta_{m_\theta}$. \square

A.4 Calculations in the proof of Theorem 2

Details on (18), in the proof of Theorem 2. By definition of i_0 , and since we assumed that $i_0 < p$

$$B \leq L \frac{\Delta T_{i_0+1}}{K_{i_0+1}^\alpha},$$

i.e., using $K_{i_0} \Delta T_{i_0} = 2^{2p+2}$,

$$B \leq 2^{1+\alpha} L \frac{\Delta T_{i_0}}{K_{i_0}^\alpha} = 2^{1+\alpha} L (\Delta T_{i_0})^{1+\alpha} 2^{-(2p+2)\alpha}.$$

From this we deduce, using $2^p \geq B$ for the second inequality,

$$(\Delta T_{i_0})^{(1+\alpha)} \geq 2^{-1-\alpha} B L^{-1} 2^{(2p+2)\alpha} \geq 2^{-1+\alpha} L^{-1} B^{2\alpha+1}.$$

Hence, using $2^{(\alpha-1)/(\alpha+1)} \geq 1/2$, we obtain $\Delta T_{i_0} \geq (1/2) L^{-1/(\alpha+1)} B^{(2\alpha+1)/(\alpha+1)}$, thus $B/\Delta T_{i_0} \leq 2 L^{1/(\alpha+1)} B^{-\alpha/(\alpha+1)}$. \square

B Anytime-MeDZO and proof

The doubling trick is the most standard way of converting non-anytime algorithms into anytime algorithms, when the regret bound is polynomial. It consists in taking fresh starts of the algorithm over a grid of dyadic times. The implementation of the trick is straightforward in our case.

Algorithm 3 Doubling trick MeDZO

- 1: **Input:** parameter $m \in [1/2, 1]$;
 - 2: **for** $i = 0, \dots$ **do**
 - 3: Run MeDZO (Alg. 2) with input $B = 2^{im}$ for 2^i rounds
 - 4: **end for**
-

Corollary (Doubling trick version). *Choose $m \in [1/2, 1]$. The doubling-trick version of MeDZO, run with m as sole input (and without the knowledge of T) ensures that for all regularity parameters $\alpha > 0$ and $L > 0$ and for $T \geq 1$*

$$\sup_{f \in \mathcal{H}(L, \alpha)} \bar{R}_T \leq 4000 (\log_2 T^m)^{3/2} \max(T^m, T L^{1/(\alpha+1)} (T^m)^{-\alpha/(\alpha+1)}) = \mathcal{O}((\log T)^{3/2} T^{\theta_m(\alpha)}).$$

As the regret bound is not exactly of the form cT^θ , we work with the polynomial version of the bound on the regret of MeDZO, equation (20), for the doubling trick to be effective. Obviously the value of the constant in the bound is not our main focus, but we still write it explicitly as it shows that there is no hidden dependence on the various parameters.

Proof. By (20), with $p_i = \lceil \log_2 2^{im} \rceil \leq 1 + \log_2 2^{im}$, in the i -th doubling trick regime, the cumulative regret is bounded by

$$2(1 + 72\sqrt{1 + \log_2 2^{im}})2^i L^{1/(\alpha+1)}(2^{im})^{-\alpha/(\alpha+1)} + (1 + \log_2 2^{im})(1 + 72\sqrt{1 + \log_2 2^{im}})2^{im}$$

Now since

$$\sum_{i=0}^{\lceil \log_2 T \rceil} 2^i = 2^{\lceil \log_2 T \rceil + 1} - 1 \geq 2T - 1 \geq T,$$

there are always less than $\lceil \log_2 T \rceil$ full regimes. Therefore, using $\log_2 2^{im} \leq \log_2 T^m$, and summing over the regimes, the first part of this sum is bounded by

$$\begin{aligned} & 2(1 + 72\sqrt{2\log_2 T^m})L^{1/(\alpha+1)} \sum_{i=0}^{\lceil \log_2 T \rceil} 2^{i(1-m\alpha/(\alpha+1))} \\ & \leq 2(1 + 72\sqrt{2\log_2 T^m})L^{1/(\alpha+1)} \frac{2^{(\lceil \log_2 T \rceil + 1)(1-m\alpha/(\alpha+1))}}{2^{1-m\alpha/(\alpha+1)} - 1} \\ & \leq 2(1 + 72\sqrt{2})\sqrt{\log_2 T^m} L^{1/(\alpha+1)} \frac{2^{2(1-m\alpha/(\alpha+1))}}{\sqrt{2} - 1} T(T^m)^{-\alpha/(\alpha+1)} \\ & \leq 2(1 + 72\sqrt{2})\sqrt{\log_2 T^m} L^{1/(\alpha+1)} \frac{4}{\sqrt{2} - 1} T(T^m)^{-\alpha/(\alpha+1)} \end{aligned}$$

where we used $2^{\lceil \log_2 T \rceil} \leq 2T$; we also used the fact that since $m \geq 1/2$, we always have the inequality $1 - m\alpha/(\alpha + 1) \geq 1/2$ to bound the denominator. Similarly, the second part is bounded by

$$2(1 + 72\sqrt{2})(\log_2 T^m)^{3/2} \sum_{i=0}^{\lceil \log_2 T \rceil} 2^{im} \leq 2(1 + 72\sqrt{2})(\log_2 T^m)^{3/2} \frac{4}{\sqrt{2} - 1} T^m.$$

All in all, we obtain the same minimax guarantees as if we had known the time horizon in advance, but with an extra multiplicative factor of $4/(\sqrt{2} - 1) \approx 9,66$. \square

C Illustration

In this section we provide a figure to illustrate the behavior of MeDZO in a schematic example.

MeDZO starts by playing on a fine discretization with a size of order \sqrt{T} , but for a short length of time, of order \sqrt{T} . At the end of the first epoch, it memorizes the empirical distribution of the arms played; then it runs a new instance of CAB1.1 with both the coarser discretization, and the memorized action. This process is repeated until the time horizon is reached.

The payoffs of the memorized actions increase until the size of the discretization reaches a critical value; after that they fluctuate. Therefore MeDZO manages to maintain a regret of order the approximation error at this critical discretization, multiplied by T .

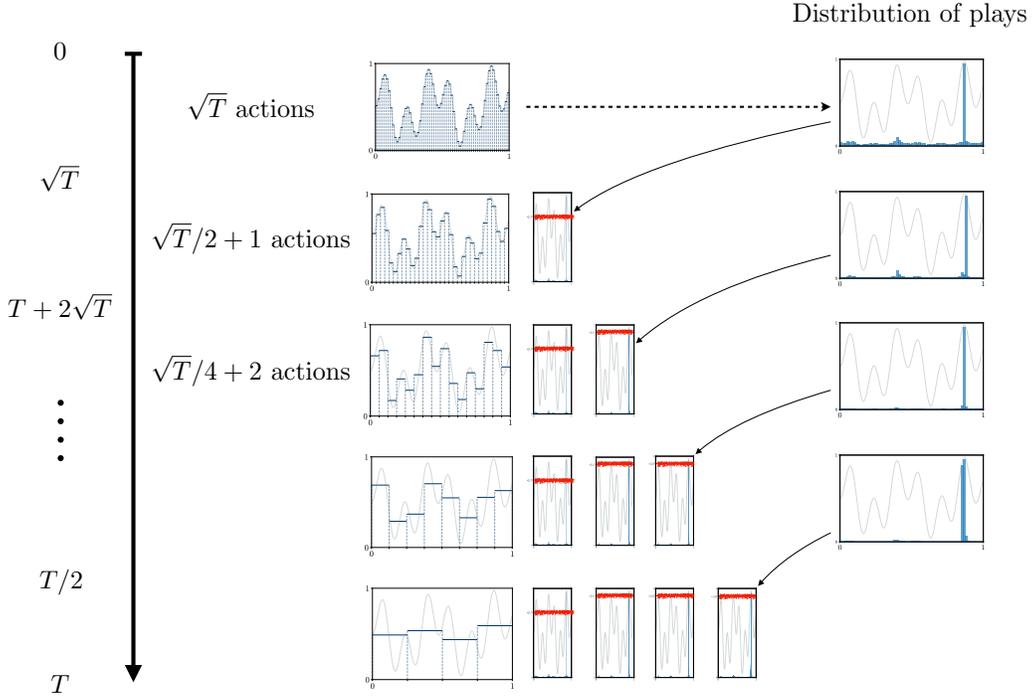


Figure 2: Behavior of MeDZO on a schematic drawing. The expected payoffs of the memorized actions are displayed in red; those from the usual discretization are in blue.

D Numerical experiments

This section contains some numerical experiments comparing the regrets of algorithms that require the knowledge of the smoothness, against MeDZO.

We examine bandit problems defined by their mean-payoff functions and gaussian $\mathcal{N}(0; 1/4)$ noise. The functions considered are $f : x \mapsto (1/2) \sin(13x) \sin(27x) + 0.5$ taken from Bubeck et al. [7], $g : x \mapsto \max(3.6x(1-x), 1 - 1/0.05|x - 0.05|)$ adapted from Coquelin and Munos [11] and the Garland function $x \mapsto x(1-x)(4 - \sqrt{|\sin(60x)|})$, which we took from Valko et al. [24]. The functions are plotted in Figure 3.

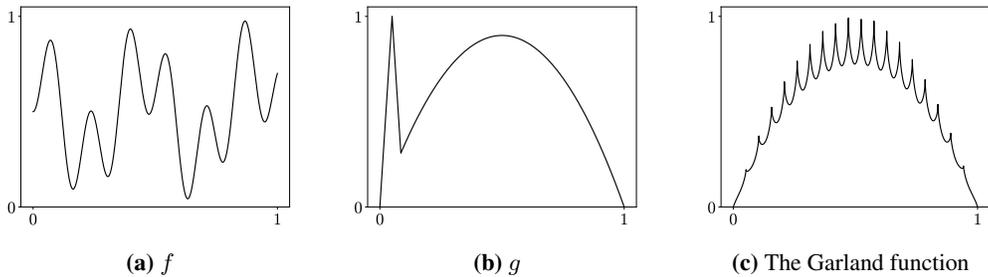


Figure 3: Problems considered

The algorithms we compare are SR from Locatelli and Carpentier [21], and CAB1 from Kleinberg [16] with MOSS as the discrete algorithm. SR takes directly the smoothness α as an input, and assumes $L = 1$. For CAB1, we compute the optimal discretization size for $L = 1$ and varying α .

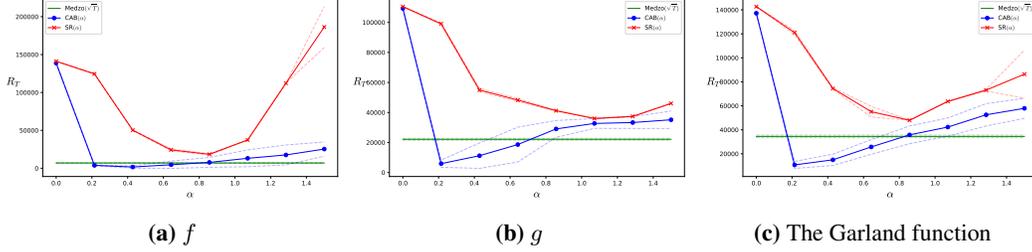


Figure 4: Regrets of MeDZO, and of SR and CAB1 run with different values of the smoothness parameter.

In Figure 4 we plot the cumulative regret of the algorithms after a time horizon $T = 300000$, for varying values of the assumed smoothness. For each problem, MeDZO was run only once, as it does not need to know the smoothness. The regret was averaged over $N = 75$ runs, and the dotted curves represent \pm one standard deviation.

We recall that minimax guarantees are worst-case guarantees, therefore comparing algorithms on a single problem can only serve as an empirical illustration.

As expected, the regrets of both SR and CAB1 depend on some careful tuning of the input parameter, determined by the smoothness. The optimal tuning is unclear, and seems to vary on the algorithm. MeDZO, on the other hand, obtains reasonable regret with no tuning. Surprisingly, CAB1 with overestimated smoothness seems to behave quite well, although the large variance sometimes makes it difficult to distinguish the results. Recall that MeDZO is the only algorithm with theoretical guarantees for high values of α .

E About simple regret

In this section, we consider the case of *simple regret*, which complements the discussion about adaptation to smoothness in sequential optimization procedures. We write out how to achieve adaptation at usual rates for simple regret under Hölder smoothness assumptions. We do not claim novelty here, as adaptive strategies have already been used for simple regret under more sophisticated regularity conditions (see, e.g., Grill et al. [14], Shang et al. [23] and a sketched out procedure in Locatelli and Carpentier [21]); however, we feel the details deserve to be written out in this simpler setting.

Let us recall the definition of simple regret. In some cases, we may only require that the algorithm outputs a recommendation \tilde{X}_T at the end of the T rounds, with the aim of minimizing the simple regret, defined as

$$\bar{r}_T = M(f) - \mathbb{E}[f(\tilde{X}_T)].$$

This setting is known under various names, e.g., pure exploration, global optimization or black-box optimization. As noted in Bubeck et al. [6], minimizing the simple regret is easier than minimizing the cumulative regret in the sense that if the decision-maker chooses a recommendation uniformly among the arms played X_1, \dots, X_T , then

$$\bar{r}_T = M(f) - \frac{1}{T} \sum_{t=1}^T \mathbb{E}[f(X_t)] = \frac{\bar{R}_T}{T}. \quad (22)$$

The minimax rates of simple regret over Hölder classes $\mathcal{H}(L, \alpha)$ are lower bounded by $\Omega(L^{1/(2\alpha+1)}T^{-\alpha/(2\alpha+1)})$, which are exactly the rates for cumulative regret divided by T (see Locatelli and Carpentier [21] for a proof of the lower bound). Consequently, at known regularity, any minimax optimal algorithm for cumulative regret automatically yields a minimax recommendation for simple regret via (22).

When the smoothness is unknown, the situation turns out to be quite different. Adapting to the Hölder parameters can be done at only a (poly-)logarithmic cost for simple regret, contrasting with the polynomial cost of adaptation of cumulative regret. This can be achieved thanks to a very general and simple cross-validation scheme defined in Shang et al. [23], named General Parallel Optimization.

Algorithm 4 GPO (General Parallel Optimization) for Hölder minimax adaptation

```

1: Input: time horizon  $T \geq 8$ 
2: Set:  $p = \lceil \log_2 T \rceil$  and define  $K_i = 2^i$  for  $i = 1, \dots, p$ 
3: for  $i = 1, \dots, p$  do // Exploration
4:   For  $\lfloor T/(2p) \rfloor$  rounds, run algorithm CAB1.1 with the discretization in  $K_i$  pieces; use MOSS
   as the discrete algorithm
5:   Define output recommendation  $\tilde{X}^{(i)}$ , uniformly chosen among the  $\lfloor T/(2p) \rfloor$  arms played
6: end for
7: for  $i = 1, \dots, p$  do // Cross-validation
8:   Play  $\lfloor T/(2p) \rfloor$  times each recommendation  $\tilde{X}^{(i)}$  and compute the average reward  $\hat{\mu}^{(i)}$ 
9: end for
10: return A recommendation  $\tilde{X}_T = \tilde{X}^{(\hat{i})}$  with  $\hat{i} \in \arg \max \hat{\mu}^{(i)}$ 

```

The next result shows that the player obtains the same simple regret bounds as when the smoothness is known (up to logarithmic factors).

Theorem 3. *GPO with CAB1.1 as a sub-algorithm (Alg. 4) achieves, given $T \geq 8$ and without the knowledge of α and L , for all $\alpha > 0$ and $L \geq 2^{\alpha+1/2} \sqrt{\lceil \log_2 T \rceil} / T$ the bound*

$$\sup_{f \in \mathcal{H}(L, \alpha)} \bar{r}_T \leq \left(54 + \frac{\sqrt{\pi}}{2} \log_2 T\right) L^{1/(2\alpha+1)} \left(\frac{\lceil \log_2 T \rceil}{T}\right)^{\alpha/(2\alpha+1)} = \tilde{\mathcal{O}}\left(L^{1/(2\alpha+1)} T^{-\alpha/(2\alpha+1)}\right).$$

The $\tilde{\mathcal{O}}$ notation hides the $\log T$ factors, and the assumption that $T \geq 8$ is needed to ensure that $T/(2p) = T/(2 \lceil \log_2 T \rceil) \geq 1$: otherwise the algorithm itself is ill-defined.

Proof. Let $f \in \mathcal{H}(L, \alpha)$ denote a mean-payoff function. Once again we decompose the error of the algorithm into two sources. The simple regret is the sum of the regret of the best recommendation among the p received, r_{\min} , and of a cross-validation error, r_{CV} ,

$$M(f) - \mathbb{E}[f(\tilde{X}_T)] = \underbrace{\min_{i=1, \dots, p} \left(M(f) - \mathbb{E}[f(\tilde{X}^{(i)})] \right)}_{r_{\min}} + \underbrace{\max_{i=1, \dots, p} \left(\mathbb{E}[f(\tilde{X}^{(i)})] - \mathbb{E}[f(\tilde{X}_T)] \right)}_{r_{\text{CV}}}. \quad (23)$$

We now show that $r_{\text{CV}} \leq p^{3/2} \sqrt{\pi/(4T)}$, by detailing an argument that is sketched in the proof of Thm. 3 in Shang et al. [23]. Denote by $\hat{\mu}^{(i)}$ the empirical reward associated to recommendation i , and $\hat{i} = \arg \max \hat{\mu}^{(i)}$, so that $\tilde{X}_T = \tilde{X}^{(\hat{i})}$. Then for any fixed i , by the tower rule,

$$\mathbb{E}[\hat{\mu}^{(i)}] = \mathbb{E}\left[\mathbb{E}[\hat{\mu}^{(i)} \mid \tilde{X}^{(i)}]\right] = \mathbb{E}\left[f(\tilde{X}^{(i)})\right]. \quad (24)$$

Therefore, by the above remarks, and since $\hat{\mu}^{(i)} \leq \hat{\mu}^{(\hat{i})}$,

$$\mathbb{E}\left[f(\tilde{X}^{(i)})\right] - \mathbb{E}\left[f(\tilde{X}_T)\right] = \mathbb{E}\left[\hat{\mu}^{(i)} - f(\tilde{X}^{(\hat{i})})\right] \leq \mathbb{E}\left[\hat{\mu}^{(\hat{i})} - f(\tilde{X}^{(\hat{i})})\right].$$

We have to be careful here, as \hat{i} is a random index that depends on the random variables $\hat{\mu}^{(i)}$'s: we cannot apply directly the tower rule as in (24). To deal with this, let us use an integrated union bound. Denote by $(\cdot)^+$ the positive part function, then

$$\mathbb{E}\left[\hat{\mu}^{(\hat{i})} - f(\tilde{X}^{(\hat{i})})\right] \leq \mathbb{E}\left[\left(\hat{\mu}^{(\hat{i})} - f(\tilde{X}^{(\hat{i})})\right)^+\right] \leq \sum_{j=1}^p \mathbb{E}\left[\left(\hat{\mu}^{(j)} - f(\tilde{X}^{(j)})\right)^+\right],$$

and we are back to handling empirical means of i.i.d. random variables. For each j , the reward given $\tilde{X}^{(j)}$ is $(1/4)$ -subgaussian. Therefore, as $\hat{\mu}^{(i)}$ is the empirical mean of $n = \lfloor T/(2p) \rfloor$ plays of the same arm $\tilde{X}^{(j)}$, this mean $\hat{\mu}^{(i)}$ is $(1/(4n))$ -subgaussian conditionally on $\tilde{X}^{(j)}$ and thus for all $\varepsilon > 0$

$$\mathbb{P}\left[\hat{\mu}^{(j)} - f(\tilde{X}^{(j)}) \geq \varepsilon\right] \leq e^{-2n\varepsilon^2}.$$

Hence by integrating over $\varepsilon \in [0, +\infty)$, using Fubini's theorem, a change of variable $x = \sqrt{4n\varepsilon}$ (and using the fact that $\lfloor T/(2p) \rfloor \geq T/(4p)$ as $T/(2p) \geq 1$):

$$\begin{aligned} \mathbb{E} \left[\left(\hat{\mu}^{(j)} - f(\tilde{X}^{(j)}) \right)^+ \right] &= \int_0^{+\infty} \mathbb{P} \left[\hat{\mu}^{(j)} - f(\tilde{X}^{(j)}) \geq \varepsilon \right] d\varepsilon \\ &\leq \int_0^{+\infty} e^{-2n\varepsilon^2} d\varepsilon = \frac{1}{\sqrt{4n}} \int_0^{+\infty} e^{-x^2/2} dx \\ &= \sqrt{\frac{\pi}{8n}} = \sqrt{\frac{\pi}{8 \lfloor T/2p \rfloor}} \leq \sqrt{\frac{\pi p}{4T}} \end{aligned}$$

Putting back the pieces together, we have shown that for any i ,

$$\mathbb{E} \left[f(\tilde{X}^{(i)}) \right] - \mathbb{E} \left[f(\tilde{X}_T) \right] \leq \sum_{j=1}^p \sqrt{\frac{\pi p}{4T}} = p^{3/2} \sqrt{\frac{\pi}{4T}}.$$

We deduce the same bound for r_{CV} by taking the maximum over i .

Let us now bound r_{\min} . By Eq. (9), using the fact that $\lfloor T/(2p) \rfloor \geq T/(4p)$ as $T/(2p) \geq 1$, for all i

$$M(f) - \mathbb{E} \left[f(\tilde{X}^{(i)}) \right] \leq \frac{L}{K_i^\alpha} + 18 \sqrt{\frac{4pK_i}{T}}.$$

We summarize a few calculations in the next lemma. These calculations come from the minimization over the K_i 's of the previous bound, with a case disjunction arising from the boundary cases.

Lemma 2. *At least one of the three following inequalities holds :*

$$L < 2^{\alpha+1/2} \sqrt{\frac{p}{T}} \quad \text{or} \quad L \geq T^\alpha \sqrt{p}$$

or

$$\min_{i=1, \dots, p} \left(\frac{L}{K_i^\alpha} + 36 \sqrt{\frac{pK_i}{T}} \right) \leq 53 L^{1/(2\alpha+1)} \left(\frac{p}{T} \right)^{\alpha/(2\alpha+1)}.$$

Let us consider these three cases separately. The first one is forbidden by the assumption that $L \geq 2^{\alpha+1/2} \sqrt{p/T}$. In the second case, the function is so irregular that the claimed bound becomes worse than $\bar{r}_T \leq 56 p^{1/2+\alpha/(2\alpha+1)}$, which is weaker than the trivial bound $\bar{r}_T \leq 1$.

Finally, in the third case, we may assume that $L \geq 2^{\alpha+1/2} \sqrt{p/T} \geq \sqrt{p/T}$. Then we have

$$L^{1/(2\alpha+1)} \geq \left(\frac{p}{T} \right)^{1/(2(2\alpha+1))} = \left(\frac{p}{T} \right)^{1/2} \left(\frac{p}{T} \right)^{-\alpha/(2\alpha+1)},$$

and thus $\sqrt{p/T} \leq L^{1/(2\alpha+1)} (p/T)^{\alpha/(2\alpha+1)}$. By injecting the bound of Lemma 2 and the bound on r_{CV} into (23):

$$\bar{r}_T \leq 53 L^{1/(2\alpha+1)} \left(\frac{p}{T} \right)^{\alpha/(2\alpha+1)} + p \sqrt{\frac{\pi}{4}} \sqrt{\frac{p}{T}} \leq (53 + p \sqrt{\pi/4}) L^{1/(2\alpha+1)} \left(\frac{p}{T} \right)^{\alpha/(2\alpha+1)}$$

and the stated bound holds, since $53 + p \sqrt{\pi/4} \leq 53 + (\log_2 T + 1) \sqrt{\pi/4} \leq 54 + \sqrt{\pi/4} \log_2 T$. \square

Proof of Lemma 2. We upper bound the minimum by comparing the two quantities

$$\frac{L}{K_i^\alpha} \quad \text{v.s.} \quad \sqrt{\frac{pK_i}{T}}.$$

As the first term is decreasing with i , and the second term is increasing with i , two extreme cases have to be dealt with. If the first term is always smaller than the second, i.e., even for $i = 1$, then:

$$\frac{L}{2^\alpha} < \sqrt{\frac{p2}{T}}.$$

This is the first case in the statement of the lemma. Otherwise, the first term might always be greater than the second one, i.e., even for $i = p$ and

$$\frac{L}{2^{\alpha p}} \geq \sqrt{\frac{p2^p}{T}}$$

which is equivalent to

$$L^2 \geq p \frac{2^{p(2\alpha+1)}}{T},$$

hence, since $2^p \geq T$,

$$L^2 \geq pT^{2\alpha}$$

which is exactly the second inequality of our statement.

Otherwise, define i^* to be an index such that

$$\frac{L}{K_{i^*-1}^\alpha} \geq \sqrt{\frac{pK_{i^*-1}}{T}} \quad \text{and} \quad \frac{L}{K_{i^*}^\alpha} \leq \sqrt{\frac{pK_{i^*}}{T}} \quad (25)$$

By the preceding discussion, i^* is well defined and $1 < i^* \leq p$. Then by definition of i^* (the first equation in (25))

$$2^{\alpha+1/2} \frac{L}{K_{i^*}^\alpha} \geq \sqrt{\frac{pK_{i^*}}{T}}.$$

Hence, by squaring and regrouping the terms

$$K_{i^*}^{2\alpha+1} \leq 2^{2\alpha+1} L^2 \frac{T}{p}$$

thus

$$K_{i^*} \leq 2L^{2/(\alpha+1)} \left(\frac{T}{p}\right)^{1/(2\alpha+1)}$$

and

$$\sqrt{\frac{pK_{i^*}}{T}} \leq \sqrt{2} L^{1/(2\alpha+1)} \left(\frac{p}{T}\right)^{\alpha/(2\alpha+1)}$$

and finally, recalling the second equation in (25)

$$\frac{L}{K_{i^*}^\alpha} + 36\sqrt{\frac{pK_{i^*}}{T}} \leq 37\sqrt{\frac{pK_{i^*}}{T}} \leq 37\sqrt{2} L^{1/(2\alpha+1)} \left(\frac{p}{T}\right)^{\alpha/(2\alpha+1)}.$$

□

F Proof of our version of the lower bound of adaptation

Here we provide the full proof of our version of the lower bound of adaptation stated in Section 2.2.

Our statement differs from that of Locatelli and Carpentier [21] on some aspects. First, and most importantly, we include the dependence on the Lipschitz constants, and we do not consider margin regularity. We also remove a superfluous requirement on B , that $B \leq cT^{(\alpha+1)/(2\alpha+1)}$, which was just an artifact of the original proof. Furthermore we believe that the additional condition that $L \leq \mathcal{O}(T^{\alpha/2})$ in our version was implicitly used in this original proof. Finally, the value of the constant differs, partly because of the analysis, and partly because we consider $(1/4)$ -subgaussian noise instead of 1-subgaussian noise.

We managed to obtain these improvements thanks to a different proof technique. In the original proof, the authors compare the empirical likelihoods of different outcomes and use the Bretagnolle-Huber inequality. We choose to build the lower bound in a slightly different way (see Garivier et al. [13]): we handle the changes of measure implicitly thanks to Pinsker's inequality (Lemma 3). Following Lattimore and Szepesvári [19], we also chose to be very precise in the definition of the bandit model, in order to make rigorous a few arguments that are often used implicitly in the literature on continuous bandits.

The main argument of the proof, that is, the sets of functions considered, are already present in Locatelli and Carpentier [21].

Before we start with the proof, let us state a technical tool. Denote by KL the Kullback-Leibler divergence. The next lemma is a generalized version of Pinsker's inequality, tailored to our needs.

Lemma 3. *Let \mathbb{P} and \mathbb{Q} be two probability measures. For any random variable $Z \in [0, 1]$,*

$$|\mathbb{E}_{\mathbb{P}}[Z] - \mathbb{E}_{\mathbb{Q}}[Z]| \leq \sqrt{\frac{\text{KL}(\mathbb{P}, \mathbb{Q})}{2}}$$

Proof. For $z \in [0, 1]$, by the classical version of Pinsker's inequality applied to the event $\{Z \geq z\}$:

$$|\mathbb{P}[Z \geq z] - \mathbb{Q}[Z \geq z]| \leq \sqrt{\frac{\text{KL}(\mathbb{P}, \mathbb{Q})}{2}}.$$

Therefore, by Fubini's theorem and the triangle inequality, and by integrating the preceding inequality:

$$|\mathbb{E}_{\mathbb{P}}[Z] - \mathbb{E}_{\mathbb{Q}}[Z]| = \left| \int_0^1 (\mathbb{P}[Z \geq z] - \mathbb{Q}[Z \geq z]) dz \right| \leq \int_0^1 |\mathbb{P}[Z \geq z] - \mathbb{Q}[Z \geq z]| dz \leq \sqrt{\frac{\text{KL}(\mathbb{P}, \mathbb{Q})}{2}}$$

□

Proof of the lower bound. For the sake of completeness, we recall in detail the construction of Locatelli and Carpentier [21], with some minor simplifications that fit our setting. Fix regularity parameters ℓ, L, α and γ satisfying $\ell \leq L$ and $\gamma \geq \alpha$, so that $\mathcal{H}(\ell, \gamma) \subset \mathcal{H}(L, \alpha)$ (remember the functions are defined on $\mathcal{X} = [0, 1]$).

Fix $M \in [1/2, 1]$. Let $K \in \mathbb{N} \setminus \{0\}$ and $\Delta \in \mathbb{R}_+$ be some parameters of the construction whose values will be determined by the analysis. We define furthermore a partition of $[0, 1]$ into $K + 1$ sets, $H_0 = [1/2, 1]$ and $H_i = [(i-1)/(2K), i/(2K)]$ for $1 \leq i \leq K$, along with their middle points $x_i \in H_i$. Finally, define the set of hypotheses ϕ_i for $i = 0, \dots, K$ as follows

$$\phi_i(x) = \begin{cases} \max(M - \Delta, M - \Delta/2 - \ell|x - x_0|^\gamma) & \text{if } x \in H_0, \\ \max(M - \Delta, M - L|x - x_i|^\alpha) & \text{if } x \in H_i \text{ and } s \neq 0, \\ M - \Delta & \text{otherwise.} \end{cases} \quad (26)$$

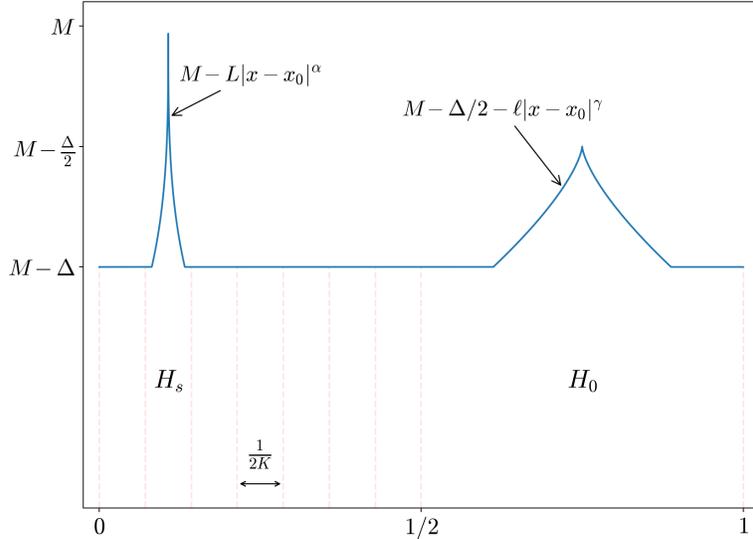


Figure 5: Mean-payoff functions for the lower bound

Figure 5 illustrates how the ϕ_i 's are defined : for $1 \leq i \leq K$, the function ϕ_i displays a peak of size Δ and of low regularity (L, α) , localized in H_i , and another peak of size $\Delta/2$, of higher

regularity (ℓ, γ) in H_0 . The function ϕ_0 only has the peak of size $\Delta/2$ and regularity (ℓ, γ) . We need to add requirements on the values of the parameters, to make sure the indeed functions belong to the appropriate regularity classes. These requirements are written in the following lemma, which we prove later.

Lemma 4. *If $(\Delta/L)^{1/\alpha} \leq 1/(4K)$ then $\phi_0 \in \mathcal{H}(\ell, \gamma)$, and if $(\Delta/(2\ell))^{1/\gamma} \leq 1/4$ then $\phi_i \in \mathcal{H}(L, \alpha)$ for $i \geq 1$.*

Fix a given algorithm. The idea of the proof of the lower bound is to use the fact that if the player has low regret, that is, less than B , when the mean-payoff function is $\phi_0 \in \mathcal{H}(L, \alpha)$, then she has to play in H_0 often. This in turn constrains the amount of exploration she can afford, and limits her ability to find the maximum when the mean-payoff functions is ϕ_i for $i > 0$.

Canonical bandit model In this paragraph, we build the necessary setting for a rigorous development. The continuous action space gives rise to measurability issues, and one should be particularly careful when handling changes of measure as we do here. Following Lattimore and Szepesvári [19, Chap. 4.7, 14 (Ex.11) and 15 (Ex.8)], we build the canonical bandit model in order to apply the chain rule for Kullback-Leibler divergences rigorously. To our knowledge, this is seldom done carefully, the two notable exceptions being the above reference and Garivier et al. [13]. We also use the notion of probability kernels in this paragraph; see Kallenberg [15, Chap. 1 and 5] for a definition and properties.

Define a sequence of measurable spaces $\Omega_t = \prod_{s=1}^t \mathcal{X} \times \mathbb{R}$, together with their Borel σ -algebra (with the usual topology on $\mathcal{X} = [0, 1]$ and on \mathbb{R}). We call $h_t = (x_1, y_1, \dots, x_t, y_t) \in \Omega_t$ a history up to time t . By an abuse of notation, we consider that $\Omega_t \subset \Omega_{t'}$ when $t \leq t'$.

An algorithm is a sequence $(K_t)_{1 \leq t \leq T}$ of (regular) probability kernels, with K_t from Ω_{t-1} to \mathcal{X} , modelling the choice of the arm at time t . By an abuse of notation, the first kernel K_1 is an arbitrary measure on \mathcal{X} , the law of the first arm picked. Define for each i another probability kernel modelling the reward obtained: $L_{i,t}$ from $\Omega_t \times \mathcal{X}$ to \mathbb{R} . We write it explicitly as :

$$L_{i,t}((x_1, y_1, \dots, x_t), B) = \sqrt{\frac{2}{\pi}} \int_B e^{-2(x - \phi_i(x_t))^2} dx$$

These kernels define probability laws $\mathbb{P}_{i,t} = L_{i,t}(K_t \mathbb{P}_{i,t-1})$ over Ω_t . Doing so, we ensured that under $\mathbb{P}_{i,t}$ the coordinate random variables $X_t : \Omega_t \rightarrow \mathcal{X}$ and $Y_t : \Omega_t \rightarrow \mathbb{R}$, defined as $X_t(x_1, \dots, x_t, y_t) = x_t$ and $Y_t(x_1, \dots, x_t, y_t) = y_t$ are such that given X_t , the reward Y_t is distributed according to $\mathcal{N}(\phi_i(X_t), 1/4)$. Denote by \mathbb{E}_i the expectation taken according to $\mathbb{P}_{i,t}$. We also index recall the pseudo-regret: $\bar{R}_{T,i} = TM(\phi_i) - \mathbb{E}_i \left[\sum_{t=1}^T \phi_i(X_t) \right]$.

A rewriting of the chain rule for Kullback-Leibler divergence with our notation would be (see Lattimore and Szepesvári [19, Exercise 11 Chap. 14] for a proof)

Proposition (Chain rule). *Let Ω and Ω' be measurable subsets of \mathbb{R}^d equipped with their natural σ -algebra. Let \mathbb{P} and \mathbb{Q} be probability distributions defined over Ω , and K and L be regular probability kernels from Ω to Ω' then*

$$\text{KL}(K\mathbb{P}, L\mathbb{Q}) = \text{KL}(\mathbb{P}, \mathbb{Q}) + \int_{\Omega} \text{KL}(K(\omega, \cdot), L(\omega, \cdot)) d\mathbb{P}(\omega)$$

The key assumptions are that Ω and Ω' are subspaces of \mathbb{R}^d , and that K and L satisfy measurability conditions, as they are regular kernels; these assumptions justify the heavy setting we introduced.

Under this setting, we may call to the chain rule twice to see that for any t :

$$\begin{aligned}
\text{KL}(\mathbb{P}_0^t, \mathbb{P}_i^t) &= \text{KL}(L_{0,t}(K_t \mathbb{P}_0^{t-1}), L_{i,t}(K_t \mathbb{P}_i^{t-1})) \\
&= \text{KL}(K_t \mathbb{P}_0^{t-1}, K_t \mathbb{P}_i^{t-1}) + \int_{\Omega_{t-1} \times \mathcal{X}} \text{KL}(L_{0,t}(h_{t-1}, x_t, \cdot), L_{i,t}(h_{t-1}, x_t, \cdot)) \, dK_t \mathbb{P}_0^{t-1}(h_{t-1}, x_t) \\
&= \text{KL}(\mathbb{P}_0^{t-1}, \mathbb{P}_i^{t-1}) + \int_{\Omega_{t-1} \times \mathcal{X}} \text{KL}(L_{0,t}(h_{t-1}, x_t, \cdot), L_{i,t}(h_{t-1}, x_t, \cdot)) \, dK_t \mathbb{P}_0^{t-1}(h_{t-1}, x_t) \\
&= \text{KL}(\mathbb{P}_0^{t-1}, \mathbb{P}_i^{t-1}) + \int_{\mathcal{X}} \text{KL}(\mathcal{N}(\phi_0(x_t), 1/4), \mathcal{N}(\phi_i(x_t), 1/4)) \, d\mathbb{P}_0^{t-1}(x_t) \\
&= \text{KL}(\mathbb{P}_0^{t-1}, \mathbb{P}_i^{t-1}) + \mathbb{E}_0[\text{KL}(\mathcal{N}(\phi_0(X_t), 1/4), \mathcal{N}(\phi_i(X_t), 1/4))]
\end{aligned}$$

where the penultimate equality comes from the fact that the density of the kernel $L_{i,t-1}$ depends only on the last coordinate x_t , and is exactly that of a gaussian variable.

We obtain the KL decomposition by iterating T times,

$$\text{KL}(\mathbb{P}_0^T, \mathbb{P}_i^T) = \mathbb{E}_0 \left[\sum_{t=1}^T \text{KL}(\mathcal{N}(\phi_0(X_t), 1/4), \mathcal{N}(\phi_i(X_t), 1/4)) \right]$$

Continuation of the proof Let us also define $N_{H_i}(T) = \sum_{t=1}^T \mathbb{1}_{\{X_t \in H_i\}}$ the number of times the algorithm selects an arm in H_i . The hypotheses ϕ_i were defined for the three following inequalities to hold. For all $i \geq 1$:

$$\bar{R}_{T,i} \geq \frac{\Delta}{2} (T - \mathbb{E}_i[N_{H_i}(T)]) = \frac{T\Delta}{2} \left(1 - \frac{\mathbb{E}_i[N_{H_i}(T)]}{T} \right), \quad (27)$$

$$\bar{R}_{T,0} \geq \frac{\Delta}{2} \sum_{i=1}^K \mathbb{E}_0[N_{H_i}(T)], \quad (28)$$

and

$$\begin{aligned}
\text{KL}(\mathbb{P}_0^T, \mathbb{P}_i^T) &= \mathbb{E}_0 \left[\sum_{t=1}^T \text{KL}(\mathcal{N}(\phi_0(X_t), 1/4), \mathcal{N}(\phi_i(X_t), 1/4)) \right] \\
&= \mathbb{E}_0 \left[\sum_{t=1}^T 2(\phi_0(X_t) - \phi_i(X_t))^2 \right] \leq 2 \mathbb{E}_0[N_{H_i}(T)] \Delta^2.
\end{aligned} \quad (29)$$

The first two inequalities come from the fact that, under \mathbb{P}_i , the player incurs an instantaneous regret of less than $\Delta/2$ whenever she picks an arm outside the optimal cell H_i . For the third inequality, first apply the chain rule to compute the Kullback-Leibler divergence, then the inequality is a consequence of the fact that ϕ_i and ϕ_0 differ only in H_i , and their difference is less than Δ .

We may now proceed with the calculations. By Lemma 3 applied to the random variable $N_{H_i}(T)/T$:

$$\frac{\mathbb{E}_i[N_{H_i}(T)]}{T} \leq \frac{\mathbb{E}_0[N_{H_i}(T)]}{T} + \sqrt{\frac{\text{KL}(\mathbb{P}_0^T, \mathbb{P}_i^T)}{2}}. \quad (30)$$

We will now show that

$$\frac{1}{K} \sum_{i=1}^K \bar{R}_{T,i} \geq \frac{T\Delta}{2} \left(1 - \frac{1}{K} - \sqrt{\frac{\Delta \bar{R}_{T,0}}{K}} \right). \quad (31)$$

Indeed by (in order) averaging (27) over $i = 1, \dots, K$, using (30), the concavity of $\sqrt{\cdot}$ and (29)

$$\begin{aligned}
\frac{1}{K} \sum_{i=1}^K \bar{R}_{T,i} &\geq \frac{T\Delta}{2} \left(1 - \frac{1}{K} \sum_{i=1}^K \frac{\mathbb{E}_i[N_{H_i}(T)]}{T} \right) \\
&\geq \frac{T\Delta}{2} \left(1 - \frac{1}{K} \sum_{i=1}^K \frac{\mathbb{E}_0[N_{H_i}(T)]}{T} - \frac{1}{K} \sum_{i=1}^K \sqrt{\frac{\text{KL}(\mathbb{P}_0^T, \mathbb{P}_i^T)}{2}} \right) \\
&\geq \frac{T\Delta}{2} \left(1 - \frac{1}{K} - \sqrt{\frac{1}{2K} \sum_{i=1}^K \text{KL}(\mathbb{P}_0^T, \mathbb{P}_i^T)} \right) \\
&\geq \frac{T\Delta}{2} \left(1 - \frac{1}{K} - \sqrt{\frac{\Delta^2}{K} \sum_{i=1}^K \mathbb{E}_0[N_{H_i}(T)]} \right).
\end{aligned}$$

This yields the claimed inequality (31) thanks to (28).

Let us assume for now that $K \geq 2$ and $\phi_0 \in \mathcal{H}(\ell, \gamma)$. Then by the assumption on the algorithm, $\bar{R}_{T,0} \leq B$, and therefore

$$\frac{1}{K} \sum_{i=1}^K \bar{R}_{T,i} \geq \frac{T\Delta}{2} \left(\frac{1}{2} - \sqrt{\frac{\Delta B}{K}} \right). \quad (32)$$

To optimize this bound, we take Δ as large as possible, while still ensuring that $\sqrt{\Delta B/K}$ is small enough, e.g., less than $1/4$. Furthermore, we impose that the ϕ_i 's belong to $\mathcal{H}(L, \alpha)$, i.e., by Lemma 4, that $(\Delta/L)^{1/\alpha} \leq 1/(4K)$. This leads to the choice

$$\Delta = c L^{1/(\alpha+1)} B^{-\alpha/(\alpha+1)} \quad \text{and} \quad K = \left\lfloor \frac{1}{4} \left(\frac{\Delta}{L} \right)^{-1/\alpha} \right\rfloor = \left\lfloor \frac{c^{-1/\alpha}}{4} (LB)^{1/(\alpha+1)} \right\rfloor,$$

with $c = 1/128$.

Conclusion, assuming that $K \geq 2$ and $\phi_0 \in \mathcal{H}(\ell, \gamma)$ With this choice of parameters, we have by definition of Δ ,

$$\Delta B = c (LB)^{1/(\alpha+1)},$$

and by definition of K , since $K \geq (c^{-1/\alpha}/8)(LB)^{1/(\alpha+1)}$,

$$\frac{\Delta B}{K} \leq 8c^{1+1/\alpha}$$

hence, using $c^{1/(2\alpha)} \leq 1$

$$\sqrt{\frac{\Delta B}{K}} \leq 2\sqrt{2}c^{1/2+1/(2\alpha)} \leq 2\sqrt{2} \cdot 2^{-7/2} = \frac{1}{4}.$$

With this in hand, we may now go back to inequality (32) to see that

$$\frac{1}{K} \sum_{i=1}^K \bar{R}_{T,i} \geq \frac{T\Delta}{2} \left(\frac{1}{2} - \frac{1}{4} \right) \geq \frac{T\Delta}{8} = \frac{c}{8} TL^{1/(\alpha+1)} B^{-\alpha/(\alpha+1)}.$$

By the definition of K , it is always true that $(\Delta/L)^{1/\alpha} \leq 1/(4K)$, and therefore, by Lemma 4, all the ϕ_i 's automatically belong to $\mathcal{H}(L, \alpha)$. Therefore, for all i , we have $\sup_{f \in \mathcal{H}(L, \alpha)} \bar{R}_T \geq \bar{R}_{T,i}$. Hence, recalling that $c = 1/128$,

$$\sup_{f \in \mathcal{H}(L, \alpha)} \bar{R}_T \geq \frac{1}{K} \sum_{i=1}^K \bar{R}_{T,i} \geq 2^{-10} TL^{1/(\alpha+1)} B^{-\alpha/(\alpha+1)}.$$

Regularity conditions on the mean-payoff functions ϕ_i We now check that $K \geq 2$, and that $\phi_0 \in \mathcal{H}(\ell, \gamma)$. Let us first focus on ϕ_0 . By Lemma 4, it is enough to impose that $(\Delta/(2\ell))^{1/\gamma} \leq 1/4$, i.e., that

$$c L^{1/(\alpha+1)} B^{-\alpha/(\alpha+1)} / (2\ell) \leq (1/4)^\gamma$$

that is,

$$L^{1/(\alpha+1)} B^{-\alpha/(\alpha+1)} \leq 2\ell (1/4)^\gamma / c = \ell 2^{1-2\gamma} c^{-1},$$

i.e., when

$$LB^{-\alpha} \leq \ell^{1+\alpha} 2^{(1-2\gamma)(1+\alpha)} c^{-(1+\alpha)}$$

hence, replacing c by its value $c = 2^{-7}$, the next condition is sufficient to ensure the regularity of the hypothesis:

$$L \leq \ell^{1+\alpha} B^\alpha c^{-(1+\alpha)} 2^{(1+\alpha)(1-2\gamma)} = \ell^{1+\alpha} B^\alpha 2^{(1+\alpha)(8-2\gamma)},$$

which is one of the two conditions in the statement of the theorem. For the bound to be valid, we must also make sure that $K \geq 2$:

$$\left\lfloor \left(\frac{c^{-1/\alpha}}{4} (LB)^{1/(\alpha+1)} \right) \right\rfloor \geq 2.$$

This condition is weaker than

$$\frac{c^{-1/\alpha}}{4} (LB)^{1/(\alpha+1)} \geq 3$$

which is equivalent to

$$L \geq c^{(\alpha+1)/\alpha} 12^{\alpha+1} B^{-1} = 2^{-7} \cdot 12 \cdot 2^{-6/\alpha} 12^\alpha B^{-1}.$$

To ensure this, we require the stronger (but more readable) condition that $L \geq 2^{-3} 12^\alpha B^{-1}$. \square

Proof of Lemma 4. A good look at Figure 6 should convince the reader of the statement. We wish to make sure that the functions ϕ_i 's satisfy (4), a Hölder condition around their maximum (and only around this maximum). Given the definition of the functions ϕ_i , we simply have to check that there is no discontinuity at the boundary of the cell H_i . We write out the details for $i > 0$ to remove any doubt; the same analysis can be carried to check that $\phi_0 \in \mathcal{H}(\ell, \gamma)$.

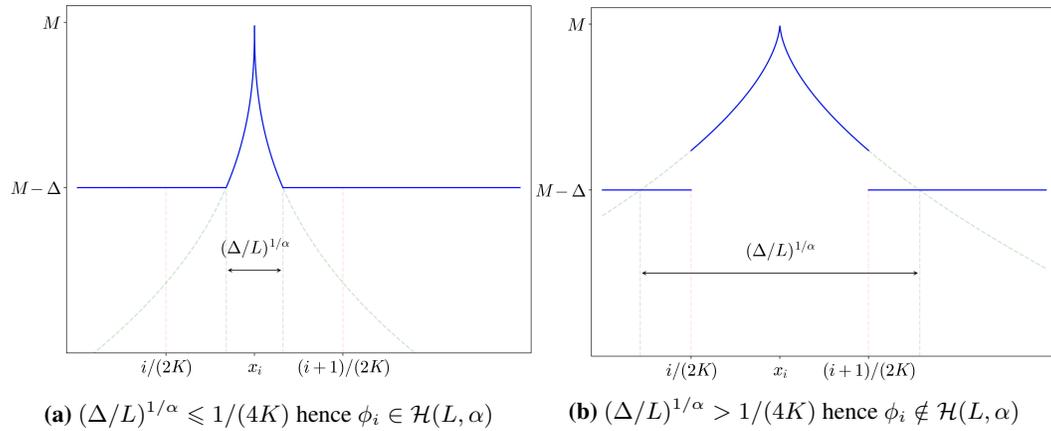


Figure 6: ϕ_i is in $\mathcal{H}(L, \alpha)$ if it is everywhere above the green dotted curve $x \mapsto M - L|x - x_i|^\alpha$, that is, if the cell H_i has enough room to contain the whole peak of size Δ

For $i > 0$, the function ϕ_i reaches its maximum at $x_i = (i - 1/2)/2K$, and the value of the maximum is M . Then for $x \in H_i$, by definition of ϕ_i :

$$\phi_i(x) = \max(M - \Delta, M - L|x_i - x|^\alpha) \geq M - L|x - x_i|^\alpha$$

thus

$$\phi_i(x_i) - \phi_i(x) = M - \phi_i(x) \leq L|x_i - x|^\alpha,$$

Now consider $x \notin H_i$. Assume, as in the statement of the lemma, that $1/(4K) \geq (\Delta/L)^{1/\alpha}$. If x is outside of H_i , then since H_i is of half-width $1/4K$,

$$|x_i - x| \geq \frac{1}{4K} \geq \left(\frac{\Delta}{L}\right)^{1/\alpha} \quad (33)$$

and, by definition of ϕ_i , for all x (even for $x \in H_0$), $\phi_i(x) \geq M - \Delta$. Therefore, by (33),

$$\phi_i(x_i) - \phi_i(x) \leq \Delta \leq L|x_i - x|^\alpha.$$

For all values of x , the Hölder condition is satisfied and $\phi_i \in \mathcal{H}(L, \alpha)$.

For ϕ_0 , the same calculations show that there is no jump at the boundary of $[1/2, 1]$, of half-width $1/4$, when the peak is of height $\Delta/2$ and regularity (ℓ, γ) if $((\Delta/2)/\ell)^{1/\gamma} \leq 1/4$. \square