



HAL
open science

Adaptation to the Range in K -Armed Bandits

Hédi Hadiji, Gilles Stoltz

► **To cite this version:**

| Hédi Hadiji, Gilles Stoltz. Adaptation to the Range in K -Armed Bandits. 2020. hal-02794382v1

HAL Id: hal-02794382

<https://hal.science/hal-02794382v1>

Preprint submitted on 5 Jun 2020 (v1), last revised 9 Jun 2022 (v3)

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.

Adaptation to the Range in K -Armed Bandits

Hédi Hadiji & Gilles Stoltz

Laboratoire de mathématiques d'Orsay

Université Paris-Saclay, CNRS

91405, Orsay, France

hedı.hadıjı,gıllıes.stoltz@math.u-psud.fr

Abstract

We consider stochastic bandit problems with K arms, each associated with a bounded distribution supported on the range $[m, M]$. We do not assume that the range $[m, M]$ is known and show that there is a cost for learning this range. Indeed, a new trade-off between distribution-dependent and distribution-free regret bounds arises, which, for instance, prevents from simultaneously achieving the typical $\ln T$ and \sqrt{T} bounds. For instance, a \sqrt{T} distribution-free regret bound may only be achieved if the distribution-dependent regret bounds are at least of order \sqrt{T} . We exhibit a strategy achieving the rates for regret indicated by the new trade-off.

1 Introduction

Virtually all articles on stochastic K -armed bandits either assume that distributions of the arms belong to some parametric family (often, one-dimensional exponential families) or to the non-parametric family of distributions supported on a known range $[m, M]$. Notable exceptions are discussed below.

We consider the second, non-parametric, framework (see Section 2) and show that the knowledge of the range $[m, M]$ is a crucial information. We do so by studying what may be achieved and what cannot be achieved anymore when this range is unknown and the strategies need to learn it. We call this problem scale-free regret minimization. Our main result (in Section 3) is a trade-off between the scale-free distribution-dependent and distribution-free regret bounds that may be achieved; it is, for instance, impossible to simultaneously achieve scale-free distribution-dependent regret bounds of order $\ln T$ and scale-free distribution-free regret bounds of order \sqrt{T} , as simple strategies like UCB strategies (by Auer et al. [5]) do in the case of a known range. Our general trade-off indicates, for instance, that if one wants to keep the same \sqrt{T} order of magnitude for the scale-free distribution-free regret bounds, then the best scale-free distribution-dependent rate that may be achieved is \sqrt{T} . We also provide (in Section 4) a strategy, based on exponential weights, that obtains optimal scale-free distribution-dependent and distribution-free regret bounds as indicated by the trade-off.

Short literature review. Optimal scale-free regret minimization under full monitoring is offered by the AdaHedge strategy by De Rooij et al. [16], which we will use as a building block in Section 4. The main difficulty in adaptation to the range is the adaptation to the upper end M (see Section 5); this is why Honda and Takemura [24] could provide optimal $\ln T$ distribution-dependent regret bounds for payoffs lying in ranges of the form $(-\infty, M]$, with a known M . Lattimore [29] considers models of distributions with a known bound on their kurtosis (a scale-free measure of the skewness of the distributions) and provides a scale-free algorithm based on the median-of-means estimators, with $\ln T$ distribution-dependent regret bounds. However, bounded bandits can have an arbitrarily high kurtosis, so our settings are not directly comparable (and we think that bounded distributions with an unknown range is a more natural assumption). Cowan and Katehakis [14] study adaptation to the range but in the restricted case of uniform distributions; see also similar results by Cowan et al. [15] for Gaussian distributions with unknown means and variances. Additional important references are discussed in Appendix A of the supplementary material.

2 Settings: stochastic bandits and bandits for oblivious individual sequences

We describe the bandit settings considered: stochastic bandits, the setting of main interest, and bandits for oblivious individual (adversarial) sequences, a setting leading to stronger regret upper bounds.

2.1 Stochastic bandits with bounded and possibly signed rewards

$K \geq 2$ arms are available. We denote by $[K]$ the set $\{1, \dots, K\}$ of arms. With each of the arm a is associated a probability distribution ν_a lying in some known model \mathcal{D} ; a model is a set of probability distributions over \mathbb{R} with a first moment. The models of interest in this article are discussed below. A bandit problem over \mathcal{D} is a K -vector of probability distributions in \mathcal{D} : we denote it by $\underline{\nu} = (\nu_a)_{a \in [K]}$. The player knows \mathcal{D} but not $\underline{\nu}$. As is standard in this setting, we denote by $\mu_a = \mathbb{E}(\nu_a)$ the mean payoff provided by an arm a . An optimal arm and the optimal mean payoff are respectively given by $a^* \in \operatorname{argmax}_{a \in [K]} \mu_a$ and $\mu^* = \max_{a \in [K]} \mu_a$. Finally, $\Delta_a = \mu^* - \mu_a$ denotes the gap of an arm a .

The online learning game goes as follows: at round $t \geq 1$, the player picks an arm $A_t \in [K]$, possibly at random according to a probability distribution $p_t = (p_{t,a})_{a \in [K]}$ based on an auxiliary randomization U_{t-1} , and then receives and observes a reward Z_t drawn independently at random according to the distribution ν_{A_t} , given A_t . More formally, a strategy of the player is a sequence of mappings from the observations to the action set, $(U_0, Z_1, U_1, \dots, Z_{t-1}, U_{t-1}) \mapsto A_t$, where U_0, U_1, \dots are i.i.d. random variables independent from all other random variables and distributed according to a uniform distribution over $[0, 1]$. At each given time $T \geq 1$, we measure the performance of a strategy through its expected regret:

$$R_T(\underline{\nu}) = T\mu^* - \mathbb{E} \left[\sum_{t=1}^T Z_t \right] = T\mu^* - \mathbb{E} \left[\sum_{t=1}^T \mu_{A_t} \right] = \sum_{t=1}^T \Delta_a \mathbb{E}[N_a(T)], \quad (1)$$

where we used the tower rule for the first equality and defined $N_a(T)$ as the number of times arm a was pulled between time rounds 1 and T .

Doob's optional skipping (see Doob [17, Chapter III, Theorem 5.2, p. 145] for the original reference, see also Chow and Teicher [13, Section 5.3] for a more recent reference) indicates that we may assume that i.i.d. sequences of rewards $(Y_{t,a})_{t \geq 1}$ are drawn beforehand, independently at random, for each arm a and that the obtained payoff at round $t \geq 1$ given the choice A_t equals $Z_t = Y_{t,A_t}$. We will use this second formulation in the rest of the paper as it is the closest to the one of oblivious individual sequences described in Section 2.2.

Models: bounded rewards with unknown range. For a given range $[m, M]$, where $m < M$ are two real numbers (not necessarily nonnegative), we denote by $\mathcal{D}_{m,M}$ the set of probability distributions supported on $[m, M]$. Then, the model corresponding to distributions with a bounded but unknown range is the union of all such $\mathcal{D}_{m,M}$:

$$\mathcal{D}_{-,+} = \bigcup_{m, M \in \mathbb{R}: m < M} \mathcal{D}_{m,M}$$

2.2 Oblivious individual sequences (oblivious adversarial bandits)

In the setting of (fully) oblivious individual sequences (see Cesa-Bianchi and Lugosi [10], Audibert and Bubeck [2]), a range $[m, M]$ is set by the environment, where m, M are real numbers (not necessarily nonnegative), and the environment picks beforehand a sequence y_1, y_2, \dots of reward vectors in $[m, M]^K$. We denote by $y_t = (y_{t,a})_{a \in [K]}$ the components of these vectors. The online learning game starts only then: at each round $t \geq 1$, the player picks an arm $A_t \in [K]$, possibly at random according to a probability distribution $p_t = (p_{t,a})_{a \in [K]}$ based on an auxiliary randomization U_{t-1} , and then receives and observes y_{t,A_t} . More formally, a strategy of the player is a sequence of mappings from the observations to the action set, $(U_0, y_{1,A_1}, U_1, \dots, y_{t-1,A_{t-1}}, U_{t-1}) \mapsto A_t$, where U_0, U_1, \dots are i.i.d. random variables independent from all other random variables and distributed according to a uniform distribution over $[0, 1]$. At each given time $T \geq 1$, denoting $y_{1:T} = (y_1, \dots, y_T)$, we measure the performance of a strategy through its expected regret:

$$R_T(y_{1:T}) = \max_{a \in [K]} \sum_{t=1}^T y_{t,a} - \mathbb{E} \left[\sum_{t=1}^T y_{t,A_t} \right], \quad (2)$$

where all randomness lies in the choice of the arms A_t only (as rewards are fixed beforehand).

Conversion of upper/lower bounds from one setting to the other. Note that (by the tower rule for the right-most equality) for all $m < M$ and for all $\underline{\nu}$ in $\mathcal{D}_{m,M}$,

$$R_T(\underline{\nu}) = \max_{a \in [K]} \mathbb{E} \left[\sum_{t=1}^T Y_{t,a} \right] - \mathbb{E} \left[\sum_{t=1}^T Y_{t,A_t} \right] \leq \mathbb{E} \left[\max_{a \in [K]} \sum_{t=1}^T Y_{t,a} - \sum_{t=1}^T Y_{t,A_t} \right] = \mathbb{E} [R_T(Y_{1:T})] \\ \leq \sup_{y_{1:T} \text{ in } [m,M]^K} R_T(y_{1:T}).$$

In particular, lower bounds on the regret for stochastic bandits are also lower bounds on the regret for oblivious adversarial bandits, and strategies designed for oblivious adversarial bandits obtain the same regret bounds for stochastic bandits when the individual payoffs y_{t,A_t} in their definition are replaced with the stochastic payoffs Y_{t,A_t} .

2.3 Scale-free regret bounds: rates for adaptation to the unknown range

Regret scales with the range length $M - m$, thus regret bounds involve a multiplicative factor $M - m$. We therefore consider such bounds divided by the scale factor $M - m$ and call them scale-free regret bounds. We denote by \mathbb{N} the set of natural integers; (rates on) regret bounds will be given by functions $\Phi : \mathbb{N} \rightarrow [0, +\infty)$.

Definition 1 (Distribution-free bounds). *A strategy for stochastic bandits, respectively, for oblivious individual sequences, is adaptive to the unknown range of payoffs with a scale-free distribution-free regret bound $\Phi : \mathbb{N} \rightarrow [0, +\infty)$ if for all real numbers $m < M$, the strategy ensures, without the knowledge of m and M :*

$$\forall \underline{\nu} \text{ in } \mathcal{D}_{m,M}, \quad \forall T \geq 1, \quad R_T(\underline{\nu}) \leq (M - m) \Phi(T), \\ \text{respectively,} \quad \forall y_1, y_2, \dots \text{ in } [m, M]^K, \quad \forall T \geq 1, \quad R_T(y_{1:T}) \leq (M - m) \Phi(T).$$

The notion of distribution-dependent regret bounds for adaptation to the range can obviously only be defined for stochastic bandits. It does not add much to the classical notion of distribution-dependent rates on regret bounds, as the scale factor $M - m$ does not appear in the definition; it merely ensures that the strategy is not informed of the range.

Definition 2 (Distribution-dependent bounds). *A strategy for stochastic bandits is adaptive to the unknown range of payoffs with a distribution-dependent rate $\Phi : \mathbb{N} \rightarrow [0, +\infty)$ if for all real numbers $m < M$, the strategy ensures, without the knowledge of m and M :*

$$\forall \underline{\nu} \text{ in } \mathcal{D}_{m,M}, \quad \limsup_{T \rightarrow +\infty} \frac{R_T(\underline{\nu})}{\Phi(T)} < +\infty.$$

Put differently, the strategy ensures that $\limsup R_T(\underline{\nu})/\Phi(T) < +\infty$ for all $\underline{\nu} \in \mathcal{D}_{-,+}$.

3 Regret lower bounds for adaptation to the range

Any scale-free distribution-free regret bound $\Phi_{\text{free}}(T)$ is larger than the optimal distribution-free regret bound on a known range. [6] provided a lower bound $(1/20) \min\{\sqrt{KT}, T\}$ on the regret of any strategy against individual sequences in $[0, 1]^K$, thus for bandit problems in $\mathcal{D}_{0,1}$. Therefore, we also have $\Phi_{\text{free}}(T) \geq (1/20) \min\{\sqrt{KT}, T\}$. We show in Section 4 a scale-free distribution-free regret bound of order $\sqrt{KT} \ln K$, which thus matches the lower bound up to a $\sqrt{\ln K}$ factor.

The situation is different for distribution-dependent bounds, where the typical $\ln T$ order of magnitude cannot be achieved when the range is unknown: all uniformly fast convergent strategies on $\mathcal{D}_{-,+}$ are such that, for all bandit problems $\underline{\nu}$ in $\mathcal{D}_{-,+}$ with at least one suboptimal arm,

$$\liminf_{T \rightarrow +\infty} \frac{R_T(\underline{\nu})}{\ln T} = +\infty. \quad (3)$$

(A strategy is said to be uniformly fast convergent on a model \mathcal{D} if for all bandit problems $\underline{\nu}$ in \mathcal{D} , it achieves a subpolynomial regret bound, that is, $R_T(\underline{\nu})/T^\alpha \rightarrow 0$ for all $\alpha \in (0, 1]$; this is a minimal requirement when studying lower bounds.) However, any rate $\varphi(T) \gg \ln T$ may be achieved thanks to a simple upper-confidence bound [UCB] strategy. Further details, including proofs of the two claims above, may be found in Appendix C of the supplementary material.

We now show that under an adaptivity assumption that is stronger than uniform fast convergence and takes finite-time guarantees into account, the distribution-dependent regret becomes polynomial in T .

3.1 Simultaneous scale-free distribution-free and distribution-dependent lower bounds

When the range $[m, M]$ of the payoffs is known, it is possible to simultaneously achieve optimal distribution-free bounds (of order \sqrt{KT}) and optimal distribution-dependent bounds (of order $\ln T$ with the optimal constant given by infima of Kullback-Leibler divergences); see the KL-UCB-switch strategy by Garivier et al. [20]. Put differently, when the range of payoffs is known, one can achieve optimal (asymptotic) distribution-dependent regret bounds while not sacrificing finite-time guarantees. Simpler strategies like UCB strategies (see Auer et al. [5]) also simultaneously achieve regret bounds of similar $\ln T$ and \sqrt{T} orders of magnitude but with suboptimal constants and/or dependencies on K .

This is not possible anymore when the range of payoffs is unknown.

To show this, we consider in this section algorithms enjoying distribution-free scale-free regret bounds and show that they suffer up to a $\Omega(\sqrt{T})$ distribution-dependent rate for adaptation to the range. Actually, the theorem below shows that there is a trade-off between the finite-time guarantees (the distribution-free scale-free regret bounds) and the asymptotic problem-dependent rates (the distribution-dependent rates for adaptation) that can be achieved. We recall that these concepts were defined in Section 2.3. The proof actually provides a finite-time (but messy) lower bound on $R_T(\underline{\nu}) / (T / \Phi_{\text{free}}(T))$.

Theorem 1. *Any strategy with a Φ_{free} distribution-free scale-free regret bound satisfying $\Phi_{\text{free}} \ll T$ may only achieve distribution-dependent rates Φ_{dep} for adaptation satisfying $\Phi_{\text{dep}}(T) \geq T / \Phi_{\text{free}}(T)$. More precisely, the regret of such a strategy is lower bounded as: for all $\underline{\nu}$ in $\mathcal{D}_{-,+}$,*

$$\liminf_{T \rightarrow \infty} \frac{R_T(\underline{\nu})}{T / \Phi_{\text{free}}(T)} \geq \frac{1}{4} \sum_{a=1}^K \Delta_a.$$

The optimal distribution-free scale-free regret bounds $\Phi_{\text{free}}(T)$ are of order \sqrt{T} (as follows from the lower bound indicated at the beginning of Section 3 and from the upper bound of Section 4). The distribution-dependent rates $\Phi_{\text{dep}}(T)$ of strategies achieving this optimal distribution-free scale-free rate are therefore larger than \sqrt{T} . More generally, there is a trade-off between the two rates: to force faster distribution-dependent rates for adaptation, one must suffer worsened distribution-free scale-free rates for adaptation. (The latter range between the optimal \sqrt{KT} rate and the trivial T rate.)

3.2 Proof of Theorem 1

We follow a standard proof technique introduced by Lai and Robbins [28] and Burnetas and Katehakis [8] and recently revisited by Garivier et al. [21]. We fix some bandit problem $\underline{\nu}$ in $\mathcal{D}_{-,+}$ and construct an alternative bandit problem $\underline{\nu}'$ in $\mathcal{D}_{-,+}$ by modifying the distribution of a single suboptimal arm a to make it optimal (which is always possible, as there is no bound on the upper end on the ranges of the payoffs in the model). We apply a fundamental inequality that links the expectations of the numbers of times $N_a(T)$ that a is pulled under $\underline{\nu}$ and $\underline{\nu}'$. We then substitute inequalities stemming from the definition of distribution-free scale-free regret bounds Φ_{dep} , and the result follows by rearranging all inequalities.

Step 1: Alternative bandit problem. The lower bound is trivial (it equals 0) when all arms of $\underline{\nu}$ are optimal. We therefore assume that at least one arm is suboptimal and fix such an arm a . For some $\varepsilon \in [0, 1]$ to be defined later by the analysis, we introduce the alternative problem $\underline{\nu}' = (\nu'_k)_{k \in [K]}$ with $\nu'_k = \nu_k$ for $j \neq a$ and $\nu'_a = (1 - \varepsilon)\nu_a + \varepsilon\delta_{\mu_a + 2\Delta_a/\varepsilon}$. This distribution ν'_a has a bounded range, so that $\underline{\nu}'$ lies indeed in $\mathcal{D}_{-,+}$. The expectation of ν'_a equals $\mu'_a = \mu_a + 2\Delta_a = \mu^* + \Delta_a > \mu^*$. Thus, a is the only optimal arm in $\underline{\nu}'$. Finally, for ε small enough, $\mu_a + 2\Delta_a/\varepsilon$ lies outside of the bounded support of ν_a . In that case, the density of ν_a with respect to ν'_ε is given by $1/(1 - \varepsilon)$ on the support of ν_a (and 0 elsewhere), so that $\text{KL}(\nu_a, \nu'_a) = \ln(1/(1 - \varepsilon))$.

Step 2: Application of a fundamental inequality. We denote by $\text{kl}(p, q)$ the Kullback-Leibler divergence between Bernoulli distributions with parameters p and q . We also index expectations in the rest of the proof by the bandit problem they are relative to: for instance, $\mathbb{E}_{\underline{\nu}}$ denotes the expectation of a random variable when the ambient randomness is given by the bandit problem $\underline{\nu}$. The fundamental

inequality for lower bounds on the regret of stochastic bandits (Garivier et al. [21], Section 2, Equation 6), which is based on the chain rule for Kullback-Leibler divergence and on a data-processing inequality for expectations of $[0, 1]$ -valued random variables, reads:

$$\text{kl}\left(\frac{\mathbb{E}_{\underline{\nu}}[N_a(T)]}{T}, \frac{\mathbb{E}_{\underline{\nu}'}[N_a(T)]}{T}\right) \leq \mathbb{E}_{\underline{\nu}}[N_a(T)] \text{KL}(\nu_a, \nu'_a) = \mathbb{E}_{\underline{\nu}}[N_a(T)] \ln(1/(1-\varepsilon)).$$

Now, since $u \in (-\infty, 1) \mapsto -u^{-1} \ln(1-u)$ is increasing, we have $\ln(1/(1-\varepsilon)) \leq \varepsilon(\ln 2)/2$ for $\varepsilon \leq 1/2$. For all $(p, q) \in [0, 1]^2$ and with the usual measure-theoretic conventions,

$$\text{kl}(p, q) = \underbrace{p \ln p + q \ln q}_{\geq -\ln 2} + \underbrace{p \ln \frac{1}{q}}_{\geq 0} + (1-p) \ln \frac{1}{1-q} \geq (1-p) \ln \frac{1}{1-q} - \ln 2,$$

so that, putting all inequalities together, we have proved

$$\left(1 - \frac{\mathbb{E}_{\underline{\nu}}[N_a(T)]}{T}\right) \ln\left(\frac{1}{1 - \mathbb{E}_{\underline{\nu}'}[N_a(T)]/T}\right) - \ln 2 \leq \frac{\ln 2}{2} \varepsilon \mathbb{E}_{\underline{\nu}}[N_a(T)]. \quad (4)$$

So far, we only imposed the constraint $\varepsilon \in [0, 1/2]$.

Step 3: Inequalities stemming from the definition of distribution-free scale-free regret bounds. We denote by $[m, M]$ a range containing the supports of all distributions of $\underline{\nu}$. By definition of Φ_{free} , given that a is a suboptimal arm (i.e., $\Delta_a > 0$):

$$\Delta_a \mathbb{E}_{\underline{\nu}}[N_a(T)] \leq R_T(\underline{\nu}) \leq (M - m) \Phi_{\text{free}}(T).$$

Because of ν'_a , the distributions of $\underline{\nu}'$ have supports within the range $[m, M_\varepsilon]$, where we denoted $M_\varepsilon = \max\{M, \mu_a + 2\Delta_a/\varepsilon\}$. For $\underline{\nu}'$, by definition of Φ_{free} , and given that all gaps Δ'_k are larger than the gap $\Delta'_a = \mu'_a - \mu^* = \Delta_a$ between the unique optimal a and the second best arms (which were the optimal arms of $\underline{\nu}$),

$$\begin{aligned} \Delta_a(T - \mathbb{E}_{\underline{\nu}'}[N_a(T)]) &= \Delta'_a(T - \mathbb{E}_{\underline{\nu}'}[N_a(T)]) \leq \sum_{j \neq a} \Delta'_j \mathbb{E}_{\underline{\nu}'}[N_j(T)] \\ &= R_T(\underline{\nu}') \leq (M_\varepsilon - m) \Phi_{\text{free}}(T). \end{aligned}$$

By rearranging the two inequalities above, we get

$$1 - \frac{\mathbb{E}_{\underline{\nu}}[N_a(T)]}{T} \geq 1 - \frac{(M - m) \Phi_{\text{free}}(T)}{T \Delta_a} \quad \text{and} \quad 1 - \frac{\mathbb{E}_{\underline{\nu}'}[N_a(T)]}{T} \leq \frac{(M_\varepsilon - m) \Phi_{\text{free}}(T)}{T \Delta_a},$$

thus, after substitution into (4),

$$\left(1 - \frac{(M - m) \Phi_{\text{free}}(T)}{T \Delta_a}\right) \ln\left(\frac{T \Delta_a}{(M_\varepsilon - m) \Phi_{\text{free}}(T)}\right) - \ln 2 \leq \frac{\ln 2}{2} \varepsilon \mathbb{E}_{\underline{\nu}}[N_a(T)]. \quad (5)$$

Step 4: Final calculations. We take $\varepsilon = \varepsilon_T = \alpha^{-1} \Phi_{\text{free}}(T)/T$ for some constant $\alpha > 0$; we will pick $\alpha = 1/8$. By the assumption $\Phi_{\text{free}}(T) \ll T$, we have $\varepsilon_T \leq 1/2$ as needed for T large enough, as well as $M_{\varepsilon_T} = \mu_a + 2\Delta_a/\varepsilon_T = \mu_a + 2\alpha\Delta_a T/\Phi_{\text{free}}(T)$. Substituting these values into (5), a finite-time lower bound on the quantity of interest is finally given by

$$\frac{\mathbb{E}_{\underline{\nu}}[N_a(T)]}{T/\Phi_{\text{free}}(T)} \geq \frac{2\alpha}{\ln 2} \left(-\ln 2 + \underbrace{\left(1 - \frac{(M - m) \Phi_{\text{free}}(T)}{T \Delta_a}\right)}_{\rightarrow 0} \ln \underbrace{\left(\frac{T \Delta_a}{2\alpha\Delta_a T + (\mu_a - m)\Phi_{\text{free}}(T)}\right)}_{\rightarrow 1/(2\alpha)} \right).$$

It entails the asymptotic lower bound

$$\liminf_{T \rightarrow +\infty} \frac{\mathbb{E}_{\underline{\nu}}[N_a(T)]}{T/\Phi_{\text{free}}(T)} \geq \frac{2\alpha}{\ln 2} (\ln(1/\alpha) - 2 \ln 2) = \frac{1}{4}$$

for the choice $\alpha = 1/8$. The claimed result follows by adding these lower bounds for each suboptimal arm a , with a factor Δ_a , following the rewriting (1) of the regret.

Algorithm 1 AdaHedge for K -armed bandits, with extra-exploration

- 1: **Input:** a sequence $(\gamma_t)_{t \geq 1}$ in $[0, 1]$ of extra-exploration rates; a payoff estimation scheme
- 2: **for** rounds $t = 1, \dots, K$ **do**
- 3: Draw arm $A_t = t$
- 4: Get and observe the payoff $y_{t,t}$
- 5: **end for**
- 6: **AdaHedge initialization:** $\eta_{K+1} = +\infty$ and $q_{K+1} = (1/K, \dots, 1/K) \stackrel{\text{def}}{=} \mathbf{1}/K$
- 7: **for** rounds $t = K + 1, \dots$ **do**
- 8: Define p_t by mixing q_t with the uniform distribution according to $p_t = (1 - \gamma_t)q_t + \gamma_t \mathbf{1}/K$
- 9: Draw an arm $A_t \sim p_t$ (independently at random according to the distribution p_t)
- 10: Get and observe the payoff y_{t,A_t}
- 11: Compute estimates $\hat{y}_{t,a}$ of all payoffs with the payoff estimation scheme considered
- 12: Compute the mixability gap $\delta_t \geq 0$ based on the distribution q_t and on these estimates:

$$\delta_t = - \sum_{a=1}^K q_{t,a} \hat{y}_{t,a} + \frac{1}{\eta_t} \ln \left(\sum_{a=1}^K q_{t,a} e^{\eta_t \hat{y}_{t,a}} \right), \quad \text{with } \delta_t = \underbrace{\sum_{a=1}^K q_{t,a} \hat{y}_{t,a} + \max_{a \in [K]} \hat{y}_{t,a}}_{\text{when } \eta_t = +\infty}$$

- 13: Compute the learning rate $\eta_{t+1} = \left(\sum_{s=K+1}^t \delta_s \right)^{-1} \ln K$
- 14: Define q_{t+1} component-wise as

$$q_{t+1,a} = \exp \left(\eta_{t+1} \sum_{s=K+1}^t \hat{y}_{a,s} \right) / \sum_{k=1}^K \exp \left(\eta_{t+1} \sum_{s=K+1}^t \hat{y}_{k,s} \right)$$

- 15: **end for**
-

4 Quasi-optimal regret bounds for range adaptation based on AdaHedge

When the range of payoffs is known, Auer et al. [6] use exponential weights (Hedge) on estimated payoffs and with extra-exploration (mixing with the uniform distribution) to achieve a regret bound of order $\sqrt{KT \ln K}$. Actually, it is folklore knowledge that the extra-exploration is unnecessary when regret bounds are considered only in expectation, as is the case in the present article.

When the range of payoffs is unknown, we consider a self-tuned version called AdaHedge (De Rooij et al. [16], see also earlier work by Cesa-Bianchi et al. [12]) and do add extra-exploration. The latter is not detrimental, given the trade-off between the distribution-free and distribution-dependent bounds discussed in the previous section; we actually achieve that trade-off. Algorithm 1 is stated in the case of adversarial oblivious learning, but to use it with stochastic payoffs, it suffices to replace y_{t,A_t} with Y_{t,A_t} . It relies on a payoff estimation scheme, which we discuss now.

In Algorithm 1, some initial exploration lasting K rounds is used to get a rough idea of the location of the payoffs and to center the estimates used at an appropriate location. Following by Auer et al. [6]), we consider, for all rounds $t \geq K + 1$ and arms $a \in [K]$,

$$\hat{y}_{t,a} = \frac{y_{t,A_t} - C}{p_{t,a}} \mathbf{1}_{\{A_t=a\}} + C \quad \text{where} \quad C \stackrel{\text{def}}{=} \frac{1}{K} \sum_{s=1}^K y_{s,s}. \quad (6)$$

Note that all $p_{t,a} > 0$ for Algorithm 1 due to the use of exponential weights. As proved by Auer et al. [6], these estimates are (conditionally) unbiased. Indeed, the distributions q_t and p_t (as well as the constant C) are measurable functions of the information $H_{t-1} = (U_0, y_{1,A_1}, U_1, \dots, y_{t-1,A_{t-1}})$ available at the beginning of round $t \geq K + 1$, and the arm A_t is drawn independently at random according to p_t based on an auxiliary randomization denoted by U_{t-1} . Therefore, given that the payoffs are oblivious, the conditional expectation of $\hat{y}_{t,a}$ with respect to H_{t-1} amounts to integrating over the randomness given by the random draw $A_t \sim p_t$: for $t \geq K + 1$,

$$\mathbb{E}[\hat{y}_{t,a} \mid H_{t-1}] = \frac{y_{t,a} - C}{p_{t,a}} \mathbb{P}(A_t = a \mid H_{t-1}) + C = \frac{y_{t,a} - C}{p_{t,a}} p_{t,a} + C = y_{t,a}. \quad (7)$$

These estimators are bounded: assuming that all $y_{t,a}$, thus also C , belong to the range $[m, M]$, and given that the distributions p_t were obtained by a mixing with the uniform distribution, with weight γ_t , we have $p_{t,a} \geq \gamma_t/K$, and therefore,

$$\forall t \geq K+1, \forall a \in [K], \quad |\hat{y}_{t,a} - C| \leq \frac{|y_{t,a} - C|}{p_{t,a}} \leq \frac{M-m}{\gamma_t/K}. \quad (8)$$

Remark. *Algorithm 1 is invariant by affine changes (translations and/or multiplications by positive factors) of the payoffs, as AdaHedge (see De Rooij et al. [16, Theorem 16]) and the payoff estimation scheme (6) so are. This is key for adaptation to the range.*

4.1 Distribution-free scale-free regret analysis

Theorem 2. *AdaHedge for K -armed bandits (Algorithm 1) with a non-increasing extra-exploration (γ_t) smaller than $1/2$ and the estimation scheme given by (6) ensures that for all bounded ranges $[m, M]$, for all oblivious individual sequences y_1, y_2, \dots in $[m, M]^K$, for all $T \geq 1$,*

$$R_T(y_{1:T}) \leq 3(M-m)\sqrt{KT \ln K} + 5(M-m)\frac{K \ln K}{\gamma_T} + (M-m) \sum_{t=K+1}^T \gamma_t.$$

Proof sketch. We provide only a sketch of proof and refer the reader to Appendix D of the supplementary material for a complete, detailed and commented proof. A direct application of the AdaHedge regret bound (Lemma 3 and Theorem 6 of De Rooij et al. [16]), bounding the variance terms of the form $\mathbb{E}[(X - \mathbb{E}[X])^2]$ by $\mathbb{E}[(X - C)^2]$, ensures that

$$\max_{k \in [K]} \sum_{t=K+1}^T \hat{y}_{t,k} - \sum_{\substack{t \geq K+1 \\ a \in [K]}} q_{t,a} \hat{y}_{t,a} \leq 2 \sqrt{\sum_{\substack{t \geq K+1 \\ a \in [K]}} q_{t,a} (\hat{y}_{t,a} - C)^2 \ln K} + \frac{M-m}{\gamma_T/K} \left(2 + \frac{4}{3} \ln K\right).$$

We take expectations, use the definition of the p_t in terms of the q_t in the left-hand side, and apply Jensen's inequality in the right-hand side to get

$$\begin{aligned} \mathbb{E} \left[\max_{k \in [K]} \sum_{t=K+1}^T \hat{y}_{t,k} - \sum_{t=K+1}^T \sum_{a=1}^K \overbrace{p_{t,a}}^{=y_{t,A_t}} \hat{y}_{t,a} + \sum_{t=K+1}^T \gamma_t \sum_{a=1}^K \overbrace{(1/K - q_{t,a})}^{\mathbb{E}[\dots] \in [m-M, M-m]} \hat{y}_{t,a} \right] \\ \leq 2 \sqrt{\sum_{t=K+1}^T \sum_{a=1}^K \mathbb{E} \left[q_{t,a} (\hat{y}_{t,a} - C)^2 \right] \ln K} + \frac{M-m}{\gamma_T/K} \left(2 + \frac{4}{3} \ln K\right). \end{aligned}$$

Since $p_{t,a} \geq (1 - \gamma_t)q_{t,a}$ with $\gamma_t \leq 1/2$ by assumption on the extra-exploration rate, we have the bound $q_{t,a} \leq 2p_{t,a}$. Together with standard calculations similar to (7), we have

$$\mathbb{E} \left[q_{t,a} (\hat{y}_{t,a} - C)^2 \right] \leq 2 \mathbb{E} \left[p_{t,a} (\hat{y}_{t,a} - C)^2 \mid H_{t-1} \right] = 2 \mathbb{E} \left[\frac{(y_{t,A_t} - C)^2}{p_{t,a}} \mathbb{1}_{\{A_t=a\}} \right] = 2 \underbrace{(y_{t,a} - C)^2}_{\leq (M-m)^2}.$$

The proof is concluded by collecting all bounds and by taking care of the first K rounds. \square

Straightforward calculations (detailed in Appendix D of the supplementary material) then lead to the following consequence of Theorem 2.

Corollary 1. *Fix a parameter $\alpha \in (0, 1)$. AdaHedge for K -armed bandits (Algorithm 1) with the extra-exploration*

$$\gamma_t = \min \left\{ 1/2, \sqrt{5(1-\alpha)K \ln K / t^\alpha} \right\}$$

and the estimation scheme given by (6) ensures that for all bounded ranges $[m, M]$, for all oblivious individual sequences y_1, y_2, \dots in $[m, M]^K$, for all $T \geq 1$,

$$R_T(y_{1:T}) \leq \left(3 + \frac{5}{\sqrt{1-\alpha}}\right) (M-m)\sqrt{K \ln K} T^{\max\{\alpha, 1-\alpha\}} + 10(M-m)K \ln K.$$

In particular, for $\alpha = 1/2$, the bound $7(M-m)\sqrt{TK \ln K} + 10(M-m)K \ln K$ holds.

This value $\alpha = 1/2$ is the best one to consider if one is only interested in a distribution-free bound (i.e., one is not interested in the distribution-dependent rates for the regret).

4.2 Distribution-dependent regret analysis, and discussion of the trade-off

For $\alpha \in [1/2, 1)$, Algorithm 1 tuned as in Corollary 1 is adaptive to the unknown range of payoffs with a distribution-free scale-free regret bound

$$\Phi_{\text{free}}^{\text{AH}}(T) = \left(3 + \frac{5}{\sqrt{1-\alpha}}\right) \sqrt{K \ln K} T^\alpha + 10K \ln K \quad (9)$$

for oblivious individual sequences thus also for stochastic bandits, with the same regret bound. (The superscript AH in $\Phi_{\text{free}}^{\text{AH}}$ stands for AdaHedge.) The trade-off stated in Theorem 1 indicates that the best possible distribution-dependent rate for adaptation to the unknown range is determined by $T/\Phi_{\text{free}}^{\text{AH}}(T)$, which is of order $T^{1-\alpha}$. It indicates, more precisely, that for all $\underline{\nu}$ in $\mathcal{D}_{-,+}$,

$$\liminf_{T \rightarrow \infty} \frac{R_T(\underline{\nu})}{T/\Phi_{\text{free}}^{\text{AH}}(T)} \geq \frac{1}{4} \sum_{a=1}^K \Delta_a.$$

The following theorem shows that this best possible distribution-dependent rate is indeed achieved and quantifies the gap between the distribution-dependent constants at hand: they differ by two multiplicative factors, a numerical factor of $4 \times 12/(1-\alpha)$ and a $\ln K$ factor.

Theorem 3. *Consider Algorithm 1 tuned as in Corollary 1, for $\alpha \in [1/2, 1)$. For all distributions ν_1, \dots, ν_K in $\mathcal{D}_{-,+}$,*

$$\limsup_{T \rightarrow \infty} \frac{R_T(\underline{\nu})}{T/\Phi_{\text{free}}^{\text{AH}}(T)} \leq \frac{12 \ln K}{1-\alpha} \sum_{a=1}^K \Delta_a. \quad (10)$$

The proof is provided in Appendix D of the supplementary material. It follows quite closely that of Theorem 3 in Seldin and Lugosi [35], where the authors study a variant of the Exp3 algorithm of Auer et al. [6] for stochastic rewards. It consists, in our setting, in showing that the number of times the algorithm chooses suboptimal arms is almost only determined by the extra-exploration. Our proof is simpler as we aim for cruder bounds. The main technical difference and issue to solve lies in controlling the learning rates η_t , which heavily depend on data in our case.

5 Extensions present in the supplementary material

Numerical experiments. They illustrate how the strategies introduced in this paper indeed adapt to the unknown range of payoffs.

One known end on the payoff range. It is folklore knowledge that there is a difference in nature between dealing with nonnegative payoffs (gains) or dealing with nonpositive payoffs (losses) for regret minimization under bandit monitoring; see Cesa-Bianchi and Lugosi [10, Remark 6.5, page 164] for an early reference and Kwon and Perchet [27] for a more complete literature review. Actually, 0 plays no special role, the issue is rather whether one end of the payoff range is known. What follows is detailed in Appendix E of the supplementary material.

Known lower end m on the payoff range. In that case we deal (up to a translation) with gains. This knowledge does not provide any advantage. Indeed, the impossibility results of Section 3 still hold, namely, no $\ln T$ rate may be achieved for scale-free distribution-dependent regret bounds, as in (3), and a trade-off exists between scale-free distribution-free and distribution-dependent regret bounds (Theorem 1 holds).

Known upper end M on the payoff range. In that case we deal (up to a translation) with losses, also known as semi-bounded rewards. The results of Section 3 do not hold anymore. The DMED strategy of Honda and Takemura [24] achieves the optimal asymptotic distribution-dependent regret bound, of order $\ln T$. We also recover some classical results: the INF strategy of Audibert and Bubeck [2] may be extended to provide a scale-free distribution-free regret bound of order \sqrt{KT} , and the AdaHedge strategy does not need any mixing with the uniform distribution to achieve the bound of Theorem 2.

Linear bandits. The techniques developed for adaptation to the range in Section 4 may be generalized to deal with (oblivious) adversarial linear bandits, see details in Appendix G of the supplementary material.

6 Broader impact

Not applicable

References

- [1] J.-Y. Audibert. Fast learning rates in statistical inference through aggregation. *Annals of Statistics*, 37(4):1591–1646, 2009.
- [2] J.-Y. Audibert and S. Bubeck. Minimax policies for adversarial and stochastic bandits. In *Proceedings of the 22nd Annual Conference on Learning Theory (COLT'09)*, pages 217–226. Omnipress, 2009.
- [3] J.-Y. Audibert, R. Munos, and C. Szepesvári. Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):1876–1902, 2009.
- [4] J.-Y. Audibert, S. Bubeck, and G. Lugosi. Regret in online combinatorial optimization. *Mathematics of Operations Research*, 39(1):31–45, 2014.
- [5] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002.
- [6] P. Auer, N. Cesa-Bianchi, Y. Freund, and R.E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.
- [7] S. Bubeck, M.B. Cohen, and Y. Li. Sparsity, variance and curvature in multi-armed bandits. In *Proceedings of the 29th International Conference on Algorithmic Learning Theory (ALT'18)*, volume 83 of PMLR, pages 111–127, 2018.
- [8] A.N. Burnetas and M.N. Katehakis. Optimal adaptive policies for sequential allocation problems. *Advances in Applied Mathematics*, 17(2):122–142, 1996.
- [9] O. Cappé, A. Garivier, O.-A. Maillard, R. Munos, and G. Stoltz. Kullback–Leibler upper confidence bounds for optimal sequential allocation. *The Annals of Statistics*, 41(3):1516–1541, 2013.
- [10] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [11] N. Cesa-Bianchi and O. Shamir. Bandit regret scaling with the effective loss range. In *Proceedings of the 29th International Conference on Algorithmic Learning Theory (ALT'18)*, volume 83 of PMLR, pages 128–151, 2018.
- [12] N. Cesa-Bianchi, Y. Mansour, and G. Stoltz. Improved second-order bounds for prediction with expert advice. *Machine Learning*, 66(2-3):321–352, 2007.
- [13] Y. Chow and H. Teicher. *Probability Theory*. Springer, 1988.
- [14] W. Cowan and M.N. Katehakis. An asymptotically optimal policy for uniform bandits of unknown support, 2015. Preprint, arXiv:1505.01918.
- [15] W. Cowan, J. Honda, and M.N. Katehakis. Normal bandits of unknown means and variances. *Journal of Machine Learning Research*, 18(154):1–28, 2018.
- [16] S. De Rooij, T. van Erven, P.D. Grünwald, and W.M. Koolen. Follow the leader if you can, hedge if you must. *Journal of Machine Learning Research*, 15(37):1281–1316, 2014.
- [17] J.L. Doob. *Stochastic Processes*. Wiley Publications in Statistics. John Wiley & Sons, 1953.
- [18] D.A. Freedman. On tail probabilities for martingales. *The Annals of Probability*, 3(1):100–118, 1975.
- [19] Y. Freund, R.E. Schapire, Y. Singer, and M.K. Warmuth. Using and combining predictors that specialize. In *Proceedings of the 29th annual ACM Symposium on Theory of Computing (STOC'97)*, pages 334–343, 1997.

- [20] A. Garivier, H. Hadji, P. Ménard, and G. Stoltz. KL-UCB-Switch: Optimal regret bounds for stochastic bandits from both a distribution-dependent and a distribution-free viewpoints, 2019. Preprint, arXiv:1805.05071.
- [21] A. Garivier, P. Ménard, and G. Stoltz. Explore first, exploit next: The true shape of regret in bandit problems. *Mathematics of Operations Research*, 44(2):377–399, 2019.
- [22] S. Gerchinovitz and T. Lattimore. Refined lower bounds for adversarial bandits. In *Advances in Neural Information Processing Systems*, pages 1198–1206, 2016.
- [23] J. Honda and A. Takemura. An asymptotically optimal policy for finite support models in the multiarmed bandit problem. *Machine Learning*, 85:361–391, 2011.
- [24] J. Honda and A. Takemura. Non-asymptotic analysis of a new bandit algorithm for semi-bounded rewards. *Journal of Machine Learning Research*, 16(113):3721–3756, 2015.
- [25] J. Kivinen and M.K. Warmuth. Averaging expert predictions. In *Proceedings of the 4th European Conference on Computational Learning Theory (EuroCOLT’99)*, pages 153–167, 1999.
- [26] W.M. Koolen. AdaFTRL. Blog post, Oct. 2016. URL <http://blog.wouterkoolen.info/AdaFTRL/post.html>.
- [27] J. Kwon and V. Perchet. Gains and losses are fundamentally different in regret minimization: The sparse case. *Journal of Machine Learning Research*, 17(227):1–32, 2016.
- [28] T.L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- [29] T. Lattimore. A scale free algorithm for stochastic bandits with bounded kurtosis. In *Advances in Neural Information Processing Systems*, pages 1584–1593, 2017.
- [30] T. Lattimore and C. Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020.
- [31] P. Massart. *Concentration Inequalities and Model Selection*, volume XXXIII of *Ecole d’Eté de Probabilités de Saint-Flour*. Springer, 2007. Lectures given in 2003, published in 2007.
- [32] H.B. McMahan. A survey of algorithms and analysis for adaptive online learning. *Journal of Machine Learning Research*, 18(1):3117–3166, 2017.
- [33] Francesco Orabona and Dávid Pál. Scale-free online learning. *Theoretical Computer Science*, 716:50–69, 2018.
- [34] M.D. Reid, R.M. Frongillo, R.C. Williamson, and N. Mehta. Generalized mixability via entropic duality. In *Proceedings of the 28th Conference on Learning Theory (COLT’15)*, volume 40 of PMLR, pages 1501–1522, 2015.
- [35] Y. Seldin and G. Lugosi. An improved parametrization and analysis of the EXP3++ algorithm for stochastic and adversarial bandits. In *Proceedings of the 30th Annual Conference on Learning Theory (COLT’17)*, volume 65 of PMLR, pages 1743–1759, 2017.
- [36] C.-Y. Wei and H. Luo. More adaptive algorithms for adversarial bandits. In *Proceedings of the 31st Conference On Learning Theory (COLT’18)*, volume 75 of PMLR, pages 1263–1291, 2018.
- [37] J. Zimmert and Y. Seldin. An optimal algorithm for stochastic and adversarial bandits. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTats’20)*, volume 89 of PMLR, pages 467–475, 2019.

Supplementary Material for “Adaptation to the Range in K -Armed Bandits”

We provide the following additions and extensions to the core results described in the main body of the article.

Appendix A offers an extended literature review on adaptation (to the range and other properties) in stochastic and adversarial K -armed bandits.

Appendix B provides numerical experiments illustrating how the strategies introduced in this paper indeed adapt to the unknown range of payoffs (and how earlier strategies do not).

Appendix C expands on the two claims stated in the introduction of Section 3, relative to distribution-dependent bounds for adaptation to the range (impossibility of a $\ln T$ rate and possibility thanks to a simple upper-confidence bound [UCB] strategy to achieve any rate $\varphi(T) \gg \ln T$).

Appendix D provides the complete proofs of the results of Section 4, namely, the ones of Theorem 2, Corollary 1, and Theorem 3.

Appendix E studies whether the adaptation results described in the main body of the article in the case of an unknown payoff range $[m, M]$ still hold when only one end of this range is unknown. It turns out that when m is known but M is unknown, achieving a distribution-dependent bound for adaptation to the range of order $\ln T$ is still impossible, and the trade-off between scale-free distribution-free and distribution-dependent regret bounds still holds (Theorem 1 holds). The picture is completely different when M is known but m is unknown, and improved scale-free distribution-free regret bounds can be provided.

Appendix F provides the statements and proofs of some technical results alluded in earlier appendices: the full-information regret bound for AdaHedge, needed in the complete proof of Theorem 2 in Appendix D, as well as the improved scale-free distribution-free regret bounds in the case the upper end M of the payoff range is known. All these results rely on a self-tuned version of a follow-the-regularized-leader (FTRL) strategy called AdaFTRL.

Appendix G deals with adaptation to the range for (oblivious) adversarial linear bandits.

A Extended literature review

We offer an extended literature review on adaptation (to the range and other properties) for stochastic and adversarial K -armed bandits.

Adaptation to the effective range in adversarial bandits. Gerchinovitz and Lattimore [22] show that it is impossible to adapt to the so-called effective range in adversarial bandits. A sequence of rewards has effective range smaller than b if for all rounds t , rewards $y_{t,a}$ at this round all lie in an interval of the form $[m_t, M_t]$ with $M_t - m_t \leq b$. The lower bound they exhibit relies on a sequence of changing intervals of fixed size. This problem is thus different from our setting. See also positive results (upper bounds) by Cesa-Bianchi and Shamir [11] for adaptation to the effective range.

Adaptation to the variance. Audibert et al. [3] consider a variant of UCB called UCB-V, which adapts to the unknown variance. Its analysis assumes that rewards lie in a known range $[0, M]$. The results crucially use Bernstein’s inequality (see, for instance, Reminder 3 in Appendix D.3 for a statement of the latter); as Bernstein’s inequality holds for random variables with supports in $[-\infty, M]$, the analysis of UCB-V might perhaps be extended to this case as well. Deviation bounds in Bernstein’s inequality contain two terms, a main term scaling with the standard deviation, and a remainder term, scaling with M . This remainder term, which seems harmless, is a true issue when M is not known, as indicated by the results of the present article.

Other criteria. Wei and Luo [36], Zimmert and Seldin [37], Bubeck et al. [7], and many more, provide strategies for adversarial bandits with rewards in a known range, say $[0, 1]$, and adapting to additional regularity in the data, like small variations or stochasticity of the data.

B Numerical experiments

We describe some numerical experiments on synthetic data to illustrate the performance of the new algorithms introduced compared to earlier approaches; we focus on how algorithms adapt to the scale of payoffs.

Five (families of) algorithms are considered. The first algorithm compared is vanilla UCB (with a $2 \ln T$ exploration factor, as in the original reference by Auer et al. [5]) and only adapt it to take the range $[m, M]$ of payoffs into account, by adding a $M - m$ factor in front of the upper confidence bound (see details below). We also compare AdaHedge for bandits and another strategy, alluded at in Section 5 and to be described in details in Appendices E.2.2 and F.4, called AdaFTRL with $1/2$ -Tsallis entropy, a generalization of the INF strategy of Audibert and Bubeck [2]. As the latter was introduced to handle losses (nonpositive payoffs), we will consider such nonpositive payoffs in our setting. It turns out that AdaHedge for bandits can be slightly improved in this case (see Appendices E.2.2 and F.3), by centering estimates at $C = 0$. Finally, we also add a simple follow-the-leader strategy (referred to as FTL; i.e., a strategy picking at each round the arm with best payoff estimate so far) and the random strategy (i.e., picking at each round an arm uniformly at random). FTL and the random strategies will exhibit undesirable performance similar to the ones of incorrectly tuned instances of UCB (respectively, with too small and too large a parameter σ).

Stochastic setting: bandit problems considered. We consider stochastic bandit problems $\underline{\nu}^{(\alpha)}$ indexed by a scale parameter $\alpha \in \{0.01, 0.1, 10, 100\}$. More precisely, $\underline{\nu}^{(\alpha)} = (\nu_a^{(\alpha)})_{a \in [K]}$ with $K = 30$ arms, each associated with a uniform distribution defined by

$$\nu_a^{(\alpha)} = \begin{cases} \text{Unif}([- \alpha, 0]) & \text{if } a = 1, \\ \text{Unif}([-1.2 \alpha, -0.2 \alpha]) & \text{if } a \neq 1, \end{cases}$$

so that all distributions are commonly supported on $[m, M] = [-1.2 \alpha, 0]$, with arm 1 being the unique optimal arm. Given the scale values $M - m = 1.2 \alpha$ obtained for the ranges $[m, M]$ as α varies, we consider four instances of UCB, with respective upper confidence bounds

$$\widehat{\mu}_a(t) + 1.2 \sigma \sqrt{\frac{2 \ln T}{N_a(t)}}, \quad \text{for } \sigma \in \{0.01, 0.1, 10, 100\},$$

where $N_a(t)$ is the number of times arm a was pulled up to round t and $\widehat{\mu}_a(t)$ denotes the empirical average of payoffs obtained for arm a when it was played.

Experimental setting. Each algorithm is run $N = 300$ times, on a time horizon $T = 100\,000$. We plot estimates of the rescaled regret $R_T(\underline{\nu}^{(\alpha)})/\alpha$ to have a meaningful comparison between the bandit problems.

These estimates are constructed as follows. We denote by

$$\mu_a^{(\alpha)} = \begin{cases} -\alpha/2 & \text{if } a = 1 \\ -0.7 \alpha & \text{if } a \neq 1 \end{cases}$$

the mean of arm a in $\underline{\nu}^{(\alpha)}$. We index the arms picked in the n -th run by an additional subscript n , so that $A_{T,n}$ refers to the arm picked by some strategy at time t in the n -th run. The expected regret of a given strategy can be rewritten as

$$R_T(\underline{\nu}^{(\alpha)}) = T \max_{a \in [K]} \mu_a^{(\alpha)} - \mathbb{E} \left[\sum_{t=1}^T \mu_{A_t}^{(\alpha)} \right] = -T\alpha/2 - \mathbb{E} \left[\sum_{t=1}^T \mu_{A_t}^{(\alpha)} \right]$$

and is estimated by

$$\widehat{R}_T(\alpha) = \frac{1}{N} \sum_{n=1}^N \widehat{R}_T(\alpha, n) \quad \text{where} \quad \widehat{R}_T(\alpha, n) = -T\alpha/2 - \sum_{t=1}^T \mu_{A_{t,n}}^{(\alpha)}.$$

On Figure 1 we therefore plot the estimates $\widehat{R}_T(\alpha)/\alpha$ of the rescaled regret as solid lines. The shaded areas correspond to ± 2 standard errors of the sequences $(\widehat{R}_T(\alpha, n)/\alpha)_{n \in [N]}$.

Table 1: Average runtimes of the (families of) algorithms considered, measured in seconds per run; as a reminder, we performed $N = 300$ runs for each algorithm.

Random play	FTL	UCB family	Bandit AdaHedge	Tsallis–AdaFTRL
$X = 1.51$ s/run	$1.7X$	$1.7X$	$7.9X$	$32.4X$

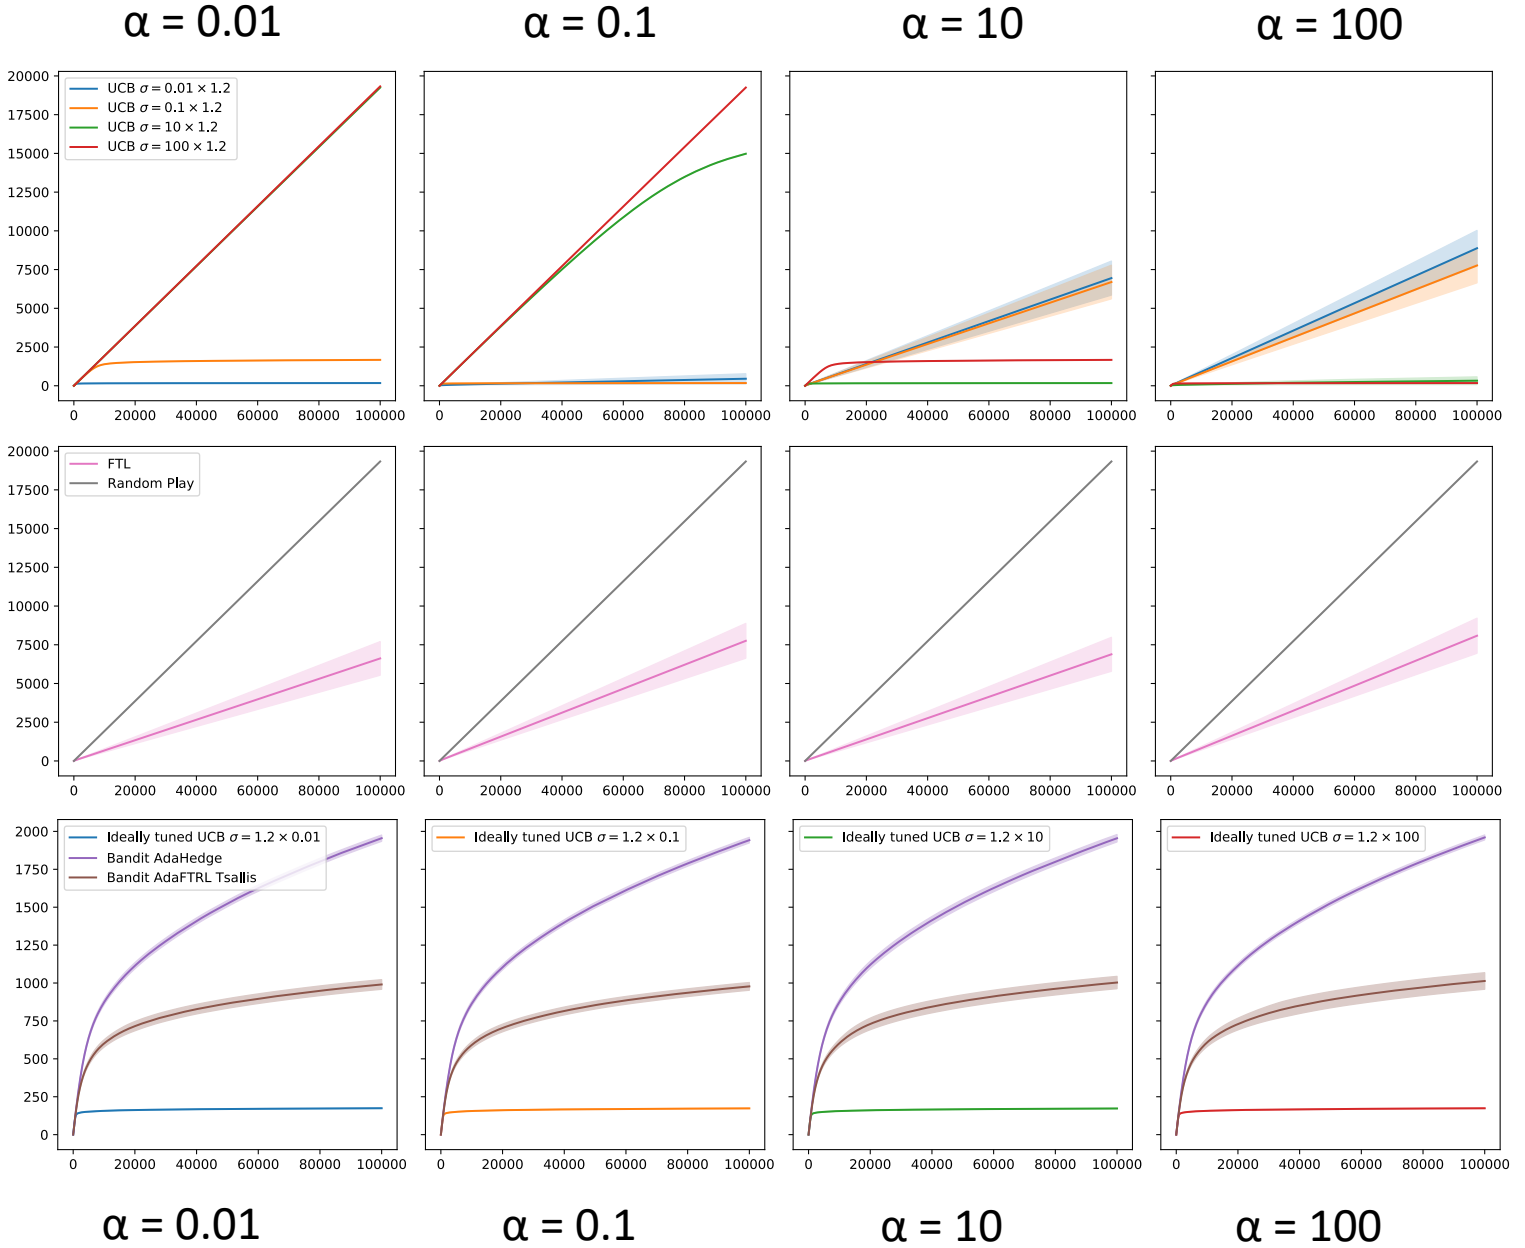


Figure 1: Comparison of the rescaled regrets of various strategies over bandit problems \mathcal{L}^α , where α ranges in $\{0.01, 0.1, 10, 100\}$. Each algorithm was run $N = 300$ times on every problem for $T = 100\,000$ time steps. Solid lines report the values of the estimated rescaled regrets, while shaded areas correspond to ± 2 standard errors of the estimates.

Complexity. The time complexity of FTL and of the instances of UCB lies only in the update of the payoff estimate of the selected arm and in the choice of the next arm. AdaHedge for bandits has a higher runtime due to the additional cost of computing the distributions q_t and p_t over the arms and the mixability gaps. AdaFTRL for bandits with $1/2$ -Tsallis entropy is the most time consuming of our algorithms, as it requires twice solving an optimization problem: once for the computation of the distributions over the arms and once for the mixability gaps; see Section F.4 for specific details on the said optimization problem and hints on an efficient solution thereof.

The memory complexity of all algorithms considered here is constant and scales linearly with K . The algorithms only need to keep in memory a vector of (cumulative or average) payoffs estimates and (for some) the cumulative mixability gaps.

All experiments were designed in Python, using the NumPy and joblib libraries, and were run on a standard laptop computer (with an Intel Core i5 processor). The code and setup for these experiments were only moderately optimized for computational efficiency. We display the average runtimes of all algorithms in Table 1; they are provided only for illustration and could certainly be significantly improved.

Discussion of the results. A first observation is that, as expected, our algorithms (see the third lines of Figure 1) are unaffected by the scale of the problems (up to a minor numerical stability issue discussed below). They yield favorable results (note that the range of the y -axis for the third line is smaller than the ranges in the first two lines), with AdaFTRL with $1/2$ -Tsallis entropy exhibiting a better performance than AdaHedge for bandits (our theoretical bounds reflect this, see Appendices E and F).

UCB tuned with the correct scale obtains the best results overall, which is consistent with the folklore knowledge that UCB performs well in practice. However, and this is our major second observation, the performance of UCB worsens dramatically when the scale is misspecified. When UCB is run with a scale parameter σ that is too small, it behaves similarly to FTL, incurring linear regret with extreme variance. When the scale parameter σ is too large, UCB is essentially playing at random and incurs linear regret too.

We conclude this section by discussing a minor issue of numerical stability: the error bars of for the expected regret of AdaFTRL with $1/2$ -Tsallis entropy seem to increase slightly with the scale α (while in theory they are independent of α). This is probably due to larger numerical errors associated with the approximate solutions of the optimization problems discussed in Section F.4.

A final note: UCB with estimated range. For the sake of completeness, we indicate that a version of UCB estimating the range, i.e., considering indices of the form

$$\hat{\mu}_a(t) + \hat{r}_t \sqrt{\frac{2 \ln T}{N_a(t)}},$$

where \hat{r}_t estimates the range $M - m$ as

$$\hat{r}_t = \max_{s \leq t} Y_{A_s, s} - \min_{s \leq t} Y_{A_s, s},$$

obtained an excellent performance on our simulations (the same as the optimally tuned version of UCB). We were however unable to provide theoretical guarantees that match our lower bounds. This is why we do not discuss this natural algorithm in the present article.

C Distribution-dependent lower bounds for adaptation to the range

In this section, we expand on the two claims stated in the introduction of Section 3, that are relative to distribution-dependent lower bounds for adaptation to the range: first, that all reasonable strategies (in the sens of Definition 3 below) ensure that for all bandit problems $\underline{\nu}$ in $\mathcal{D}_{-,+}$ with at least one suboptimal arm,

$$\liminf_{T \rightarrow +\infty} \frac{R_T(\underline{\nu})}{\ln T} = +\infty,$$

while, second, any rate $\varphi(T) \gg \ln T$ may be achieved thanks to a simple upper-confidence bound [UCB] strategy. Before we do so, we remind the reader of the “classical” results, for an abstract model \mathcal{D} and then, for the model $\mathcal{D}_{m,M}$ corresponding to payoff distributions with a known range $[m, M]$.

Definition 3. *A strategy is uniformly fast convergent on a model \mathcal{D} if for all bandit problems $\underline{\nu}$ in \mathcal{D} , it achieves a subpolynomial regret bound, that is, $R_T(\underline{\nu})/T^\alpha \rightarrow 0$ for all $(\alpha, 1]$.*

A lower bound on the distribution-dependent rates that such a strategy may achieve is provided by a general result of Lai and Robbins [28] and Burnetas and Katehakis [8] (see also its rederivation by Garivier et al. [21]). It involves a quantity defined as an infimum of Kullback-Leibler divergences: we recall that for two probability distributions ν, ν' defined on the same probability space (Ω, \mathcal{F}) ,

$$\text{KL}(\nu, \nu') = \begin{cases} \int_{\Omega} \ln\left(\frac{d\nu}{d\nu'}\right) d\nu & \text{if } \nu \ll \nu', \\ +\infty & \text{otherwise,} \end{cases}$$

where $\nu \ll \nu'$ means that ν is absolutely continuous with respect to ν' and $d\nu/d\nu'$ then denotes the Radon-Nikodym derivative. Now, for any probability distribution ν , any real number x , and any model \mathcal{D} , we define

$$\mathcal{K}_{\text{inf}}(\nu, x, \mathcal{D}) = \inf\{\text{KL}(\nu, \nu') : \nu' \in \mathcal{D} \text{ and } E(\nu') > x\},$$

where by convention, the infimum of an empty set equals $+\infty$ and where we denoted by $E(\nu')$ the expectation of ν' . The quantity $\mathcal{K}_{\text{inf}}(\nu, x, \mathcal{D})$ can be null. With the usual measure-theoretic conventions, in particular, $0/0 = 0$, we then have the following lower bound.

Reminder 1. *For all models \mathcal{D} , for all uniformly fast convergent strategies on \mathcal{D} , for all bandit problems $\underline{\nu}$ in \mathcal{D} ,*

$$\liminf_{T \rightarrow +\infty} \frac{R_T(\underline{\nu})}{\ln T} \geq \sum_{a \in [K]} \frac{\Delta_a}{\mathcal{K}_{\text{inf}}(\nu_a, \mu^*, \mathcal{D})}.$$

When the range $[m, M]$ of payoffs is known, i.e., when the model is $\mathcal{D}_{m,M}$, there exist strategies achieving the lower bound of Reminder 1, like the DMED strategy of Honda and Takemura [23, 24] or the KL-UCB strategy of Cappé et al. [9] and Garivier et al. [20]. (This can even be extended to the case of semi-bounded only rewards with a known upper bound on the payoffs, as is discussed in details in Appendix E.2.1.)

No logarithmic regret distribution-dependent regret bound under adaptation to the range.

Now, the lower bound in Reminder 1 cannot be achieved any more when the range is not known, that is, when we consider the model $\mathcal{D}_{-,+}$ of bounded distributions with unknown range. Actually, the proof reveals that the important fact is that the upper end of the payoff range is unknown. The impossibility result also holds for models $\mathcal{D}_{m,+}$ of bounded distributions with unknown upper end on the range and known lower end m on the range, for some fixed $m \in \mathbb{R}$:

$$\mathcal{D}_{m,+} = \bigcup_{\substack{M \in \mathbb{R} \\ m < M}} \mathcal{D}_{m,M}. \quad (11)$$

Theorem 4. *All uniformly fast convergent strategies on $\mathcal{D}_{-,+}$ are such that, for all bandit problems $\underline{\nu}$ in $\mathcal{D}_{-,+}$ with at least one suboptimal arm,*

$$\liminf_{T \rightarrow +\infty} \frac{R_T(\underline{\nu})}{\ln T} = +\infty.$$

The same result holds for all models $\mathcal{D}_{m,+}$, where $m \in \mathbb{R}$.

Strategies that are adaptive to the range thus cannot get rates Φ for distribution-dependent regret bounds on the regret of the order of $\ln T$ in Definition 2. A similar phenomenon was discussed by Lattimore [29] in the case of stochastic bandits with sub-Gaussian distributions. It turns out that any rate Φ such that $\Phi(T) \gg \ln T$ may be achieved, through a simple upper-confidence bound [UCB] strategy, as also discussed by Lattimore [29]; see further details after the proof.

Proof. We fix $m \in \mathbb{R}$ and provide the proof for $\mathcal{D}_{m,+}$. Given Reminder 1 and since we assumed that at least one arm a is suboptimal, i.e., is associated with a gap $\Delta_a = \mu^* - \mu_a > 0$, it is necessary and sufficient to show that $\mathcal{K}_{\text{inf}}(\nu_a, \mu^*, \mathcal{D}_{m,+}) = 0$, where $\nu_a \in \mathcal{D}_{m,+}$.

We have in particular $\mu_a \geq m$. We use the same construction as in the proof of Theorem 1. Let $\nu'_\varepsilon = (1 - \varepsilon)\nu_a + \varepsilon\delta_{\mu_a + 2\Delta_a/\varepsilon}$ for $\varepsilon \in (0, 1)$: it is a bounded probability distribution, with lower end of support larger than m , that is, $\nu'_\varepsilon \in \mathcal{D}_{m,+}$. For ε small enough, $\mu_a + 2\Delta_a/\varepsilon$ lies outside of the bounded support of ν_a . In that case, the density of ν_a with respect to ν'_ε is given by $1/(1 - \varepsilon)$ on the support of ν_a (and 0 elsewhere), so that

$$\text{KL}(\nu_a, \nu'_\varepsilon) = \ln\left(\frac{1}{1 - \varepsilon}\right).$$

Moreover, $\mathbb{E}(\nu'_\varepsilon) = (1 - \varepsilon)\mu_a + \varepsilon(\mu_a + 2\Delta_a/\varepsilon) = \mu_a + 2\Delta_a = \mu^* + \Delta_a > \mu^*$. Therefore, by definition of \mathcal{K}_{inf} as an infimum,

$$\mathcal{K}_{\text{inf}}(\nu_a, \mu^*, \mathcal{D}_{m,+}) \leq \text{KL}(\nu_a, \nu'_\varepsilon) = \ln\left(\frac{1}{1 - \varepsilon}\right).$$

This upper bound holds for all $\varepsilon > 0$ small enough and thus shows that we actually have $\mathcal{K}_{\text{inf}}(\nu_a, \mu^*, \mathcal{D}_{m,+}) = 0$.

The exact same construction and proof can be performed in the case of $\mathcal{D}_{-,+}$, without the need of indicating that the lower end of the support of ν'_ε is larger than m . \square

UCB with an increased exploration rate adapts to the range The lower bound of Theorem 4 does not prevent distribution-dependent rates for adaptation that are arbitrarily larger than a logarithm. Consider UCB with indexes of the form

$$\widehat{\mu}_a(t) + \sqrt{\frac{\varphi(t)}{N_a(t)}} \quad \text{where} \quad \frac{\varphi(t)}{\ln t} \rightarrow +\infty$$

and where $\widehat{\mu}_a(t)$ denotes the empirical average of payoffs obtained till round t when playing arm a . Following the analysis of Lattimore [29] in the case of Gaussian bandits with unknown variances, it can be shown that such a UCB is adaptive to the unknown range of payoffs with a distribution-dependent rate φ . However the trick used here is purely asymptotic and gives up on finite-time guarantees.

D Complete proofs of the results of Section 4

We provide here complete proofs for Corollary 1, Theorem 2, and Theorem 3, in this order.

D.1 Proof of Corollary 1

Proof of Corollary 1. We have, first,

$$\sum_{t=K+1}^T \gamma_t \leq \sqrt{5(1-\alpha)K \ln K} \sum_{t=K+1}^T t^{-\alpha} \leq \sqrt{5(1-\alpha)K \ln K} \int_0^T \frac{1}{t^\alpha} dt = \sqrt{\frac{5K \ln K}{1-\alpha}} T^{1-\alpha},$$

second, using the definition of γ_T as a minimum,

$$\frac{K \ln K}{\gamma_T} \leq \frac{K \ln K}{1/2} + \frac{T^\alpha K \ln K}{\sqrt{5(1-\alpha)K \ln K}} = 2K \ln K + \sqrt{\frac{K \ln K}{5(1-\alpha)}} T^\alpha,$$

and third, $\sqrt{T} \leq T^{\max\{\alpha, 1-\alpha\}}$, so that the regret bound of Theorem 2 may be further bounded by

$$R_T(y_{1:T}) \leq (M-m)\sqrt{K \ln K} \left(3 + 2\sqrt{\frac{5}{1-\alpha}} \right) T^{\max\{\alpha, 1-\alpha\}} + 10(M-m)K \ln K.$$

The claimed bound is obtained by bounding $2\sqrt{5}$ by 5. \square

D.2 Proof of Theorem 2

In Algorithm 1, for time steps $t \geq K+1$, the weights q_t are obtained by using the AdaHedge algorithm of De Rooij et al. [16] on the payoff estimates $\hat{y}_{t,a}$. AdaHedge is designed for the case of a full monitoring (not a bandit monitoring), but the use of these estimates emulates a full monitoring. Section 2.2 of De Rooij et al. [16] (see also an earlier analysis by Cesa-Bianchi et al. [12]) ensures the bound stated next in Reminder 2. For the sake of completeness, we rederive this bound in Appendix F.2. We call pre-regret the quantity at hand in Reminder 2: it corresponds to some regret defined in terms of the payoff estimates.

Reminder 2 (Application of Lemma 3 and Theorem 6 of De Rooij et al. [16]). *For all sequences of payoff estimates $\hat{y}_{t,a}$ lying in some bounded real-valued interval, denoted by $[b, B]$, for all $T \geq K+1$, the pre-regret of AdaHedge satisfies*

$$\max_{k \in [K]} \sum_{t=K+1}^T \hat{y}_{t,k} - \sum_{t=K+1}^T \sum_{a=1}^K q_{t,a} \hat{y}_{t,a} \leq 2 \sum_{t=K+1}^T \delta_t$$

where

$$\sum_{t=K+1}^T \delta_t \leq \underbrace{\sqrt{\sum_{t=K+1}^T \sum_{a=1}^K q_{t,a} \left(\hat{y}_{t,a} - \sum_{k \in [K]} q_{t,k} \hat{y}_{t,k} \right)^2 \ln K}}_{\leq \sqrt{\sum_{t=K+1}^T \sum_{a=1}^K q_{t,a} (\hat{y}_{t,a} - c)^2 \ln K} \text{ for any } c \in \mathbb{R}} + (B-b) \left(1 + \frac{2}{3} \ln K \right)$$

and AdaHedge does not require the knowledge of $[b, B]$ to achieve this bound.

The bound of Reminder 2 will prove itself particularly handy for three reasons: first, it is valid for signed payoffs (payoffs in \mathbb{R}); second, it is adaptive to the range of payoffs; third, the right-hand side looks at first sight not intrinsic enough a bound (as it also depends on the weights q_t) but we will see later that this dependency is particularly useful.

We recall that we start the summation in Reminder 2 at $t = K+1$ because the AdaHedge algorithm is only started at this time, after the initial exploration. The bound holding “for any $c \in \mathbb{R}$ ” is obtained by a classical bound on the variance.

Proof of Theorem 2. We deal with the contribution of the initial exploration by using the inequality $\max(u + v) \leq \max u + \max v$, together with the fact that $y_{t,a} - y_{t,A_t} \leq M - m$ for any $a \in [K]$:

$$R_T(y_{1:T}) \leq \underbrace{\max_{a \in [K]} \sum_{t=1}^K y_{t,a} - \mathbb{E} \left[\sum_{t=1}^K y_{t,A_t} \right]}_{\leq K(M-m)} + \max_{a \in [K]} \sum_{t=K+1}^T y_{t,a} - \mathbb{E} \left[\sum_{t=K+1}^T y_{t,A_t} \right]. \quad (12)$$

We now transform the pre-regret bound of Reminder 2, which is stated with the distributions q_t , into a pre-regret bound with the distributions p_t ; we do so while substituting the bounds $B = C + KM/\gamma_T$ and $b = C + Km/\gamma_T$ implied by (8) and the fact that (γ_t) is non-increasing, and by using the definition $q_{t,a} = p_{t,a} - \gamma_t(1/K - q_{t,a})$ for all $a \in [K]$:

$$\max_{k \in [K]} \sum_{t=K+1}^T \hat{y}_{t,k} - \sum_{t=K+1}^T \sum_{a=1}^K p_{t,a} \hat{y}_{t,a} + \sum_{t=K+1}^T \gamma_t \sum_{a=1}^K (1/K - q_{t,a}) \hat{y}_{t,a} \leq 2 \sum_{t=K+1}^T \delta_t \quad (13)$$

$$\text{where } \sum_{t=K+1}^T \delta_t \leq \sqrt{\sum_{t=K+1}^T \sum_{a=1}^K q_{t,a} (\hat{y}_{t,a} - C)^2 \ln K} + \frac{(M-m)K}{\gamma_T} \left(1 + \frac{2}{3} \ln K \right).$$

As noted by Auer et al. [6], by the very definition (6) of the estimates,

$$\sum_{a=1}^K p_{t,a} \hat{y}_{t,a} = y_{t,A_t}.$$

By (7), the tower rule and the fact that q_t is H_{t-1} -measurable, on the one hand, and the fact that the expectation of a maximum is larger than the maximum of expectations, on the other hand, the left-hand side of the first inequality in (13) thus satisfies

$$\begin{aligned} & \mathbb{E} \left[\max_{k \in [K]} \sum_{t=K+1}^T \hat{y}_{t,k} - \sum_{t=K+1}^T \sum_{a=1}^K p_{t,a} \hat{y}_{t,a} + \sum_{t=K+1}^T \gamma_t \sum_{a=1}^K (1/K - q_{t,a}) \hat{y}_{t,a} \right] \\ & \geq \max_{k \in [K]} \sum_{t=K+1}^T y_{t,k} - \mathbb{E} \left[\sum_{t=K+1}^T y_{t,A_t} \right] + \sum_{t=K+1}^T \gamma_t \left(\underbrace{\sum_{a=1}^K y_{t,a}/K}_{\in [m,M]} - \underbrace{\sum_{a=1}^K \mathbb{E}[q_{t,a}] y_{t,a}}_{\in [m,M]} \right) \\ & \geq \max_{k \in [K]} \sum_{t=K+1}^T y_{t,k} - \mathbb{E} \left[\sum_{t=K+1}^T y_{t,A_t} \right] - (M-m) \sum_{t=1}^T \gamma_t. \end{aligned}$$

As for the right-hand side of the second inequality in (13), we first note that by definition (see line 4 in Algorithm 1), $p_{t,a} \geq (1 - \gamma_t)q_{t,a}$ with $\gamma_t \leq 1/2$ by assumption on the extra-exploration rate, so that $q_{t,a} \leq 2p_{t,a}$; therefore, by substituting first this inequality and then by using Jensen's inequality,

$$\begin{aligned} \mathbb{E} \left[\sqrt{\sum_{t=K+1}^T \sum_{a=1}^K q_{t,a} (\hat{y}_{t,a} - C)^2 \ln K} \right] & \leq \sqrt{2} \mathbb{E} \left[\sqrt{\sum_{t=K+1}^T \sum_{a=1}^K p_{t,a} (\hat{y}_{t,a} - C)^2 \ln K} \right] \\ & \leq \sqrt{2} \sqrt{\sum_{t=K+1}^T \sum_{a=1}^K \mathbb{E} [p_{t,a} (\hat{y}_{t,a} - C)^2] \ln K}. \end{aligned} \quad (14)$$

Standard calculations (see Auer et al. [6] again) show, similarly to (7), that for all $a \in [K]$,

$$\mathbb{E} \left[p_{t,a} (\hat{y}_{t,a} - C)^2 \mid H_{t-1} \right] = \mathbb{E} \left[\frac{(y_{t,A_t} - C)^2}{p_{t,a}} \mathbb{1}_{\{A_t=a\}} \right] = (y_{t,a} - C)^2 \leq (M-m)^2,$$

where the last inequality comes from (8). By the tower rule, the same upper bound holds for the (unconditional) expectation. Therefore, taking the expectation of both sides of (13) and collecting all bounds together, we proved so far

$$R_T(y_{1:T}) \leq \underbrace{2\sqrt{2}}_{\leq 3} (M-m) \sqrt{KT \ln K} + (M-m) \frac{K \ln K}{\gamma_T} \underbrace{\left(\frac{2 + \gamma_T}{\ln K} + \frac{4}{3} \right)}_{\leq 5} + (M-m) \sum_{t=K+1}^T \gamma_t,$$

where we used $\gamma_T \leq 1/2$ and $\ln K \geq \ln 2$ as $K \geq 2$. \square

D.3 Proof of Theorem 3

Proof of Theorem 3. Given the decomposition (1) of the regret, it is necessary and sufficient to upper bound the expected number of times $\mathbb{E}[N_a(t)]$ any suboptimal arm a is drawn, where by definition of Algorithm 1,

$$\mathbb{E}[N_a(t)] = 1 + \mathbb{E} \left[\sum_{t=K+1}^T \left((1 - \gamma_t) q_{t,a} + \frac{\gamma_t}{K} \right) \right] \leq 1 + \sum_{t=K+1}^T \mathbb{E}[q_{t,a}] + \frac{1}{K} \sum_{t=K+1}^T \gamma_t.$$

We show below (and this is the main part of the proof) that

$$\sum_{t=K+1}^T \mathbb{E}[q_{t,a}] = \mathcal{O}(\ln T). \quad (15)$$

The proof of Corollary 1] shows in particular that

$$\frac{1}{K} \sum_{t=K+1}^T \gamma_t \leq \sqrt{\frac{5 \ln K}{(1 - \alpha)K}} T^{1-\alpha}.$$

Substituting the value (9) of $\Phi_{\text{free}(T)}^{\text{AH}}$ and using the decomposition (1) of $R_T(\nu)$ into $\sum \Delta_a \mathbb{E}[N_a(t)]$ then yield

$$\frac{R_T(\nu)}{T/\Phi_{\text{free}(T)}^{\text{AH}}} \leq \sum_{a \in [K]} \Delta_a \sqrt{\frac{5 \ln K}{(1 - \alpha)K}} \left(3 + \frac{5}{\sqrt{1 - \alpha}} \right) \sqrt{K \ln K} (1 + o(1)) + \mathcal{O} \left(\frac{\ln T}{T^{1-\alpha}} \right),$$

from which the stated bound follows, via the crude inequality $3\sqrt{5}\sqrt{1 - \alpha} + 5 \leq 12$.

Structure of the proof of (15). Let a^* denote an optimal arm. By definition of $q_{t,a}$ and by lower bounding a sum of exponential terms by any of the summands, we get

$$q_{t,a} = \frac{\exp \left(\eta_t \sum_{s=K+1}^{t-1} \hat{y}_{t,a} \right)}{\sum_{k=1}^K \exp \left(\eta_t \sum_{s=K+1}^{t-1} \hat{y}_{t,k} \right)} \leq \exp \left(\eta_t \sum_{t=K+1}^{t-1} (\hat{y}_{t,a} - \hat{y}_{t,a^*}) \right).$$

Then, by separating cases, depending on whether $\sum_{t=K+1}^{t-1} (\hat{y}_{t,a} - \hat{y}_{t,a^*})$ is smaller or larger than $-(t-1-K)\Delta_a/2$, and by remembering that the probability $q_{t,a}$ is always smaller than 1, we get

$$\begin{aligned} \sum_{t=K+1}^T \mathbb{E}[q_{t,a}] &\leq \sum_{t=K+1}^T \mathbb{E} \left[\exp \left(-\eta_t \frac{(t-1-K)\Delta_a}{2} \right) \right] \\ &\quad + \sum_{t=K+1}^T \mathbb{P} \left[\sum_{s=K+1}^{t-1} (\hat{y}_{s,a} - \hat{y}_{s,a^*}) \geq -\frac{(t-1-K)\Delta_a}{2} \right]. \end{aligned} \quad (16)$$

We show that the sums in the right-hand side of (16) are respectively $\mathcal{O}(1)$ and $\mathcal{O}(\ln T)$.

First sum in the right-hand side of (16). Given the definition of the learning rates (see the statement of Algorithm 1), namely,

$$\eta_t = \ln K \Big/ \sum_{s=K+1}^{t-1} \delta_s, \quad (17)$$

we are interested in upper bounds on the sum of the δ_s . Such upper bounds were already derived in the proof of Theorem 2; the second inequality in (13) together with the bound $q_{t,a} \leq 2p_{t,a}$ stated in the middle of the proof immediately yield

$$\begin{aligned} \sum_{s=K+1}^{t-1} \delta_s &\leq \sqrt{\sum_{s=K+1}^t \sum_{a=1}^K q_{s,a} (\hat{y}_{s,a} - C)^2 \ln K + \frac{(M-m)K}{\gamma_t} \left(1 + \frac{2}{3} \ln K \right)} \\ &\leq \sqrt{2} \sqrt{\sum_{s=K+1}^t \sum_{a=1}^K p_{s,a} (\hat{y}_{s,a} - C)^2 \ln K + \frac{(M-m)K}{\gamma_t} \left(1 + \frac{2}{3} \ln K \right)}. \end{aligned}$$

Unlike what we did to complete the proof of Theorem 2, we do not take expectations and rather proceed with deterministic bounds. By the definition (6) of the estimated payoffs for the equality below, by (8) for the first inequality below, and by the fact that the exploration rates are non-increasing for the second inequality below, we have, for all $s \geq K + 1$,

$$\sum_{a=1}^K p_{s,a} (\hat{y}_{s,a} - C)^2 = \frac{(y_{s,A_s} - C)^2}{p_{s,A_s}} \leq \frac{(M - m)^2}{\gamma_s/K} \leq \frac{(M - m)^2}{\gamma_t/K}. \quad (18)$$

Therefore,

$$\sum_{s=K+1}^{t-1} \delta_s \leq \sqrt{2}(M - m) \sqrt{\frac{t K \ln K}{\gamma_t}} + \frac{(M - m)K}{\gamma_t} \left(1 + \frac{2}{3} \ln K\right) \stackrel{\text{def}}{=} D_t = \Theta\left(\sqrt{t/\gamma_t} + 1/\gamma_t\right).$$

For the sake of concision, we denoted by D_t the obtained bound. Via the definition (17) of η_t , the sum of interest is in turn bounded by

$$\sum_{t=K+1}^T \exp\left(-\eta_t(t-1-K)\frac{\Delta_a}{2}\right) \leq \sum_{t=K+1}^T \exp\left(-\frac{\Delta_a \ln K}{2} \frac{t-1-K}{D_t}\right) = \mathcal{O}(1),$$

where the equality to $\mathcal{O}(1)$, i.e., the fact that the considered series is bounded, follows from the fact that

$$-(t-1-K)/D_t = \Theta\left(\sqrt{t\gamma_t} + t\gamma_t\right) = \Theta\left(t^{(1-\alpha)/2} + t^{1-\alpha}\right).$$

Second sum in the right-hand side of (16). We will use Bernstein's inequality for martingales, and more specifically, the formulation of the inequality by Freedman [18, Thm. 1.6] (see also Massart [31, Section 2.2]), as stated next.

Reminder 3. Let $(X_n)_{n \geq 1}$ be a martingale difference sequence with respect to a filtration $(\mathcal{F}_n)_{n \geq 0}$, and let $N \geq 1$ be a summation horizon. Assume that there exist real numbers b and v_N such that, almost surely,

$$\forall n \leq N, \quad X_n \leq b \quad \text{and} \quad \sum_{n=1}^N \mathbb{E}[X_n^2 | \mathcal{F}_{n-1}] \leq v_N.$$

Then for all $\delta \in (0, 1)$,

$$\mathbb{P}\left[\sum_{n=1}^N X_n \geq \sqrt{2v_N \ln \frac{1}{\delta}} + \frac{b}{3} \ln \frac{1}{\delta}\right] \leq \delta.$$

For $s \geq K + 1$, we consider the increments $X_s = \Delta_a - \hat{y}_{s,a^*} + \hat{y}_{s,a}$, which are adapted to the filtration $\mathcal{F}_s = \sigma(A_1, Z_1, \dots, A_s, Z_s)$, where we recall that Z_1, \dots, Z_s denote the payoffs obtained in rounds $1, \dots, s$. Also, as p_s is measurable with respect to past information \mathcal{F}_{s-1} and since payoffs are drawn independently from everything else (see Section 2), we have, by the definition (6) of the estimated payoffs (where we rather denote by $Y_{s,a}$ the payoffs drawn at random according to ν_a , to be in line with the notation of Section 2 for stochastic bandits): for all $a \in [K]$,

$$\mathbb{E}[\hat{y}_{s,a} | \mathcal{F}_{s-1}] = \frac{\mathbb{E}[Y_{s,a} | \mathcal{F}_{s-1}] - C}{p_{s,a}} \mathbb{1}_{\{A_s=a\}} + C = \frac{\mu_a - C}{p_{s,a}} \mathbb{1}_{\{A_s=a\}} + C = \mu_a.$$

As a consequence, $\mathbb{E}[X_s | \mathcal{F}_{s-1}] = \mathbb{E}[\Delta_a - \hat{y}_{s,a^*} + \hat{y}_{s,a} | \mathcal{F}_{s-1}] = 0$. Put differently, $(X_s)_{s \geq K+1}$ is indeed a martingale difference sequence with respect to the filtration $(\mathcal{F}_s)_{s \geq K}$.

We now check that the additional assumptions of Reminder are satisfied. Manipulations and arguments similar to the ones used in (8) and (18) show that for all $s \geq K + 1$,

$$\begin{aligned} \Delta_a - \hat{y}_{s,a^*} + \hat{y}_{s,a} &\leq \Delta_a - \frac{Y_{s,a^*} - C}{p_{s,a}} \mathbb{1}_{\{A_s=a^*\}} + \frac{Y_{s,a} - C}{p_{s,a}} \mathbb{1}_{\{A_s=a\}} \\ &\leq (M - m)(1 + K/\gamma_s) \leq b \stackrel{\text{def}}{=} (M - m)(1 + K/\gamma_t). \end{aligned}$$

For the variance bound, we first note that for all $s \leq t-1$, we have $(\hat{y}_{s,a} - C)(\hat{y}_{s,a^*} - C) = 0$ because of the indicator functions, and therefore,

$$\begin{aligned} \mathbb{E}\left[(\Delta_a - \hat{y}_{s,a^*} + \hat{y}_{s,a})^2 \mid \mathcal{F}_{s-1}\right] &\leq \mathbb{E}\left[(\hat{y}_{s,a^*} + \hat{y}_{s,a})^2 \mid \mathcal{F}_{s-1}\right] \\ &\leq \mathbb{E}\left[(\hat{y}_{s,a^*} - C)^2 \mid \mathcal{F}_{s-1}\right] + \mathbb{E}\left[(\hat{y}_{s,a} - C)^2 \mid \mathcal{F}_{s-1}\right]; \end{aligned}$$

in addition, for all $a \in [K]$ (including a^*),

$$\mathbb{E}\left[(\hat{y}_{s,a} - C)^2 \mid \mathcal{F}_{s-1}\right] = \mathbb{E}\left[\frac{(Y_{s,A_s} - C)^2}{p_{s,a}^2} \mathbb{1}_{\{A_s=a\}} \mid \mathcal{F}_{s-1}\right] \leq \frac{(M-m)^2}{p_{s,a}} \leq \frac{(M-m)^2 K}{\gamma_t}.$$

Therefore

$$\sum_{s=K+1}^{t-1} \mathbb{E}\left[(\Delta_a - \hat{y}_{s,a^*} + \hat{y}_{s,a})^2 \mid \mathcal{F}_{s-1}\right] \leq \frac{2K(M-m)^2(t-1-K)}{\gamma_t} \leq v_t \stackrel{\text{def}}{=} \frac{2(M-m)^2 t K}{\gamma_t}.$$

Bernstein's inequality (Reminder 3) may thus be applied; the choice $\delta = 1/t$ therein leads to

$$\mathbb{P}\left[\sum_{s=K+1}^{t-1} (\Delta_a - (\hat{y}_{s,a^*} - \hat{y}_{s,a})) \geq \underbrace{2(M-m)\sqrt{\frac{tK}{\gamma_t} \ln t + \frac{M-m}{3} \left(1 + \frac{K}{\gamma_t}\right) \ln t}}_{\stackrel{\text{def}}{=} D'_t}\right] \leq \frac{1}{t}.$$

As $\sqrt{t/\gamma_t} = \mathcal{O}(t^{(1+\alpha)/2})$ and $1/\gamma_t = \mathcal{O}(t^\alpha)$ as $t \rightarrow \infty$, where $\alpha < 1$, and as $\Delta_a > 0$ (given that we are considering a suboptimal arm a), there exists $t_0 \in \mathbb{N}$ such that for all $t \geq t_0$,

$$D'_t \leq \frac{(t-1-K)\Delta_a}{2}$$

thus

$$\begin{aligned} \mathbb{P}\left[\sum_{s=K+1}^{t-1} (\hat{y}_{s,a} - \hat{y}_{s,a^*}) \geq -\frac{(t-1-K)\Delta_a}{2}\right] &= \mathbb{P}\left[\sum_{s=K+1}^{t-1} (\Delta_a - (\hat{y}_{s,a^*} - \hat{y}_{s,a})) \geq \frac{(t-1-K)\Delta_a}{2}\right] \\ &\leq \mathbb{P}\left[\sum_{s=K+1}^{t-1} (\Delta_a - (\hat{y}_{s,a^*} - \hat{y}_{s,a})) \geq D'_t\right] \leq \frac{1}{t}. \end{aligned}$$

Therefore, as $T \rightarrow \infty$

$$\sum_{t=1}^T \mathbb{P}\left[\sum_{t=K+1}^{t-1} (\hat{y}_{t,a} - \hat{y}_{t,a^*}) \geq -\frac{(t-1-K)\Delta_a}{2}\right] = \mathcal{O}(\ln T),$$

as claimed. This concludes the proof. \square

E The case of one known end of the payoff range (bandits with gains or with losses)

In this section, we only discuss distribution-free and distribution-dependent upper bounds on the regret, as well as distribution-dependent lower bounds on the regret. This is because the $(M - m)\sqrt{KT}$ distribution-free regret lower bound of Auer et al. [6] holds even in the case when both ends m and M of the range are known.

We identified two difficulties in this paper when the range of bounded payoffs is unknown. First, no $\ln T$ rate for distribution-dependent bounds may be achieved, see (3) and Theorem 4. Second, there exists a trade-off between distribution-free and distribution-dependent rates for range adaptation, see Theorem 1. It turns out that when the upper end M on the payoff range is known, these difficulties (should) disappear. On the contrary, they remain when only the lower end m on the payoff range is known. These statements are detailed and proved below. We therefore contribute to enlightening the difference in nature between bandits with gains and bandits with losses, a topic that was already discussed by Cesa-Bianchi and Lugosi [10, Remark 6.5, page 164] and Kwon and Perchet [27].

E.1 Known lower end m but unknown upper end M on the payoff range

This case corresponds to considering the model $\mathcal{D}_{m,+}$ defined in (11) as

$$\mathcal{D}_{m,+} = \bigcup_{\substack{M \in \mathbb{R}, \\ m < M}} \mathcal{D}_{m,M}.$$

What is discussed below actually also holds for the larger model $\mathcal{D}_{m,+\infty}$ consisting of probability distributions with a first moment supported $[m, +\infty)$. Note that we have the strict inclusion $\mathcal{D}_{m,+} \subset \mathcal{D}_{m,+\infty}$ as distributions in $\mathcal{D}_{m,+\infty}$ are not bounded in general.

Definitions 1 and 2 handle the case of $\mathcal{D}_{-,+}$ but can be adapted in an obvious way to $\mathcal{D}_{m,+}$ by fixing m , by having the strategy know m , and require the bounds to hold for all $M \in [m, +\infty)$ and all bandit problems in $\mathcal{D}_{m,M}$. We then refer to scale-free distribution-free regret bounds and distribution-dependent rates for adaptation to the upper end of the range.

We already explained that the construction used to prove Theorem 4 not only works for $\mathcal{D}_{-,+}$ but also for $\mathcal{D}_{m,+}$. It turns out that the exact same construction was considered in Theorem 1: defining $\nu'_a = (1 - \varepsilon)\nu_a + \varepsilon\delta_{\mu_a + 2\Delta_a/\varepsilon}$ from a distribution ν_a . When $\nu_a \in \mathcal{D}_{m,+}$, we also have that ν'_a is a bounded distribution, with support lower bounded by m , that is $\nu'_a \in \mathcal{D}_{m,+}$. The proof and thus the result of Theorem 1 thus also holds for the case of $\mathcal{D}_{m,+}$.

E.2 Known upper end M but unknown lower end m on the payoff range

When the upper end M of the payoff range is known, $\ln T$ distribution-dependent regret rates are possible and there exists an algorithm achieving the optimal problem-dependent constant (Section E.2.1). Also, \sqrt{KT} scale-free distribution-free regret upper bounds may be achieved (Section E.2.2), which exactly match the distribution-free lower bound. We could not exhibit a strategy that would simultaneously achieve both optimal distribution-dependent and distribution-free regret bounds, unlike what is known in the case of a known payoff range (the KL-UCB-switch strategy by Garivier et al. [20]). We however conjecture that this should be possible and that, at least, no trade-off exists between the two bounds (unlike the one imposed by Theorem 1).

The case considered in this subsection corresponds to the models $\mathcal{D}_{-,M}$, for $M \in \mathbb{R}$, defined as

$$\mathcal{D}_{-,M} = \bigcup_{\substack{m \in \mathbb{R}, \\ m < M}} \mathcal{D}_{m,M}.$$

Some of the results actually also hold more generally for semi-bounded payoffs, which correspond to the models $\mathcal{D}_{-\infty,M}$, for $M \in \mathbb{R}$, defined as the sets of probability distributions with a first moment supported on $(-\infty, M]$. Note that we have the strict inclusion $\mathcal{D}_{-,M} \subset \mathcal{D}_{-\infty,M}$ as distributions in $\mathcal{D}_{-\infty,M}$ are not bounded in general.

E.2.1 Known M but unknown m , part 1: distribution-dependent bounds

We may again adapt Definitions 1 and 2 to define the concepts of distribution-free and distribution-dependent rates for adaptation to the lower end of the range, by considering the models $\mathcal{D}_{-,M}$ or $\mathcal{D}_{-\infty,M}$ therein. The DMED strategy of Honda and Takemura [24] does achieve a $\ln T$ distribution-dependent rate for adaptation to the lower end of the range and is even competitive against all bandit problems in $\mathcal{D}_{-\infty,M}$. The achieved upper bound is asymptotically optimal as indicated by Reminder 1.

Reminder 4 (Honda and Takemura [24], main theorem). *The regret of the DMED strategy is bounded, for all bandit problems $\underline{\nu}$ in $\mathcal{D}_{-\infty,M}$, as*

$$\limsup_{T \rightarrow \infty} \frac{R_T(\underline{\nu})}{\ln T} \leq \sum_{a=1}^K \frac{\Delta_a}{\mathcal{K}_{\text{inf}}(\nu_a, \mu^*, \mathcal{D}_{-\infty,M})}.$$

The nice and deep result of Reminder 4 implies that from the distribution-dependent point of view, adaptation to the lower end m of the range is automatic (if such a lower end exists: result holds also when there is no lower bound on the payoffs). Our intuition and understanding for this situation is the following. When the model is $\mathcal{D}_{m,M}$ for known ends m and M , the optimal constant for the $\ln T$ regret is given (see again Reminder 1) for all bandit problems $\underline{\nu}$ in $\mathcal{D}_{m,M}$ by

$$C(\underline{\nu}, m, M) = \sum_{a=1}^K \frac{\Delta_a}{\mathcal{K}_{\text{inf}}(\nu_a, \mu^*, \mathcal{D}_{m,M})}.$$

But it actually turns out, as indicated by Proposition 1 below, that $C(\underline{\nu}, m, M)$ is independent of m and equals $C(\underline{\nu}, -\infty, M)$.

Proposition 1. *Fix $M \in \mathbb{R}$. For all $m \leq M$, for all $\nu \in \mathcal{D}_{m,M}$ and all $\mu > E(\nu)$,*

$$\mathcal{K}_{\text{inf}}(\nu, \mu, \mathcal{D}_{m,M}) = \mathcal{K}_{\text{inf}}(\nu, \mu, \mathcal{D}_{-\infty,M}).$$

Proof. The inequality \geq is immediate, as the right-hand side of the equality is an infimum over the larger set $\mathcal{D}_{-\infty,M}$. For the inequality \leq , we may assume with no loss of generality that $\mu < M$, as otherwise, there is no distribution ν' neither in $\mathcal{D}_{m,M}$ nor in $\mathcal{D}_{-\infty,M}$ with $E(\nu') > \mu \geq M$, so that both \mathcal{K}_{inf} quantities equal $+\infty$.

We fix M, m, ν and μ as in the statement of the proposition. It suffices to show that in the case $\mu < M$, for all $\nu' \in \mathcal{D}_{-\infty,M}$ with $E(\nu') > \mu$ and $\nu \ll \nu'$, there exists $\nu'' \in \mathcal{D}_{m,M}$ with $E(\nu'') > \mu$ and $\text{KL}(\nu, \nu'') \leq \text{KL}(\nu, \nu')$. (If ν is not absolutely continuous with respect to ν' , then $\text{KL}(\nu, \nu') = +\infty$ and taking ν'' as the Dirac mass δ_M at M is a suitable choice.) To do so, given such a distribution ν' , we first note that $\nu \ll \nu'$ and $\nu \in \mathcal{D}_{m,M}$, i.e., $\nu([m, M]) = 1$, entail that $\nu'([m, M]) > 0$, so that we may define the restriction $\nu'' = \nu'|_{[m, M]}$ of ν' to $[m, M]$; its density with respect to ν' is given by

$$\frac{d\nu''}{d\nu'}(x) = \nu'([m, M])^{-1} \mathbf{1}_{\{x \in [m, M]\}} \quad \nu'\text{-a.s. for all } x \in \mathbb{R}.$$

We have the absolute-continuity chain $\nu \ll \nu'' \ll \nu'$, and the Radon-Nykodym derivatives thus defined satisfy

$$\frac{d\nu}{d\nu'}(x) = \frac{d\nu}{d\nu''}(x) \frac{d\nu''}{d\nu'}(x) = \nu'([m, M])^{-1} \frac{d\nu}{d\nu''}(x) \mathbf{1}_{\{x \in [m, M]\}} \quad \nu'\text{-a.s. for all } x \in \mathbb{R}. \quad (19)$$

Moreover $E(\nu'') \geq E(\nu')$, and thus $E(\nu'') > \mu$, as

$$\begin{aligned} E(\nu') &= \int_{(-\infty, m)} x d\nu'(x) + \int_{[m, M]} x d\nu'(x) \\ &\leq \left(1 - \nu'([m, M])\right)m + \nu'([m, M]) E(\nu'') \leq E(\nu''). \end{aligned}$$

Finally, by (19), which also holds ν -almost surely, and the definition of Kullback-Leibler divergences,

$$\begin{aligned} \text{KL}(\nu, \nu') &= \int_{(-\infty, M]} \ln\left(\frac{d\nu}{d\nu'}\right) d\nu = -\ln \nu'([m, M]) + \int_{[m, M]} \ln\left(\frac{d\nu}{d\nu''}\right) d\nu \\ &= -\ln \nu'([m, M]) + \text{KL}(\nu, \nu'') \geq \text{KL}(\nu, \nu''). \end{aligned}$$

This concludes the proof. \square

E.2.2 Known M but unknown m , part 2: distribution-free bounds

A first observation is that (as in the case of a fully known payoff range) AdaHedge does not require any extra-exploration (i.e., any mixing with the uniform distribution) to achieve a scale-free distribution-free regret bound of order $(M - m)\sqrt{KT \ln K}$. This is formally detailed in Appendix F.3. Both this result and the one described next rely on the AdaFTRL methodology of Orabona and Pál [33], which we recall in Appendix F.1.

The INF strategy of Audibert and Bubeck [2] can be seen as an instance of FTRL with $1/2$ -Tsallis entropy, as essentially noted by Audibert et al. [4]. The INF strategy provides a distribution-free regret bound of order \sqrt{KT} in case of a known payoff range. Up to some technical issues, which we could solve, it may be extended to provide a similar scale-free distribution regret bound, which is optimal as it does not contain any superfluous $\sqrt{\ln K}$ factor. The exact statement to be proved in Appendix F.4 is the following: AdaFTRL with $1/2$ -Tsallis entropy relying on an upper bound M on the payoffs ensures that for all $m \in \mathbb{R}$ with $m \leq M$, for all oblivious individual sequences y_1, y_2, \dots in $[m, M]^K$, for all $T \geq 1$,

$$R_T(y_{1:T}) \leq 4(M - m)\sqrt{KT} + 2(M - m).$$

We conclude this section by providing a high-level idea of the technical issues that were solved to obtain the latter bound. We consider estimates $\hat{y}_{t,a}$ obtained from (6) by replacing the constant C therein by the known upper end M . We however could not simply derive the regret bound from some generic full-information regret guarantee for AdaFTRL with $1/2$ -Tsallis entropy, as to the best of our knowledge, there are no meaningful full-information regret bounds for Tsallis entropy in the first place, and as these would anyway scale with the effective range of the estimates. We instead provide a more careful analysis exploiting special properties of the estimates: $\hat{y}_{t,a} = M$ for all $a \neq A_t$ and $\hat{y}_{t,A_t} \leq M$.

Open problem 1. We however were unable so far to provide a non-trivial scale-free distribution-dependent regret bound for our strategy AdaFTRL with $1/2$ -Tsallis entropy. Note that there exist $\mathcal{O}(\ln T)$ bounds for FTRL with $1/2$ -Tsallis entropy, i.e., with a different tuning of the learning rates (namely, η_t of order $1/\sqrt{t}$, but then, the range adaptive distribution-free guarantees are lost); see Zimmert and Seldin [37]. We would have liked to prove such a $\mathcal{O}(\ln T)$ scale-free distribution-dependent regret bound for AdaFTRL with $1/2$ -Tsallis entropy (or even achieve a more modest aim like a poly-logarithmic bound), as this seems possible and would have shown with certainty that the trade-off imposed by Theorem 1 does not hold anymore when the upper end M on the payoff range is known. The techniques of Seldin and Lugosi [35], which consist in a precise tuning of the extra-exploration in their variant of the Exp3 algorithm of Auer et al. [6] together with a gap estimation scheme, or the ones of Zimmert and Seldin [37] might be helpful to that end. We leave this problem for future research.

Open problem 2. For the sake of completeness, we underline here that either getting rid of the $\sqrt{\ln K}$ factor in the scale-free distribution-free regret bound of AdaHedge for K -armed bandits in the general case of an unknown upper end M on the payoff range, or, alternatively, exhibiting a larger lower bound of order $\sqrt{KT \ln K}$ for this scale-free distribution-free regret, is also a problem that we could not solve yet.

F Known results on AdaFTRL and AdaHedge in full information and applications thereof in the bandit setting

The aim of this section is two-fold: first, we provide, for the sake of self-completeness, a proof of the full-information bound for AdaHedge (Reminder 2 in Appendix D); second, we state and prove the improved bandit regret bounds alluded at in Appendix E.2.2, in the case of a known upper end M but unknown lower end m of the payoff range. We do respectively so in Appendices F.2 (for the full-information bound for AdaHedge) and in Appendices F.3 and F.4 (for the improved bandit regret bounds).

All these bounds can be put under the umbrella of the AdaFTRL methodology of Orabona and Pál [33] (AdaFTRL stands for adaptive follow-the-regularized-leader), which we recall, again for the sake of self-completeness, in Section F.1. This AdaFTRL methodology was partially built on and inspired the analysis for AdaHedge, which is a special case of AdaFTRL with entropic regularizer (see De Rooij et al. [16] for AdaHedge, as well as the earlier analysis by Cesa-Bianchi et al. [12]). Koolen [26] actually proposes an alternative analysis of AdaFTRL, closer to the AdaHedge formulation, namely, using directly some mixability gaps instead of upper bounds thereon; this is the analysis we recall below in Section F.1.

F.1 AdaFTRL for full information (reminder of known results)

To avoid confusion with the notation used in the main body of the paper, we first describe the considered setting of prediction of oblivious individual sequences with full information.

Full-information setting. The game between the player and the environment is actually the same as the one described in Section 2.2, except that the player observes at each step the entire payoff vector, not just the obtained payoff. More formally (and with a different piece of notation z instead of y , to better distinguish the two settings), the environment first picks a sequence of payoff vectors $z_t \in \mathbb{R}^K$, for all $t \geq 1$. Then, in a sequential manner, at every time step t , the player picks an action A_t , distributed according to a probability p_t over the action set $[K]$, obtains the payoff z_{t,A_t} , and observes the entire vector z_t (i.e., also the payoffs $z_{t,a}$ corresponding to the actions $a \neq A_t$).

In the sequel, we denote by \mathcal{S} the simplex of probability distributions over $[K]$ and we use the short-hand notation, for $p \in \mathcal{S}$ and $z \in \mathbb{R}^K$,

$$\langle p, z \rangle = \sum_{a \in [K]} p_a z_a.$$

FTRL (follow-the-regularized-leader). The FTRL method consists in choosing p_t according to

$$p_t \in \operatorname{argmin}_{p \in \mathcal{S}: F(p) < +\infty} \left\{ \frac{F(p)}{\eta_t} - \sum_{s=1}^{t-1} \langle p, z_s \rangle \right\},$$

where $F : \mathbb{R}^K \rightarrow \mathbb{R} \cup \{+\infty\}$ is a convex function, called the regularizer, and η_t is a non-negative learning rate in $(0, +\infty]$, which may depend on past observations. The condition $F(p) < +\infty$ will always be satisfied for some $p \in \mathcal{S}$ by the considered regularizers (see below) and is only meant to avoid the undefined $+\infty / +\infty$ in the case $\eta_t = +\infty$. For the sake of concision we will however omit it in the sequel.

Let us give a succinct account of the convex analysis results we use here, following the exposition of Lattimore and Szepesvári [30, Chapter 26]. Using their terminology, the domain $\operatorname{Dom} L$ of a convex function $L : \mathbb{R}^K \rightarrow \mathbb{R} \cup \{+\infty\}$ is the set $\{x \in \mathbb{R}^K : L(x) < +\infty\}$ of those points where it takes finite values. A convex function $L : \mathbb{R}^K \rightarrow \mathbb{R} \cup \{+\infty\}$ is said to be Legendre if the interior of its domain $\operatorname{Int}(\operatorname{Dom} L)$ is non-empty, if L is strictly convex and differentiable on $\operatorname{Int}(\operatorname{Dom} L)$, and if its gradient ∇L blows up on the boundary of $\operatorname{Dom} L$. The minimizers of Legendre functions may be seen to satisfy the following properties.

Proposition 2 (Special case of Lattimore and Szepesvári [30, Proposition 26.14]). *Let L be a Legendre function and $A \subseteq \mathbb{R}^d$ be a convex set that intersects $\operatorname{Int}(\operatorname{Dom} L)$. Then L possesses a unique minimizer x^* over A , which belongs to $\operatorname{Int}(\operatorname{Dom} L)$, therefore ensuring that L is differentiable at x^* . Furthermore,*

$$\forall x \in A \cap \operatorname{Dom} L, \quad \langle \nabla L(x^*), x - x^* \rangle \geq 0.$$

Finally, for $x, y \in \mathbb{R}^d$, if $F : \mathbb{R}^K \rightarrow \mathbb{R} \cup \{+\infty\}$ is differentiable at y , we define the Bregman divergence between x and y as

$$B_F(x, y) = F(x) - F(y) - \langle \nabla F(y), x - y \rangle; \quad (20)$$

when F is convex, we have $B_F(x, y) \geq 0$ for all $x \in \mathbb{R}^d$.

We are now ready to state our first reminder, which is a classical regret bound for FTRL (see, e.g., Lattimore and Szepesvári [30, Chapter 28, Exercise 28.12] for references, and McMahan [32] for more general versions). It involves the diameter D_F of the action set (the K -dimensional simplex \mathcal{S} in our case):

$$D_F = \max_{p, q \in \mathcal{S}} \{F(p) - F(q)\}.$$

Reminder 5 (Generic full-information FTRL bound over the simplex). *The FTRL method with a Legendre regularizer F (of finite diameter D_F) and with any rule for picking the learning rates so that they form a non-increasing sequence satisfies the following guarantee: for all sequences z_1, z_2, \dots of vector payoffs in \mathbb{R}^K , the regret is bounded by*

$$\begin{aligned} \max_{a \in [K]} \sum_{t=1}^T z_{t,a} - \sum_{t=1}^T \langle p_t, z_t \rangle &\leq \frac{D_F}{\eta_T} + \sum_{t=1}^{T-1} \left(\langle p_t - p_{t+1}, -z_t \rangle - \frac{B_F(p_{t+1}, p_t)}{\eta_t} \right) \\ &\quad + \left(\langle p_T - p^*, -z_T \rangle - \frac{B_F(p^*, p_T)}{\eta_T} \right), \end{aligned} \quad (21)$$

$$\text{where } p^* \in \operatorname{argmax}_{p \in \mathcal{S}} \sum_{t=1}^T \langle p, z_t \rangle$$

and where the regret bound is well defined, thanks to the following observations and conventions: for rounds $t \geq 1$ where $\eta_t < +\infty$, the function F is indeed differentiable at p_t so that $B_F(p_{t+1}, p_t)$ is well defined; for rounds $t \geq 1$ where $\eta_t = +\infty$, we set $B_F(p_{t+1}, p_t)/\eta_t = 0$ irrespectively of the fact whether F is differentiable at p_t .

Proof of Reminder 5. Denote by S_t the cumulative vector payoff up to time $t \geq 1$. Fix $T \geq 1$. For the sake of concision of the equations, we define $p_{T+1} = p^*$, which is a Dirac mass at some arm (that is, p_{T+1} is not given by FTRL). The regret can therefore be rewritten as

$$\begin{aligned} \max_{a \in [K]} \sum_{t=1}^T z_{t,a} - \sum_{t=1}^T \langle p_t, z_t \rangle &= \max_{p \in \mathcal{S}} \sum_{t=1}^T \langle p, z_t \rangle - \sum_{t=1}^T \langle p_t, z_t \rangle \\ &= \sum_{t=1}^T \langle p_{T+1}, z_t \rangle - \sum_{t=1}^T \langle p_t, z_t \rangle = \sum_{t=1}^T \langle p_t - p_{T+1}, -z_t \rangle. \end{aligned}$$

By summation by parts,

$$\begin{aligned} &\sum_{t=1}^T \langle p_t - p_{T+1}, -z_t \rangle \\ &= \sum_{t=1}^T \sum_{s=t}^T \langle p_s - p_{s+1}, -z_t \rangle = \sum_{s=1}^T \sum_{t=1}^s \langle p_s - p_{s+1}, -z_t \rangle = \sum_{s=1}^T \langle p_s - p_{s+1}, -S_s \rangle \\ &= \sum_{t=1}^T \langle p_t - p_{t+1}, -z_t \rangle + \sum_{t=1}^T \langle p_t - p_{t+1}, -S_{t-1} \rangle. \end{aligned} \quad (22)$$

If $\eta_t < +\infty$, then by the optimality condition from Proposition 2 applied to the Legendre function $L : x \mapsto \eta_t^{-1} F(x) - \langle S_{t-1}, x \rangle$, we know that L thus F are differentiable at p_t and that

$$\langle \eta_t^{-1} \nabla F(p_t) - S_{t-1}, p_{t+1} - p_t \rangle \geq 0,$$

$$\text{that is, } \langle p_t - p_{t+1}, -S_{t-1} \rangle \leq \langle \eta_t^{-1} \nabla F(p_t), p_{t+1} - p_t \rangle.$$

If $\eta_t = +\infty$, the previous inequality holds too, as by definition of p_t , we have $\langle p_t - p_{t+1}, -S_{t-1} \rangle \leq 0$ and as we set by convention $\eta_t^{-1} \nabla F(p_t) = 0$ regardless of whether F is differentiable at p_t or not. Substituting in (22), we proved so far

$$\sum_{t=1}^T \langle p_t - p_{T+1}, -z_t \rangle \leq \sum_{t=1}^T \langle p_t - p_{t+1}, -z_t \rangle + \langle \eta_t^{-1} \nabla F(p_t), p_{t+1} - p_t \rangle. \quad (23)$$

This inequality can be rewritten in terms of Bregman divergences:

$$\sum_{t=1}^T \langle p_t - p^*, -z_t \rangle \leq \sum_{t=1}^T \left(\langle p_t - p_{t+1}, -z_t \rangle - \frac{B_F(p_{t+1}, p_t)}{\eta_t} \right) + \sum_{t=1}^T \frac{F(p_{t+1}) - F(p_t)}{\eta_t}$$

We now upper bound the second sum in the right-hand side: again by summation by parts, with the convention $\eta_0 = +\infty$ and $1/\eta_0 = 0$:

$$\begin{aligned} \sum_{t=1}^T \frac{F(p_{t+1}) - F(p_t)}{\eta_t} &= \sum_{t=1}^T (F(p_{t+1}) - F(p_t)) \sum_{s=1}^t \left(\frac{1}{\eta_s} - \frac{1}{\eta_{s-1}} \right) \\ &= \sum_{s=1}^T \sum_{t=s}^T (F(p_{t+1}) - F(p_t)) \left(\frac{1}{\eta_s} - \frac{1}{\eta_{s-1}} \right) = \sum_{s=1}^T \underbrace{(F(p_{T+1}) - F(p_s))}_{\leq D_F} \underbrace{\left(\frac{1}{\eta_s} - \frac{1}{\eta_{s-1}} \right)}_{\geq 0} \leq \frac{D_F}{\eta_T}, \end{aligned}$$

where the final equality is obtained by a telescoping sum, using that the sequence of learning rates is non-increasing. \square

AdaFTRL, an adaptive version of FTRL. The AdaFTRL approach consists in tuning the learning rate in a way that scales with the observed data. More precisely, it relies on a quantity called the (generalized) mixability gap, which naturally appears as an upper bound on the summands in the FTRL bound of Reminder 5:

$$\delta_t^F \stackrel{\text{def}}{=} \max_{p \in \mathcal{S}} \left\{ \langle p_t - p, -z_t \rangle - \frac{B_F(p, p_t)}{\eta_t} \right\} \geq 0. \quad (24)$$

That mixability gaps are always nonnegative can be seen by taking $p = p_t$ in the definition. We may further upper bound (21) when it holds by using this mixability gap:

$$\max_{a \in [K]} \sum_{t=1}^T z_{t,a} - \sum_{t=1}^T \langle p_t, z_t \rangle \leq \frac{D_F}{\eta_T} + \sum_{t=1}^T \delta_t^F. \quad (25)$$

The AdaFTRL learning rate balances the two terms in the above regret bound by taking

$$\eta_t = D_F \left/ \sum_{s=1}^{t-1} \delta_s^F \right. \in (0, +\infty] \quad (26)$$

Note that this rule for picking learning rates indeed leads to non-increasing sequences thereof, as the mixability gaps are non-negative. We summarize the discussion above in the theorem stated next, from which subsequent (closed-from) regret bounds will be derived by using the specific properties of the regularizer F at hand to upper bound the mixability gaps.

Theorem 5 (AdaFTRL tool box). *Under the assumptions of Reminder 5 and with its conventions, the regret of the FTRL method based on the learning rates (26) satisfies*

$$\max_{a \in [K]} \sum_{t=1}^T z_{t,a} - \sum_{t=1}^T \langle p_t, z_t \rangle \leq 2 \sum_{t=1}^T \delta_t^F \quad (27)$$

where, moreover,

$$\left(\sum_{t=1}^T \delta_t^F \right)^2 = 2D_F \sum_{t=1}^T \frac{\delta_t^F}{\eta_t} + \sum_{t=1}^T (\delta_t^F)^2. \quad (28)$$

Proof. Inequality (27) follows from (25) and (26). The equality (28) is obtained by expanding the squared sum,

$$\left(\sum_{t=1}^T \delta_t^F \right)^2 = \sum_{t=1}^T (\delta_t^F)^2 + 2 \sum_{t=1}^T \sum_{s=1}^{t-1} \delta_t^F \delta_s^F = \sum_{t=1}^T (\delta_t^F)^2 + 2 \sum_{t=1}^T \delta_t^F \frac{D_F}{\eta_t}$$

where the final equality is obtained by substituting the definition (26) of η_t . \square

F.2 AdaHedge for full information (reminder of known results)

The content of this section is extracted from various sources, out of which the most important is Koolen [26]. We claim no novelty. This section recalls how the bound for AdaHedge (Reminder 2, for which a direct proof was provided by De Rooij et al. [16]) can also be seen as a special case of the results of Section F.1.

It is well-known (see Freund et al. [19], Kivinen and Warmuth [25], Audibert [1]) and can be found again by a simple optimization under a linear constraint that the Hedge weight update corresponds to FTRL with the negentropy as a regularizer:

$$H_{\text{neg}}(p) = \sum_{a=1}^K p_a \ln p_a,$$

with value $+\infty$ whenever $p_a = 0$ for some $a \in [K]$. That is,

$$\operatorname{argmin}_{p \in \mathcal{S}} \left\{ \frac{H_{\text{neg}}(p)}{\eta_t} - \sum_{s=1}^{t-1} \langle p, z_s \rangle \right\} = \{p_t\}$$

with $p_{t,a} = \exp \left(\eta_t \sum_{s=1}^{t-1} z_{s,a} \right) / \sum_{k=1}^K \exp \left(\eta_t \sum_{s=1}^{t-1} z_{k,s} \right).$ (29)

Straightforward calculation show that the regularizer H_{neg} is indeed Legendre (see Lattimore and Szepesvári [30], Example 26.11) and the H_{neg} -diameter of the simplex equals $D_{H_{\text{neg}}} = \ln K$. Reminder 5 and Theorem 5 can therefore be applied.

AdaHedge is exactly AdaFTRL with H_{neg} as a regularizer. Indeed, the mixability gap (24) can be computed in closed form (as noted by Reid et al. [34, Lemma 5]) and reads in this case:

$$\delta_t^{\text{neg}} = \begin{cases} -\langle p_t, z_t \rangle + \eta_t^{-1} \ln \left(\sum_{a=1}^K p_{t,a} e^{\eta_t z_{t,a}} \right) & \text{if } \eta_t < +\infty, \\ -\langle p_t, z_t \rangle + \max_{a \in [K]} z_{t,a} & \text{if } \eta_t = +\infty. \end{cases} \quad (30)$$

Proof of the rewriting (30). When $\eta_t = +\infty$, the mixability gap equals, by definition,

$$\delta_t^F = \max_{p \in \mathcal{S}} \{ \langle p_t - p, -z_t \rangle \} = -\langle p_t, z_t \rangle + \max_{p \in \mathcal{S}} \langle p, z_t \rangle = -\langle p_t, z_t \rangle + \max_{a \in [K]} z_{t,a}.$$

For the case $\eta_t < +\infty$, the following formula, which is at the heart of the closed-form formula for the Hedge updates (29), will be useful: for any $S \in \mathbb{R}^d$,

$$\min_{p \in \mathcal{S}} \left\{ H_{\text{neg}}(p) - \langle p, S \rangle \right\} = \sum_{i=1}^K \frac{e^{S_i}}{\sum_{j=1}^K e^{S_j}} \left(\ln \left(\frac{e^{S_i}}{\sum_{j=1}^K e^{S_j}} \right) - S_i \right) = -\ln \left(\sum_{i=1}^K e^{S_i} \right). \quad (31)$$

When $\eta_t < +\infty$, Equation (29) shows that p_t lies in the interior $\text{Int}(\mathcal{S})$ of \mathcal{S} . The Bregman divergence at hand in the definition (24) of the mixability gaps may be simplified into

$$B_F(p, p_t) = H_{\text{neg}}(p) - H_{\text{neg}}(p_t) - \langle \nabla H_{\text{neg}}(p_t), p - p_t \rangle = H_{\text{neg}}(p) - \langle \nabla H_{\text{neg}}(p_t), p \rangle + 1,$$

where the second inequality holds by taking into account the fact that H_{neg} is twice differentiable at any $p \in \text{Int}(\mathcal{S})$, with

$$\nabla H_{\text{neg}}(p) = (1 + \ln p_i)_{i \in [K]} \quad \text{so that} \quad \langle \nabla H_{\text{neg}}(p), p \rangle = 1 + \sum_{i=1}^K p_i \ln p_i = 1 + H_{\text{neg}}(p).$$

The mixability gaps can therefore be rewritten

$$\begin{aligned}\delta_t^F &= \max_{p \in \mathcal{S}} \left\{ \langle p_t - p, -z_t \rangle - \frac{B_F(p, p_t)}{\eta_t} \right\} \\ &= -\langle p_t, z_t \rangle - \frac{1}{\eta_t} + \frac{1}{\eta_t} \max_{p \in \mathcal{S}} \{ \eta_t \langle p, z_t \rangle - H_{\text{neg}}(p) + \langle \nabla H_{\text{neg}}(p_t), p \rangle \} \\ &= -\langle p_t, z_t \rangle - \frac{1}{\eta_t} - \frac{1}{\eta_t} \min_{p \in \mathcal{S}} \{ H_{\text{neg}}(p) - \langle p, \eta_t z_t + \nabla H_{\text{neg}}(p_t) \rangle \}\end{aligned}$$

Now by (31), specialized with $S = \eta_t z_t + \nabla H_{\text{neg}}(p_t)$, we can compute the value of the minimum:

$$\min_{p \in \mathcal{S}} \{ H_{\text{neg}}(p) - \langle p, \eta_t z_t + \nabla H_{\text{neg}}(p_t) \rangle \} = -\ln \left(\sum_{i=1}^K e^{\eta_t z_i + 1 + \ln p_i} \right) = -1 - \ln \left(\sum_{i=1}^K p_i e^{\eta_t z_i} \right).$$

Collecting all equalities together concludes the proof. \square

Reminder 2 is thus a special case of the following bound.

Theorem 6 (See Lemma 3 and Theorem 6 of De Rooij et al. [16]). *For all sequences of payoffs $z_{t,a}$ lying in some bounded real-valued interval, denoted by $[b, B]$, for all $T \geq 1$, the regret of the AdaHedge algorithm with full information, as defined by (29) and (30), satisfies*

$$\begin{aligned}\max_{k \in [K]} \sum_{t=1}^T z_{t,k} - \sum_{t=1}^T \sum_{a=1}^K p_{t,a} z_{t,a} &\leq 2 \sum_{t=1}^T \delta_t^{\text{neg}} \\ \text{where } \sum_{t=1}^T \delta_t^{\text{neg}} &\leq \sqrt{\sum_{t=1}^T \sum_{a=1}^K p_{t,a} \left(z_{t,a} - \sum_{k \in [K]} q_{t,k} z_{t,k} \right)^2} \ln K + (B - b) \left(1 + \frac{2}{3} \ln K \right),\end{aligned}$$

and AdaHedge does not require the knowledge of $[b, B]$ to achieve this bound.

The quantities

$$v_t \stackrel{\text{def}}{=} \sum_{a=1}^K p_{t,a} \left(z_{t,a} - \sum_{k \in [K]} q_{t,k} z_{t,k} \right)^2$$

in the bound correspond to the variance of the random variables taking values $z_{t,a}$ with probability $p_{t,a}$; the variational formula for variances indicates that

$$\sum_{a=1}^K p_{t,a} \left(z_{t,a} - \sum_{k \in [K]} q_{t,k} z_{t,k} \right)^2 = \min_{c \in \mathbb{R}} \sum_{a=1}^K p_{t,a} (z_{t,a} - c)^2,$$

which entails the final bound given as a note in the statement of Reminder 2.

The following formulation of Bernstein's inequality will be useful in the proof of Theorem 6.

Lemma 1 (Bernstein's inequality tailored to our needs). *Let X be a random variable in $[0, 1]$, with variance denoting by $\text{Var}(X)$. Then for all $\eta > 0$,*

$$\frac{\ln \left(\mathbb{E} \left[e^{\eta(X - \mathbb{E}[X])} \right] \right)}{\eta^2} \leq \frac{1}{2} \text{Var}(X) + \frac{1}{3} \frac{\ln \left(\mathbb{E} \left[e^{\eta(X - \mathbb{E}[X])} \right] \right)}{\eta}.$$

Proof. Denote by $\psi_X(\eta) = \ln \left(\mathbb{E} \left[e^{\eta(X - \mathbb{E}[X])} \right] \right)$ the log-moment generating function of X . A version of Bernstein's inequality with an appropriate control of the moments (as stated by Massart [31, Section 2.2.3] and applied to X with $c = 1/3$) indicates that for all $\eta \in (0, 3)$,

$$\left(1 - \frac{\eta}{3} \right) \psi_X(\eta) \leq \frac{\eta^2}{2} \text{Var}(X).$$

Actually, this inequality also holds for $\eta \geq 3$ as its left-hand side is non-positive while its right-hand side is nonnegative. The claimed result is derived by rearranging the terms

$$\psi_X(\eta) \leq \frac{\eta^2}{2} \text{Var}(X) + \frac{\eta}{3} \psi_X(\eta)$$

and by dividing both sides by η^2 . \square

Proof of Theorem 6. We apply Theorem 5. To that end, we first bound the mixability gaps. The rewriting (30) (and Jensen's inequality) directly shows that $0 \leq \delta_t^{\text{neg}} \leq B - b$. We may also prove the bound

$$\frac{\delta_t^{\text{neg}}}{\eta_t} \leq \frac{v_t}{2} + \frac{1}{3}(B - b)\delta_t^{\text{neg}}. \quad (32)$$

It suffices to do so for $\eta_t < +\infty$. Consider the random variable X taking values $(z_{t,a} - b)/(B - b)$ with probability $p_{t,a}$, for $a \in \{1, \dots, K\}$. The mixability gap can be rewritten as

$$\delta_t^{\text{neg}} = \frac{1}{\eta_t} \psi_X(\eta_t(B - b))$$

with the notation of the proof of Lemma 1. The variance of X equals $v_t/(B - b)^2$. Lemma 1 with $\eta = \eta_t(B - b)$ yields

$$\frac{\delta_t^{\text{neg}}}{\eta_t(B - b)^2} \leq \frac{v_t}{2(B - b)^2} + \frac{\delta_t^{\text{neg}}}{3(B - b)}.$$

from which we obtain (32) by rearranging.

From (28) and (32), we deduce, together with the bound $(\delta_t^{\text{neg}})^2 \leq (B - b)\delta_t^{\text{neg}}$, that

$$\left(\sum_{t=1}^T \delta_t^{\text{neg}} \right)^2 \leq (\ln K) \sum_{t=1}^T v_t + (B - b) \left(\frac{2}{3} \ln K + 1 \right) \sum_{t=1}^T \delta_t^{\text{neg}}.$$

Therefore, using the fact that $x^2 \leq a + bx$ implies $x \leq \sqrt{a} + b$ for all $a, b, x \geq 0$,

$$\sum_{t=1}^T \delta_t^{\text{neg}} \leq \sqrt{\ln K \sum_{t=1}^T v_t} + (B - b) \left(\frac{2}{3} \ln K + 1 \right),$$

which thanks to (27) concludes the proof of Theorem 6. \square

F.3 AdaHedge with known upper bound M on the payoffs (application of Section F.2)

We show how to obtain a scale-free distribution-free regret bound of order $(M - m)\sqrt{KT \ln K}$ with no extra-exploration (including no initial exploration) when an upper bound M on the payoffs is given to the player. We consider Algorithm 2, where no mixing takes place (unlike in Algorithm 1) and where the probability distributions p_t are directly computed via an AdaHedge update (no need for intermediate probabilities q_t). Note also that we use the estimates (6) with the choice $C_t = M$, that is,

$$\hat{y}_{t,a} = \frac{y_{t,a} - M}{p_{t,a}} \mathbb{1}_{\{A_t=a\}} + M. \quad (33)$$

The following observation is key in the analysis below: $\hat{y}_{t,a} = M$ for all $a \neq A_t$ and $\hat{y}_{t,A_t} \leq M$. We will also use, as in the proof of Theorem 2,

$$\sum_{a=1}^K p_{t,a} \hat{y}_{t,a} = y_{t,A_t}.$$

The performance bound for this simpler algorithm is stated next.

Theorem 7. *AdaHedge for K -armed bandits relying on an upper bound M on the payoffs (Algorithm 2) ensures that for all $m \in \mathbb{R}$ with $m \leq M$, for all oblivious individual sequences y_1, y_2, \dots in $[m, M]^K$, for all $T \geq 1$,*

$$R_T(y_{1:T}) \leq 2(M - m)\sqrt{KT \ln K} + 2(M - m).$$

Algorithm 2 AdaHedge for K -armed bandits, when an upper bound on the payoffs is given

- 1: **Input:** an upper bound M on the payoffs
- 2: **AdaHedge initialization:** $\eta_1 = +\infty$ and $p_1 = (1/K, \dots, 1/K)$
- 3: **for** rounds $t = 1, 2, \dots$ **do**
- 4: Draw an arm $A_t \sim p_t$ (independently at random according to the distribution p_t)
- 5: Get and observe the payoff y_{t,A_t}
- 6: Compute the estimates of all payoffs

$$\hat{y}_{t,a} = \frac{y_{t,a} - M}{p_{t,a}} \mathbb{1}_{\{A_t=a\}} + M$$

- 7: Compute the mixability gap δ_t based on the distribution p_t and on these estimates:

$$\delta_t = \begin{cases} -\sum_{a=1}^K p_{t,a} \hat{y}_{t,a} + \frac{1}{\eta_t} \ln \left(\sum_{a=1}^K p_{t,a} e^{\eta_t \hat{y}_{t,a}} \right) & \text{if } \eta_t < +\infty \\ -\sum_{a=1}^K p_{t,a} \hat{y}_{t,a} + \max_{a \in [K]} \hat{y}_{t,a} & \text{if } \eta_t = +\infty \end{cases}$$

- 8: Compute the learning rate $\eta_{t+1} = \left(\sum_{s=1}^t \delta_s \right)^{-1} \ln K$

- 9: Define p_{t+1} component-wise as

$$p_{t+1,a} = \exp \left(\eta_{t+1} \sum_{s=1}^t \hat{y}_{a,s} \right) / \sum_{k=1}^K \exp \left(\eta_{t+1} \sum_{s=1}^t \hat{y}_{k,s} \right)$$

- 10: **end for**

The main technical difference with respect to the analysis of Algorithm 1 is that the mixability gaps are directly bounded by the range $M - m$. We no longer need to artificially control the size of the estimates (which we did via extra-exploration) to get, in turn, a control of the mixability gaps.

Lemma 2 (Improved mixability gap bound). *The mixability gaps of AdaHedge for K -armed bandits relying on an upper bound M on the payoffs (Algorithm 2) are bounded, for all $m \in \mathbb{R}$ with $m \leq M$, for all oblivious individual sequences y_1, y_2, \dots in $[m, M]^K$, for all $t \geq 1$, by*

$$0 \leq \delta_t \leq M - m \quad \text{and} \quad \frac{\delta_t}{\eta_t} \leq \frac{1}{2} p_{t,A_t}^{-1} (M - y_{t,A_t})^2.$$

Proof. The fact that $\delta_t \geq 0$ holds by definition of the gaps and Jensen's inequality. For $\delta_t \leq M - m$, the observations after (33) indicate that when $\eta_t = +\infty$,

$$\delta_t = -\sum_{a=1}^K p_{t,a} \hat{y}_{t,a} + \max_{a \in [K]} \hat{y}_{t,a} = M - \hat{y}_{t,A_t},$$

while for $\eta_t < +\infty$,

$$\begin{aligned} \delta_t &= -y_{t,A_t} + \frac{1}{\eta_t} \ln \left((1 - p_{t,A_t}) e^{\eta_t M} + p_{t,A_t} e^{\eta_t M} e^{\eta_t (y_{t,A_t} - M)/p_{t,A_t}} \right) \\ &\leq M - y_{t,A_t} + \frac{1}{\eta_t} \ln \left((1 - p_{t,A_t}) + p_{t,A_t} \underbrace{e^{\eta_t (y_{t,A_t} - M)/p_{t,A_t}}}_{\leq 1} \right), \end{aligned}$$

which entails $\delta_t \leq M - y_{t,A_t} \leq M - m$.

Furthermore, in the case $\eta_t < +\infty$, using the inequality $e^{-x} \leq 1 - x + x^2/2$ valid for $x \geq 0$, followed by the inequality $\ln(1 + u) \leq u$, valid for all $u > -1$, we get

$$\delta_t \leq M - \hat{y}_{t,A_t} + \frac{1}{\eta_t} \ln \left(\underbrace{1 - p_{A_t,t} + p_{A_t,t}}_{=1} \underbrace{-\eta_t(M - y_{t,A_t}) + \eta_t^2 \frac{(M - y_{t,A_t})^2}{2p_{A_t,t}}}_{=u} \right) \leq \eta_t \frac{(M - y_{t,A_t})^2}{2p_{t,A_t}}.$$

The second inequality is trivial in case $\eta_t = +\infty$, as $\delta_t/\eta_t = 0$. \square

We are now ready to prove Theorem 7.

Proof of Theorem 7. As indicated in Section F.2, AdaHedge is a special case of AdaFTRL and the bound of Theorem 5 is applicable.

Equation (28) and Lemma 2, which entails in particular that $\delta_t^2 \leq (M - m)\delta_t$, yield

$$\left(\sum_{t=1}^T \delta_t \right)^2 = 2(\ln K) \sum_{t=1}^T \frac{\delta_t}{\eta_t} + \sum_{t=1}^T (\delta_t)^2 \leq (\ln K) \sum_{t=1}^T p_{t,A_t}^{-1} (M - y_{t,A_t})^2 + (M - m) \sum_{t=1}^T \delta_t,$$

which, through the fact that $x^2 \leq a + bx$ implies $x \leq \sqrt{a} + b$ for all $a, b, x \geq 0$, leads in turn to

$$\sum_{t=1}^T \delta_t \leq \sqrt{\sum_{t=1}^T p_{t,A_t}^{-1} (M - \hat{y}_{t,A_t})^2 \ln K} + (M - m).$$

Therefore, Equation (27) guarantees that

$$\max_{k \in [K]} \sum_{t=1}^T \hat{y}_{t,k} - \underbrace{\sum_{t=1}^T \sum_{a=1}^K p_{t,a} \hat{y}_{t,a}}_{=y_{t,A_t}} \leq 2 \sqrt{\sum_{t=1}^T p_{t,A_t}^{-1} (M - \hat{y}_{t,A_t})^2 \ln K} + 2(M - m). \quad (34)$$

We conclude the proof by integrating the inequality above and using Jensen's inequality, exactly as in the proof of Theorem 2. Indeed, Equation (12) therein indicates that

$$R_T(y_{1:T}) = \max_{k \in [K]} \sum_{t=1}^T y_{t,k} - \mathbb{E} \left[\sum_{t=1}^T y_{t,A_t} \right] \leq \mathbb{E} \left[\max_{k \in [K]} \sum_{t=1}^T \hat{y}_{t,k} - \sum_{t=1}^T y_{t,A_t} \right]$$

and, by the same manipulations as in (14) and in the equation that follows it,

$$\begin{aligned} \mathbb{E} \left[\sqrt{\sum_{t=1}^T p_{t,A_t}^{-1} (M - \hat{y}_{t,A_t})^2 \ln K} \right] &\leq \sqrt{\mathbb{E} \left[\sum_{t=1}^T p_{t,A_t}^{-1} (M - y_{t,A_t})^2 \ln K \right]} \\ &= \sqrt{\mathbb{E} \left[\sum_{t=1}^T \sum_{a=1}^K (M - y_{t,a})^2 \ln K \right]} \leq (M - m) \sqrt{KT \ln K} \end{aligned}$$

The claimed result is obtained by collecting all bounds together. \square

F.4 AdaFTRL with Tsallis entropy in the case of a known upper bound M on the payoffs

In this section we describe how the AdaHedge learning rate scheme can be used in the FTRL framework with a different regularizer, namely Tsallis entropy, to improve the scale-free distribution-free regret bound into a bound of optimal order $(M - m)\sqrt{KT}$, i.e., without any superfluous $\sqrt{\ln K}$ factor.

Tsallis entropy. We focus on the (rescaled) $1/2$ -Tsallis entropy, which is defined by

$$H_{1/2}(p) = - \sum_{a=1}^K 2\sqrt{p_a}.$$

This regularizer is Legendre over the domain $[0, +\infty)^K$ (see Lattimore and Szepesvári [30, Example 26.10]). Its diameter equals

$$D_{H_{1/2}} = \max_{p \in \mathcal{S}} H_{1/2}(p) - \min_{q \in \mathcal{S}} H_{1/2}(q) = -2 - (-2\sqrt{K}) = 2(\sqrt{K} - 1), \quad (35)$$

as for all $p \in \mathcal{S}$, we have (by concavity of the square root for the right-most inequality)

$$1 \leq \sum_{a=1}^K p_a \leq \sum_{a=1}^K \sqrt{p_a} \leq \sqrt{K},$$

where 1 is achieved with $p = (1, 0, \dots, 0)$ and \sqrt{K} with the uniform distribution.

The function $H_{1/2}$ is differentiable at all $q \in (0, +\infty)^K$, with $\nabla H_{1/2}(q) = (-1/\sqrt{q_a})_{a \in [K]}$. The Bregman divergence associated with $H_{1/2}$ equals, for $p, q \in \mathcal{S}$ such that $q_a > 0$ for all a :

$$\begin{aligned} B_{H_{1/2}}(p, q) &= -2 \sum_{a=1}^K \sqrt{p_a} + 2 \sum_{a=1}^K \sqrt{q_a} + \sum_{a=1}^K \frac{1}{\sqrt{q_a}} (p_a - q_a) \\ &= -2 \sum_{a=1}^K \frac{\sqrt{p_a} - \sqrt{q_a}}{2\sqrt{q_a}} (2\sqrt{q_a} - (\sqrt{p_a} + \sqrt{q_a})) = \sum_{a=1}^K \frac{(\sqrt{p_a} - \sqrt{q_a})^2}{\sqrt{q_a}}. \end{aligned}$$

AdaFTRL with $1/2$ -Tsallis entropy. We consider FTRL with the $1/2$ -Tsallis entropy on the estimated losses (33):

$$p_t \in \operatorname{argmin}_{p \in \mathcal{S}} \left\{ \frac{H_{1/2}(p)}{\eta_t} - \sum_{s=1}^{t-1} \langle p, \hat{y}_s \rangle \right\} = \operatorname{argmin}_{p \in \mathcal{S}} \left\{ -\frac{1}{\eta_t} \sum_{a=1}^K 2\sqrt{p_a} - \sum_{a=1}^K p_a \sum_{s=1}^{t-1} \hat{y}_{s,a} \right\}.$$

FTRL with the $1/2$ -Tsallis entropy was essentially introduced by Audibert and Bubeck [2] to get rid of a $\sqrt{\ln K}$ factor in the distribution-free regret bound of K -armed adversarial bandits (with known payoff range). It was later noted by Audibert et al. [4] that it actually is an instance of mirror descent with Tsallis entropy as a regularizer. More recently, Zimmert and Seldin [37] showed that this regularizer can obtain quasi-optimal regret bounds for both stochastic and adversarial rewards.

We more precisely consider AdaFTRL with the $1/2$ -Tsallis, that is, we compute the learning rates η_t based on the mixability gaps (24); see Algorithm 3. We denote by δ_t^{Ts} the mixability gaps (24).

On the implementation. For Tsallis entropy, the optimization problems involved in the computation of the updates p_t and of the mixability gaps δ_t^{Ts} admit a (semi-)explicit formula. Indeed, p_t can be computed thanks to the formula, for all $z \in \mathbb{R}^K$,

$$\operatorname{argmin}_{p \in \mathcal{S}} \{ H_{1/2}(p) - \langle p, z \rangle \} = \operatorname{argmax}_{p \in \mathcal{S}} \left\{ \langle p, z \rangle + \sum_{a=1}^K 2\sqrt{p_a} \right\} = \left(\frac{1}{(c(z) - z_a)^2} \right)_{a \in [K]}, \quad (36)$$

where $c(z)$ is an implicit normalization constant, such that the vector lies in the simplex \mathcal{S} and $c(z) > z_a$ for all $a \in [K]$. This constant $c(z)$ is in fact the Lagrange multiplier associated with the constraint $p_1 + \dots + p_K = 1$. See Zimmert and Seldin [37] for more details on how to compute $c(z)$ efficiently, see also Audibert et al. [4]. To compute the mixability gap, rewrite

$$\begin{aligned} \delta_t^{\text{Ts}} &= \max_{p \in \mathcal{S}} \left\{ \langle p_t - p, -\hat{y}_t \rangle - \frac{H_{1/2}(p) - H_{1/2}(p_t) - \langle \nabla H_{1/2}(p_t), p - p_t \rangle}{\eta_t} \right\} \\ &= \langle p_t, -\hat{y}_t \rangle + \frac{H_{1/2}(p_t)}{\eta_t} - \frac{\langle \nabla H_{1/2}(p_t), p_t \rangle}{\eta_t} + \frac{1}{\eta_t} \max_{p \in \mathcal{S}} \left\{ \langle p, \nabla H_{1/2}(p_t) + \eta_t \hat{y}_t \rangle - H_{1/2}(p) \right\}, \end{aligned} \quad (37)$$

where the maximum in the left-most side of these equalities can be computed efficiently, thanks to (36).

Algorithm 3 AdaFTRL with Tsallis entropy for K -armed bandits, when an upper bound on the payoffs is given

- 1: **Input:** an upper bound M on the payoffs
- 2: **Initialization:** $\eta_1 = +\infty$ and $p_1 = (1/K, \dots, 1/K)$
- 3: **for** rounds $t = 1, 2, \dots$ **do**
- 4: Draw an arm $A_t \sim p_t$ (independently at random according to the distribution p_t)
- 5: Get and observe the payoff y_{t,A_t}
- 6: Compute the estimates of all payoffs

$$\hat{y}_{t,a} = \frac{y_{t,a} - M}{p_{t,a}} \mathbb{1}_{\{A_t=a\}} + M$$

- 7: Compute the mixability gap δ_t^{Ts} based on the distribution p_t and on these estimates, e.g., using the efficient implementation stated around (37):

$$\delta_t^{\text{Ts}} = \max_{p \in \mathcal{S}} \left\{ \langle p_t - p, -\hat{y}_t \rangle - \frac{B_{H_{1/2}}(p, p_t)}{\eta_t} \right\}$$

- 8: Compute the learning rate $\eta_{t+1} = 2 \left(\sum_{s=1}^t \delta_s^{\text{Ts}} \right)^{-1} (\sqrt{K} - 1)$
- 9: Define p_{t+1} as

$$p_{t+1} \in \operatorname{argmin}_{p \in \mathcal{S}} \left\{ - \sum_{a=1}^K p_a \sum_{s=1}^t \hat{y}_{s,a} - \frac{1}{\eta_{t+1}} \sum_{a=1}^K 2\sqrt{p_a} \right\},$$

where an efficient implementation is provided by, e.g., (36)

10: **end for**

Analysis of the algorithm. We provide the following performance bound.

Theorem 8. *AdaFTRL with $1/2$ -Tsallis entropy for K -armed bandits relying on an upper bound M on the payoffs (Algorithm 3) ensures that for all $m \in \mathbb{R}$ with $m \leq M$, for all oblivious individual sequences y_1, y_2, \dots in $[m, M]^K$, for all $T \geq 1$,*

$$R_T(y_{1:T}) \leq 4(M - m)\sqrt{KT} + 2(M - m).$$

As in Section F.3, the proof scheme is a combination of the AdaFTRL bound of Theorem 5 (which is indeed applicable), together with an improved bound on the mixability gap that exploits the specific shape of the estimates. This bound is stated in the next lemma, which is much similar to Lemma 2.

Lemma 3. *The mixability gaps of AdaFTRL with Tsallis entropy for K -armed bandits relying on an upper bound M on the payoffs (Algorithm 3) are bounded, for all $m \in \mathbb{R}$ with $m \leq M$, for all oblivious individual sequences y_1, y_2, \dots in $[m, M]^K$, for all $t \geq 1$, by*

$$0 \leq \delta_t^{\text{Ts}} \leq M - m \quad \text{and} \quad \frac{\delta_t^{\text{Ts}}}{\eta_t} \leq p_{t,A_t}^{-1/2} (M - y_{t,A_t})^2.$$

The proof of Lemma 3 is postponed to the end of this section and we now proceed with the proof of Theorem 8.

Proof of Theorem 8. The structure of the proof is much similar to the one of Theorem 7, which is why we only sketch our arguments. The bound of Theorem 5 is applicable. We use Lemma 3 with (28) to see that

$$\left(\sum_{t=1}^T \delta_t^{\text{Ts}} \right)^2 \leq 2D_{H_{1/2}} \sum_{t=1}^T p_{t,A_t}^{-1/2} (M - y_{t,A_t})^2 + (M - m) \sum_{t=1}^T \delta_t^{\text{Ts}}. \quad (38)$$

Again, using the fact that for all $a, b, x \geq 0$, the inequality $x^2 \leq a + bx$ implies $x \leq \sqrt{a + b}$:

$$\sum_{t=1}^T \delta_t^{\text{Ts}} \leq \sqrt{2D_{H_{1/2}} \sum_{t=1}^T p_{t,A_t}^{-1/2} (M - y_{t,A_t})^2 + (M - m)} \quad (39)$$

By (27), by taking expectations, and by Jensen's inequality:

$$R_T(y_{1:T}) \leq 2\mathbb{E} \left[\sum_{t=1}^T \delta_t^{\text{Ts}} \right] \leq 2\sqrt{2D_{H_{1/2}} \sum_{t=1}^T \mathbb{E} \left[p_{t,A_t}^{-1/2} (M - y_{t,A_t})^2 \right] + 2(M - m)}. \quad (40)$$

We conclude by observing that for all t , by definition of the payoff estimates,

$$\begin{aligned} \mathbb{E} \left[p_{t,A_t}^{-1/2} (M - y_{t,A_t})^2 \right] &= \mathbb{E} \left[\sum_{a=1}^K p_{t,a} p_{t,a}^{-1/2} (M - y_{t,a})^2 \right] \leq (M - m)^2 \mathbb{E} \left[\sum_{a=1}^K \sqrt{p_{t,a}} \right] \\ &\leq (M - m)^2 \sqrt{K}, \end{aligned}$$

where the last inequality follows from the concavity of the square root. The final claim is obtained by bounding the diameter $D_{H_{1/2}}$ by $2\sqrt{K}$. \square

We conclude this section by providing a proof of Lemma 3.

Proof of Lemma 3. The fact that $\delta_t^{\text{Ts}} \geq 0$ holds actually for all regularizers and can be seen from the definition (24) with $p = p_t$. For the inequality $\delta_t^{\text{Ts}} \leq M - m$, we start with elementary manipulations of the definition of the mixability gap (24). Denoting by \vec{M} the vector with coordinates (M, \dots, M) and noting that $\langle p_t - q, \vec{M} \rangle = 0$ for all $q \in \mathcal{S}$, we have

$$\delta_t^{\text{Ts}} = \max_{q \in \mathcal{S}} \left\{ \langle p_t - q, -\hat{y}_t \rangle - \frac{B_{H_{1/2}}(q, p_t)}{\eta_t} \right\} = \max_{q \in \mathcal{S}} \left\{ \langle p_t - q, \vec{M} - \hat{y}_t \rangle - \frac{B_{H_{1/2}}(q, p_t)}{\eta_t} \right\}. \quad (41)$$

Since all the coordinates of $\vec{M} - \hat{y}_t$ are non-negative and by non-negativity of the Bregman divergence, this implies that

$$\delta_t^{\text{Ts}} \leq \langle p_t, \vec{M} - \hat{y}_t \rangle = M - y_{A_t,t} \leq M - m.$$

We now prove the second inequality; we may assume that $\eta_t < +\infty$, as the bound holds trivially otherwise. By Proposition 2 (and by calculations similar to the ones performed in the proof of Remark 5) the maximum in the rewriting (41) of δ_t^{Ts} is achieved on the interior of the domain of $H_{1/2}$, which equals $(0, +\infty)^K$, thus in the interior of \mathcal{S} . We therefore only need to prove that

$$\forall q \in \text{Int}(\mathcal{S}), \quad \langle p_t - q, \vec{M} - \hat{y}_t \rangle - \frac{B_{H_{1/2}}(q, p_t)}{\eta_t} \leq \eta_t p_{t,A_t}^{-1/2} (M - y_{t,A_t})^2. \quad (42)$$

We fix such a $q \in \text{Int}(\mathcal{S})$, i.e., such that $q_a > 0$ for all a . We consider two cases. First, if $q_{A_t} \geq p_{t,A_t}$, then, given the observations made after (33),

$$\langle p_t - q, \vec{M} - \hat{y}_t \rangle - \frac{B_{H_{1/2}}(q, p_t)}{\eta_t} = \underbrace{\left(\frac{M - y_{t,A_t}}{p_{t,A_t}} \right)}_{\geq 0} \underbrace{(p_{t,A_t} - q_{A_t})}_{\leq 0} - \frac{B_{H_{1/2}}(q, p_t)}{\eta_t} \leq 0.$$

Otherwise, when $q_{A_t} < p_{t,A_t}$, a standard way of bounding the mixability gap, detailed below, indicates that

$$\langle p_t - q, \vec{M} - \hat{y}_t \rangle - \frac{B_{H_{1/2}}(q, p_t)}{\eta_t} \leq \frac{\eta_t}{2} \left\langle \vec{M} - \hat{y}_t, \nabla^2 H_{1/2}(z)^{-1} (\vec{M} - \hat{y}_t) \right\rangle, \quad (43)$$

where z is some probability distribution of the open segment $\text{Seg}(q, p_t)$ between q and p_t , and where $\nabla^2 H_{1/2}(z)^{-1}$ denotes the inverse of the positive definite Hessian of $H_{1/2}$ at z . Since at $w \in (0, +\infty)^K$, the function $H_{1/2}$ is indeed twice differentiable, with

$$\nabla H_{1/2}(w) = (-w_a^{-1/2})_{a \in [K]} \quad \text{and} \quad \nabla^2 H_{1/2}(w) = \text{Diag}(w_a^{-3/2}/2)_{a \in [K]},$$

we have $\nabla^2 H_{1/2}(z)^{-1} = \text{Diag}(2z_a^{3/2})_{a \in [K]}$. We substitute this value into (43) and recall that the vector $\vec{M} - \hat{y}_t$ has null coordinates except for its A_t -th coordinate:

$$\frac{\eta_t}{2} \left\langle \vec{M} - \hat{y}_t, \nabla^2 H_{1/2}(z)^{-1} (\vec{M} - \hat{y}_t) \right\rangle = \eta_t z_{A_t}^{3/2} (M - \hat{y}_{t,A_t})^2.$$

Finally, remember that z lies in the open segment $\text{Seg}(q, p_t)$ and that we assumed $q_{A_t} < p_{t,A_t}$; we thus also have $z_{A_t} < p_{t,A_t}$. As a consequence, using the very definition of \hat{y}_{t,A_t} ,

$$\eta_t z_{A_t}^{3/2} (M - \hat{y}_{t,A_t})^2 \leq \eta_t p_{t,A_t}^{3/2} (M - \hat{y}_{t,A_t})^2 = \eta_t p_{t,A_t}^{-1/2} (M - y_{t,A_t})^2.$$

Therefore, in all cases, that is, whether $q_{A_t} \geq p_{t,A_t}$ or $q_{A_t} < p_{t,A_t}$, the bound (42) is obtained. It only remains to prove the standard inequality (43).

This inequality is essentially stated as Theorem 26.13 in Lattimore and Szepesvári [30] but we provide a proof for the sake of completeness. As we assumed that $\eta_t < +\infty$, we have (as above, by Proposition 2) that p_t lies in the interior of \mathcal{S} . In particular, as both p_t and q are in the interior of \mathcal{S} , the function $H_{1/2}$ is \mathcal{C}^2 over the closed segment $\text{Seg}(q, p_t)$ between q and p_t . Therefore, by the mean-value theorem, there exists z in the open segment $\text{Seg}(q, p_t)$ such that

$$\underbrace{H_{1/2}(q) - H_{1/2}(p_t)}_{=B_{H_{1/2}}(q,p_t)} - \left\langle \nabla H_{1/2}(p_t), q - p_t \right\rangle = \frac{1}{2} \left\langle q - p_t, \nabla^2 H_{1/2}(z) (q - p_t) \right\rangle.$$

It is useful to introduce the standard notation from convex analysis for the local norm (which is indeed a norm because the Hessian is positive definite):

$$\|q - p_t\|_{\nabla^2 H_{1/2}(z)}^2 \stackrel{\text{def}}{=} \left\langle q - p_t, \nabla^2 H_{1/2}(z) (q - p_t) \right\rangle.$$

We therefore have so far the rewriting:

$$-\frac{B_{H_{1/2}}(q,p_t)}{\eta_t} = -\frac{1}{2\eta_t} \left\langle q - p_t, \nabla^2 H_{1/2}(z) (q - p_t) \right\rangle.$$

Now, by the Cauchy-Schwarz inequality,

$$\begin{aligned} \langle p_t - q, \vec{M} - \hat{y}_t \rangle &= \left\langle \nabla^2 H_{1/2}(z)^{1/2} (p_t - q), \nabla^2 H_{1/2}(z)^{-1/2} (\vec{M} - \hat{y}_t) \right\rangle \\ &\leq \|p_t - q\|_{\nabla^2 H_{1/2}(z)} \|\vec{M} - \hat{y}_t\|_{\nabla^2 H_{1/2}(z)^{-1}}. \end{aligned}$$

Combining the rewriting and the bound above, we get

$$\begin{aligned} \langle p_t - q, M - \hat{y}_t \rangle &= \frac{B_{H_{1/2}}(q,p_t)}{\eta_t} \\ &\leq \|p_t - q\|_{\nabla^2 H_{1/2}(z)} \|\vec{M} - \hat{y}_t\|_{\nabla^2 H_{1/2}(z)^{-1}} - \frac{1}{2\eta_t} \|q - p_t\|_{\nabla^2 H_{1/2}(z)}^2 \\ &\leq \frac{\eta_t}{2} \|\vec{M} - \hat{y}_t\|_{\nabla^2 H_{1/2}(z)^{-1}}^2, \end{aligned}$$

where we used $ab - b^2/2 \leq a^2/2$ to get the second inequality. This is exactly (43). \square

G Adaptation to the range for linear bandits

To illustrate the generality of the techniques discussed in this paper, we quickly describe how these can be used to obtain range adaptive algorithms for linear bandits. This section is meant for illustration and not for completeness. In particular, we focus on the case of (oblivious) adversarial linear bandits, for which we refer the reader to Lattimore and Szepesvári [30, Chapter 27], which we follow closely, for a more thorough description of the setting; we do not describe the application of our techniques to stochastic linear bandits.

Learning protocol. A finite action set $\mathcal{A} \subset \mathbb{R}^d$, of cardinality K , is given. (The setting of vanilla K -armed bandits considered in the rest of the article corresponds to \mathcal{A} formed by the vertices of the probability simplex of \mathbb{R}^K .) The environment selects beforehand a sequence $(y_t)_{t \geq 1}$ of vectors in \mathbb{R}^d satisfying a boundedness assumption: there exists an interval $[m, M]$ such that

$$\forall t \geq 1, \forall x \in \mathcal{A}, \quad x^\top y_t \in [m, M]. \quad (44)$$

We assume that the player does not know in advance m nor M . To simplify the exposition, we also assume that $m \leq 0 \leq M$.

At every time step, the player chooses an action $X_t \in \mathcal{A}$ and receives and only observes the payoff $X_t^\top y_t$. It does not observe y_t nor the payoffs $x^\top y_t$ associated with choices $x \neq X_t$. The action X_t is chosen independently at random according to a distribution over \mathcal{A} denoted by $p_t = (p_t(a))_{a \in \mathcal{A}}$.

The expected regret is defined as

$$R_T(y_{1:T}) = \max_{x \in \mathcal{A}} \sum_{t=1}^T x^\top y_t - \mathbb{E} \left[\sum_{t=1}^T X_t^\top y_t \right].$$

Estimating the unobserved payoffs. As in the case of vanilla K -armed bandits, the key is to estimate unobserved payoffs. We may actually build an estimate \hat{y}_t of the vectors y_t , from which we form the estimates $x^\top \hat{y}_t$. This estimate takes advantage of the linear structure of the problem.

Fix a distribution π such that the non-negative symmetric matrix

$$M(\pi) \stackrel{\text{def}}{=} \sum_{x \in \mathcal{A}} \pi(x) x x^\top$$

is invertible: such a distribution exists whenever \mathcal{A} spans \mathbb{R}^d , which we may assume with no loss of generality; see Lemma 4 below. This distribution π will be used to explore the arms; it is in general not uniform over the arms. For all distributions q over \mathcal{A} and all $\gamma \in (0, 1]$, the distribution $p = (1 - \gamma)q + \gamma\pi$ is such that the non-negative symmetric matrix $M(p)$ is invertible as well (as it is larger than $\gamma M(\pi)$, in the sense of the partial inequality \succcurlyeq over non-negative symmetric matrices). We then define

$$\hat{y}_t = M(p_t)^{-1} X_t X_t^\top y_t \quad (45)$$

and note that

$$\mathbb{E}[\hat{y}_t \mid p_t] = M(p_t)^{-1} \underbrace{\left(\sum_{x \in \mathcal{A}} p_t(x) x x^\top y_t \right)}_{=M(p_t)} = y_t; \quad (46)$$

indeed, conditioning on p_t amounts to integrating over the random choice of X_t according to p_t .

An algorithm adaptive to the unknown range. When the range is given, a well-known strategy is to use plain exponential weights over actions in \mathcal{A} with the estimates $x^\top \hat{y}_t$ to obtain distributions q_t that are then mixed with π to form the final distributions p_t . When the range is unknown, we suggest to simply replace plain exponential weights with AdaHedge (the difference lies in the tuning of the rates η_t), which leads to Algorithm 4. In this algorithm, we refer to rates γ_t as exploration rates (and not as extra-exploration rates as in Algorithm 1) and similarly, to π as the exploration distribution. This is because for adversarial linear bandits, exploration was always required even to get expected results (unlike for K -armed bandits, see the introduction of Section 4).

Algorithm 4 AdaHedge for adversarial linear bandits

- 1: **Input:** an exploration distribution π over \mathcal{A} and exploration rates $(\gamma_t)_{t \geq 1}$ in $[0, 1]$
- 2: **Initialization:** $\eta_1 = +\infty$ and q_1 is the uniform distribution over \mathcal{A}
- 3: **for** rounds $t = 1, \dots$ **do**
- 4: Define p_t by mixing q_t with π according to

$$p_t = (1 - \gamma_t)q_t + \gamma_t\pi$$

- 5: Draw an arm $X_t \sim p_t$ (independently at random according to the distribution p_t)
- 6: Get and observe the payoff $X_t^\top y_t$
- 7: Compute estimates $x^\top \hat{y}_t$ of all payoffs according to (45)
- 8: Compute the mixability gap δ_t based on the distribution q_t and on these estimates:

$$\delta_t = \begin{cases} -\sum_{x \in \mathcal{A}} q_t(x) x^\top \hat{y}_t + \frac{1}{\eta_t} \ln \left(\sum_{x \in \mathcal{A}} q_t(x) e^{\eta_t x^\top \hat{y}_t} \right) & \text{if } \eta_t < +\infty \\ -\sum_{x \in \mathcal{A}} q_t(x) x^\top \hat{y}_t + \max_{x \in \mathcal{A}} x^\top \hat{y}_t & \text{if } \eta_t = +\infty \end{cases}$$

- 9: Compute the learning rate $\eta_{t+1} = \left(\sum_{s=1}^t \delta_s \right)^{-1} \ln K$
 - 10: Define q_{t+1} component-wise as
- $$q_{t+1}(a) = \exp \left(\eta_{t+1} \sum_{s=1}^t a^\top \hat{y}_s \right) / \sum_{x \in \mathcal{A}} \exp \left(\eta_{t+1} \sum_{s=1}^t x^\top \hat{y}_s \right)$$
- 11: **end for**
-

The analysis of this algorithm relies on the same ingredients as the ones already encountered in Section 4.1, with the addition of the following lemma, that quantifies the quality of the exploration. This lemma requires that \mathcal{A} spans \mathbb{R}^d , which we may assume with no loss of generality (otherwise, we just replace \mathbb{R}^d by the vector space generated by \mathcal{A}).

Lemma 4 (Lattimore and Szepesvári [30, Theorem 21.1]). *There exists a distribution π over \mathcal{A} such that*

$$M(\pi) = \sum_{x \in \mathcal{A}} \pi(x) x x^\top \text{ is invertible} \quad \text{and} \quad \max_{x \in \mathcal{A}} x^\top M(\pi)^{-1} x = d.$$

We are now ready to state the main result of this section. It is the counterpart of Corollary 1; for the sake of simplicity, we only state it for the value $\alpha = 1/2$.

Theorem 9. *AdaHedge for adversarial linear bandits (Algorithm 4) with the extra-exploration*

$$\gamma_t = \min \left\{ 1/2, \sqrt{2.5 d (\ln K) t^{-1/2}} \right\}$$

ensures that for all bounded ranges $[m, M]$, for all oblivious individual sequences y_1, y_2, \dots satisfying the boundedness condition (44),

$$R_T(y_{1:T}) \leq 12(M - m) \sqrt{dT \ln K} + 18(M - m)d \ln K.$$

The proof starts by following closely the ones of Theorem 2 and Corollary 1; the differences are underlined and dealt with in the second part of the proof.

Proof. By Reminder 2, since the player plays the AdaHedge strategy over the payoff estimates $x^\top \hat{y}_t$, the pre-regret satisfies

$$\max_{x \in \mathcal{A}} \sum_{t=1}^T x^\top \hat{y}_t - \sum_{t=1}^T \sum_{a \in \mathcal{A}} q_t(a) a^\top \hat{y}_t \leq 2\sqrt{VT \ln K} + M_T \left(2 + \frac{4}{3} \ln K \right)$$

with $V_T = \sum_{t=1}^T \sum_{x \in \mathcal{A}} q_t(x) (x^\top \hat{y}_t)^2$ and

$$M_T = \max\{x^\top \hat{y}_t : t \leq T \text{ and } x \in \mathcal{A}\} - \min\{x^\top \hat{y}_t : t \leq T \text{ and } x \in \mathcal{A}\}.$$

As in Theorem 2, since $\gamma_t \leq 1/2$, we have $q_t(x) \leq 2p_t(x)$ for all $x \in \mathcal{A}$. We therefore define

$$V'_T = \sum_{t=1}^T \sum_{x \in \mathcal{A}} p_t(x) (x^\top \hat{y}_t)^2$$

and have $V_t \leq 2V'_T$. By the tower rule, based on the equality (46), and given that the expectation of a maximum is larger than the maximum of the expectations (for the first inequality), and by the definition of the p_t (for the second inequality), we have proved so far that

$$\begin{aligned} R_T(y_{1:T}) &\leq \mathbb{E} \left[\max_{x \in \mathcal{A}} \sum_{t=1}^T x^\top \hat{y}_t - \sum_{t=1}^T \sum_{a \in \mathcal{A}} p_t(a) a^\top \hat{y}_t \right] \\ &\leq \mathbb{E} \left[\max_{x \in \mathcal{A}} \sum_{t=1}^T x^\top \hat{y}_t - \sum_{t=1}^T \sum_{a \in \mathcal{A}} q_t(a) a^\top \hat{y}_t \right] + \mathbb{E} \left[\sum_{t=1}^T \gamma_t \sum_{a \in \mathcal{A}} (\pi(a) - q_t(a)) a^\top \hat{y}_t \right] \\ &\leq \mathbb{E} \left[2\sqrt{2V'_T \ln K} + M_T \left(2 + \frac{4}{3} \ln K \right) \right] + \underbrace{\sum_{t=1}^T \gamma_t \sum_{a \in \mathcal{A}} (\pi(a) - q_t(a)) a^\top y_t}_{\leq (M-m)} . \end{aligned}$$

Hence by Jensen's inequality and by the bounds $\mathbb{E}[V'_T] \leq (M-m)^2 dT$ and $M_T \leq 2(M-m)d/\gamma_T$ proved below, we finally get

$$\begin{aligned} R_T(y_{1:t}) &\leq 2\sqrt{2\mathbb{E}[V'_T] \ln K} + \mathbb{E}[M_T] \left(2 + \frac{4}{3} \ln K \right) + (M-m) \sum_{t=1}^T \gamma_t \\ &\leq 2\sqrt{2}(M-m)\sqrt{dT \ln K} + \left(2 + \frac{4}{3} \ln K \right) \frac{2(M-m)d}{\gamma_T} + (M-m) \sum_{t=1}^T \gamma_t \\ &\leq 3(M-m)\sqrt{dT \ln K} + 9(M-m) \frac{d \ln K}{\gamma_T} + (M-m) \sum_{t=1}^T \gamma_t . \end{aligned}$$

Replacing the γ_t by their values and using the same bounds as in Corollary 1 yields the claimed result; the factor 12 in the bound comes from

$$3 + \sqrt{10} + 9\sqrt{\frac{2}{5}} \leq 12.$$

We only need to prove the two claimed bounds to complete the proof; they can be extracted from the proof of Theorem 27.1 by Lattimore and Szepesvári [30] but we provide derivations for the sake of completeness.

Proof of $M_T \leq 2(M-m)d/\gamma_T$. We fix $x \in \mathcal{A}$ and $t \leq T$. We recall that $M(p_t)$ and thus $M(p_t)^{-1}$ are positive definite symmetric matrices. By the Cauchy-Schwarz inequality applied with the norm induced by the positive $M(p_t)^{-1}$,

$$|x^\top M(p_t)^{-1} X_t| \leq \sqrt{x^\top M(p_t)^{-1} x} \sqrt{X_t^\top M(p_t)^{-1} X_t} \leq \max_{x \in \mathcal{A}} \{x^\top M(p_t)^{-1} x\}.$$

As indicated right before (45), we have $M(p_t) \succcurlyeq \gamma_t M(\pi)$ and therefore $M(p_t)^{-1} \preccurlyeq M(\pi)^{-1}/\gamma_t$. This entails

$$|x^\top M(p_t)^{-1} X_t| \leq \frac{1}{\gamma_t} \max_{x \in \mathcal{A}} \{x^\top M(\pi)^{-1} x\} = \frac{d}{\gamma_t} \leq \frac{d}{\gamma_T},$$

where the equality follows from Lemma 4 and where we used $\gamma_T \leq \gamma_t$ for the second inequality. Finally, keeping in mind that we assumed $m \leq 0 \leq M$,

$$x^\top \hat{y}_t = \underbrace{x^\top M(p_t)^{-1} X_t}_{\in [-d/\gamma_t, d/\gamma_t]} \underbrace{X_t^\top y_t}_{\in [m, M]} \in \left[-\frac{d \max\{-m, M\}}{\gamma_T}, \frac{d \max\{-m, M\}}{\gamma_T} \right],$$

from which the bound

$$M_t = 2 \frac{d \max\{-m, M\}}{\gamma_T} \leq \frac{2d(M - m)}{\gamma_T}$$

follows, as desired.

Proof of $\mathbb{E}[V'_T] \leq (M - m)^2 dT$. Since $|X_t^\top y_t| \leq \max\{-m, M\} \leq M - m$, the definition (45) leads to

$$\begin{aligned} (x^\top \hat{y}_t)^2 &= \left(x^\top M(p_t)^{-1} X_t X_t^\top y_t \right)^2 \leq (M - m)^2 \left(x^\top M(p_t)^{-1} X_t \right)^2 \\ &= (M - m)^2 X_t^\top M(p_t)^{-1} x x^\top M(p_t)^{-1} X_t. \end{aligned}$$

Therefore, summing over $x \in \mathcal{A}$ and using the very definition of $M(p_t)$, we get

$$\begin{aligned} \sum_{x \in \mathcal{A}} p_t(x) (x^\top \hat{y}_t)^2 &\leq (M - m)^2 X_t^\top M(p_t)^{-1} \left(\sum_{x \in \mathcal{A}} p_t(x) x x^\top \right) M(p_t)^{-1} X_t \\ &= (M - m)^2 X_t^\top M(p_t)^{-1} X_t = (M - m)^2 \text{Tr} \left(M(p_t)^{-1} X_t X_t^\top \right). \end{aligned}$$

Now, by the linearity of the trace,

$$\mathbb{E} \left[\text{Tr} \left(M(p_t)^{-1} X_t X_t^\top \right) \right] = \mathbb{E} \left[\sum_{x \in \mathcal{A}} p_t(x) \text{Tr} \left(M(p_t)^{-1} x x^\top \right) \right] = \mathbb{E} [\text{Tr}(I_d)] = d,$$

where I_d is the d -dimensional identity matrix. Collecting all bounds together and summing over t yields the claimed inequality $\mathbb{E}[V'_T] \leq (M - m)^2 dT$. \square