

Tsallis-INF: An Optimal Algorithm for Stochastic and Adversarial Bandits

Julian Zimmert

Yevgeny Seldin

University of Copenhagen, Copenhagen, Denmark

ZIMMERT@DI.KU.DK

SELDIN@DI.KU.DK

Editor: Peter Auer

Abstract

We derive an algorithm that achieves the optimal (within constants) pseudo-regret in both adversarial and stochastic multi-armed bandits without prior knowledge of the regime and time horizon.¹ The algorithm is based on online mirror descent (OMD) with Tsallis entropy regularization with power $\alpha = 1/2$ and reduced-variance loss estimators. More generally, we define an adversarial regime with a self-bounding constraint, which includes stochastic regime, stochastically constrained adversarial regime (Wei and Luo, 2018), and stochastic regime with adversarial corruptions (Lykouris et al., 2018) as special cases, and show that the algorithm achieves logarithmic regret guarantee in this regime and all of its special cases simultaneously with the optimal regret guarantee in the adversarial regime. The algorithm also achieves adversarial and stochastic optimality in the utility-based dueling bandit setting. We provide empirical evaluation of the algorithm demonstrating that it significantly outperforms UCB1 and EXP3 in stochastic environments. We also provide examples of adversarial environments, where UCB1 and THOMPSON SAMPLING exhibit almost linear regret, whereas our algorithm suffers only logarithmic regret. To the best of our knowledge, this is the first example demonstrating vulnerability of THOMPSON SAMPLING in adversarial environments. Last but not least, we present a general stochastic analysis and a general adversarial analysis of OMD algorithms with Tsallis entropy regularization for $\alpha \in [0, 1]$ and explain the reason why $\alpha = 1/2$ works best.

Keywords: Bandits, Online Learning, Best of Both Worlds, Online Mirror Descent, Tsallis Entropy, Multi-armed Bandits, Stochastic, Adversarial, I.I.D.

1. Introduction

Stochastic (i.i.d.) and adversarial multi-armed bandits are two fundamental sequential decision making problems in online learning (Thompson, 1933; Robbins, 1952; Lai and Robbins, 1985; Auer et al., 2002a,b). When prior information about the nature of environment is available, it is possible to achieve $\mathcal{O}\left(\sum_{i:\Delta_i>0} \frac{\log(T)}{\Delta_i}\right)$ pseudo-regret in the stochastic case (Lai and Robbins, 1985; Auer et al., 2002a) and $\mathcal{O}(\sqrt{KT})$ pseudo-regret in the adversarial case (Audibert and Bubeck, 2009, 2010), where T is the time horizon, K is the number of actions (a.k.a. arms), and Δ_i are suboptimality gaps. Both results match the lower bounds within constants, see Bubeck and Cesa-Bianchi (2012) for a survey.² The challenge in recent

1. The paper expands and improves our earlier work (Zimmert and Seldin, 2019).

2. To be precise, the $\mathcal{O}\left(\sum_{i:\Delta_i>0} \frac{\log(T)}{\Delta_i}\right)$ stochastic regret rate is optimal when the means of the rewards are close to $\frac{1}{2}$, see Lai and Robbins (1985), Cappé et al. (2013), and Kaufmann et al. (2012) for refined

years has been to achieve the optimal regret rates without prior knowledge about the nature of the problem.

One approach pursued by Bubeck and Slivkins (2012) and later refined by Auer and Chiang (2016) is to start playing under the assumption that the environment is i.i.d. and constantly monitor whether the assumption is satisfied. If a deviation from the i.i.d. assumption is detected, the algorithm performs an irreversible switch into an adversarial operation mode. This approach recovers the optimal bound in the stochastic case, but suffers from a multiplicative logarithmic factor in the regret in the adversarial case. Furthermore, the time horizon needs to be known in advance. The best known doubling schemes lead to extra multiplicative logarithmic factors in either the stochastic or the adversarial regime (Besson and Kaufmann, 2018).

Another approach pioneered by Seldin and Slivkins (2014) alters algorithms designed for adversarial bandits to achieve improved regret in the stochastic setting without losing the adversarial guarantees. They have introduced EXP3++, a modification of the EXP3 algorithm for adversarial bandits, which was later improved by Seldin and Lugosi (2017) to achieve an anytime regret of $\mathcal{O}\left(\sum_{i:\Delta_i>0} \frac{\log(T)^2}{\Delta_i}\right)$ in the stochastic case while preserving optimality in the adversarial case. A related approach by Wei and Luo (2018) uses log-barrier regularization instead of entropic regularization behind the EXP3. Their stochastic regret bound scales with $\log(T)$, although the constants are not spelled out explicitly and by empirical evaluation seem to be very large. Their adversarial regret guarantee scales with a square root of the cumulative loss of the best action in hindsight rather than a square root of the time horizon, but has an extra $\log T$ factor.

Seldin and Slivkins (2014), Lykouris et al. (2018), and Wei and Luo (2018) also define a number of intermediate regimes between stochastic and adversarial bandits and provide improved regret guarantees for them.

The question of whether it is at all possible to achieve simultaneous optimality in both worlds with no prior knowledge about the regime has remained open since the work of Bubeck and Slivkins (2012). Auer and Chiang (2016) have shown that no algorithm obtaining the optimal stochastic pseudo-regret bound can simultaneously achieve the optimal high-probability adversarial regret bound. Neither can an algorithm obtain the optimal stochastic pseudo-regret guarantee simultaneously with the optimal expected regret guarantee for adaptive adversaries.³ In addition, Abbasi-Yadkori et al. (2018) have shown that in the pure exploration setting it is also impossible to obtain the optimal rates in both stochastic and adversarial regimes.

We show that for pseudo-regret it is possible to achieve optimality in both regimes with a surprisingly simple algorithm. Moreover, we define a more general *adversarial regime with a self-bounding constraint*, which includes the stochastic, stochastically constrained adversarial (Wei and Luo, 2018), and adversarially corrupted stochastic (Lykouris et al., 2018) regimes as special cases. We propose an algorithm that achieves logarithmic pseudo-regret guarantee in the adversarial regime with a self-bounding constraint simultaneously with the adversarial regret guarantee. The algorithm is based on online mirror descent

lower and upper bounds otherwise. However, the refined analysis applies to stochastic bandits, whereas we consider a more general setting, see Section 2 for details.

3. This does not contradict our result because we bound the pseudo-regret, which is weaker than the expected regret.

with regularization by Tsallis entropy with power α . We name it α -TSALLIS-INF, or simply TSALLIS-INF for $\alpha = \frac{1}{2}$, where INF stands for Implicitly Normalized Forecaster (Audibert and Bubeck, 2009). The proposed algorithm is anytime: it requires neither the knowledge of the time horizon nor doubling schemes.

The main contributions of the paper are summarized in the following bullet points:

1. We propose the TSALLIS-INF algorithm, which is based on online mirror descent with regularization by Tsallis entropy with power $\alpha = \frac{1}{2}$. The algorithm achieves the optimal logarithmic pseudo-regret rate in the stochastic regime simultaneously with the optimal square-root adversarial regret guarantee with no prior knowledge of the regime. This resolves an open question of Bubeck and Slivkins (2012).
2. When combined with reduced-variance loss estimators proposed by Zimmert and Lattimore (2019), the leading constant of the stochastic regret bound for the TSALLIS-INF algorithm matches the asymptotic lower bound of Lai and Robbins (1985) within a multiplicative factor of 2.
3. The leading constant of the adversarial regret bound for the same combination matches the minimax lower bound of Cesa-Bianchi and Lugosi (2006, Theorem 6.1) within a multiplicative factor of less than 15. To the best of our knowledge, this is the best leading constant in an adversarial regret bound known today, matching the result of Zimmert and Lattimore (2019).
4. We introduce an adversarial regime with a self-bounding constraint, which includes stochastic, stochastically constrained adversarial, and adversarially corrupted stochastic regimes as special cases. We show that TSALLIS-INF achieves logarithmic regret in the new regime.
5. We improve the regret bound for adversarially corrupted stochastic regimes.
6. We use TSALLIS-INF in a SPARRING framework (Ailon et al., 2014) to obtain an algorithm that achieves stochastic and adversarial optimality in utility-based dueling bandits.
7. We provide a general analysis of OMD with Tsallis-Entropy regularization with power $\alpha \in [0, 1]$ and provide an intuition on why $\alpha = \frac{1}{2}$ works best.
8. We provide an empirical comparison of TSALLIS-INF with standard algorithms from the literature, UCB1, THOMPSON SAMPLING, EXP3, EXP3++, BROAD. We show that in stochastic environments with expected losses close to 0.5, TSALLIS-INF is only slightly worse than THOMPSON SAMPLING and significantly outperforms all other competitors, whereas in stochastically constrained adversarial environments TSALLIS-INF significantly outperforms all the competitors.
9. In one of the empirical comparisons, we design a stochastically constrained adversarial environment, where THOMPSON SAMPLING suffers almost linear regret. To the best of our knowledge, this is the first evidence that THOMPSON SAMPLING is not suitable for adversarial environments.

	Regime	$\frac{\text{Upper Bound}}{\text{Lower Bound}}$
<p>BROAD (Wei and Luo, 2018)</p> <p><i>Corresponds to Tsallis entropy regularization with $\alpha = 0$.</i></p> <p><i>Doubling is used for tuning the learning rate.</i></p>	<p>Sto.</p> <p>Adv.</p>	<p>$\mathcal{O}(K)$</p> <p>$\mathcal{O}(\sqrt{\log T})$</p>
<p>$\alpha = \frac{1}{2}$ (This paper)</p> <p><i>Anytime. No need for gap estimation, doubling, or mixing.</i></p>	<p>Sto. & Adv.</p>	<p>$\mathcal{O}(1)$</p>
<p>EXP3++ (Seldin and Lugosi, 2017)</p> <p><i>Corresponds to Tsallis entropy regularization with $\alpha = 1$.</i></p> <p><i>Anytime. Mixed-in exploration is used for gap estimation.</i></p>	<p>Sto.</p> <p>Adv.</p>	<p>$\mathcal{O}(\log T)$</p> <p>$\mathcal{O}(\sqrt{\log K})$</p>

Table 1: Ratio of regret upper to lower bound for TSALLIS-INF and the closest prior work, BROAD and EXP3++.

The paper is structured in the following way: In Section 2, we provide a formal definition of the problem setting, including the adversarial environment and the adversarial environment with a self-bounding constraint. Stochastic environments are a special case of the latter. In Section 3, we briefly review the framework of online mirror descent. We follow the techniques of Bubeck (2010) to derive an anytime version of the family of algorithms based on regularization by α -Tsallis Entropy (Tsallis, 1988; Abernethy et al., 2015). Section 4 contains the main theorems. We show that $\alpha = \frac{1}{2}$ provides an algorithm that is optimal in both adversarial regime and adversarial regime with a self-bounding constraint. The latter implies optimality in the stochastic regime. Interestingly, it is the same regularization power $\alpha = \frac{1}{2}$ that has been used by Audibert and Bubeck (2009, 2010) in POLY-INF algorithm to achieve the optimal regret rate in the adversarial regime. We analyse the algorithm with standard importance-weighted loss estimators and with reduced-variance loss estimators proposed by Zimmert and Lattimore (2019). The latter further reduces the constants and gets within a multiplicative factor of less than 15 from the minimax lower bound in the adversarial case and a multiplicative factor of 2 from the asymptotic lower bound in the stochastic case. Table 1 relates our results to the closest prior work on best-of-both-worlds algorithms. Wei and Luo (2018) use logarithmic regularization, which corresponds to Tsallis entropy with power $\alpha = 0$, and apply doubling for tuning the learning rate. Seldin and Lugosi (2017) use entropic regularization, which corresponds to Tsallis entropy with power $\alpha = 1$, and mix in additional exploration for estimation of the gaps. TSALLIS-INF with $\alpha = \frac{1}{2}$ requires neither doubling nor mixing nor estimation of the gaps. At the end of Section 4, we also provide a general analysis of the regret of α -TSALLIS-INF with $\alpha \in [0, 1]$ in adversarial environments and a general analysis of the regret of α -TSALLIS-INF with $\alpha \in [0, 1]$ in stochastic environments. We show that for $\alpha \neq \frac{1}{2}$ the optimal form of regularization and learning rate for the adversarial regime and for the stochastic regime differ. Thus, for $\alpha \neq \frac{1}{2}$ the algorithm does not achieve simultaneous optimality in both. Furthermore, for $\alpha \neq \frac{1}{2}$ the optimal regularizer for the stochastic regime requires oracle access to the unknown gaps. Prior work (Seldin and Slivkins, 2014; Seldin and Lugosi, 2017; Wei and Luo, 2018) used additional techniques, such as mixed-in exploration or dou-

bling, to control the regret, but as we show in Table 1 the results were suboptimal. In Section 5 we show that the stochastic regime with adversarial corruptions (Lykouris et al., 2018) is a special case of the adversarial regime with a self-bounding constraint and that TSALLIS-INF achieves the optimal regret rate there as well. In Section 6, we apply TSALLIS-INF to dueling bandits. Section 7 contains proofs of our main theorems. In Section 8, we provide an empirical comparison of TSALLIS-INF with baseline stochastic and adversarial bandit algorithms from the literature. We show that in stochastic environments with loss means close to 0.5, TSALLIS-INF with reduced-variance loss estimators significantly outperforms UCB1, EXP3, EXP3++, and BROAD, and follows closely behind THOMPSON SAMPLING, whereas in certain adversarial environments it significantly outperforms UCB1 and THOMPSON SAMPLING, which suffer almost linear regret, and also significantly outperforms EXP3, EXP3++, and BROAD. To the best of our knowledge, this is also the first evidence that THOMPSON SAMPLING is vulnerable in adversarial environments. We conclude with a summary in Section 9.

2. Problem Setting

At time $t = 1, 2, \dots$, the agent chooses an arm $I_t \in \{1, \dots, K\}$ out of a set of K arms. The environment picks a loss vector $\ell_t \in [0, 1]^K$ and the agent observes and suffers *only* the loss of the arm played, ℓ_{t, I_t} . The performance of an algorithm is measured in terms of pseudo-regret:

$$\overline{Reg}_T = \mathbb{E} \left[\sum_{t=1}^T \ell_{t, I_t} \right] - \min_i \mathbb{E} \left[\sum_{t=1}^T \ell_{t, i} \right] = \mathbb{E} \left[\sum_{t=1}^T (\ell_{t, I_t} - \ell_{t, i_t^*}) \right],$$

where $i_t^* \in \arg \min_i \mathbb{E} \left[\sum_{t=1}^T \ell_{t, i} \right]$ is defined as a best arm in expectation in hindsight and the expectation is taken over internal randomization of the algorithm and the environment.

In the (*adaptive*) *adversarial setting*, the adversary selects the losses arbitrarily, potentially based on the history of the agent's actions (I_1, \dots, I_{t-1}) and the adversary's own internal randomization. For deterministic oblivious adversaries, the definition of pseudo-regret coincides with the expected regret defined as $\mathbb{E}[Reg_T] = \mathbb{E} \left[\min_i \sum_{t=1}^T (\ell_{t, I_t} - \ell_{t, i}) \right]$.

We further define an *adversarial regime with a* (Δ, C, T) *self-bounding constraint*, where $\Delta \in [0, 1]^K$ and $C \geq 0$. In this regime, the adversary selects losses such that at time T the regret of any algorithm satisfies

$$\overline{Reg}_T \geq \sum_{t=1}^T \sum_i \Delta_i \mathbb{P}(I_t = i) - C. \quad (1)$$

The above condition should be satisfied at time T , but there is no requirement that it should be satisfied for all $t < T$.

A simple instance of an adversarial regime with a self-bounding constraint is the *stochastic* regime. In the stochastic regime, the losses $\ell_{t, i}$ are drawn from distributions with fixed means, $\mathbb{E}[\ell_{t, i}] = \mu_i$ independently of t , and the pseudo-regret can be written as

$$\overline{Reg}_T = \sum_{t=1}^T \sum_i \Delta_i \mathbb{P}(I_t = i), \quad (2)$$

where $\Delta_i = \mathbb{E}[\ell_{t,i}] - \min_i \mathbb{E}[\ell_{t,i}]$ is the suboptimality gap of action i . Thus, (1) is satisfied with Δ being the vector of suboptimality gaps and $C = 0$. In the stochastic regime, the best arm $i^* = \arg \min_i \mu_i$ is the same for all the rounds, $i_T^* = i^*$ for all T (if there is more than one best arm we can pick one arbitrarily).

Another instance of an adversarial regime with a self-bounding constraint is the *stochastically constrained adversarial setting* (Wei and Luo, 2018). In this setting, the losses $\ell_{t,i}$ are drawn from distributions with fixed gaps, $\mathbb{E}[\ell_{t,i} - \ell_{t,j}] = \tilde{\Delta}_{i,j}$ independently of t , but the means, as well as other parameters of the distributions of all arms, are allowed to change with time and may depend on the agent's past actions I_1, \dots, I_{t-1} . Obviously, the stochastic regime is a special case of a stochastically constrained adversary. By using $i^* = \arg \min_i \tilde{\Delta}_{i,1}$ to denote an optimal arm (if there is more than one, we can pick one arbitrarily) we define a vector of suboptimality gaps Δ by taking $\Delta_i = \tilde{\Delta}_{i,i^*}$, and then the pseudo-regret satisfies the identity in (2) and the condition in equation (1) is satisfied with the vector Δ and $C = 0$. In the stochastically constrained adversarial setting, the best arm is also the same for all rounds, $i_T^* = i^*$ for all T .

In Section 5, we show that *stochastic bandits with adversarial corruptions* (Lykouris et al., 2018) are also a special case of an adversarial regime with a self-bounding constraint.

The motivation behind the definition of the adversarial regime with a self-bounding constraint will become clear when we explain the analysis. For simplified intuition, the reader can think about its special case, the stochastic regime, where the constraint (1) is satisfied by the identity in (2).

3. Online Mirror Descent

We recall a number of basic definitions and facts from convex analysis. The convex conjugate (a.k.a. Fenchel conjugate) of a function $f : \mathbb{R}^K \rightarrow \mathbb{R}$ is defined by

$$f^*(y) = \sup_{x \in \mathbb{R}^K} \{\langle x, y \rangle - f(x)\}.$$

We use

$$\mathcal{I}_{\mathcal{A}}(x) := \begin{cases} 0, & \text{if } x \in \mathcal{A} \\ \infty, & \text{otherwise} \end{cases}$$

to denote the characteristic function of a closed and convex set $\mathcal{A} \subset \mathbb{R}^K$. Hence, $(f + \mathcal{I}_{\mathcal{A}})^*(y) = \max_{x \in \mathcal{A}} \{\langle x, y \rangle - f(x)\}$. By standard results from convex analysis (Rockafellar, 2015), for differentiable and convex f with invertible gradient $(\nabla f)^{-1}$, it holds that

$$\nabla(f + \mathcal{I}_{\mathcal{A}})^*(y) = \arg \max_{x \in \mathcal{A}} \{\langle x, y \rangle - f(x)\} \in \mathcal{A}.$$

3.1 General Framework

The traditional online mirror descent (OMD) framework uses a fixed regularizer Ψ with certain regularity constraints (Shalev-Shwartz, 2012). The update rule is

$$w_1 = \min_{w \in \mathcal{A}} \Psi(w), \quad w_{t+1} = \min_{w \in \mathcal{A}} a_t \langle w, \ell_t \rangle + D_{\Psi}(w, w_t),$$

where ℓ_t is the observed loss at time t , \mathcal{A} is the convex body of the action set, a_t is a weight parameter, and D_Ψ is the Bregman divergence $D_\Psi(x, y) = \Psi(x) - \Psi(y) - \langle x - y, \nabla \Psi(y) \rangle$. If the norm of the gradient of the regularizer $\|\nabla \Psi(x)\|$ is unbounded at the boundary of \mathcal{A} , then the update rule is equivalent to $w_{t+1} = \nabla(\Psi + \mathcal{I}_{\mathcal{A}})^*(-\sum_{s=1}^t a_s \ell_s)$, where $\sum_{s=1}^t a_s \ell_s$ is a weighted sum of past losses. This setting has been generalized to time-varying regularizers Ψ_t (Orabona et al., 2015), where the updates are given by $w_{t+1} = \nabla(\Psi_t + \mathcal{I}_{\mathcal{A}})^*(-\sum_{s=1}^t \ell_s)$. Note that this formulation uses no weighting a_s of the losses.

In the bandit setting, we do not observe the complete loss vector ℓ_t . Instead, an unbiased estimator $\hat{\ell}_t$ satisfying $\mathbb{E}_{I_t \sim w_t} [\hat{\ell}_t] = \ell_t$ is used for updating the cumulative losses. The common way of constructing unbiased loss estimators is through importance-weighted sampling:

$$\hat{\ell}_{t,i} = \frac{\mathbb{1}_t(i)\ell_{t,i}}{w_{t,i}}, \text{ where } \mathbb{1}_t(i) = \mathbb{1}(I_t = i) \text{ is the indicator function.} \quad (\text{IW})$$

We use (IW) to denote these estimators. Zimmert and Lattimore (2019) proposed reduced-variance importance-weighted loss estimators, which we call for brevity reduced-variance estimators or (RV)-estimators, and they are defined by

$$\hat{\ell}_{t,i} = \frac{\mathbb{1}_t(i)(\ell_{t,i} - \mathbb{B}_t(i))}{w_{t,i}} + \mathbb{B}_t(i), \text{ where } \mathbb{B}_t(i) := \frac{1}{2}\mathbb{1}(w_{t,i} \geq \eta_t^2). \quad (\text{RV})$$

For any $\mathbb{B}_t(i) \in [0, 1]$ the loss estimators remain unbiased, but their second moment $\mathbb{E}[\hat{\ell}_{t,i}^2]$ and variance are reduced. The value $\mathbb{B}_t(i) = \frac{1}{2}$ minimizes the worst-case variance of $\hat{\ell}_{t,i}$. However, the reduced-variance estimators can take negative values, $\hat{\ell}_{t,i} \geq -\frac{1}{2}\left(\frac{1}{w_{t,i}} - 1\right)$, while the analysis relies on non-negativity of the loss estimators. Zimmert and Lattimore (2019) show that negative loss estimators can be dealt with, as long as they satisfy $\hat{\ell}_{t,i} \geq -\frac{1}{2}\eta_t^{-2}$. We achieve this by only reducing variance of the estimators with $w_{t,i} \geq \eta_t^2$.

The algorithm is provided in Algorithm 1. At every step, we choose a probability distribution over arms w_t . We add $\mathcal{I}_{\Delta^{K-1}}$ to the regularizers Ψ_t , thereby ensuring that $w_t \in \Delta^{K-1}$, where Δ^{K-1} is the probability simplex. Note that the framework is equivalent to what Abernethy et al. (2014) call GRADIENT-BASED PREDICTION (GBP), where they replace $\nabla(\Psi_t + \mathcal{I}_{\Delta^{K-1}})^*$ with suitable functions $\nabla \Phi_t : \mathbb{R}^K \rightarrow \Delta^{K-1}$. We adopt the notation $\Phi_t := (\Psi_t + \mathcal{I}_{\Delta^{K-1}})^*$.

Algorithm 1: Online Mirror Descent for bandits

- Input:** $(\Psi_t)_{t=1,2,\dots}$
- 1 **Initialize:** $\hat{L}_0 = \mathbf{0}_K$ (where $\mathbf{0}_K$ is a vector of K zeros)
 - 2 **for** $t = 1, \dots$ **do**
 - 3 choose $w_t = \nabla(\Psi_t + \mathcal{I}_{\Delta^{K-1}})^*(-\hat{L}_{t-1})$ % see Alg. 2 for an explicit calculation
 - 4 sample $I_t \sim w_t$
 - 5 observe ℓ_{t,I_t}
 - 6 use (IW) or (RV) to construct $\hat{\ell}_t$
 - 7 update $\hat{L}_t = \hat{L}_{t-1} + \hat{\ell}_t$
-

3.2 OMD with Tsallis Entropy Regularization

We now consider a family of algorithms, which are regularized by the (negative) α -Tsallis entropy $H_\alpha(x) := \frac{1}{1-\alpha} (1 - \sum_i x_i^\alpha)$ (Tsallis, 1988). We change the scaling and add linear terms, resulting in the following regularizer with learning rate η_t :

$$\begin{aligned}\Psi(w) &:= - \sum_i \frac{w_i^\alpha - \alpha w_i}{\alpha(1-\alpha)\xi_i}, \\ \Psi_t(w) &:= \frac{1}{\eta_t} \Psi(w).\end{aligned}$$

Unless stated otherwise, we assume that $\xi_i = 1$ for all i , which leads to *symmetric regularization*. In the stochastic analysis of α -TSALLIS-INF with $\alpha \neq \frac{1}{2}$, we take $\xi_i = \Delta_i^{1-2\alpha}$ which leads to *asymmetric regularization*. Since the gaps are unknown, the latter is mainly interesting from a theoretical point of view.

The resulting family of algorithms is a subset of INF (Audibert and Bubeck, 2009), which we call α -TSALLIS-INF. α -TSALLIS-INF with symmetric regularization is related to the POLY-INF algorithm of Audibert and Bubeck (2009, 2010) and equivalent to the GBP algorithm proposed by Abernethy et al. (2015).

As has been observed earlier (Abernethy et al., 2015; Agarwal et al., 2017), α -TSALLIS-INF includes EXP3, which is based on the negative Shannon entropy $\sum_{i=1}^K w_i \log(w_i)$ (Cover and Thomas, 2006), and algorithms based on the log-barrier potential $\sum_{i=1}^K -\log(w_i)$ (Foster et al., 2016) as special cases.⁴ This can be seen by adding a constant term to the regularizer, so that $\Psi(w) = - \sum_i \frac{w_i^\alpha - \alpha w_i - (1-\alpha)}{\alpha(1-\alpha)\xi_i}$, and taking the respective limits $\alpha \rightarrow 0$ and $\alpha \rightarrow 1$. It gives:

$$\begin{aligned}\lim_{\alpha \rightarrow 0} - \frac{w_i^\alpha - \alpha w_i - (1-\alpha)}{\alpha(1-\alpha)\xi_i} &= \lim_{\alpha \rightarrow 0} - \frac{\log(w_i)w_i^\alpha - w_i + 1}{(1-2\alpha)\xi_i} = -\xi_i^{-1}(\log(w_i) - w_i + 1), \\ \lim_{\alpha \rightarrow 1} - \frac{w_i^\alpha - \alpha w_i - (1-\alpha)}{\alpha(1-\alpha)\xi_i} &= \lim_{\alpha \rightarrow 1} - \frac{\log(w_i)w_i^\alpha - w_i + 1}{(1-2\alpha)\xi_i} = \xi_i^{-1}(\log(w_i)w_i - w_i + 1),\end{aligned}$$

which are within linear and constant terms identical to the log-barrier potential and the negative Shannon entropy, respectively. Note that for symmetric regularization, neither the constant nor the linear terms influence the algorithm's choice of w , since it is normalized.

3.3 Implementation Details

The weights $w_{t,i}$ in TSALLIS-INF are given implicitly through a solution of a constrained optimization problem:

$$w_t = \arg \max_{w \in \Delta^{K-1}} \langle w, -\hat{L}_t \rangle + \frac{4}{\eta_t} \sum_i \sqrt{w_i}.$$

The solution takes the form

$$w_{t,i} = 4 \left(\eta_t \left(\hat{L}_{t,i} - x \right) \right)^{-2},$$

4. We use \log to denote the natural logarithm throughout the paper.

where the normalization factor x is defined implicitly through the constraint

$$\sum_i 4 \left(\eta_t \left(\hat{L}_{t,i} - x \right) \right)^{-2} = 1.$$

The normalization factor can be efficiently approximated by Newton's Method, reaching a sufficient precision in very few iterations. Details of the computation are provided in Algorithm 2.

Algorithm 2: Newton's Method approximation of w_t in TSALLIS-INF ($\alpha = \frac{1}{2}$)

Input: x, \hat{L}_t, η_t %we use x from the previous iteration as a warmstart

- 1 **repeat**
- 2 $\forall i : w_{t,i} \leftarrow 4(\eta_t(\hat{L}_{t,i} - x))^{-2}$
- 3 $x \leftarrow x - (\sum_i w_{t,i} - 1) / (\eta_t \sum_i w_{t,i}^{\frac{3}{2}})$
- 4 **until** convergence

4. Main Results

In this section we present our main result, the TSALLIS-INF algorithm with $\alpha = \frac{1}{2}$ that achieves the optimal regret bounds in both adversarial and stochastic bandits. We show that it also achieves a logarithmic regret guarantee in the more general adversarial regime with a self-bounding constraint. In fact, the stochastic regret bound follows as a special case of the more general analysis. We then present a general analysis of α -TSALLIS-INF with $\alpha \in [0, 1]$ and explain the intuition of why $\alpha = \frac{1}{2}$ works best.

4.1 Analysis of Tsallis-INF with $\alpha = 1/2$

We show that TSALLIS-INF with $\alpha = \frac{1}{2}$ and symmetric regularizer achieves the optimal \sqrt{T} regret scaling in the adversarial regime and simultaneously $\log(T)$ regret scaling in the adversarial regime with a self-bounding constraint. The latter ensures the same regret scaling in stochastic and stochastically constrained adversarial environments as special cases. We analyse the algorithm with (IW) and (RV) loss estimators. Both estimators achieve the optimal regret scaling in both regimes, but the (RV) estimator yields better constants. The results for the two estimators are presented alongside each other using cases brackets and marked by (IW) and (RV), respectively.

Theorem 1 *The pseudo-regret of TSALLIS-INF with $\alpha = \frac{1}{2}$, symmetric regularization ($\xi_i = 1$), and learning rate*

$$\eta_t = \begin{cases} 2\sqrt{\frac{1}{t}}, & \text{for (IW) estimators,} \\ 4\sqrt{\frac{1}{t}}, & \text{for (RV) estimators,} \end{cases}$$

in any adversarial bandit problem satisfies:

$$\overline{\text{Reg}}_T \leq \begin{cases} 4\sqrt{KT} + 1, & \text{with (IW) estimators,} \\ 2\sqrt{KT} + 10K \log(T) + 16, & \text{with (RV) estimators.} \end{cases} \quad (3)$$

If there exists a vector $\Delta \in [0, 1]^K$ with a unique zero entry i^* (i.e., $\Delta_{i^*} = 0$ and $\Delta_i > 0$ for all $i \neq i^*$) and a constant C , such that the pseudo-regret at time T satisfies

$$\overline{\text{Reg}}_T \geq \mathbb{E} \left[\sum_{t=1}^T \sum_{i \neq i^*} w_{t,i} \Delta_i \right] - C, \quad (4)$$

then the pseudo-regret further satisfies

$$\overline{\text{Reg}}_T \leq \begin{cases} \left(\sum_{i \neq i^*} \frac{4 \log(T) + 12}{\Delta_i} \right) + 4 \log(T) + \frac{1}{\Delta_{\min}} + \frac{3}{2} \sqrt{K} + 8 + C, & \text{with (IW),} \\ \left(\sum_{i \neq i^*} \frac{\log(T) + 3}{\Delta_i} \right) + 28K \log(T) + \frac{1}{\Delta_{\min}} + \frac{3}{2} \sqrt{K} + 32 + C, & \text{with (RV),} \end{cases}$$

where $\Delta_{\min} := \min_{\Delta_i > 0} \Delta_i$. If C satisfies

$$\begin{aligned} C &> \left(\sum_{i \neq i^*} \frac{4 \log(T) + 12}{\Delta_i} \right) + \frac{1}{\Delta_{\min}}, & \text{for (IW),} \\ C &> \left(\sum_{i \neq i^*} \frac{\log(T) + 3}{\Delta_i} \right) + \frac{1}{\Delta_{\min}}, & \text{for (RV),} \end{aligned}$$

then the regret additionally satisfies

$$\overline{\text{Reg}}_t \leq \begin{cases} 2 \sqrt{\left(\left(\sum_{i \neq i^*} \frac{4 \log(T) + 12}{\Delta_i} \right) + \frac{1}{\Delta_{\min}} \right) C + 4 \log(T) + \frac{3}{2} \sqrt{K} + 8}, & \text{with (IW),} \\ 2 \sqrt{\left(\left(\sum_{i \neq i^*} \frac{\log(T) + 3}{\Delta_i} \right) + \frac{1}{\Delta_{\min}} \right) C + 28K \log(T) + \frac{3}{2} \sqrt{K} + 32}, & \text{with (RV).} \end{cases}$$

The proof is postponed to Section 7. We call the condition in equation (4) a (Δ, C, T) *self-bounding property* of the regret. As we have mentioned in Section 2, in the stochastically constrained adversarial environments and stochastic bandits as their special case $\overline{\text{Reg}}_T = \mathbb{E} \left[\sum_{t=1}^T \sum_{i \neq i^*} w_{t,i} \Delta_i \right]$, where Δ is the vector of suboptimality gaps, and under the assumption that the best arm is unique, the condition in equation (4) is satisfied with $C = 0$. Thus, in the above regimes, the regret of TSALLIS-INF with $\alpha = \frac{1}{2}$ and RV loss estimators is

$$\overline{\text{Reg}}_T \leq \left(\sum_{i \neq i^*} \frac{\log(T) + 3}{\Delta_i} \right) + 28K \log(T) + \frac{1}{\Delta_{\min}} + \frac{3}{2} \sqrt{K} + 32.$$

The worst case lower bound for stochastic multiarmed bandits (MAB) with Bernoulli losses is achieved when the expectations of the losses are close to $\frac{1}{2}$. Let Δ denote the vector of gaps and let $\mathbb{E}[\ell_{t,i}] = \frac{1}{2} + \Delta_i$. By adapting the well known divergence-dependent lower bound of Lai and Robbins (1985), we can show that for any consistent algorithm

$$\lim_{\|\Delta\| \rightarrow 0} \left(\left(\sum_{i: \Delta_i > 0} \frac{1}{\Delta_i} \right)^{-1} \liminf_{t \rightarrow \infty} \frac{\mathbb{E}[\overline{\text{Reg}}_t]}{\log(t)} \right) \geq \frac{1}{2}.$$

See Appendix A for details. Therefore, the asymptotic regret upper bound of TSALLIS-INF with RV-estimators in the stochastic regime is optimal within a multiplicative factor of 2,

which is arguably a small price for a significant gain in robustness against adversaries. We leave it to future work to close the gap or prove that it is impossible to do so without compromising on the adversarial guarantees.

To the best of our knowledge, the leading constant 2 in the adversarial regret bound of TSALLIS-INF with RV estimators (the bound in equation (3)) provides the tightest adversarial regret guarantee known today. It matches the minimax adversarial lower bound in Cesa-Bianchi and Lugosi (2006, Theorem 6.1) within a multiplicative factor of less than 15. Under the assumption of known time horizon, Zimmert and Lattimore (2019) provide an adversarial regret bound with a leading constant of $\sqrt{2}$. The $\sqrt{2}$ multiplicative difference between their result and ours is the standard conversion rate between fixed-horizon and anytime regret bounds.

Remark 2 *The assumption that Δ has a unique zero entry and the corresponding assumption on uniqueness of the best arm in the stochastically constrained adversarial setting is a technical assumption we had to use in our proofs, but our experiments suggest that this is an artifact of the analysis. We conjecture that it can be removed, but explain the challenges in achieving the goal in Section 7.*

4.2 A General Analysis of α -Tsallis-INF with $\alpha \in [0, 1]$

Now we provide a general analysis of α -TSALLIS-INF with $\alpha \in [0, 1]$ and then explain the intuition of why $\alpha = \frac{1}{2}$ works best. Since $\alpha \neq \frac{1}{2}$ anyway leads to suboptimal regret rates and in order to keep things simple, we restrict the general analysis to IW estimators. We note that in Theorem 1, the RV estimators helped improve the constants, but they did not change the rates. Therefore, we save the effort of optimising the constants in a priori suboptimal bounds. To keep things even simpler, we derive logarithmic bounds for stochastically constrained adversarial environments rather than the more general adversarial regime with a self-bounding constraint (technically speaking, we work with $C = 0$).

Note that the adversarial analysis in Theorem 3 and stochastic analysis in Theorem 4 consider different versions of α -TSALLIS-INF. The adversarial analysis uses *symmetric* regularization, whereas stochastic analysis uses *asymmetric* regularization. We get back to this point after we present the results.

4.2.1 ADVERSARIAL REGIME

α -TSALLIS-INF with symmetric regularization has been previously analyzed in the adversarial setting by Abernethy et al. (2015) and Agarwal et al. (2017). Abernethy et al. provide a finite-time analysis for $\alpha \in (0, 1]$, while Agarwal et al. analyze the case of $\alpha = 0$. The main contribution of the following theorem is that it provides a unified and anytime treatment for all $\alpha \in [0, 1]$. The bound recovers the constants from Abernethy et al. without the need of tuning the learning rate by the time horizon T .

Theorem 3 *For any $\alpha \in [0, 1]$ and any adversarial bandit problem, the pseudo-regret of α -TSALLIS-INF with symmetric regularizer ($\xi_i = 1$), learning rate $\eta_t = \sqrt{\frac{K^{1-2\alpha} - K^{-\alpha}}{1-\alpha} \frac{1-t^{-\alpha}}{\alpha t}}$,*

and IW loss estimators at any time T satisfies

$$\overline{\text{Reg}}_T \leq 2\sqrt{\min\left\{\frac{1}{\alpha - \alpha^2}, \frac{\log(K)}{\alpha}, \frac{\log(T)}{1 - \alpha}\right\}KT} + 1.$$

(At the boundaries $\alpha = 0$ and $\alpha = 1$, the learning rates are defined by $\lim_{\alpha \rightarrow 0} \eta_t = \sqrt{\frac{(K-1)\log(t)}{t}}$ and $\lim_{\alpha \rightarrow 1} \eta_t = \sqrt{\frac{\log(K)(1-t^{-1})}{t}}$, respectively.)

The proof is postponed to Section 7.

4.2.2 STOCHASTICALLY CONSTRAINED ADVERSARIAL REGIME

Now we present an analysis of α -TSALLIS-INF with $\alpha \in [0, 1]$ and *asymmetric* regularization in the stochastically constrained adversarial setting. We let $\bar{t} = \max\{e, t\}$. For learning rates $\eta_t = \frac{16^\alpha}{4} \frac{1-\bar{t}^{-1+\alpha}}{(1-\alpha)t^\alpha}$ and asymmetric regularizer with $\xi_i = \Delta_i^{1-2\alpha}$ for $i \neq i^*$ and $\xi_{i^*} = \Delta_{\min}^{1-2\alpha}$, where $\Delta_{\min} = \min_{i \neq i^*} \Delta_i$, we prove the following theorem:

Theorem 4 *For any $\alpha \in [0, 1]$ and any stochastically constrained adversarial regime with a unique best arm (i.e., $\Delta_i > 0$ for all i except a unique index i^* for which $\Delta_{i^*} = 0$), the pseudo-regret of α -TSALLIS-INF with learning rate $\eta_t = \frac{16^\alpha}{4} \frac{1-\bar{t}^{-1+\alpha}}{(1-\alpha)t^\alpha}$ and asymmetric regularizer with parameters $\xi_i = \Delta_i^{1-2\alpha}$ for $i \neq i^*$ and $\xi_{i^*} = \Delta_{\min}^{1-2\alpha}$ at any time T satisfies*

$$\overline{\text{Reg}}_T \leq \sum_{i \neq i^*} \left(\frac{(8 \min\{\frac{1}{1-\alpha}, \log(T)\} + 64) \log(T)}{\Delta_i} \right) + \frac{16 \log^4\left(\frac{16}{\Delta_{\min}^2} \log^2\left(\frac{16}{\Delta_{\min}^2}\right)\right)}{\Delta_{\min}} + 4.$$

The proof is provided in Appendix D.

Remark 5 *We emphasize that for $\alpha \neq \frac{1}{2}$, the result in Theorem 4 requires knowledge of the gaps Δ_i for tuning the regularization parameters ξ_i . For $\alpha = \frac{1}{2}$, this knowledge is not required. Therefore, Theorem 4 is primarily interesting from the theoretical perspective of characterization of behavior of α -TSALLIS-INF in stochastically constrained adversarial environments, whereas $\alpha = \frac{1}{2}$ is the only practically interesting value with the refined analysis in Theorem 1.*

Remark 6 *For $\alpha \neq \frac{1}{2}$, the version α -TSALLIS-INF in Theorem 4 uses asymmetric regularization, whereas α -TSALLIS-INF in Theorem 3 uses symmetric regularization. The corresponding learning rates also differ. Therefore, for $\alpha \neq \frac{1}{2}$, neither of the two versions of α -TSALLIS-INF achieves simultaneous optimality in the stochastic and adversarial setting. In fact, the time dependence of the adversarial regret guarantee for α -TSALLIS-INF in Theorem 4 is in the order of $T^\alpha + T^{1-\alpha}$.*

Remark 7 *We note that while Tsallis entropy with $\alpha = 0$ corresponds to log-barrier potential used in BROAD, and Tsallis entropy with $\alpha = 1$ corresponds to entropic regularization used in EXP3++, the two algorithms (BROAD and EXP3++) use symmetric regularization, whereas α -TSALLIS-INF in Theorem 4 uses asymmetric regularization. Therefore, there is no direct relation between the result of Theorem 4 and these two algorithms. In particular, BROAD and EXP3++ use other techniques to achieve slightly suboptimal, but simultaneous stochastic and adversarial regret guarantees (as described in Table 1), which is not the case for α -TSALLIS-INF with asymmetric regularization in Theorem 4.*

4.3 Intuition Behind the Success of Tsallis-Inf with $\alpha = \frac{1}{2}$

It has been previously shown that regularization by Tsallis entropy with power $\alpha = 1/2$ leads to the minimax optimal regret rate in the adversarial regime (Audibert and Bubeck, 2009). Here we provide some basic intuition on why the same value of α works well in the stochastic case. We also highlight the key breakthroughs that allow us to overcome challenges faced in prior work.

We start with a simple “back of the envelope” approximation of the form of the weights w_t played by TSALLIS-INF. By definition of Algorithm 1, at round t we have

$$w_t = \arg \max_{w \in \Delta^{K-1}} \left\{ \left\langle w, -\hat{L}_{t-1} \right\rangle + \frac{1}{\eta_t} \sum_i \frac{w_i^\alpha - \alpha w_i}{\alpha(1-\alpha)\xi_i} \right\}.$$

Taking a derivative of the Lagrangian of the above expression with respect to w_i and equating it to zero, we obtain

$$-\hat{L}_{t-1,i} + \frac{1}{\eta_t(1-\alpha)\xi_i} (w_{t,i}^{\alpha-1} - 1) - \nu = 0,$$

where ν is a Lagrange multiplier corresponding to the constraint that w is a probability distribution. We can express ν as

$$\nu = \frac{1}{\eta_t(1-\alpha)\xi_{i^*}} (w_{t,i^*}^{\alpha-1} - 1) - \hat{L}_{t-1,i^*}.$$

For $i \neq i^*$ this gives

$$\begin{aligned} w_{t,i} &= \left(\eta_t(1-\alpha)\xi_i \left(\hat{L}_{t-1,i} + \nu \right) + 1 \right)^{\frac{1}{\alpha-1}} \\ &= \left(\eta_t(1-\alpha)\xi_i \left(\hat{L}_{t-1,i} - \hat{L}_{t-1,i^*} + \frac{1}{\eta_t(1-\alpha)\xi_{i^*}} (w_{t,i^*}^{\alpha-1} - 1) \right) + 1 \right)^{\frac{1}{\alpha-1}} \\ &= \left(\eta_t(1-\alpha)\xi_i \left(\hat{L}_{t-1,i} - \hat{L}_{t-1,i^*} \right) + \frac{\xi_i}{\xi_{i^*}} (w_{t,i^*}^{\alpha-1} - 1) + 1 \right)^{\frac{1}{\alpha-1}} \\ &\approx \left(\eta_t(1-\alpha)\xi_i \left(\hat{L}_{t-1,i} - \hat{L}_{t-1,i^*} \right) \right)^{\frac{1}{\alpha-1}}, \end{aligned}$$

where the approximation holds because asymptotically the first term dominates the sum. A bit more explicitly, in order for the algorithm to deliver non-trivial regret guarantee, w_{t,i^*} should be close to 1. Thus, the last two terms in the brackets are roughly a constant. At the same time, as we discuss below, the whole expression in the brackets must grow roughly as $(\Delta_i^2 t)^{1-\alpha}$. Thus, the first term must dominate. In the stochastic regime $\mathbb{E} \left[\hat{L}_{t,i} - \hat{L}_{t,i^*} \right] = \Delta_i t$. If we use this in our back-of-the-envelope calculation, we obtain that for $i \neq i^*$ in the stochastic regime $\mathbb{E}[w_{t,i}] \approx \mathbb{E} \left[\left(\eta_t(1-\alpha)\xi_i (\hat{L}_{t-1,i} - \hat{L}_{t-1,i^*}) \right)^{\frac{1}{\alpha-1}} \right] \propto (\eta_t \xi_i \Delta_i t)^{\frac{1}{\alpha-1}}$. (Strictly speaking, when we take the expectation inside the power we obtain an inequality, but we ignore this detail in the high-level discussion. We also ignore the $(1-\alpha)$ factor, which can be seen as a constant for $\alpha < 1$.)

In order to achieve a regret rate of $\Theta(\sum_{i \neq i^*} \frac{\log t}{\Delta_i})$ in the stochastic regime, the suboptimal arms should be explored at a rate of $\Theta(\frac{1}{\Delta_i^2 t})$ per round (if $\mathbb{E}[w_{t,i}] = \Theta(\frac{1}{\Delta_i^2 t})$, then $\Delta_i \mathbb{E}[\sum_{s=1}^t w_{s,i}] = \Theta(\frac{\log t}{\Delta_i})$, as desired). Exploring more than that leads to excessive regret from the exploration alone. Exploring less is also prohibitive, because it leads to an overly high probability of misidentifying the best arm. By looking at the approximation of $\mathbb{E}[w_{t,i}]$ from the previous paragraph, we obtain that we should have $(\eta_t \xi_i \Delta_i t)^{\frac{1}{\alpha-1}} \propto \frac{1}{\Delta_i^2 t}$ or, equivalently, $\eta_t \xi_i \propto t^{-\alpha} \Delta_i^{1-2\alpha}$. The learning rate takes care of the time-dependent quantities, i.e., $\eta_t \propto t^{-\alpha}$, and ξ_i should take care of the arm-dependent quantities, i.e., we should have $\xi_i \propto \Delta_i^{1-2\alpha}$. Note that $\alpha = \frac{1}{2}$ leads to a symmetric regularizer Ψ (i.e., $\xi_i = 1$), whereas for $\alpha \neq \frac{1}{2}$ the regularizer must be tuned using unknown gaps Δ_i . The necessity to tune the regularizer based on unknown gaps has hindered progress in the work of Wei and Luo (2018), who used the log-barrier regularizer corresponding to $\alpha = 0$.

Another crucial novelty behind the success of our analysis is basing it on the self-bounding property of the regret in equation (4). The new proof technique uses the same mechanism for controlling the regret in stochastic and adversarial regimes and we explain the intuition behind it in Section 7.1. The earlier approach by Seldin and Slivkins (2014) and Seldin and Lugosi (2017) has controlled the regret in stochastic and adversarial regimes through separate mechanisms. The stochastic analysis was based on using empirical estimates of the gaps and high-probability control of the weights $w_{t,i}$. However, gap estimation is challenging, because the variance of $\hat{L}_{t,i}$ is of the order of $\sum_{s=1}^t \frac{1}{w_{s,i}}$. If the arms are played according to the target probabilities of $w_{t,i} \approx \frac{1}{t \Delta_i^2}$, then the variance of $(\hat{L}_{t,i} - \hat{L}_{t,i^*})$ is of the order of $\Theta(\Delta_i^2 t^2)$. This is prohibitively large, because the square root of the variance is of the same order as the expected cumulative gap and standard tools, such as Bernstein's inequality, cannot guarantee concentration of $(\hat{L}_{t,i} - \hat{L}_{t,i^*})$ around $\Delta_i t$. Seldin and Slivkins (2014) have coped with this by mixing in additional exploration, but this has led to a regret growth rate of the order of $(\log T)^3$ in the stochastic regime. Seldin and Lugosi (2017) have mixed in less exploration and used unweighted losses for the gap estimates, which has decreased the regret growth rate down to $(\log T)^2$. It is currently unknown whether direct gap estimation can be further improved to support the desired $\log T$ stochastic regret rates. Additionally, existing oracle analysis in Seldin and Slivkins (2014, Theorem 2) and Theorem 4 here only support $(\log T)^2$ regret rate for EXP3-based algorithms (corresponding to $\alpha = 1$) in the stochastic regime. It is also unknown whether this rate can be improved. To summarize, the main breakthrough compared to this line of work is moving from $\alpha = 1$ to $\alpha = \frac{1}{2}$ and shifting from an analysis based on gap estimation to an analysis based on self-boundedness of the regret. The proposed algorithm does not mix in any additional exploration.

5. Additional Intermediate Regimes Between Stochastic and Adversarial

In this section, we show that stochastic bandits with adversarial corruptions proposed by Lykouris et al. (2018) are also a special case of an adversarial environment with a self-bounding constraint. We further propose an extension of their regime by combining it with a stochastically constrained adversary. We show that the combination is also a special case

of an adversarial environment with a $(\Delta, 2C, T)$ self-bounding constraint, where TSALLIS-INF achieves logarithmic regret. We finish the section with an open question on whether TSALLIS-INF can achieve logarithmic regret guarantees in the intermediate regimes defined by Seldin and Slivkins (2014).

5.1 Stochastic Bandits with Adversarial Corruptions

Lykouris et al. (2018) have proposed a regime in which an adversary is allowed to make corruptions to an otherwise stochastic environment. Let $\bar{\mathcal{L}}_T = (\bar{\ell}_1, \dots, \bar{\ell}_T)$ and $\mathcal{L}_T = (\ell_1, \dots, \ell_T)$ be two sequences of losses, then the amount of corruption is measured by $\sum_{t=1}^T \|\bar{\ell}_t - \ell_t\|_\infty$.

Let $\bar{\mathcal{L}}_T$ be a sequence of losses generated by a stochastically constrained adversary with best arm i^* and gaps Δ_i , and let \mathcal{L}_T be its adaptively corrupted version with corruption amount bounded by C . The regret of an algorithm executed on \mathcal{L}_T satisfies

$$\begin{aligned} \overline{Reg}_T &= \max_i \mathbb{E} \left[\sum_{t=1}^T \ell_{t,I_t} - \ell_{t,i} \right] \geq \mathbb{E} \left[\sum_{t=1}^T \ell_{t,I_t} - \ell_{t,i^*} \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \bar{\ell}_{t,I_t} - \bar{\ell}_{t,i^*} \right] + \mathbb{E} \left[\sum_{t=1}^T \ell_{t,I_t} - \bar{\ell}_{t,I_t} \right] + \mathbb{E} \left[\sum_{t=1}^T \bar{\ell}_{t,i^*} - \ell_{t,i^*} \right] \\ &\geq \sum_{t=1}^T \sum_{i \neq i^*} \Delta_i \mathbb{E}[w_{t,i}] - 2C. \end{aligned} \tag{5}$$

Thus, a stochastically constrained adversary with adversarial corruptions is an adversarial regime with a $(\Delta, 2C, T)$ self-bounding constraint. This leads to a direct corollary of Theorem 1, which improves upon the pseudo-regret bounds of Lykouris et al. (2018) and Gupta et al. (2019), the latter providing an $\mathcal{O} \left(\sum_{i \neq i^*} \frac{\log(T)}{\Delta_i} + KC \right)$ guarantee. We note that Lykouris et al. (2018) and Gupta et al. (2019) do not assume uniqueness of the best arm and also provide high-probability regret guarantees, but they only consider the more restricted stochastic setting with adversarial corruptions rather than stochastically constrained adversarial setting with adversarial corruptions.

Corollary 8 *The regret of TSALLIS-INF in a stochastically constrained adversarial environment with a unique best arm i^* , adaptively corrupted with corruption amount bounded by C satisfies*

$$\overline{Reg}_T = \mathcal{O} \left(\sum_{i \neq i^*} \frac{\log(T)}{\Delta_i} + \sqrt{\sum_{i \neq i^*} \frac{\log(T)}{\Delta_i} C} \right).$$

Remark 9 *We emphasize that the assumption of best arm uniqueness is on the stochastically constrained adversary before corruption. After the adaptive corruption, it is allowed to have multiple best arms and the identity of the best arm is allowed to change.*

Proof By equation (5), the self-bounding condition (4) of Theorem 1 is satisfied with Δ being the vector of gaps of the underlying stochastically constrained adversary and the

constant being $2C$. Thus, with RV loss estimators for $2C \leq \left(\sum_{i \neq i^*} \frac{\log(T)+3}{\Delta_i}\right) + \frac{1}{\Delta_{\min}}$ TSALLIS-INF achieves

$$\overline{Reg}_T \leq \left(\sum_{i \neq i^*} \frac{\log(T)+3}{\Delta_i}\right) + 28K \log(T) + \frac{1}{\Delta_{\min}} + \frac{3}{2}\sqrt{K} + 32 + 2C$$

and otherwise

$$\overline{Reg}_t \leq 2 \sqrt{\left(\left(\sum_{i \neq i^*} \frac{\log(T)+3}{\Delta_i}\right) + \frac{1}{\Delta_{\min}}\right) 2C + 28K \log(T) + \frac{3}{2}\sqrt{K} + 32}.$$

■

5.2 Open Problem: The Performance in Seldin and Slivkins’ Environments

Seldin and Slivkins (2014) define *moderately contaminated stochastic regime* and *an adversarial regime with a gap*. In the moderately contaminated stochastic regime, the adversary is allowed to change up to $\frac{t\Delta_i}{4}$ arbitrarily selected observations for a suboptimal arm i and up to $\frac{t\Delta_{\min}}{4}$ observations for the optimal arm i^* (where $\Delta_{\min} = \min_{\Delta_i > 0} \Delta_i$). The logic behind the definition is that in expectation, the adversary can reduce the gap Δ_i by a factor of 2, but cannot eliminate it completely. The adversarial regime with a gap is an adversarial regime, where starting from a certain time τ (unknown to the algorithm) the cumulative loss of an optimal arm maintains a certain gap Δ_τ to all other arms until the end of the game. Seldin and Slivkins show that their EXP3++ algorithm achieves “logarithmic” regret in both regimes. Note that in the moderately contaminated stochastic regime, the amount of contamination is allowed to grow linearly with time. While the regime could be seen as a special case of stochastic bandits with adversarial corruptions discussed earlier, the regret bound in Corollary 8 only supports “logarithmic” regret for “logarithmic” amount of corruption C . So far we have been unable to obtain “logarithmic” regret guarantees for TSALLIS-INF in the intermediate regimes of Seldin and Slivkins (the analysis proposed in Zimmert and Seldin (2019) is incorrect). The challenge is that the gaps are defined through cumulative rather than instantaneous quantities. Deriving “logarithmic” regret guarantees for TSALLIS-INF in these regimes is an interesting open problem.

6. Dueling Bandits

In the sparring approach to stochastic utility-based dueling bandits, (Ailon et al., 2014) each side in the sparring can be modeled as a stochastically constrained adversarial environment. This makes it a perfect application domain for TSALLIS-INF. The problem is defined by K arms with utilities $u_i \in [0, 1]$. At each round, an agent has to select two arms, I_t and J_t , to “duel”. The feedback is the winner W_t of the “duel”, which is chosen according to

$\mathbb{P}[W_t = I_t] = \frac{1+u_{I_t}-u_{J_t}}{2}$. The regret is defined by the distance to the optimal utility:

$$\overline{Reg}_T = \sum_{t=1}^T 2u_{i_t^*} - \mathbb{E} \left[\sum_{t=1}^T (u_{I_t} + u_{J_t}) \right].$$

In the adversarial version of the problem, the utilities u_i are not constant but time dependent, $u_{t,i}$, and selected by an adversary. The regret in this case is the difference to the optimal utility in hindsight:

$$\overline{Reg}_T = \max_i \mathbb{E} \left[\sum_{t=1}^T 2u_{t,i} \right] - \mathbb{E} \left[\sum_{t=1}^T (u_{t,I_t} + u_{t,J_t}) \right].$$

Ailon et al. (2014) have proposed the SPARRING algorithm, in which two black-box MAB algorithms spar with each other. The first algorithm selects I_t and receives the loss $\ell_{t,I_t} = \mathbb{1}(W_t \neq I_t)$. The second algorithm selects J_t and receives the loss $\ell_{t,J_t} = \mathbb{1}(W_t \neq J_t)$. They have shown that the regret is the sum of individual regret values for both MABs, thereby recovering $\mathcal{O}(\sqrt{KT})$ regret in the adversarial case if MABs with $\mathcal{O}(\sqrt{KT})$ adversarial regret bound are used. In the stochastic case, each black-box MAB plays in a stochastically constrained adversarial environment because the relative winning probability of the arms stays fixed, but depending on the arm choice of the sparring partner, the baseline shifts up and down. Since no algorithm has been known to achieve $\log(T)$ regret in stochastically constrained adversarial setting, Ailon et al. (2014) provide no analysis of SPARRING in the stochastic case. Indeed, as we demonstrate in our experiments, standard algorithms for stochastic multi-armed bandits, such as UCB or THOMPSON SAMPLING, may exhibit almost linear regret in stochastically constrained adversarial setting and, therefore, are not suitable for sparring.

By applying Theorem 1, we directly obtain the following corollary.

Corollary 10 *In a utility-based dueling bandit problem SPARRING with two independent versions of TSALLIS-INF suffers a regret of*

$$\overline{Reg}_T \leq \mathcal{O} \left(\sum_{i:\Delta_i>0} \frac{\log(T)}{\Delta_i} \right)$$

in the stochastic case with a unique best arm and

$$\overline{Reg}_T \leq \mathcal{O}(\sqrt{KT})$$

in the adversarial case.

7. Proofs

In this section, we first revise the general proof framework of OMD and provide a compact summary of how to modify it to obtain stochastic guarantees. Afterward, we provide proofs of Theorems 1 and 3. A proof of Theorem 4 along with proofs of all the lemmas in this section are provided in the appendix.

7.1 High-Level Overview of OMD Modification for Stochastic Analysis

We follow the standard OMD analysis (Lattimore and Szepesvári, 2019, Chapter 28) and introduce the potential function $\Phi_t(-L) = \max_{w \in \Delta^{K-1}} \{\langle w, -L \rangle - \Psi_t(w)\}$ to decompose the regret into *stability* and *penalty* terms.

$$\begin{aligned} \overline{Reg}_T &= \mathbb{E} \left[\sum_{t=1}^T \left(\ell_{t, I_t} - \ell_{t, i_T^*} \right) \right] \\ &= \underbrace{\mathbb{E} \left[\sum_{t=1}^T \ell_{t, I_t} + \Phi_t(-\hat{L}_t) - \Phi_t(-\hat{L}_{t-1}) \right]}_{\text{stability}} + \underbrace{\mathbb{E} \left[\sum_{t=1}^T -\Phi_t(-\hat{L}_t) + \Phi_t(-\hat{L}_{t-1}) - \ell_{t, i_T^*} \right]}_{\text{penalty}}. \end{aligned} \quad (6)$$

The OMD analysis bounds the *stability* and *penalty* terms separately. For Tsallis-entropy regularizers, Abernethy et al. (2015) have proven the following bounds:

$$\begin{aligned} \text{stability} &\leq \sum_{t=1}^T \eta_t \sum_{i=1}^K f(\mathbb{E}[w_{t,i}]), \\ \text{penalty} &\leq \sum_{t=1}^T (\eta_{t+1}^{-1} - \eta_t^{-1}) \sum_{i=1}^K g(\mathbb{E}[w_{t,i}]), \end{aligned}$$

where $f(x)$ and $g(x)$ are proportional to $x^{1-\alpha}$ and x^α , respectively. Adversarial bounds that scale with \sqrt{T} are obtained by applying $\sum_{i=1}^K f(\mathbb{E}[w_{t,i}]) \leq \max_{w \in \Delta^{K-1}} \sum_{i=1}^K f(w)$, $\sum_{i=1}^K g(\mathbb{E}[w_{t,i}]) \leq \max_{w \in \Delta^{K-1}} \sum_{i=1}^K g(w)$, and choosing an appropriate learning rate. In particular, for $\alpha = 1/2$ we have $f(x) \propto \sqrt{x}$ and $g(x) \propto \sqrt{x}$ and we use $\eta_t \propto 1/\sqrt{t}$, for which $\eta_{t+1}^{-1} - \eta_t^{-1} = \Theta(1/\sqrt{t})$. This gives

$$\overline{Reg}_T \leq \sum_{t=1}^T c \frac{1}{\sqrt{t}} \sum_{i=1}^K \sqrt{\mathbb{E}[w_{t,i}]} \leq \sum_{t=1}^T c \frac{1}{\sqrt{t}} \max_{z \in \Delta^{K-1}} \sum_{i=1}^K \sqrt{z_i} \leq \sum_{t=1}^T c \frac{1}{\sqrt{t}} \sqrt{K} \leq 2c\sqrt{KT},$$

where c is a small constant and we replace $\mathbb{E}[w_{t,i}]$ with z_i in the maximization.

The main insight of the paper is that the same framework can be used to obtain logarithmic bounds in the stochastic case. The key novelty is that if we constrain the maximization of $\mathbb{E}[w_{t,i}]$ by the self-bounding property of the regret (4), the space of solutions excludes the worst-case scenario, where the regret grows with the square root of the time horizon. For simplicity, we first explain the approach with $C = 0$. By the self-bounding property (4), we then have $\overline{Reg}_T \geq \sum_{t=1}^T \sum_{i \neq i^*} \Delta_i \mathbb{E}[w_{t,i}] = \sum_{t=1}^T \sum_i \Delta_i \mathbb{E}[w_{t,i}]$ (since $\Delta_{i^*} = 0$ by definition), which we can use to write

$$\overline{Reg}_T \leq 2\overline{Reg}_T - \sum_{t=1}^T \sum_i \Delta_i \mathbb{E}[w_{t,i}]. \quad (7)$$

The negative contributions $-\Delta_i \mathbb{E}[w_{t,i}]$ are used to achieve better control of the growth of $\mathbb{E}[w_{t,i}]$, but they are only helpful for i with $\Delta_i > 0$, i.e., only for $i \neq i^*$. Therefore, we derive

refined bounds for the stability and penalty terms:

$$\begin{aligned} \text{stability} &\leq \sum_{t=1}^T \eta_t \sum_{i \neq i^*} \tilde{f}(\mathbb{E}[w_{t,i}]), \\ \text{penalty} &\leq \sum_{t=1}^T (\eta_{t+1}^{-1} - \eta_t^{-1}) \sum_{i \neq i^*} g(\mathbb{E}[w_{t,i}]), \end{aligned}$$

where the summation excludes the best arm i^* , which has no negative contribution in (7). The cost of excluding the best arm is an addition of a linear term to f : $\tilde{f}(x) = f(x) + c'x \leq (1 + c')f(x)$, where c' is a small constant. In particular, for $\alpha = \frac{1}{2}$ and learning rate $\eta_t \propto 1/\sqrt{t}$ we have

$$\begin{aligned} \overline{\text{Reg}}_T &\leq 2\overline{\text{Reg}}_T - \sum_{t=1}^T \sum_{i \neq i^*} \Delta_i \mathbb{E}[w_{t,i}] \\ &\leq \sum_{t=1}^T \sum_{i \neq i^*} \left(2c_1 \frac{1}{\sqrt{t}} \sqrt{\mathbb{E}[w_{t,i}]} - \Delta_i \mathbb{E}[w_{t,i}] \right) \\ &\leq \sum_{t=1}^T \sum_{i \neq i^*} \max_z \left(2c_1 \frac{1}{\sqrt{t}} \sqrt{z} - \Delta_i z \right) \\ &\leq \sum_{t=1}^T \sum_{i \neq i^*} \frac{c_2}{\Delta_i t} \\ &= O \left(\sum_{i \neq i^*} \frac{\log T}{\Delta_i} \right), \end{aligned}$$

where c_1 and c_2 are small constants and in the second line we used the refined stability and penalty bounds to bound $2\overline{\text{Reg}}_T$. The negative contribution is exploited in the maximization in the third line, which is now done coordinate-wise and the constraint that w_t is a probability distribution is dropped.

We assume uniqueness of the zero-entry in Δ , because currently we are only able to exclude one arm from the summation in the refined bound on stability. Had there been multiple arms with $\Delta_i = 0$, they would have no negative contributions to control $\mathbb{E}[w_{t,i}]$. The challenge in excluding more than one arm from the summation is explained in Lemma 11, where we derive the refined bound.

In the more general analysis with $C > 0$, we introduce a parameter λ and write $\overline{\text{Reg}}_T \leq (1 + \lambda)\overline{\text{Reg}}_T - \lambda \left(\sum_{t=1}^T \sum_{i \neq i^*} \Delta_i \mathbb{E}[w_{t,i}] - C \right)$. We use λ for optimizing the dependence on C . The parameter λ can also be seen as a Lagrange multiplier in a constrained optimization problem of maximizing the regret bound (stability bound + penalty bound) under the self-bounding constraint that (stability bound + penalty bound) $\geq \sum_{t=1}^T \sum_{i \neq i^*} \Delta_i \mathbb{E}[w_{t,i}] - C$.

7.2 Key Lemmas

The proofs of Theorems 1, 3, and 4 are based on the following two lemmas that bound the *stability* and *penalty* terms. The proofs of the lemmas are provided in Appendix C.

Lemma 11 *For a positive learning rate, the instantaneous stability of α -TSALLIS-INF satisfies at any time t*

$$\mathbb{E} \left[\ell_{t, I_t} + \Phi_t(-\hat{L}_t) - \Phi_t(-\hat{L}_{t-1}) \right] \leq \begin{cases} \min \left\{ \sum_{i=1}^K \frac{\eta_t \xi_i}{2} \mathbb{E} [w_{t,i}]^{1-\alpha}, 1 \right\}, & \text{if 1.} \\ \frac{\eta_t^2}{2} + \sum_{i=1}^K \frac{\eta_t}{2} \mathbb{E} [w_{t,i}]^{\frac{1}{2}} (1 - \mathbb{E} [w_{t,i}]), & \text{if 2.} \\ \frac{7\eta_t^2}{8} K + \sum_{i=1}^K \frac{\eta_t}{8} \mathbb{E} [w_{t,i}]^{\frac{1}{2}} (1 - \mathbb{E} [w_{t,i}]), & \text{if 3.} \\ \sum_{i \neq j} \left(\frac{\eta_t \xi_i}{2} \mathbb{E} [w_{t,i}]^{1-\alpha} + \frac{\eta_t (\xi_i + 2\xi_j)}{2} \mathbb{E} [w_{t,i}] \right), & \text{if 4.,} \end{cases}$$

where

1. \hat{L}_t is based on IW estimators. The inequality holds for any $\eta_t > 0$ and $\alpha \in [0, 1]$.
2. \hat{L}_t is based on IW estimators, $1 \geq \eta_t > 0$, and $\alpha = \frac{1}{2}$.
3. \hat{L}_t is based on RV estimators, $1 \geq \eta_t > 0$, and $\alpha = \frac{1}{2}$.
4. \hat{L}_t is based on IW estimators and $\eta_t \xi_i \leq \frac{1}{4}$ for all i . The inequality holds for any j and $\alpha \in [0, 1]$.

The first part of the Lemma is due to Abernethy et al. (2015). The remaining parts are non-trivial refinements that are crucial for our analysis, as outlined in the previous section. The first inequality is used in the proof of Theorem 3, the second and third inequalities are used for the two results in Theorem 1, and the last inequality is used in the proof of Theorem 4. In the proof of Theorem 1 we use $\mathbb{E} [w_{t,i}]^{\frac{1}{2}} (1 - \mathbb{E} [w_{t,i}]) \leq \mathbb{E} [w_{t,i}]^{\frac{1}{2}}$ for $i \neq i^*$ and $\mathbb{E} [w_{t,i^*}]^{\frac{1}{2}} (1 - \mathbb{E} [w_{t,i^*}]) \leq (1 - \mathbb{E} [w_{t,i^*}]) = \sum_{i \neq i^*} \mathbb{E} [w_{t,i}]$ for i^* . This eliminates $\mathbb{E} [w_{t,i^*}]$ from the regret bound and allows to exploit the self-bounding property. The approach only allows to eliminate one arm from the regret bound, which is the reason we rely on the assumption of uniqueness of the best arm.

Lemma 12 *For any $\alpha \in [0, 1]$ and any unbiased loss estimators the penalty term of α -TSALLIS-INF satisfies:*

1. *For the symmetric regularizer and a non-increasing sequence of positive learning rates η_1, η_2, \dots*

$$\mathbb{E} \left[\sum_{t=1}^T \left(\Phi_t(-\hat{L}_{t-1}) - \Phi_t(-\hat{L}_t) - \ell_{t, i_T^*} \right) \right] \leq \frac{(K^{1-\alpha} - 1)(1 - T^{-\alpha})}{(1 - \alpha)\alpha\eta_T} + 1.$$

2. *For an arbitrary regularizer, a non-increasing sequence of positive learning rates η_1, η_2, \dots , and any $x \in [1, \infty]$*

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T \left(\Phi_t(-\hat{L}_{t-1}) - \Phi_t(-\hat{L}_t) - \ell_{t,i_T^*} \right) \right] \\ & \leq \frac{1 - T^{-\alpha x}}{\alpha} \sum_{i \neq i_T^*} \left(\frac{\mathbb{E}[w_{1,i}]^\alpha - \alpha \mathbb{E}[w_{1,i}]}{\eta_1 \xi_i (1 - \alpha)} + \sum_{t=2}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \frac{\mathbb{E}[w_{t,i}]^\alpha - \alpha \mathbb{E}[w_{t,i}]}{\xi_i (1 - \alpha)} \right) + T^{1-x}. \end{aligned}$$

The first part of the Lemma is a straightforward improvement of the penalty bound in Abernethy et al. (2015) with the techniques from Agarwal et al. (2017). The second part is again a crucial refinement. It is obtained by exploiting the negative contribution of $\Psi_T(\mathbf{e}_{i^*})$ in an intermediate step of the proof, which Abernethy et al. (2015) trivially bounded by 0.

7.3 Proofs of Theorems 1 and 3

Now we are ready to present proofs of the main theorems.

Proof of Theorem 1 We provide a proof of regret bounds for TSALLIS-INF with RV estimators. The analysis of TSALLIS-INF with IW estimators in the adversarial case is analogous to the proof of Theorem 3 and under the self-bounding constraint (4), it is analogous to the analysis of RV estimators with the bound in Part 3 of Lemma 11 replaced by the bound in Part 2. Therefore, the proofs of both results for the IW estimators are omitted.

To analyze the regret, we start by bounding the *stability* term. We use Lemma 11. For $t < 16$ we have $\eta_t > 1$ and the RV estimators are equivalent to IW estimators. Thus, we can apply the first part of the lemma to bound the instantaneous stability by 1. For $t \geq 16$, we use the third part of the lemma.

$$\begin{aligned} \text{stability} &= \mathbb{E} \left[\sum_{t=1}^T \ell_{t,I_t} + \Phi_t(-\hat{L}_t) - \Phi_t(-\hat{L}_{t-1}) \right] \\ &\leq 15 + \sum_{t=16}^T \left(\frac{7\eta_t^2}{8} K + \sum_{i=1}^K \frac{\eta_t}{8} \sqrt{\mathbb{E}[w_{t,i}]} (1 - \mathbb{E}[w_{t,i}]) \right) \\ &\leq 15 + 14K \log(T) + \sum_{t=16}^T \sum_{i=1}^K \frac{\sqrt{\mathbb{E}[w_{t,i}]} (1 - \mathbb{E}[w_{t,i}])}{2\sqrt{t}}. \end{aligned} \quad (8)$$

Adversarial bound. We bound $\sum_{i=1}^K \sqrt{\mathbb{E}[w_{t,i}]} (1 - \mathbb{E}[w_{t,i}]) \leq \sum_{i=1}^K \sqrt{\mathbb{E}[w_{t,i}]} \leq \sqrt{K}$, where the last step holds by simple maximization. Then we have

$$\begin{aligned} \text{stability} &\leq 15 + 14K \log(T) + \sum_{t=16}^T \frac{\sqrt{K}}{2\sqrt{t}} \\ &\leq 15 + 14K \log(T) + \sqrt{KT}. \end{aligned}$$

For the *penalty* term, we use the first part of Lemma 12 to obtain

$$\text{penalty} \leq \sqrt{KT} + 1.$$

Combining *stability* and *penalty* completes the proof.

Bound under the self-bounding constraint (4). We continue bounding the *stability* up from equation (8). For $i \neq i^*$, we use $\sqrt{\mathbb{E}[w_{t,i}]}(1 - \mathbb{E}[w_{t,i}]) \leq \sqrt{\mathbb{E}[w_{t,i}]}$. For i^* , we use $\sqrt{\mathbb{E}[w_{t,i^*}]}(1 - \mathbb{E}[w_{t,i^*}]) \leq (1 - \mathbb{E}[w_{t,i^*}]) = \sum_{i \neq i^*} \mathbb{E}[w_{t,i}]$. For a constant $0 < \lambda \leq 1$ that will be specified at a later stage of the proof and $t \leq T_0 = \left\lceil \left(\frac{1}{2\lambda\Delta_{\min}}\right)^2 \right\rceil$, we further bound the last expression as $\sum_{i \neq i^*} \mathbb{E}[w_{t,i}] \leq 1$. Altogether, this gives

$$\begin{aligned} \sum_{t=16}^T \sum_{i=1}^K \frac{\sqrt{\mathbb{E}[w_{t,i}]}(1 - \mathbb{E}[w_{t,i}])}{2\sqrt{t}} &\leq \sum_{t=16}^{T_0} \frac{1}{2\sqrt{t}} + \sum_{i \neq i^*} \left(\sum_{t=T_0+1}^T \frac{\mathbb{E}[w_{t,i}]}{2\sqrt{t}} + \sum_{t=16}^T \frac{\sqrt{\mathbb{E}[w_{t,i}]}}{2\sqrt{t}} \right) \\ &\leq \sqrt{T_0} + \sum_{i \neq i^*} \left(\sum_{t=1}^T \frac{\sqrt{\mathbb{E}[w_{t,i}]}}{2\sqrt{t}} + \sum_{t=T_0+1}^T \frac{\mathbb{E}[w_{t,i}]}{2\sqrt{t}} \right) \end{aligned}$$

and

$$\text{stability} \leq 15 + 14K \log(T) + \sqrt{T_0} + \sum_{i \neq i^*} \left(\sum_{t=1}^T \frac{\sqrt{\mathbb{E}[w_{t,i}]}}{2\sqrt{t}} + \sum_{t=T_0+1}^T \frac{\mathbb{E}[w_{t,i}]}{2\sqrt{t}} \right).$$

In order to bound the *penalty* term, we use the second part of Lemma 12 with $x = \infty$. At the end of the derivation we use Lemma 13, by which $\sum_{t=2}^{\infty} \left(\sqrt{t} - \sqrt{t-1} - \frac{1}{2\sqrt{t}} \right) \leq \frac{1}{4}$.

$$\begin{aligned} \text{penalty} &= \mathbb{E} \left[\sum_{t=1}^T -\Phi_t(-\hat{L}_t) + \Phi_t(-\hat{L}_{t-1}) - \ell_{t,i^*} \right] \\ &\leq 4 \sum_{i \neq i^*} \left(\frac{\sqrt{\mathbb{E}[w_{1,i}]} - \frac{1}{2} \mathbb{E}[w_{1,i}]}{\eta_1} + \sum_{t=2}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \left(\sqrt{\mathbb{E}[w_{t,i}]} - \frac{1}{2} \mathbb{E}[w_{t,i}] \right) \right) \\ &= \sum_{i \neq i^*} \left(\left(\sqrt{\mathbb{E}[w_{1,i}]} - \frac{1}{2} \mathbb{E}[w_{1,i}] \right) + \sum_{t=2}^T \left(\sqrt{t} - \sqrt{t-1} \right) \left(\sqrt{\mathbb{E}[w_{t,i}]} - \frac{1}{2} \mathbb{E}[w_{t,i}] \right) \right) \\ &= \sum_{t=1}^T \left(\sum_{i \neq i^*} \frac{\sqrt{\mathbb{E}[w_{t,i}]} - \frac{1}{2} \mathbb{E}[w_{t,i}]}{2\sqrt{t}} \right) + \sum_{i \neq i^*} \left(\frac{\sqrt{\mathbb{E}[w_{1,i}]} - \frac{1}{2} \mathbb{E}[w_{1,i}]}{2} \right. \\ &\quad \left. + \sum_{t=2}^T \left(\sqrt{t} - \sqrt{t-1} - \frac{1}{2\sqrt{t}} \right) \left(\sqrt{\mathbb{E}[w_{t,i}]} - \frac{1}{2} \mathbb{E}[w_{t,i}] \right) \right) \\ &\leq \sum_{t=1}^T \left(\sum_{i \neq i^*} \frac{\sqrt{\mathbb{E}[w_{t,i}]} - \frac{1}{2} \mathbb{E}[w_{t,i}]}{2\sqrt{t}} \right) + \left(\frac{1}{2} + \sum_{t=2}^T \left(\sqrt{t} - \sqrt{t-1} - \frac{1}{2\sqrt{t}} \right) \right) \sqrt{K} \\ &\leq \sum_{t=1}^T \left(\sum_{i \neq i^*} \frac{\sqrt{\mathbb{E}[w_{t,i}]} - \frac{1}{2} \mathbb{E}[w_{t,i}]}{2\sqrt{t}} \right) + \frac{3}{4} \sqrt{K}. \end{aligned}$$

Combining *penalty* and *stability* gives the bound

$$\overline{\text{Reg}}_T \leq \sum_{i \neq i^*} \left(\sum_{t=1}^T \frac{\sqrt{\mathbb{E}[w_{t,i}]} + \sum_{t=T_0+1}^T \frac{\mathbb{E}[w_{t,i}]}{4\sqrt{t}} \right) + \sqrt{T_0} + \underbrace{\frac{3}{4}\sqrt{K} + 15 + 14K \log(T)}_{=:M}.$$

By using the self-bounding property (4) and $(1 + \lambda) \leq 2$ we obtain

$$\begin{aligned} \overline{\text{Reg}}_T &\leq \overline{\text{Reg}}_T + \lambda \left(\overline{\text{Reg}}_T - \sum_{t=1}^T \sum_{i \neq i^*} \Delta_i \mathbb{E}[w_{t,i}] + C \right) \\ &\leq \sum_{i \neq i^*} \left(\sum_{t=1}^T \frac{2\sqrt{\mathbb{E}[w_{t,i}]} + \sum_{t=T_0+1}^T \frac{\mathbb{E}[w_{t,i}]}{2\sqrt{t}} \right) + 2\sqrt{T_0} + 2M - \lambda \sum_{t=1}^T \sum_{i \neq i^*} \Delta_i \mathbb{E}[w_{t,i}] + \lambda C \\ &= \sum_{i \neq i^*} \left(\sum_{t=1}^{T_0} \left(\frac{2\sqrt{\mathbb{E}[w_{t,i}]} + \frac{1}{2} \mathbb{E}[w_{t,i}]}{\sqrt{t}} - \lambda \Delta_i \mathbb{E}[w_{t,i}] \right) + \sum_{t=T_0+1}^T \left(\frac{2\sqrt{\mathbb{E}[w_{t,i}] + \frac{1}{2} \mathbb{E}[w_{t,i}]} + \frac{1}{2} \mathbb{E}[w_{t,i}]}{\sqrt{t}} - \lambda \Delta_i \mathbb{E}[w_{t,i}] \right) \right) \\ &\quad + 2\sqrt{T_0} + 2M + \lambda C \\ &\leq \sum_{i \neq i^*} \left(\sum_{t=1}^{T_0} \max_{z \geq 0} \left(\frac{2\sqrt{z}}{\sqrt{t}} - \lambda \Delta_i z \right) + \sum_{t=T_0+1}^T \max_{z \geq 0} \left(\frac{2\sqrt{z} + \frac{1}{2}z}{\sqrt{t}} - \lambda \Delta_i z \right) \right) \\ &\quad + 2\sqrt{T_0} + 2M + \lambda C. \end{aligned}$$

Simple optimization shows that $\max_{z > 0} 2\gamma\sqrt{z} - \beta z = \frac{\gamma^2}{\beta}$. Thus, we have

$$\max_{z \geq 0} \frac{2\sqrt{z}}{\sqrt{t}} - \lambda \Delta_i z = \frac{1}{\lambda \Delta_i t}$$

and

$$\begin{aligned} \max_{z \geq 0} \frac{2(\sqrt{z} + \frac{1}{4}z)}{\sqrt{t}} - \lambda \Delta_i z &= \frac{1}{(\lambda \Delta_i - \frac{1}{4\sqrt{t}})t} \\ &= \frac{1}{\lambda \Delta_i t} + \frac{1}{(\lambda \Delta_i - \frac{1}{4\sqrt{t}})t} - \frac{1}{\lambda \Delta_i t} \\ &= \frac{1}{\lambda \Delta_i t} + \frac{1}{4\lambda^2 \Delta_i^2 t^{\frac{3}{2}} - \lambda \Delta_i t}. \end{aligned}$$

In order to bound the summation of the above terms, we use the following bound from Lemma 15 in the appendix:

$$\sum_{t=T_0+1}^T \frac{1}{bt^{\frac{3}{2}} - ct} \leq \frac{2}{b\sqrt{T_0} - c}.$$

By definition of T_0 we have $\frac{1}{2\lambda\Delta_{\min}} \leq \sqrt{T_0} \leq \frac{1}{2\lambda\Delta_{\min}} + 1$ and

$$\frac{1}{2\lambda^2 \Delta_i^2 \sqrt{T_0} - \frac{1}{2} \lambda \Delta_i} = \frac{2}{\lambda \Delta_i (4\lambda \Delta_i \sqrt{T_0} - 1)} \leq \frac{2}{\lambda \Delta_i \left(2\frac{\Delta_i}{\Delta_{\min}} - 1 \right)} \leq \frac{2}{\lambda \Delta_i}.$$

By plugging the calculations into the regret bound above we obtain:

$$\begin{aligned}
 \overline{Reg}_T &\leq \sum_{i \neq i^*} \left(\sum_{t=1}^T \frac{1}{\lambda \Delta_i t} + \sum_{t=T_0+1}^T \frac{1}{4\lambda^2 \Delta_i^2 t^{\frac{3}{2}} - \lambda \Delta_i t} \right) + 2\sqrt{T_0} + 2M + \lambda C \\
 &\leq \sum_{i \neq i^*} \left(\left(\sum_{t=1}^T \frac{1}{\lambda \Delta_i t} \right) + \frac{1}{2\lambda^2 \Delta_i^2 \sqrt{T_0} - \frac{1}{2}\lambda \Delta_i} \right) + 2\sqrt{T_0} + 2M + \lambda C \\
 &\leq \sum_{i \neq i^*} \frac{\log(T) + 3}{\lambda \Delta_i} + \frac{1}{\lambda \Delta_{\min}} + 2(M + 1) + \lambda C.
 \end{aligned}$$

Finally, choosing $\lambda = \min \left\{ 1, \sqrt{\left(\sum_{i \neq i^*} \frac{\log(T)+3}{\Delta_i} + \frac{1}{\Delta_{\min}} \right) / C} \right\}$ completes the proof. \blacksquare

Proof of Theorem 3 We start from equation (6). Since the regularization is symmetric, we have $\xi_i = 1$ for all i . Using Lemma 11, we bound the *stability* term as

$$\begin{aligned}
 \text{stability} &= \mathbb{E} \left[\sum_{t=1}^T \ell_{t, I_t} + \Phi_t(-\hat{L}_t) - \Phi_t(-\hat{L}_{t-1}) \right] \leq \sum_{t=1}^T \sum_{i=1}^K \frac{\eta_t}{2} \mathbb{E} [w_{t,i}]^{1-\alpha} \\
 &\leq \left(\sum_{t=1}^T \frac{\eta_t}{2} \right) \max_{z \in \Delta^{K-1}} \sum_{i=1}^K z_i^{1-\alpha} = \left(\sum_{t=1}^T \sqrt{\frac{K^{1-2\alpha} - K^{-\alpha} (1-t^{-\alpha})}{1-\alpha}} \frac{K^\alpha}{\alpha t} \right) \frac{K^\alpha}{2} \\
 &\leq \left(\sum_{t=1}^T \sqrt{\frac{1 - K^{\alpha-1} (1-T^{-\alpha})}{1-\alpha}} \frac{\sqrt{K}}{\alpha t} \right) \frac{\sqrt{K}}{2} \leq \sqrt{\frac{1 - K^{\alpha-1} (1-T^{-\alpha})}{1-\alpha}} \frac{KT}{\alpha}.
 \end{aligned}$$

The *penalty* is bounded according to Lemma 12

$$\begin{aligned}
 \text{penalty} &= \mathbb{E} \left[\sum_{t=1}^T -\Phi_t(-\hat{L}_t) + \Phi_t(-\hat{L}_{t-1}) - \ell_{t, i_T^*} \right] \\
 &\leq \frac{(K^{1-\alpha} - 1)(1 - T^{-\alpha})}{(1-\alpha)\alpha\eta_T} + 1 = \sqrt{\frac{1 - K^{\alpha-1} (1 - T^{-\alpha})}{1-\alpha}} \frac{KT}{\alpha} + 1.
 \end{aligned}$$

The proof is completed by noting that the first factor is bounded by $\sqrt{\frac{1}{1-\alpha}}$ and monotonically increasing in α with the limit $\lim_{\alpha \rightarrow 1} \sqrt{\frac{1 - K^{\alpha-1}}{1-\alpha}} = \sqrt{\log(K)}$ (details in Lemma 14 in the appendix). By the same argument, the second factor is bounded by $\sqrt{\frac{1}{\alpha}}$ and monotonically decreasing in α with the limit $\lim_{\alpha \rightarrow 0} \sqrt{\frac{1 - T^{-\alpha}}{\alpha}} = \sqrt{\log(T)}$. \blacksquare

8. Experiments

We provide an empirical comparison of TSALLIS-INF with IW and with RV loss estimators with the classical algorithms for stochastic bandits, UCB1 (Auer et al., 2002a, with parameter $\alpha = 1.5$) and THOMPSON SAMPLING (Thompson, 1933)⁵, and the classical algorithm for adversarial bandits, EXP3, implemented for the losses (Bubeck and Cesa-Bianchi, 2012). We also compare with the state-of-the-art algorithms for stochastic and adversarial bandits, EXP3++ with parametrization proposed by Seldin and Lugosi (2017) and BROAD (Wei and Luo, 2018). The pseudo-regret is estimated by 100 repetitions of the corresponding experiments and two standard deviations of the empirical pseudo-regret, $\sum_{t=1}^T \Delta_{I_t}$, over the 100 repetitions are depicted by the shaded areas on the plots. We always show the first 10000 time steps on a linear plot and then the time steps from 10^4 to 10^7 on a separate log-log plot.

The first experiment, shown in Figures 1 and 2, is a standard stochastic MAB, where the mean rewards are $(1 + \Delta)/2$ for the single optimal arm and $(1 - \Delta)/2$ for all the suboptimal arms. The number of arms K and the gaps Δ are varied as described in the figures. Unsurprisingly, THOMPSON SAMPLING exhibits the lowest regret, but TSALLIS-INF with RV estimators follows closely behind and outperforms all other competitors by a large margin. TSALLIS-INF with IW estimators takes a confident third place, while UCB1, EXP3, and EXP3++ fall roughly in the same league. BROAD suffers from extremely large constant factors and is out of question for practical applications.

The second experiment, shown in Figures 3 and 4, simulates stochastically constrained adversaries. The mean loss of (optimal arm, all sub-optimal arms) switches between $(1 - \Delta, 1)$ and $(0, \Delta)$, while staying unchanged for phases that are increasing exponentially in length. Both UCB1 and THOMPSON-SAMPLING suffer almost linear regret. To the best of our knowledge, this is the first empirical evidence clearly demonstrating that THOMPSON SAMPLING is unsuitable for adversarial regimes. All other algorithms are almost unaffected by the shifting of the means.

Both experiments confirm that TSALLIS-INF with IW estimators achieves logarithmic regret in stochastic and stochastically constrained adversarial environments and that RV estimators significantly improve the constants.

8.1 Multiple Optimal Arms

Since our theoretical results for the stochastic setting do not include multiple optimal arms, we explore this setting empirically. We use a single suboptimal arm with a mean loss of $9/16$. All other arms are optimal with a mean loss of $7/16$. We run the experiment with 1000 repetitions and increase the number of arms. Figure 5 clearly shows that the regret does not suffer if the optimal arm is not unique. On the opposite, we observe that the regret decreases with the growth of the number of suboptimal arms. Therefore, we conjecture that the requirement of uniqueness is merely an artifact of the analysis.

5. Another leading stochastic algorithm, KL-UCB (Cappé et al., 2013), has performed comparably to THOMPSON SAMPLING in our experiments and, therefore, is not reported in the figures.




Figure 1: Comparison of TSALLIS-INF with IW and with RV estimators with THOMPSON SAMPLING, UCB1, EXP3, EXP3++, and BROAD in a stochastic environment with fixed mean losses of $\frac{1-\Delta}{2}$ for the optimal arm and $\frac{1+\Delta}{2}$ for all sub-optimal arms. The experiment is repeated for different number of arms K and different gaps Δ . The figure shows the first 10000 time steps on a linear plot. The pseudo-regret is estimated by 100 repetitions and we depict 2 standard deviations of the empirical pseudo-regret by the shaded areas.




Figure 2: Comparison of TSALLIS-INF with IW and with RV estimators with THOMPSON SAMPLING, UCB1, EXP3, EXP3++, and BROAD in a stochastic environment with fixed mean losses of $\frac{1-\Delta}{2}$ for the optimal arm and $\frac{1+\Delta}{2}$ for all sub-optimal arms. The experiment is repeated for different number of arms K and different gaps Δ . The figure shows the time steps from 10^4 to 10^7 on a log-log plot. The pseudo-regret is estimated by 100 repetitions and we depict 2 standard deviations of the empirical pseudo-regret by the shaded areas.




Figure 3: Comparison of TSALLIS-INF with IW and with RV estimators with THOMPSON SAMPLING, UCB1, EXP3, EXP3++, and BROAD in a stochastically constrained adversarial environment. The environment (unknown to the agent) alternates between two stochastic settings. In the first setting the expected loss of the optimal arm is 0 and Δ for sub-optimal arms. In the second the expected losses are $1 - \Delta$ and 1, respectively. The time between alternations increases exponentially (with factor 1.6) after each switch. The experiment is repeated for different number of arms K and different gaps Δ . The figure shows the first 10000 time steps on a linear plot. The pseudo-regret is estimated by 100 repetitions and we depict 2 standard deviations of the empirical pseudo-regret by the shaded areas.




Figure 4: Comparison of TSALLIS-INF with IW and with RV estimators with THOMPSON SAMPLING, UCB1, EXP3, EXP3++, and BROAD in a stochastically constrained adversarial environment. The environment (unknown to the agent) alternates between two stochastic settings. In the first setting the expected loss of the optimal arm is 0 and Δ for sub-optimal arms. In the second the expected losses are $1 - \Delta$ and 1 respectively. The time between alternations increases exponentially (with factor 1.6) after each switch. The experiment is repeated for different number of arms K and different gaps Δ . The figure shows the time steps from 10^4 to 10^7 on a log-log plot. The pseudo-regret is estimated by 100 repetitions and we depict 2 standard deviations of the empirical pseudo-regret by the shaded areas.

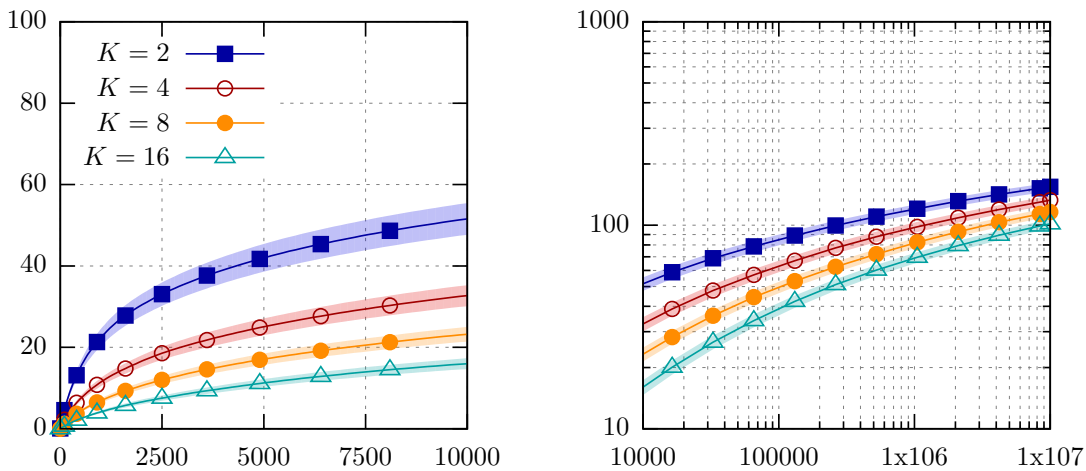


Figure 5: Increasing number of copies of the best arm

9. Discussion

We have presented a general analysis of online mirror descent algorithms regularized by Tsallis entropy with $\alpha \in [0, 1]$. As the main contribution, we have shown that the special case of $\alpha = \frac{1}{2}$ achieves optimality in both adversarial and stochastic regimes, while being oblivious to the environment at hand. Thereby, we have closed logarithmic gaps to lower bounds, which were present in existing best-of-both-worlds algorithms. We introduced a novel proof technique based on the self-bounding property of the regret, circumventing the need of controlling the variance of loss estimates. We have provided an empirical evidence that our algorithm outperforms UCB1 in stochastic environments and is significantly more robust than UCB1 and THOMPSON SAMPLING in non-i.i.d. settings. We have introduced an adversarial regime with a self-bounding constraint, which includes stochastically constrained adversaries and adversarially corrupted stochastic bandits as special cases and improved regret bounds for the latter two regimes. We have also shown that TSALLIS-INF can be applied to achieve stochastic and adversarial optimality in utility-based dueling bandits.

A weak point of the current analysis is the assumption on uniqueness of the zero entry in a vector of suboptimality gaps in the adversarial regime with a self-bounding constraint. In stochastic and stochastically constrained adversarial settings, it corresponds to assumption of uniqueness of the best arm. Our experiments suggest that this is most likely an artifact of the analysis and we aim to address this shortcoming in future work.

Another open question is whether it is possible to close the remaining factor 2 gap between the upper and lower gap-dependent asymptotic regret bounds, either by improving the upper bound in the stochastic regime or deriving a tighter lower bound for the adversarial regime with a self-bounding constraint.

One more open question is whether logarithmic regret is achievable by TSALLIS-INF in the intermediate regimes defined by Seldin and Slivkins (2014). We have discussed this question in more detail in Section 5.2.

An additional direction for future research is the application of TSALLIS-INF to other problems. The fact that the algorithm relies solely on importance weighted losses makes it a suitable candidate for partial monitoring games. One step in this direction has already been taken by Zimmert et al. (2019).

Acknowledgments

We would like to thank Chloé Rouyer for pointing out several bugs in the previous version of the work (Zimmert and Seldin, 2019) and Haipeng Luo for the idea on how to improve our regret bounds for stochastic bandits with adversarial corruptions in the large C case. We are also grateful to the anonymous reviewers for their comments. We acknowledge partial support by the Independent Research Fund Denmark, grant number 9040-00361B.

Appendix A. Asymptotic Lower Bound

If the optimal arm has mean reward $\frac{1}{2}$ and suboptimal arms have the gaps Δ_i then the following lower bound for any consistent algorithm follows from Lai and Robbins (1985, Theorem 2)

$$\lim_{t \rightarrow \infty} \frac{\overline{Reg}_t}{\log(t)} \geq \sum_{i: \Delta_i > 0} \frac{\Delta_i}{\text{kl}(\frac{1}{2} + \Delta_i, \frac{1}{2})}.$$

For any $\Delta_i \in [0, 0.5]$ the kl term can be upper bounded as

$$\text{kl}(\frac{1}{2} + \Delta_i, \frac{1}{2}) \leq 2\Delta_i^2 + 3\Delta_i^3,$$

which can be verified by taking Taylor's expansion at $\Delta_i = 0$. Therefore,

$$\begin{aligned} \sum_{i: \Delta_i > 0} \frac{\Delta_i}{\text{kl}(\frac{1}{2} + \Delta_i, \frac{1}{2})} &\geq \sum_{i: \Delta_i > 0} \frac{1}{2\Delta_i + 3\Delta_i^2} \\ &= \frac{1}{2} \sum_{i: \Delta_i > 0} \frac{1}{\Delta_i} - \sum_{i: \Delta_i > 0} \frac{\frac{3}{2}\Delta_i}{2\Delta_i + 3\Delta_i^2} \geq \frac{1}{2} \left(\sum_{i: \Delta_i > 0} \frac{1}{\Delta_i} - \frac{3}{2}K \right). \end{aligned}$$

Thus, for any consistent algorithm we obtain

$$\lim_{\|\Delta\| \rightarrow 0} \left(\left(\sum_{i: \Delta_i > 0} \frac{1}{\Delta_i} \right)^{-1} \liminf_{t \rightarrow \infty} \frac{\mathbb{E}[\overline{Reg}_t]}{\log(t)} \right) \geq \lim_{\|\Delta\| \rightarrow 0} \left(\frac{1}{2} - \frac{3}{4} \left(\sum_{i: \Delta_i > 0} \frac{1}{\Delta_i} \right)^{-1} K \right) = \frac{1}{2},$$

since $\lim_{\|\Delta\| \rightarrow 0} \left(\sum_{i: \Delta_i > 0} \Delta_i^{-1} \right)^{-1} K = 0$.

Appendix B. Technical Lemmas

Lemma 13

$$\sum_{t=2}^{\infty} \left(\sqrt{t} - \sqrt{t-1} - \frac{1}{2\sqrt{t}} \right) \leq \frac{1}{4}.$$

Proof We have $\sqrt{t} - \sqrt{t-1} = \frac{(\sqrt{t}-\sqrt{t-1})(\sqrt{t}+\sqrt{t-1})}{\sqrt{t}+\sqrt{t-1}} = \frac{1}{\sqrt{t}+\sqrt{t-1}}$ and $\frac{1}{\sqrt{t}+\sqrt{t-1}} - \frac{1}{2\sqrt{t}} = \frac{1}{2\sqrt{t}(\sqrt{t}+\sqrt{t-1})^2} \leq \frac{1}{8(t-1)^{3/2}}$. Summation of the latter is related to the Riemann zeta function $\zeta\left(\frac{3}{2}\right) = \sum_{t=1}^{\infty} \frac{1}{t^{3/2}} \leq 2.62$. To get a slightly tighter bound, we count the first two terms explicitly and bound the rest using Riemann zeta function:

$$\begin{aligned} \sum_{t=2}^{\infty} \left(\sqrt{t} - \sqrt{t-1} - \frac{1}{2\sqrt{t}} \right) &= \sqrt{2} - 1 - \frac{1}{2\sqrt{2}} + \sqrt{3} - \sqrt{2} - \frac{1}{2\sqrt{3}} + \sum_{t=4}^{\infty} \left(\sqrt{t} - \sqrt{t-1} - \frac{1}{2\sqrt{t}} \right) \\ &\leq \sqrt{3} - 1 - \frac{1}{2\sqrt{2}} - \frac{1}{2\sqrt{3}} + \frac{1}{8} \sum_{t=3}^{\infty} \frac{1}{t^{3/2}} \\ &= \sqrt{3} - 1 - \frac{1}{2\sqrt{2}} - \frac{1}{2\sqrt{3}} + \frac{1}{8} \left(\zeta\left(\frac{3}{2}\right) - 1 - \frac{1}{2\sqrt{2}} \right) \\ &\leq \frac{1}{4}. \end{aligned}$$

■

Lemma 14 For any $y > 0$ and $x > 0$ the function $\frac{1-y^{-x}}{x}$ is non-increasing in x and has the limit

$$\lim_{x \rightarrow 0} \frac{1-y^{-x}}{x} = \log(y),$$

therefore, $\frac{1-y^{-x}}{x} \leq \min\{x^{-1}, \log(y)\}$.

Proof Taking the derivative and using the inequality $z \leq e^z - 1$:

$$\frac{\partial}{\partial x} \left(\frac{1-y^{-x}}{x} \right) = \frac{\log(y)y^{-x}x - (1-y^{-x})}{x^2} \leq \frac{(e^{\log(y)x} - 1)y^{-x} - (1-y^{-x})}{x^2} = 0.$$

The limit by L'Hôpital's rule is

$$\lim_{x \rightarrow 0} \frac{1-y^{-x}}{x} = \lim_{x \rightarrow 0} \frac{\log(y)y^{-x}}{1} = \log(y).$$

■

Lemma 15 For any $b > 0$ and $c > 0$ and $T_0, T \in \mathbb{N}$, such that $T_0 < T$ and $b\sqrt{T_0} > c$, it holds that

$$\sum_{t=T_0+1}^T \frac{1}{bt^{\frac{3}{2}} - ct} \leq \frac{2}{b\sqrt{T_0} - c}.$$

Proof In the domain $(c^2/b^2, \infty)$, the function $f(t) = \frac{1}{bt^{\frac{3}{2}} - ct}$ is positive, monotonically decreasing and has the antiderivative

$$F(t) = \frac{1}{c} \left(2 \log(b\sqrt{t} - c) - \log(t) \right),$$

which can be verified by taking the respective derivatives. Therefore, we can bound

$$\begin{aligned} \sum_{t=T_0+1}^T f(t) &\leq \int_{T_0}^{\infty} f(t) dt = \lim_{t \rightarrow \infty} F(t) - F(T_0) = \frac{1}{c} \left(2 \log(b) - 2 \log(b\sqrt{T_0} - c) + \log(T_0) \right) \\ &= \frac{2}{c} \log \left(\frac{b\sqrt{T_0}}{b\sqrt{T_0} - c} \right) \leq \frac{2}{c} \left(\frac{b\sqrt{T_0}}{b\sqrt{T_0} - c} - 1 \right) = \frac{2}{b\sqrt{T_0} - c}. \end{aligned}$$

■

Lemma 16 For any $\alpha \in [0, 1]$ and $z \geq 1$, it holds that

$$\frac{1 - z^{-1+\alpha}}{1 - \alpha} \leq (\log z)^\alpha,$$

where for $\alpha = 1$ we consider it as the limit case $\lim_{\alpha \rightarrow 1} \frac{1 - z^{-1+\alpha}}{1 - \alpha} \leq \log(z)$ and for $z = 1$ and $\alpha = 0$ we use the convention $0^0 = 1$.

Proof For $\alpha = 0$ we have

$$\frac{1 - z^{-1+0}}{1 - 0} = 1 - z^{-1} < (\log z)^0 = 1,$$

where for $z = 1$ we use the convention $0^0 = 1$. For $\alpha = 1$ we consider the limit case

$$\lim_{\alpha \rightarrow 1} \frac{1 - z^{-1+\alpha}}{1 - \alpha} = \log(z) = (\log z)^1.$$

It is left to verify the statement for $\alpha \in (0, 1)$. Consider the function

$$f(z) = (\log z)^\alpha - \frac{1 - z^{-1+\alpha}}{1 - \alpha}.$$

The function is continuous for $z \geq 1$, takes the value 0 at $z = 1$ and goes to infinity for $z \rightarrow \infty$. If there is a point where the function is negative, there must also exist an extreme point. Setting the derivative to 0 shows that all extreme points z^* satisfy

$$z^* = \log(z^*) \alpha^{\frac{1}{\alpha-1}}.$$

The function values at the extreme points are therefore lower bounded by

$$f(z^*) \geq \min_{z \geq 1} \left((\log z)^\alpha - \frac{1 - (\log(z) \alpha^{\frac{1}{\alpha-1}})^{-1+\alpha}}{1 - \alpha} \right) = \min_{\tilde{z} \geq 0} \left(\tilde{z}^\alpha - \frac{1 - \tilde{z}^{-1+\alpha} \alpha}{1 - \alpha} \right),$$

where we apply the substitution $\tilde{z} = \log(z)$. The RHS goes to infinity for $\tilde{z} \rightarrow 0$ and $\tilde{z} \rightarrow \infty$, which means that the only extreme point (can be verified by taking the derivative) at $\tilde{z} = 1$ is the minimum. Since $1^\alpha - \frac{1 - 1^{-1+\alpha} \alpha}{1 - \alpha} = 0$, the function f is always positive, which concludes the proof. ■

Appendix C. Support Lemmas for Section 7

We use $v = (v_i)_{i=1,\dots,K}$ to denote a column vector $v \in \mathbb{R}^K$ with elements v_1, \dots, v_K . We use $\text{diag}(v)$ to denote a $K \times K$ matrix with v_1, \dots, v_K on the diagonal and 0 elsewhere. For a positive semidefinite matrix M we use $\|\cdot\|_M = \sqrt{\langle \cdot, M \cdot \rangle}$ to denote the canonical norm with respect to M . We also use the following properties of the potential function.

$$\begin{aligned}\Psi_t(w) &= -\sum_i \frac{w_i^\alpha - \alpha w_i}{\alpha(1-\alpha)\eta_t \xi_i}, \\ \nabla \Psi_t(w) &= -\left(\frac{w_i^{\alpha-1} - 1}{(1-\alpha)\eta_t \xi_i} \right)_{i=1,\dots,K}, \\ \nabla^2 \Psi_t(w) &= \text{diag} \left(\left(\frac{w_i^{\alpha-2}}{\eta_t \xi_i} \right)_{i=1,\dots,K} \right).\end{aligned}$$

For $Y \leq 0$:

$$\begin{aligned}\Psi_t^*(Y) &= \max_w \langle w, Y \rangle + \frac{1}{\eta_t} \sum_i \frac{w_i^\alpha - \alpha w_i}{\alpha(1-\alpha)\xi_i}, \\ \nabla \Psi_t^*(Y) &= \arg \max_w \langle w, Y \rangle + \frac{1}{\eta_t} \sum_i \frac{w_i^\alpha - \alpha w_i}{\alpha(1-\alpha)\xi_i} = \left((-\eta_t(1-\alpha)\xi_i Y_i + 1)^{\frac{1}{\alpha-1}} \right)_{i=1,\dots,K}.\end{aligned}\quad (9)$$

As we have shown in Section 4.3, there exists a Lagrange multiplier ν , such that the algorithm picks the probabilities

$$w_t = \nabla \Phi_t(-\hat{L}_{t-1}) = \nabla \Psi_t^*(-\hat{L}_{t-1} + \nu \mathbf{1}_K).\quad (10)$$

Ψ_t is a Legendre function, which implies that its gradient is invertible and $\nabla \Psi_t^{-1} = \nabla \Psi_t^*$ (Rockafellar, 2015). Furthermore, by the Inverse Function theorem,

$$\nabla^2 \Psi_t^*(\nabla \Psi_t(w)) = \nabla^2 \Psi_t(w)^{-1} = \text{diag}(\eta_t \xi_i w_i^{2-\alpha})_{i=1,\dots,K}.\quad (11)$$

The Bregman divergence associated with a Legendre function f is defined by

$$D_f(x, y) = f(x) - f(y) - \langle \nabla f(y), x - y \rangle.$$

By Taylor's theorem, it satisfies for some $z \in \text{conv}(x, y)$

$$D_f(x, y) \leq \frac{1}{2} \|x - y\|_{\nabla^2 f(z)}^2.\quad (12)$$

C.1 Controlling the *stability* Term

Equation (12) gives a way of bounding the *stability* term. The following lemma allows to control the eigenvalues of the Hessian $\nabla^2 \Psi_t^*$.

Lemma 17 *Let $w \in \Delta^{K-1}$ and $\tilde{w} = \nabla \Psi_t^*(\nabla \Psi_t(w) - \ell)$. If $\eta_t \xi_i \leq \frac{1}{4}$ for all i , then for all $\ell \in \mathbb{R}^K$ with $\ell_i \geq -1$ for all i , it holds that $\tilde{w}_i^{2-\alpha} \leq 2w_i^{2-\alpha}$ for all i .*

Proof Since $\nabla\Psi_t$ is the inverse of $\nabla\Psi_t^*$, we have

$$\begin{aligned} \nabla\Psi_t(w)_i - \nabla\Psi_t(\tilde{w})_i &= \ell_i \geq -1, \\ \frac{w_i^{\alpha-1} - 1}{(1-\alpha)\eta_t\xi_i} - \frac{\tilde{w}_i^{\alpha-1} - 1}{(1-\alpha)\eta_t\xi_i} &\leq 1, \\ \tilde{w}_i^{1-\alpha} &\leq \frac{w_i^{1-\alpha}}{1 - \eta_t\xi_i(1-\alpha)w_i^{1-\alpha}} \leq \frac{w_i^{1-\alpha}}{1 - \eta_t\xi_i(1-\alpha)}, \\ \tilde{w}_i^{2-\alpha} &\leq \frac{w_i^{2-\alpha}}{(1 - \eta_t\xi_i(1-\alpha))^{\frac{2-\alpha}{1-\alpha}}}. \end{aligned}$$

It remains to bound $(1 - \eta_t\xi_i(1-\alpha))^{-\frac{2-\alpha}{1-\alpha}}$. Note that this function is monotonically decreasing in α , which can be verified by confirming that the derivative is negative in $[0, 1]$. Using the fact that $\eta_t\xi_i \leq \frac{1}{4}$, we have

$$(1 - \eta_t\xi_i(1-\alpha))^{-\frac{2-\alpha}{1-\alpha}} \leq (1 - \eta_t\xi_i)^{-2} \leq \frac{4^2}{3^2} \leq 2. \quad \blacksquare$$

For the RV estimators and $\alpha = 1/2$, we provide a tighter bound for the stability term by using the following two lemmas. For $\alpha = 1/2$ and $\xi_i = 1$ we have

$$\begin{aligned} \Psi_t(w) &= -4\eta_t^{-1} \sum_{i=1}^K (w_i^{\frac{1}{2}} - \frac{1}{2}w_i), \\ \nabla\Psi_t(w) &= \left(-2\eta_t^{-1}(w_i^{-\frac{1}{2}} - 1) \right)_{i=1, \dots, K}. \end{aligned} \quad (13)$$

Lemma 18 *The convex conjugate of $\Psi_t(w) = -4\eta_t^{-1} \sum_{i=1}^K (w_i^{\frac{1}{2}} - \frac{1}{2}w_i)$ is*

$$\Psi_t^*(Y) = \begin{cases} \sum_{i=1}^K \frac{2\eta_t^{-1}}{1-\eta_t Y_i/2}, & \text{if } Y_i < 2\eta_t^{-1} \text{ for all } i, \\ \infty, & \text{otherwise.} \end{cases}$$

Proof

$$\begin{aligned} \Psi_t^*(Y) &= \sup_{w \in \mathbb{R}^K} \langle w, Y \rangle - \Psi_t(w) \\ &= \sum_{i=1}^K \sup_{w \in \mathbb{R}} w(Y_i - 2\eta_t^{-1}) + 4\eta_t^{-1}w^{\frac{1}{2}}. \end{aligned}$$

For $Y_i \geq 2\eta_t^{-1}$, the term goes to infinity as $w \rightarrow \infty$. Otherwise, the maximum is obtained by $w = \frac{1}{(1-\frac{1}{2}\eta_t Y_i)^2}$, which concludes the proof. \blacksquare

Using the explicit form of the convex conjugate, we can show a general bound on the stability.

Lemma 19 *Let $\alpha = 1/2$ and $\xi_i = 1$. Then for any x , such that $\min_i \eta_t (\hat{\ell}_{t,i} - x) w_{t,i}^{\frac{1}{2}} \geq -1$, the instantaneous stability satisfies*

$$\langle w_t, \hat{\ell}_t \rangle + \Phi_t(-\hat{L}_t) - \Phi_t(-\hat{L}_{t-1}) \leq \sum_{i=1}^K \frac{\eta_t}{2} w_{t,i}^{\frac{3}{2}} (\hat{\ell}_{t,i} - x)^2 + \frac{\eta_t^2}{2} w_{t,i}^2 |x - \hat{\ell}_{t,i}|_+^3,$$

where $|z|_+ = \max\{z, 0\}$.

Proof By equation (10) and since $\nabla \Psi_t^{-1} = \nabla \Psi_t^*$ and $w_t = \nabla \Phi_t(-\hat{L}_{t-1}) = \nabla(\Psi_t + \mathcal{I}_{\Delta^{K-1}})^*(-\hat{L}_{t-1})$, there exists a Lagrange multiplier ν such that

$$-\hat{L}_{t-1} = \nabla \Psi_t(w_t) - \nu \mathbf{1}_K.$$

Furthermore, $\Phi_t(-L - \nu \mathbf{1}_K) = \Phi_t(-L) - \nu$, since the maximization over w is restricted to the probability simplex. Using these two properties, we have for any $x \in \mathbb{R}$

$$\begin{aligned} \langle w_t, \hat{\ell}_t \rangle + \Phi_t(-\hat{L}_t) - \Phi_t(-\hat{L}_{t-1}) &= \langle w_t, \hat{\ell}_t \rangle + \Phi_t(-\hat{\ell}_t + \nabla \Psi_t(w_t) - \nu \mathbf{1}_K) - \Phi_t(\nabla \Psi_t(w_t) - \nu \mathbf{1}_K) \\ &= \langle w_t, \hat{\ell}_t \rangle + \Phi_t(-\hat{\ell}_t + \nabla \Psi_t(w_t)) - \Phi_t(\nabla \Psi_t(w_t)) \\ &= \langle w_t, \hat{\ell}_t - x \mathbf{1}_K \rangle + \Phi_t(x \mathbf{1}_K - \hat{\ell}_t + \nabla \Psi_t(w_t)) - \Phi_t(\nabla \Psi_t(w_t)) \\ &\leq \langle w_t, \hat{\ell}_t - x \mathbf{1}_K \rangle + \Psi_t^*(x \mathbf{1}_K - \hat{\ell}_t + \nabla \Psi_t(w_t)) - \Psi_t^*(\nabla \Psi_t(w_t)), \end{aligned}$$

where the last line uses $\Phi_t(\nabla \Psi_t(w)) = \Psi_t^*(\nabla \Psi_t(w))$ for any $w \in \Delta^{K-1}$, which holds because the argmax in both terms is w , and the inequality $\Phi_t(L) \leq \Psi_t^*(L)$, which holds because Φ_t is a constrained version of Ψ_t^* .

Using the explicit expression for the convex conjugate in Lemma 18 and the explicit expression for $\nabla \Psi_t(w)$ in equation (13) and assuming that x is in the range defined in the statement of Lemma 19, which ensures that the convex conjugate is bounded, we have

$$\begin{aligned} \langle w_t, \hat{\ell}_t - x \mathbf{1}_K \rangle + \Psi_t^*(x \mathbf{1}_K - \hat{\ell}_t + \nabla \Psi_t(w_t)) - \Psi_t^*(\nabla \Psi_t(w_t)) &= \sum_{i=1}^K w_{t,i} (\hat{\ell}_{t,i} - x) + \frac{2}{\eta_t} \left(w_{t,i}^{-\frac{1}{2}} + \frac{\eta_t}{2} (\hat{\ell}_{t,i} - x) \right)^{-1} - \frac{2}{\eta_t} w_{t,i}^{\frac{1}{2}} \\ &= \sum_{i=1}^K \frac{2}{\eta_t} w_{t,i}^{\frac{1}{2}} \left(\frac{\eta_t}{2} (\hat{\ell}_{t,i} - x) w_{t,i}^{\frac{1}{2}} + \left(1 + \frac{\eta_t}{2} (\hat{\ell}_{t,i} - x) w_{t,i}^{\frac{1}{2}} \right)^{-1} - 1 \right) \\ &= \sum_{i=1}^K \frac{\eta_t}{2} w_{t,i}^{\frac{3}{2}} (\hat{\ell}_{t,i} - x)^2 \left(1 + \frac{\eta_t}{2} (\hat{\ell}_{t,i} - x) w_{t,i}^{\frac{1}{2}} \right)^{-1}. \end{aligned} \quad (14)$$

From $\min_i \eta_t (\hat{\ell}_{t,i} - x) w_{t,i}^{\frac{1}{2}} \geq -1$ it follows that $\forall i : \left(1 + \frac{\eta_t}{2} (\hat{\ell}_{t,i} - x) w_{t,i}^{\frac{1}{2}} \right)^{-1} \leq 2$, so

$$\begin{aligned} \left(1 + \frac{\eta_t}{2} (\hat{\ell}_{t,i} - x) w_{t,i}^{\frac{1}{2}} \right)^{-1} &= 1 - \frac{\eta_t}{2} (\hat{\ell}_{t,i} - x) w_{t,i}^{\frac{1}{2}} \left(1 + \frac{\eta_t}{2} (\hat{\ell}_{t,i} - x) w_{t,i}^{\frac{1}{2}} \right)^{-1} \\ &\leq 1 + \eta_t |x - \hat{\ell}_{t,i}|_+ w_{t,i}^{\frac{1}{2}}. \end{aligned} \quad (15)$$

Combining equations (14) and (15) completes the proof. \blacksquare

Now we have all the tools to prove the main *stability* lemma.

Proof of Lemma 11 We begin by proving the first and the last part of the lemma followed by the second and third.

First part of the lemma. First, we bound the *stability* by 1. By convexity of Φ_t , we have

$$\ell_{t,I_t} + \Phi_t(-\hat{L}_t) - \Phi_t(-\hat{L}_{t-1}) \leq \ell_{t,I_t} - \langle \nabla \Phi_t(-\hat{L}_t), \hat{\ell}_t \rangle \leq 1,$$

where the second inequality uses the non-negativity of the IW estimators.

Recall that $w_t = \nabla \Phi_t(-\hat{L}_{t-1})$ and $\ell_{t,I_t} = \langle w_t, \hat{\ell}_t \rangle$. Furthermore, $\Phi_t(L + x\mathbf{1}_K) = \Phi_t(L) + x$, where $\mathbf{1}_K$ is a vector of K ones, since we take the argmax over probability distributions. Finally, from equation (10) follows the existence of a constant c_t , such that $\nabla \Psi_t(w_t) = -\hat{L}_{t-1} + c_t\mathbf{1}_K$. Hence, for any $x \in \mathbb{R}$

$$\begin{aligned} & \mathbb{E} \left[\ell_{t,I_t} + \Phi_t(-\hat{L}_t) - \Phi_t(-\hat{L}_{t-1}) \right] \\ &= \mathbb{E} \left[\langle w_t, \hat{\ell}_t \rangle + \Phi_t(-\hat{L}_t) - \Phi_t(-\hat{L}_{t-1}) \right] \\ &= \mathbb{E} \left[\langle w_t, \hat{\ell}_t \rangle + \Phi_t(\nabla \Psi_t(w_t) - \hat{\ell}_t) - \Phi_t(\nabla \Psi_t(w_t)) \right] \\ &= \mathbb{E} \left[\langle w_t, \hat{\ell}_t - x\mathbf{1}_K \rangle + \Phi_t(\nabla \Psi_t(w_t) - \hat{\ell}_t + x\mathbf{1}_K) - \Phi_t(\nabla \Psi_t(w_t)) \right] \\ &\leq \mathbb{E} \left[\langle w_t, \hat{\ell}_t - x\mathbf{1}_K \rangle + \Psi_t^*(\nabla \Psi_t(w_t) - \hat{\ell}_t + x\mathbf{1}_K) - \Psi_t^*(\nabla \Psi_t(w_t)) \right] \quad (16) \\ &= \mathbb{E} \left[D_{\Psi_t^*}(\nabla \Psi_t(w_t) - \hat{\ell}_t + x\mathbf{1}_K, \nabla \Psi_t(w_t)) \right] \end{aligned}$$

$$\leq \mathbb{E} \left[\max_{z \in \text{conv}(\nabla \Psi_t(w_t), \nabla \Psi_t(w_t) - \hat{\ell}_t + x\mathbf{1}_K)} \frac{1}{2} \|\hat{\ell}_t - x\mathbf{1}_K\|_{\nabla^2 \Psi_t^*(z)}^2 \right] \quad (17)$$

$$= \mathbb{E} \left[\max_{w \in \text{conv}(w_t, \nabla \Psi_t^*(\nabla \Psi_t(w_t) - \hat{\ell}_t + x\mathbf{1}_K))} \frac{1}{2} \|\hat{\ell}_t - x\mathbf{1}_K\|_{\nabla^2 \Psi_t(w)^{-1}}^2 \right] \quad (18)$$

$$\leq \mathbb{E} \left[\sum_{i=1}^K \max_{w_i \in [w_{t,i}, \nabla \Psi_t^*(\nabla \Psi_t(w_t) - \hat{\ell}_t + x\mathbf{1}_K)]_i} \frac{\eta_t \xi_i}{2} (\hat{\ell}_{t,i} - x)^2 w_i^{2-\alpha} \right],$$

where in equation (16) we have $\Phi_t(x) \leq \Psi_t^*(x)$, because Φ_t is a constrained version of Ψ_t^* , and $\Phi_t(\nabla \Psi_t(w_t)) = \Psi_t^*(\nabla \Psi_t(w_t))$, because $\arg \max_{w \in \mathbb{R}^K} \langle w, \nabla \Psi_t(w_t) \rangle - \Psi(w) = w_t$ is in the probability simplex and the constraint is inactive. Inequality (17) follows by equation (12), and (18) by equation (11).

In order to prove the first part of the Lemma, we set $x = 0$ and observe that $\nabla \Psi_t^*(\nabla \Psi_t(w_t) - \hat{\ell}_t)_i \leq \nabla \Psi_t^*(\nabla \Psi_t(w_t))_i = w_{t,i}$ because of non-negativity of the losses and the fact that $\nabla \Psi_t^*$ is monotonically increasing, see (9). (The observation implies that the highest value of $w_i \in [w_{t,i}, \nabla \Psi_t^*(\nabla \Psi_t(w_t) - \hat{\ell}_t)_i]$ is $w_{t,i}$.) Since the importance weighted losses are 0 for the

arms that were not played, we have

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^K \max_{w_i \in [w_{t,i}, \nabla \Psi_t^*(\nabla \Psi_t(w_t) - \hat{\ell}_t)_i]} \frac{\eta_t \xi_i}{2} \hat{\ell}_{t,i}^2 w_i^{2-\alpha} \right] &= \mathbb{E} \left[\sum_{i=1}^K \frac{\eta_t \xi_i}{2} \hat{\ell}_{t,i}^2 w_{t,i}^{2-\alpha} \right] \\ &= \frac{\eta_t \xi_i}{2} \mathbb{E} \left[\sum_{i=1}^K \frac{\ell_{t,i}^2}{w_{t,i}^2} w_{t,i}^{2-\alpha} \mathbf{1}_t(i) \right] = \frac{\eta_t \xi_i}{2} \mathbb{E} \left[\sum_{i=1}^K \frac{\ell_{t,i}^2}{w_{t,i}^2} w_{t,i}^{3-\alpha} \right] \leq \sum_{i=1}^K \frac{\eta_t \xi_i}{2} \mathbb{E} [w_{t,i}]^{1-\alpha}, \end{aligned}$$

where we use that $\mathbb{E}[\mathbf{1}_t(i) | \ell_1, \dots, \ell_{t-1}, I_1, \dots, I_{t-1}] = w_{t,i}$. The last inequality follows by Jensen's inequality.

Fourth part of the lemma. We set $x = \mathbf{1}_t(j) \ell_{t,j}$. In the calculation below, for the events $I_t \in \{1, \dots, K\} \setminus j$, we have $x = 0$ and use the same derivation as in the previous case. When $I_t = j$, for $i \neq j$ we have $\hat{\ell}_{t,i} - x = -x \geq -1$ and for j we have $\hat{\ell}_{t,j} - x \geq 0$. For $i \neq j$ we use Lemma 17 to bound $(\nabla \Psi_t^*(\nabla \Psi_t(w_t) - \hat{\ell}_t + x \mathbf{1}_K)_i)^{2-\alpha} \leq 2w_{t,i}^{2-\alpha}$ and for j we use $\nabla \Psi_t^*(\nabla \Psi_t(w_t) - \hat{\ell}_t)_j \leq \nabla \Psi_t^*(\nabla \Psi_t(w_t))_j = w_{t,j}$.

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^K \max_{\tilde{w}_i \in [w_{t,i}, \nabla \Psi_t^*(\nabla \Psi_t(w_t) - \hat{\ell}_t + x \mathbf{1}_K)_i]} \frac{\eta_t \xi_i}{2} (\hat{\ell}_{t,i} - x)^2 \tilde{w}_i^{2-\alpha} \right] \\ \leq \sum_{i \neq j} \frac{\eta_t \xi_i}{2} \mathbb{E} [w_{t,i}]^{1-\alpha} + \mathbb{E} \left[\mathbf{1}_t(j) \left(\frac{\eta_t \xi_j}{2} \left(\frac{\ell_{t,j}}{w_{t,j}} - \ell_{t,j} \right)^2 w_{t,j}^{2-\alpha} + \sum_{i \neq j} \frac{\eta_t \xi_i}{2} \ell_{t,j}^2 2w_{t,i}^{2-\alpha} \right) \right] \\ \leq \sum_{i \neq j} \frac{\eta_t \xi_i}{2} \mathbb{E} [w_{t,i}]^{1-\alpha} + \mathbb{E} \left[\frac{\eta_t \xi_j}{2} (1 - w_{t,j})^2 w_{t,j}^{1-\alpha} + \sum_{i \neq j} \eta_t \xi_i w_{t,i}^{2-\alpha} w_{t,j} \right] \\ \leq \sum_{i \neq j} \left(\frac{\eta_t \xi_i}{2} \mathbb{E} [w_{t,i}]^{1-\alpha} + \frac{\eta_t (\xi_j + 2\xi_i)}{2} \mathbb{E} [w_{t,i}] \right), \end{aligned}$$

where in the last step for the middle term we use $(1 - w_{t,j})^2 w_{t,j}^{1-\alpha} \leq 1 - w_{t,j} = \sum_{i \neq j} w_{t,i}$ and for the last term $w_{t,i}^{2-\alpha} \leq 1$.

Second part of the lemma. We set $x = \ell_{t,I_t}$ and first verify that Lemma 19 can be applied. We have for any i :

$$\eta_t (\hat{\ell}_{t,i} - x) w_{t,i}^{\frac{1}{2}} \geq -\eta_t w_{t,i}^{\frac{1}{2}} \geq -\eta_t \geq -1,$$

where the last inequality is by the assumption of the lemma. Since $\mathbb{E}[\ell_{t,I_t}] = \mathbb{E}[\langle w_t, \ell_t \rangle] = \mathbb{E}[\langle w_t, \hat{\ell}_t \rangle]$, by applying Lemma 19 we have

$$\begin{aligned} \mathbb{E} \left[\ell_{t,I_t} + \Phi_t(-\hat{L}_t) - \Phi_t(-\hat{L}_{t-1}) \right] &\leq \mathbb{E} \left[\sum_{i=1}^K \frac{\eta_t}{2} w_{t,i} \frac{3}{2} (\hat{\ell}_{t,i} - \ell_{t,I_t})^2 + \frac{\eta_t^2}{2} w_{t,i}^2 \left| \ell_{t,I_t} - \hat{\ell}_{t,i} \right|_+^3 \right] \\ &\leq \frac{\eta_t}{2} \mathbb{E} \left[w_{t,I_t}^{-\frac{1}{2}} (1 - w_{t,I_t})^2 + \sum_{i \neq I_t} w_{t,i} \frac{3}{2} \right] + \frac{\eta_t^2}{2} \end{aligned} \quad (19)$$

$$\begin{aligned} &= \frac{\eta_t}{2} \mathbb{E} \left[w_{t,I_t}^{-\frac{1}{2}} (1 - w_{t,I_t})^2 + \left(\sum_{i=1}^K w_{t,i} \frac{3}{2} \right) - w_{t,I_t} \frac{3}{2} \right] + \frac{\eta_t^2}{2} \\ &= \frac{\eta_t}{2} \mathbb{E} \left[\sum_{i=1}^K w_{t,i} \frac{1}{2} (1 - w_{t,i})^2 + (1 - w_{t,i}) w_{t,i} \frac{3}{2} \right] + \frac{\eta_t^2}{2} \end{aligned} \quad (20)$$

$$\begin{aligned} &= \frac{\eta_t}{2} \mathbb{E} \left[\sum_{i=1}^K w_{t,i} \frac{1}{2} (1 - w_{t,i}) \right] + \frac{\eta_t^2}{2} \\ &\leq \frac{\eta_t}{2} \sum_{i=1}^K \mathbb{E}[w_{t,i}] \frac{1}{2} (1 - \mathbb{E}[w_{t,i}]) + \frac{\eta_t^2}{2}, \end{aligned} \quad (21)$$

where equation (19) uses non-negativeness of IW estimators and boundedness of the losses in $[0, 1]$, by which $\left| \ell_{t,I_t} - \hat{\ell}_{t,i} \right|_+ \leq 1$, and explicit form of $\hat{\ell}_{t,i}$ for the first term in the summation; equation (20) uses the conditional probability of $I_t = i$, which is $w_{t,i}$; and equation (21) follows by concavity of the function $f(z) = z^{\frac{1}{2}}(1 - z)$ and Jensen's inequality.

Third part of the lemma. We set $x = \ell_{t,I_t}$ and first verify that Lemma 19 can be applied. Recall that $\mathbb{B}_t(i) := \frac{1}{2} \mathbf{1}(w_{t,i} \geq \eta_t^2)$. For $i \neq I_t$ we have $\hat{\ell}_{t,i} = \mathbb{B}_t(i)$ and

$$\eta_t (\hat{\ell}_{t,i} - x) w_{t,i}^{\frac{1}{2}} = \eta_t (\mathbb{B}_t(i) - \ell_{t,I_t}) w_{t,i}^{\frac{1}{2}} \geq -\eta_t \geq -1,$$

while for I_t we have $\hat{\ell}_{t,I_t} = \frac{\ell_{t,I_t} - \mathbb{B}_t(I_t)}{w_{t,I_t}} + \mathbb{B}_t(I_t)$ and

$$\eta_t (\hat{\ell}_{t,I_t} - \ell_{t,I_t}) w_{t,I_t}^{\frac{1}{2}} = \eta_t (\ell_{t,I_t} - \mathbb{B}_t(I_t)) \left(\frac{1}{w_{t,I_t}} - 1 \right) w_{t,I_t}^{\frac{1}{2}} \geq -\eta_t \mathbb{B}_t(I_t) w_{t,I_t}^{-\frac{1}{2}} \geq -1.$$

Since $\mathbb{E}[\ell_{t,I_t}] = \mathbb{E}[\langle w_t, \ell_t \rangle] = \mathbb{E}[\langle w_t, \hat{\ell}_t \rangle]$, applying Lemma 19 we have

$$\begin{aligned} \mathbb{E} \left[\ell_{t,I_t} + \Phi_t(-\hat{L}_t) - \Phi_t(-\hat{L}_{t-1}) \right] &\leq \mathbb{E} \left[\sum_{i=1}^K \frac{\eta_t}{2} w_{t,i} \frac{3}{2} (\hat{\ell}_{t,i} - \ell_{t,I_t})^2 + \frac{\eta_t^2}{2} w_{t,i}^2 \left| \ell_{t,I_t} - \hat{\ell}_{t,i} \right|_+^3 \right] \\ &\leq \mathbb{E} \left[\sum_{i=1}^K \frac{\eta_t}{2} w_{t,i} \frac{3}{2} (\hat{\ell}_{t,i} - \ell_{t,I_t})^2 + \frac{\eta_t^2}{2} w_{t,i}^2 \left| \hat{\ell}_{t,i} - \ell_{t,I_t} \right|_+^3 \right]. \end{aligned} \quad (22)$$

For any $\ell \in [0, 1]$ and any i , we have $|\ell - \mathbb{B}_t(i)| \leq |1 - \mathbb{B}_t(i)|$. For I_t , we have $|\hat{\ell}_{t,I_t} - \ell_{t,I_t}| = |\ell_{t,I_t} - \mathbb{B}_t(I_t)| \frac{1-w_{t,I_t}}{w_{t,I_t}} \leq |1 - \mathbb{B}_t(I_t)| \frac{1-w_{t,I_t}}{w_{t,I_t}}$, while for $i \neq I_t$ we have $|\hat{\ell}_{t,i} - \ell_{t,I_t}| \leq |1 - \mathbb{B}_t(i)|$. Let $\bar{\mathbb{B}}_t(i) = \frac{1}{2} \mathbb{1}(w_{t,i} < \eta_t^2) = \frac{1}{2} - \mathbb{B}_t(i)$, then $|1 - \mathbb{B}_t(i)| = |\frac{1}{2} + \bar{\mathbb{B}}_t(i)|$. We have $|\frac{1}{2} + \bar{\mathbb{B}}_t(i)|^2 = \frac{1}{4} + \frac{3}{2}\bar{\mathbb{B}}_t(i)$ and $|\frac{1}{2} + \bar{\mathbb{B}}_t(i)|^3 \leq 1$. Thus, for I_t we have

$$\begin{aligned} (\hat{\ell}_{t,I_t} - \ell_{t,I_t})^2 &\leq \left(\frac{1}{4} + \frac{3}{2}\bar{\mathbb{B}}_t(I_t) \right) \frac{(1-w_{t,I_t})^2}{w_{t,I_t}^2}, \\ |\hat{\ell}_{t,I_t} - \ell_{t,I_t}|^3 &\leq \frac{(1-w_{t,I_t})^3}{w_{t,I_t}^3} \end{aligned}$$

and for $i \neq I_t$ we have

$$\begin{aligned} (\hat{\ell}_{t,i} - \ell_{t,I_t})^2 &\leq \left(\frac{1}{4} + \frac{3}{2}\bar{\mathbb{B}}_t(I_t) \right), \\ |\hat{\ell}_{t,i} - \ell_{t,I_t}|^3 &\leq 1. \end{aligned}$$

Plugging this into equation (22) leads to

$$\begin{aligned} &\mathbb{E} \left[\sum_{i=1}^K \frac{\eta_t}{2} w_{t,i}^{\frac{3}{2}} (\hat{\ell}_{t,i} - \ell_{t,I_t})^2 + \frac{\eta_t^2}{2} w_{t,i}^2 |\hat{\ell}_{t,i} - \ell_{t,I_t}|^3 \right] \\ &\leq \mathbb{E} \left[\left(\frac{1}{4} + \frac{3}{2}\bar{\mathbb{B}}_t(I_t) \right) \frac{\eta_t}{2} w_{t,I_t}^{-\frac{1}{2}} (1-w_{t,I_t})^2 + \frac{\eta_t^2}{2} w_{t,I_t}^{-1} (1-w_{t,I_t})^3 \right. \\ &\quad \left. + \sum_{i \neq I_t} \left(\frac{1}{4} + \frac{3}{2}\bar{\mathbb{B}}_t(i) \right) \frac{\eta_t}{2} w_{t,i}^{\frac{3}{2}} + \frac{\eta_t^2}{2} w_{t,i}^2 \right] \\ &= \mathbb{E} \left[\sum_{i=1}^K \left(\left(\frac{1}{4} + \frac{3}{2}\bar{\mathbb{B}}_t(i) \right) \frac{\eta_t}{2} w_{t,i}^{\frac{1}{2}} (1-w_{t,i})^2 + \frac{\eta_t^2}{2} (1-w_{t,i})^3 \right. \right. \\ &\quad \left. \left. + (1-w_{t,i}) \left(\frac{1}{4} + \frac{3}{2}\bar{\mathbb{B}}_t(i) \right) \frac{\eta_t}{2} w_{t,i}^{\frac{3}{2}} + (1-w_{t,i}) \frac{\eta_t^2}{2} w_{t,i}^2 \right) \right] \\ &\leq \mathbb{E} \left[\sum_{i=1}^K \left(\frac{\eta_t}{8} w_{t,i}^{\frac{1}{2}} (1-w_{t,i}) + \frac{3}{2}\bar{\mathbb{B}}_t(i) \frac{\eta_t}{2} w_{t,i}^{\frac{1}{2}} + \frac{\eta_t^2}{2} \right) \right] \\ &\leq \frac{7\eta_t^2}{8} K + \mathbb{E} \left[\sum_{i=1}^K \frac{\eta_t}{8} w_{t,i}^{\frac{1}{2}} (1-w_{t,i}) \right] \\ &\leq \frac{7\eta_t^2}{8} K + \sum_{i=1}^K \frac{\eta_t}{8} \mathbb{E}[w_{t,i}]^{\frac{1}{2}} (1 - \mathbb{E}[w_{t,i}]), \end{aligned}$$

where in the penultimate step we use $\bar{\mathbb{B}}_t(i) w_{t,i}^{\frac{1}{2}} = \frac{1}{2} \mathbb{1}(w_{t,i} < \eta_t^2) w_{t,i}^{\frac{1}{2}} \leq \frac{1}{2} \eta_t$ and the last step follows by concavity of $f(z) = z^{\frac{1}{2}}(1-z)$ and Jensen's inequality. \blacksquare

C.2 Controlling the penalty Term

We begin with a standard lemma to simplify the *penalty* term.

Lemma 20 *For any $\alpha \in [0, 1]$, any positive learning rate, and any fixed $v, u \in \Delta^{K-1}$, the penalty term satisfies*

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T \left(\Phi_t(-\hat{L}_{t-1}) - \Phi_t(-\hat{L}_t) - \ell_{t, i_T^*} \right) \right] \\ & \leq \mathbb{E} \left[\frac{\Psi(v) - \Psi(w_1)}{\eta_1} + \sum_{t=2}^T (\eta_t^{-1} - \eta_{t-1}^{-1}) (\Psi(v) - \Psi(w_t)) + \frac{\Psi(u) - \Psi(v)}{\eta_T} \right] + \langle u - \mathbf{e}_{i_T^*}, L_T \rangle. \end{aligned}$$

Proof First, note that all the terms involving $\Psi(v)$ in the lemma sum up to 0. Then, recall that w_t is defined as $\arg \max_{w \in \Delta^{K-1}} \left\{ \langle w, -\hat{L}_{t-1} \rangle - \frac{\Psi(w)}{\eta_t} \right\}$. Therefore,

$$\Phi_t(-\hat{L}_{t-1}) = - \left\langle w_t, \hat{L}_{t-1} \right\rangle - \frac{\Psi(w_t)}{\eta_t}.$$

Furthermore, by definition of the potential function, for any $\tilde{w} \in \Delta^{K-1}$ it holds that:

$$-\Phi_t(-\hat{L}_t) = - \max_{w \in \Delta^{K-1}} \left\{ \langle w, -\hat{L}_t \rangle - \frac{\Psi(w)}{\eta_t} \right\} \leq \langle \tilde{w}, \hat{L}_t \rangle + \frac{\Psi(\tilde{w})}{\eta_t}.$$

Setting \tilde{w} to w_{t+1} for $t < T$ and to u for $t = T$, and using $\hat{L}_0 = \mathbf{0}_K$, where $\mathbf{0}_K$ is a vector of K zeros, the sum of potential differences can be bounded as follows:

$$\begin{aligned} & \sum_{t=1}^T \left(\Phi_t(-\hat{L}_{t-1}) - \Phi_t(-\hat{L}_t) \right) \\ & \leq \sum_{t=1}^T \left(- \left\langle w_t, \hat{L}_{t-1} \right\rangle - \frac{\Psi(w_t)}{\eta_t} \right) + \sum_{t=1}^{T-1} \left(\left\langle w_{t+1}, \hat{L}_t \right\rangle + \frac{\Psi(w_{t+1})}{\eta_t} \right) + \langle u, \hat{L}_T \rangle + \frac{\Psi(u)}{\eta_T} \\ & = - \frac{\Psi(w_1)}{\eta_1} - \sum_{t=2}^T (\eta_t^{-1} - \eta_{t-1}^{-1}) \Psi(w_t) + \frac{\Psi(u)}{\eta_T} + \langle u, \hat{L}_T \rangle. \end{aligned}$$

The proof is finalized by taking the expectation and subtracting the optimal loss. Due to unbiasedness of the loss estimators, for a fixed u we have $\mathbb{E}[\langle u, \hat{L}_T \rangle] = \langle u, L_T \rangle$. \blacksquare

Proof of Lemma 12 The proof of both parts of the lemma is based on Lemma 20.

Part 1: We set $v = w_1$. Since $w_1 = \arg \max_{w \in \Delta^{K-1}} -\Psi_1(w) = \arg \max_{w \in \Delta^{K-1}} -\Psi(w)$, we have $\Psi(w_1) - \Psi(w_t) \leq 0$ for any t . Since the learning rate is non-increasing, the terms

$(\eta_t^{-1} - \eta_{t-1}^{-1})$ are all positive, so

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T \left(\Phi_t(-\hat{L}_{t-1}) - \Phi_t(-\hat{L}_t) - \ell_{t,i_T^*} \right) \right] \\ & \leq \mathbb{E} \left[\frac{\Psi(w_1) - \Psi(w_1)}{\eta_1} + \sum_{t=2}^T (\eta_t^{-1} - \eta_{t-1}^{-1}) (\Psi(w_1) - \Psi(w_t)) + \frac{\Psi(u) - \Psi(w_1)}{\eta_T} \right] \\ & \qquad \qquad \qquad + \left\langle u - \mathbf{e}_{i_T^*}, L_T \right\rangle \\ & \leq \mathbb{E} \left[\frac{\Psi(u) - \Psi(w_1)}{\eta_T} \right] + \left\langle u - \mathbf{e}_{i_T^*}, L_T \right\rangle. \end{aligned}$$

Following the trick of Agarwal et al. (2017), we set $u_{i_T^*} = 1 - T^{-1}$ and $u_i = \frac{T^{-1}}{K-1}$ for $i \neq i_T^*$. The losses are bounded in $[0, T]$, so this choice of u implies $\langle u - \mathbf{e}_{i_T^*}, L_T \rangle \leq 1$. Since we assume that the regularizer is symmetric, the explicit form of w_1 is $w_{1,i} = K^{-1}$ and

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T \left(\Phi_t(-\hat{L}_{t-1}) - \Phi_t(-\hat{L}_t) - \ell_{t,i_T^*} \right) \right] \\ & \leq \frac{K^{1-\alpha}}{\alpha(1-\alpha)\eta_T} - \frac{(K-1)^{1-\alpha}T^{-\alpha} + (1-T^{-1})^\alpha}{\alpha(1-\alpha)\eta_T} + 1. \end{aligned}$$

It remains to bound $K^{1-\alpha} - (K-1)^{1-\alpha}T^{-\alpha} - (1-T^{-1})^\alpha$. Since x^α and $x^{1-\alpha}$ are concave functions, by Taylor's expansion around $X-1$ we have $X^{1-\alpha} \leq (X-1)^{1-\alpha} + (1-\alpha)(X-1)^{-\alpha}$ and $X^\alpha \leq (X-1)^\alpha + \alpha(X-1)^{\alpha-1}$ for any $X > 1$, thus

$$\begin{aligned} K^{1-\alpha} + T^\alpha & \leq (K-1)^{1-\alpha} + (1-\alpha)(K-1)^{-\alpha} + (T-1)^\alpha + \alpha(T-1)^{\alpha-1} \\ & \leq (K-1)^{1-\alpha} + (T-1)^\alpha + 1, \end{aligned}$$

where the last line uses $(T-1)^{\alpha-1}, (K-1)^{-\alpha} \leq 1$. Therefore,

$$\begin{aligned} K^{1-\alpha} - (K-1)^{1-\alpha}T^{-\alpha} - (1-T^{-1})^\alpha & = K^{1-\alpha} + T^{-\alpha}(-K^{1-\alpha} - (T-1)^\alpha) \\ & \leq K^{1-\alpha} + T^{-\alpha}(-K^{1-\alpha} - T^\alpha + 1) \\ & = (K^{1-\alpha} - 1)(1 - T^{-\alpha}). \end{aligned}$$

Part 2: Set

$$\begin{aligned} \tilde{w} & = \arg \max_{w \in \Delta^{K-1}} -\Psi \left((1 - T^{-x})\mathbf{e}_{i_T^*} + T^{-x}w \right), \\ v & = u = (1 - T^{-x})\mathbf{e}_{i_T^*} + T^{-x}\tilde{w}, \\ v_t & = (1 - T^{-x})\mathbf{e}_{i_T^*} + T^{-x}w_t. \end{aligned}$$

By definition, $\Psi(v) \leq \Psi(v_t)$ for all t . So

$$\Psi(v) - \Psi(w_t) \leq \Psi(v_t) - \Psi(w_t) \leq \sum_{i \neq i_T^*} \frac{(w_{t,i}^\alpha - \alpha w_{t,i})(1 - T^{-\alpha x})}{\alpha(1-\alpha)\xi_i}.$$

In the last inequality, we have used the fact that the contribution of the optimal arm i_T^* is non-positive, since $v_{t,i_T^*} \geq w_{t,i_T^*}$ and $w^\alpha - \alpha w$ are monotonically increasing in w over $[0, 1]$. The choice of u ensures that $\langle u - \mathbf{e}_{i_T^*}, L_T \rangle \leq T^{1-x}$. Starting again with Lemma 20, we have:

$$\begin{aligned}
 & \mathbb{E} \left[\sum_{t=1}^T \left(\Phi_t(-\hat{L}_{t-1}) - \Phi_t(-\hat{L}_t) - \ell_{t,i_T^*} \right) \right] \\
 & \leq \mathbb{E} \left[\frac{\Psi(v) - \Psi(w_1)}{\eta_1} + \sum_{t=2}^T (\eta_t^{-1} - \eta_{t-1}^{-1}) (\Psi(v) - \Psi(w_t)) + \frac{\Psi(u) - \Psi(v)}{\eta_T} \right] \\
 & \qquad \qquad \qquad + \langle u - \mathbf{e}_{i_T^*}, L_T \rangle \\
 & \leq \frac{1 - T^{-\alpha x}}{\alpha} \sum_{i \neq i_T^*} \mathbb{E} \left[\frac{w_{1,i}^\alpha - \alpha w_{1,i}}{(1-\alpha)\eta_1 \xi_i} + \sum_{t=2}^T (\eta_t^{-1} - \eta_{t-1}^{-1}) \frac{w_{t,i}^\alpha - \alpha w_{t,i}}{(1-\alpha)\xi_i} \right] + T^{1-x} \\
 & \leq \frac{1 - T^{-\alpha x}}{\alpha} \sum_{i \neq i_T^*} \left(\frac{\mathbb{E}[w_{1,i}^\alpha] - \alpha \mathbb{E}[w_{1,i}]}{(1-\alpha)\eta_1 \xi_i} + \sum_{t=2}^T (\eta_t^{-1} - \eta_{t-1}^{-1}) \frac{\mathbb{E}[w_{t,i}^\alpha] - \alpha \mathbb{E}[w_{t,i}]}{(1-\alpha)\xi_i} \right) + T^{1-x}.
 \end{aligned}$$

■

Appendix D. Proof of Theorem 4

We follow the same strategy as outlined in Section 7. In order to cover the limit cases $\alpha \in \{0, 1\}$, the proof is significantly more technical than the proof of Theorem 1.

Proof of Theorem 4 Recall that the learning rate is $\eta_t = \frac{16^\alpha}{4} \frac{1 - \bar{t}^{-1+\alpha}}{(1-\alpha)\bar{t}^\alpha}$, where $\bar{t} = \max\{e, t\}$, and regularization parameters are $\xi_i = \Delta_i^{1-2\alpha}$ for $i \neq i^*$ and $\xi_{i^*} = \Delta_{\min}^{1-2\alpha}$.

Bounding the *stability* term We start by bounding the *stability* term. For $t \leq T_0$ we use the first part of Lemma 11 and otherwise the second with $j = i^*$. The value of T_0 is chosen so that $\eta_{T_0} \xi_i \leq \frac{1}{4}$.

$$\begin{aligned}
 \text{stability} &= \mathbb{E} \left[\sum_{t=1}^T \ell_{t,I_t} + \Phi_t(-\hat{L}_t) - \Phi_t(-\hat{L}_{t-1}) \right] \\
 &\leq \underbrace{\sum_{t=1}^T \sum_{i \neq i^*} \frac{\eta_t \xi_i \mathbb{E}[w_{t,i}]^{1-\alpha}}{2}}_{\text{concave}} + \underbrace{\sum_{t=1}^{T_0} \frac{\eta_t \xi_{i^*} \mathbb{E}[w_{t,i^*}]^{1-\alpha}}{2}}_{\text{constant}} + \underbrace{\sum_{t=T_0+1}^T \sum_{i \neq i^*} \frac{\eta_t (\xi_i + 2\xi_{i^*})}{2} \mathbb{E}[w_{t,i}]}_{\text{linear}}.
 \end{aligned}$$

Bounding the concave part Since $w^{1-\alpha}$ is a concave function of w , it can be upper bounded by the first order Taylor's approximation. For any w^* :

$$\begin{aligned}
 w_{t,i}^{1-\alpha} &\leq w^{*1-\alpha} + (1-\alpha)w^{*-\alpha}(w_{t,i} - w^*) \\
 &= \alpha w^{*1-\alpha} + (1-\alpha)w^{*-\alpha}w_{t,i}.
 \end{aligned}$$

Taking $w^* = \frac{16}{\Delta_i^2 t}$ (with $\eta_t = \frac{16^\alpha}{4} \frac{1-\bar{t}^{-1+\alpha}}{(1-\alpha)t^\alpha}$, $\xi_i = \Delta_i^{1-2\alpha}$):

$$\begin{aligned}
 \sum_{t=1}^T \frac{\eta_t \xi_i}{2} \mathbb{E}[w_{t,i}]^{1-\alpha} &\leq \sum_{t=1}^T \frac{\Delta_i^{1-2\alpha} \frac{16^\alpha}{4} \frac{1-\bar{t}^{-1+\alpha}}{(1-\alpha)t^\alpha}}{2} \left(\alpha \left(\frac{16}{\Delta_i^2 t} \right)^{1-\alpha} + (1-\alpha) \left(\frac{16}{\Delta_i^2 t} \right)^{-\alpha} \mathbb{E}[w_{t,i}] \right) \\
 &= \sum_{t=1}^T \frac{1-\bar{t}^{-1+\alpha}}{1-\alpha} \left(\frac{2\alpha}{\Delta_i t} + \frac{1-\alpha}{8} \Delta_i \mathbb{E}[w_{t,i}] \right) \\
 &\leq \frac{1-T^{-1+\alpha}}{1-\alpha} \frac{2(\log(T)+1)}{\Delta_i} + \sum_{t=1}^T \frac{\Delta_i \mathbb{E}[w_{t,i}]}{8}. \tag{23}
 \end{aligned}$$

Finally, we bound the leading factor of the log term with Lemma 14:

$$\frac{1-T^{-1+\alpha}}{1-\alpha} \leq \min\left\{ \frac{1}{1-\alpha}, \log(T) \right\}.$$

Bounding the linear part We first show that all $t > T_0 = \frac{16}{\Delta_{\min}^2} \log^2\left(\frac{16}{\Delta_{\min}^2}\right)$ satisfy $\eta_t \xi_i \leq \frac{\Delta_i}{4}$.

$$\begin{aligned}
 \eta_t \xi_i = \Delta_i^{1-2\alpha} \frac{16^\alpha}{4} \frac{1-\bar{t}^{-1+\alpha}}{(1-\alpha)t^\alpha} &< \frac{\Delta_i}{4} \left(\frac{16}{\Delta_{\min}^2 T_0} \right)^\alpha \frac{1-\bar{T}_0^{-1+\alpha}}{1-\alpha} \\
 &\leq \frac{\Delta_i}{4} \frac{1 - \left(\frac{16}{\Delta_{\min}^2} \log^2\left(\frac{16}{\Delta_{\min}^2}\right)\right)^{-1+\alpha}}{(1-\alpha) \left(\log\left(\frac{16}{\Delta_{\min}^2}\right)\right)^{2\alpha}}.
 \end{aligned}$$

It remains to show that $\frac{1 - \left(\frac{16}{\Delta_{\min}^2} \log^2\left(\frac{16}{\Delta_{\min}^2}\right)\right)^{-1+\alpha}}{(1-\alpha) \left(\log\left(\frac{16}{\Delta_{\min}^2}\right)\right)^{2\alpha}} \leq 1$. By Lemma 16 we have

$$\begin{aligned}
 \frac{1 - \left(\frac{16}{\Delta_{\min}^2} \log^2\left(\frac{16}{\Delta_{\min}^2}\right)\right)^{-1+\alpha}}{1-\alpha} &\leq \left(\log \left(\frac{16}{\Delta_{\min}^2} \log^2\left(\frac{16}{\Delta_{\min}^2}\right) \right) \right)^\alpha \\
 &\leq \left(2 \log\left(\frac{16}{\Delta_{\min}^2}\right) \right)^\alpha \leq \left(\log\left(\frac{16}{\Delta_{\min}^2}\right) \right)^{2\alpha},
 \end{aligned}$$

which concludes the proof. Therefore,

$$\sum_{i \neq i^*} \sum_{t=T_0+1}^T \frac{\eta_t (\xi_i + 2\xi_{i^*})}{2} \mathbb{E}[w_{t,i}] \leq \sum_{i \neq i^*} \sum_{t=1}^T \frac{\Delta_i + 2\Delta_{\min}}{8} \mathbb{E}[w_{t,i}] \leq \sum_{i \neq i^*} \sum_{t=1}^T \frac{3\Delta_i \mathbb{E}[w_{t,i}]}{8}. \tag{24}$$

Bounding the constant part Recall $T_0 = \frac{16}{\Delta_{\min}^2} \log^2\left(\frac{16}{\Delta_{\min}^2}\right) \geq 16$. We can use the estimation $\sum_{t=1}^{T_0} t^{-\alpha} \leq 1 + \int_1^{T_0} t^{-\alpha} dt = 1 + \frac{T_0^{1-\alpha}-1}{1-\alpha} \leq 2 \frac{T_0^{1-\alpha}-1}{1-\alpha}$:

$$\begin{aligned}
 \sum_{t=1}^{T_0} \frac{\eta_t \xi_{i^*}}{2} &\leq \sum_{t=1}^{T_0} \frac{\Delta_{\min}^{1-2\alpha} 16^\alpha (1 - T_0^{-1+\alpha})}{8(1-\alpha)t^\alpha} = \frac{\Delta_{\min}^{1-2\alpha} 16^\alpha (1 - T_0^{-1+\alpha})}{8(1-\alpha)} \sum_{t=1}^{T_0} \frac{1}{t^\alpha} \\
 &\leq \frac{\Delta_{\min}^{1-2\alpha} 16^\alpha (1 - T_0^{-1+\alpha})(T_0^{1-\alpha} - 1)}{4(1-\alpha)^2} \\
 &= \frac{16^\alpha T_0^{1-\alpha} (1 - T_0^{-1+\alpha})^2}{4\Delta_{\min}^{2\alpha-1} (1-\alpha)^2} = \frac{4 \log^{2-2\alpha}\left(\frac{16}{\Delta_{\min}^2}\right)}{\Delta_{\min}} \left(\frac{1 - T_0^{-1+\alpha}}{1-\alpha}\right)^2 \\
 &\leq \frac{4 \log^2\left(\frac{16}{\Delta_{\min}^2}\right) \log^2(T_0)}{\Delta_{\min}} \leq \frac{4 \log^4(T_0)}{\Delta_{\min}}, \tag{25}
 \end{aligned}$$

where the last line uses Lemma 14. Combining (23), (24), and (25) we obtain:

$$\text{stability} \leq \sum_{i \neq i^*} \left(\min \left\{ \frac{1}{1-\alpha}, \log(T) \right\} \frac{2(\log(T) + 1)}{\Delta_i} + \sum_{t=1}^T \frac{\Delta_i \mathbb{E}[w_{t,i}]}{2} \right) + \frac{4 \log^4(T_0)}{\Delta_{\min}}. \tag{26}$$

Bounding the penalty term For the *penalty* term, we start with the second part of Lemma 12 with $x = 1$. We have

$$\begin{aligned}
 \text{penalty} &= \mathbb{E} \left[\sum_{t=1}^T -\Phi_t(-\hat{L}_t) + \Phi_t(-\hat{L}_{t-1}) - \ell_{t,i^*} \right] \\
 &\leq \frac{1 - T^{-\alpha}}{\alpha} \sum_{i \neq i^*} \left(\frac{\mathbb{E}[w_{1,i}]^\alpha - \alpha \mathbb{E}[w_{1,i}]}{\eta_1 \xi_i (1-\alpha)} + \sum_{t=2}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \frac{\mathbb{E}[w_{t,i}]^\alpha - \alpha \mathbb{E}[w_{t,i}]}{(1-\alpha)\xi_i} \right) + 1 \\
 &\leq \sum_{i \neq i^*} \left(\underbrace{\frac{\mathbb{E}[w_{1,i}]^\alpha - \alpha \mathbb{E}[w_{1,i}]}{\eta_1 \xi_i (1-\alpha)}}_{\text{constant}} \log(T) + \underbrace{\sum_{t=2}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \frac{\mathbb{E}[w_{t,i}]^\alpha - \alpha \mathbb{E}[w_{t,i}]}{(1-\alpha)\alpha \xi_i}}_{\text{concave}} \right) + 1,
 \end{aligned}$$

where in the first term we use $\frac{1-T^{-\alpha}}{\alpha} \leq \log(T)$ and in the second $\frac{1-T^{-\alpha}}{\alpha} \leq \frac{1}{\alpha}$, both bounds following from Lemma 14.

Bounding the concave term Since w^α is a concave function of w it can be upper bounded by the first order Taylor's approximation:

$$w_{t,i}^\alpha \leq w^{*\alpha} + \alpha w^{*\alpha-1} (w_{t,i} - w^*) = (1-\alpha)w^{*\alpha} + \alpha w^{*\alpha-1} w_{t,i}.$$

Taking $w^* = \frac{16}{\Delta_i^2 t}$ (with $\eta_t = \frac{16^\alpha}{4} \frac{1-\bar{t}^{-1+\alpha}}{(1-\alpha)t^\alpha}$ and $\xi_i = \Delta_i^{1-2\alpha}$):

$$\begin{aligned}
 & \sum_{t=1}^{T-1} \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) \frac{\mathbb{E}[w_{t+1,i}]^\alpha - \alpha \mathbb{E}[w_{t+1,i}]}{(1-\alpha)\alpha\xi_i} & (27) \\
 & \leq \sum_{t=1}^{T-1} \left(\frac{(t+1)^\alpha}{1-(t+1)^{-1+\alpha}} - \frac{t^\alpha}{1-t^{-1+\alpha}} \right) \\
 & \quad \cdot \frac{4\Delta_i^{2\alpha-1}}{\alpha 16^\alpha} \left((1-\alpha) \left(\frac{16}{\Delta_i^2 t} \right)^\alpha + \alpha \left(\frac{16}{\Delta_i^2 t} \right)^{\alpha-1} \mathbb{E}[w_{t+1,i}] - \alpha \mathbb{E}[w_{t+1,i}] \right) \\
 & \leq \sum_{t=1}^{T-1} \left(\frac{(t+1)^\alpha - t^\alpha}{1-t^{-1+\alpha}} \right) \left(\frac{4(1-\alpha)}{\alpha\Delta_i t^\alpha} + \frac{\Delta_i \mathbb{E}[w_{t+1,i}]}{4t^{\alpha-1}} \left(1 - \left(\frac{16}{\Delta_i^2 t} \right)^{1-\alpha} \right) \right) \\
 & \text{(by Taylor's approximation } (t+1)^\alpha \leq t^\alpha + \alpha t^{\alpha-1} \text{ and also use } \frac{16}{\Delta_i^2} > 1) \\
 & \leq \sum_{t=1}^{T-1} \left(\frac{\alpha t^{\alpha-1}}{1-t^{-1+\alpha}} \right) \left(\frac{4(1-\alpha)}{\alpha\Delta_i t^\alpha} + \frac{\Delta_i \mathbb{E}[w_{t+1,i}]}{4t^{\alpha-1}} \left(1 - \left(\frac{1}{t} \right)^{1-\alpha} \right) \right) \\
 & \leq \sum_{t=1}^{T-1} \left(\frac{1-\alpha}{1-e^{-1+\alpha}} \frac{4}{\Delta_i t} + \frac{\Delta_i \mathbb{E}[w_{t+1,i}]}{4} \frac{1-t^{-1+\alpha}}{1-t^{-1+\alpha}} \right) \\
 & \leq \frac{1-\alpha}{1-e^{-1+\alpha}} \frac{4(\log(T)+1)}{\Delta_i} + \sum_{t=1}^T \frac{\Delta_i \mathbb{E}[w_{t,i}]}{4} \\
 & \leq \frac{8(\log(T)+1)}{\Delta_i} + \sum_{t=1}^T \frac{\Delta_i \mathbb{E}[w_{t,i}]}{4}. & (28)
 \end{aligned}$$

The last step follows by the leading factor $\frac{1-\alpha}{1-e^{-1+\alpha}}$ being bounded by 2.

Bounding the constant term Since $\frac{w^\alpha - \alpha w}{1-\alpha}$ is monotonically decreasing in α and ξ_i is monotonically increasing in α , we have

$$\frac{\mathbb{E}[w_{1,i}]^\alpha - \alpha \mathbb{E}[w_{1,i}]}{(1-\alpha)\eta_1 \xi_i} \leq \frac{4}{\Delta_i(1-e^{-1})} \leq \frac{8}{\Delta_i}. \quad (29)$$

Combining (28) and (29) we obtain

$$\text{penalty} \leq \sum_{i \neq i^*} \left(\frac{16(\log(T)+1)}{\Delta_i} + \sum_{t=1}^T \frac{\Delta_i \mathbb{E}[w_{t,i}]}{4} \right) + 1. \quad (30)$$

Finishing the proof Finally, we combine (26), (30), and rearrange the terms to get

$$\begin{aligned} \overline{Reg}_T &\leq \sum_{t=1}^T \sum_{i \neq i^*} \frac{3\Delta_i \mathbb{E}[w_{t,i}]}{4} + \sum_{i \neq i^*} \left(\frac{\left(2 \min\{\frac{1}{1-\alpha}, \log(T)\} + 16\right) (\log(T) + 1)}{\Delta_i} \right) \\ &\qquad\qquad\qquad + \frac{4 \log^4(T_0)}{\Delta_{\min}} + 1. \\ &= \frac{3\overline{Reg}_T}{4} + \sum_{i \neq i^*} \left(\frac{\left(2 \min\{\frac{1}{1-\alpha}, \log(T)\} + 16\right) (\log(T) + 1)}{\Delta_i} \right) + \frac{4 \log^4(T_0)}{\Delta_{\min}} + 1. \end{aligned}$$

Rearranging and multiplying by 4 finishes the proof. ■

References

- Yasin Abbasi-Yadkori, Peter Bartlett, Victor Gabillon, Alan Malek, and Michal Valko. Best of both worlds: Stochastic & adversarial best-arm identification. In *Proceedings of the International Conference on Computational Learning Theory (COLT)*, 2018.
- Jacob Abernethy, Chansoo Lee, Abhinav Sinha, and Ambuj Tewari. Online linear optimization via smoothing. In *Proceedings of the International Conference on Computational Learning Theory (COLT)*, 2014.
- Jacob D Abernethy, Chansoo Lee, and Ambuj Tewari. Fighting bandits with a new kind of smoothness. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.
- Alekh Agarwal, Haipeng Luo, Behnam Neyshabur, and Robert E Schapire. Corraling a band of bandit algorithms. In *Proceedings of the International Conference on Computational Learning Theory (COLT)*, 2017.
- Nir Ailon, Zohar Karnin, and Thorsten Joachims. Reducing dueling bandits to cardinal bandits. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2014.
- Jean-Yves Audibert and Sébastien Bubeck. Minimax policies for adversarial and stochastic bandits. In *Proceedings of the International Conference on Computational Learning Theory (COLT)*, 2009.
- Jean-Yves Audibert and Sébastien Bubeck. Regret bounds and minimax policies under partial monitoring. *Journal of Machine Learning Research*, 11, 2010.
- Peter Auer and Chao-Kai Chiang. An algorithm with nearly optimal pseudo-regret for both stochastic and adversarial bandits. In *Proceedings of the International Conference on Computational Learning Theory (COLT)*, 2016.
- Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3), 2002a.

- Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1), 2002b.
- Lilian Besson and Emilie Kaufmann. What doubling tricks can and can't do for multi-armed bandits. *arXiv preprint arXiv:1803.06971*, 2018.
- Sébastien Bubeck. *Bandits games and clustering foundations*. PhD thesis, Université des Sciences et Technologie de Lille-Lille I, 2010.
- Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1), 2012.
- Sébastien Bubeck and Aleksandrs Slivkins. The best of both worlds: Stochastic and adversarial bandits. In *Proceedings of the International Conference on Computational Learning Theory (COLT)*, 2012.
- Olivier Cappé, Aurélien Garivier, Odalric-Ambrym Maillard, Rémi Munos, and Gilles Stoltz. Kullback-Leibler upper confidence bounds for optimal sequential allocation. *The Annals of Statistics*, 41, 2013.
- Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley Series in Telecommunications and Signal Processing, 2nd edition, 2006.
- Dylan J Foster, Zhiyuan Li, Thodoris Lykouris, Karthik Sridharan, and Eva Tardos. Learning in games: Robustness of fast convergence. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 4734–4742, 2016.
- Anupam Gupta, Tomer Koren, and Kunal Talwar. Better algorithms for stochastic bandits with adversarial corruptions. In *Proceedings of the International Conference on Computational Learning Theory (COLT)*, 2019.
- Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: An optimal finite time analysis. In *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, 2012.
- Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1), 1985.
- Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press (preprint), 2019.
- Thodoris Lykouris, Vahab Mirrokni, and Renato Paes Leme. Stochastic bandits robust to adversarial corruptions. In *Proceedings of the Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, 2018.
- Francesco Orabona, Koby Crammer, and Nicolò Cesa-Bianchi. A generalized online mirror descent with applications to classification and regression. *Machine Learning*, 99(3), 2015.

- Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5), 1952.
- Ralph Tyrell Rockafellar. *Convex analysis*. Princeton university press, 2015.
- Yevgeny Seldin and Gábor Lugosi. An improved parametrization and analysis of the EXP3++ algorithm for stochastic and adversarial bandits. In *Proceedings of the International Conference on Computational Learning Theory (COLT)*, 2017.
- Yevgeny Seldin and Aleksandrs Slivkins. One practical algorithm for both stochastic and adversarial bandits. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2014.
- Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2), 2012.
- William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4), 1933.
- Constantino Tsallis. Possible generalization of Boltzmann-Gibbs statistics. *Journal of statistical physics*, 52(1-2), 1988.
- Chen-Yu Wei and Haipeng Luo. More adaptive algorithms for adversarial bandits. In *Proceedings of the International Conference on Computational Learning Theory (COLT)*, 2018.
- Julian Zimmert and Tor Lattimore. Connections between mirror descent, Thompson sampling, and the information ratio. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Julian Zimmert and Yevgeny Seldin. An optimal algorithm for stochastic and adversarial bandits. In *Proceedings on the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.
- Julian Zimmert, Haipeng Luo, and Chen-Yu Wei. Beating stochastic and adversarial semi-bandits optimally and simultaneously. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019.