
Near-optimal Per-Action Regret Bounds for Sleeping Bandits

Quan Nguyen, Nishant A. Mehta

Department of Computer Science, University of Victoria
manhquan233@gmail.com, nmehta@uvic.ca

Abstract

We derive near-optimal per-action regret bounds for sleeping bandits, in which both the sets of available arms and their losses in every round are chosen by an adversary. In a setting with K total arms and at most A available arms in each round over T rounds, the best known upper bound is $O(K\sqrt{TA\ln K})$, obtained indirectly via minimizing internal sleeping regrets. Compared to the minimax $\Omega(\sqrt{TA})$ lower bound, this upper bound contains an extra multiplicative factor of $K \ln K$. We address this gap by directly minimizing the per-action regret using generalized versions of EXP3, EXP3-IX and FTRL with Tsallis entropy, thereby obtaining near-optimal bounds of order $O(\sqrt{TA\ln K})$ and $O(\sqrt{T\sqrt{AK}})$. We extend our results to the setting of bandits with advice from sleeping experts, generalizing EXP4 along the way. This leads to new proofs for a number of existing adaptive and tracking regret bounds for standard non-sleeping bandits. Extending our results to the bandit version of experts that report their confidences leads to new bounds for the confidence regret that depends primarily on the sum of experts' confidences. We prove a lower bound, showing that for any minimax optimal algorithms, there exists an action whose regret is sublinear in T but linear in the number of its active rounds.

1 INTRODUCTION

The multi-armed bandit (MAB) framework and its variants have been widely used for practical applica-

tions in various domains such as clinical trials, finance and recommender systems (see e.g. [Bouneffouf et al., 2020](#)). In the standard MAB framework, a learner interacts with K arms over T rounds. In each round, the learner chooses to observe the loss of one of the arms. While the losses of the arms in each round are unknown to the learner, the number of arms K is assumed to be fixed in every round. However, this assumption does not always hold in practice. For example, in drug testing where each arm is a drug type, certain types of drugs can only be tested in some certain rounds, or new and more effective drugs might be available only in later rounds. In such scenarios, it is important to have algorithms capable of learning with time-varying sets of available arms. This is the *sleeping bandits* setting ([Kleinberg et al., 2010](#)), where in each round $t = 1, \dots, T$, only a subset $\mathbb{A}_t \subseteq \{1, \dots, K\}$ of active arms are accessible to the learner.

The *sleeping experts* setting (also known as the specialist setting) ([Blum, 1997](#); [Freund and Schapire, 1997](#)) is the full-information feedback variant of this problem, in which the losses of the active arms are revealed at the end of each round. Prior works on sleeping experts have mainly used two different notions of regret to measure the performance of a learner, namely per-action regret ([Blum and Mansour, 2007](#); [Gaillard et al., 2014](#); [Luo and Schapire, 2015](#)) and ordering regret ([Kleinberg et al., 2010](#); [Kanade and Steinke, 2014](#); [Neu and Valko, 2014](#)). Besides the notions of regrets, an important characteristic of the setting is the stochastic or adversarial nature of the sets \mathbb{A}_t and the arms' losses. [Kanade and Steinke \(2014\)](#) indicated that obtaining a sublinear ordering regret bound is computationally hard when both \mathbb{A}_t and losses are adversarial. As a result, subsequent works on sleeping bandits usually assume at least one component to be stochastic ([Slivkins, 2013](#); [Neu and Valko, 2014](#); [Slivkins, 2014](#); [Saha et al., 2020](#)). Recently, [Gaillard et al. \(2023\)](#) developed a new notion of regret for sleeping bandits called *sleeping internal regret*, which can be minimized efficiently in the fully adversarial setting with adversarial \mathbb{A}_t and adversarial losses.

Our work focuses on minimizing the per-action regret in the fully adversarial setting. This notion of regret compares the cumulative loss of the learner to that of a single best arm in hindsight during the rounds in which that arm was active. To the best of our knowledge, in the fully adversarial setting, no prior work has focused on directly deriving optimal per-action regret bounds. We are interested in obtaining more fine-grained bounds that depend on the maximum number of active arms in any round A , where $A = \max_{t=1, \dots, T} |\mathbb{A}_t| \leq K$. The smallest existing bound is the $O(K\sqrt{TA \ln K})$ bound by Gaillard et al. (2023), obtained indirectly from minimizing the internal sleeping regret. This bound can be much larger than an $\Omega(\sqrt{TA})$ minimax lower bound implied by suitably adapting an existing minimax lower bound construction for standard bandits (Auer et al., 2002). Moreover, as we show in this work, the factor of K outside the square root can be eliminated entirely.

Another motivation for bounding the per-action regret in sleeping bandits is its implication on the adaptive and tracking regrets in standard non-sleeping bandits. Adaptive regret (also known as interval regret) (Hazan and Seshadhri, 2009; Luo et al., 2018) is the regret against a fixed arm on a time interval, while tracking regret (also known as shifting or switching regret) (Herbster and Warmuth, 1998) is the regret against a sequence of arms over T rounds. Previous work obtained adaptive and tracking regret bounds for standard non-sleeping experts via a reduction to regret bounds for sleeping experts (Freund et al., 1997; Adamskiy et al., 2016). In bandits, instead of the reduction to sleeping bandits, $\tilde{O}(\sqrt{T})$ bounds¹ on tracking and adaptive regret have been obtained via Fixed Share (Auer et al., 2002; Herbster and Warmuth, 1998; Luo et al., 2018). Our work shows that the reduction to sleeping bandits also leads to $\tilde{O}(\sqrt{T})$ adaptive and tracking bounds.

Overview of Main Results and Techniques

We extend the EXP3 (Auer et al., 2002), EXP3-IX (Neu, 2015), Follow-The-Regularized-Leader (FTRL) with Tsallis entropy (Audibert and Bubeck, 2009; Abernethy et al., 2015) and EXP4 (Auer et al., 2002) algorithms for standard bandits to sleeping bandits, obtaining new bounds that strictly generalize the existing bounds. Our results lead to new proofs for $\tilde{O}(\sqrt{T})$ adaptive and tracking regret bounds for standard bandits. The generalized algorithms and analyses are adapted to the bandit-feedback version of the experts that report their confidences setting (Blum and Mansour, 2007). A summary of

our contributions in comparison to prior works is in Table 1. All of our results hold for both pseudo-regret and high probability regret bounds. Our paper is organized as follows (all proofs are in the appendix):

- Section 3 introduces the $O(\sqrt{TA \ln K})$ and $O(\sqrt{T\sqrt{AK}})$ regret bounds for sleeping bandits. These bounds improve the best existing $O(K\sqrt{TA \ln K})$ bound, as well as recover the near-optimal $O(\sqrt{TK \ln K})$ and minimax $O(\sqrt{TK})$ bounds in non-sleeping bandits. Section 3.1 shows a novel algorithm called SB-EXP3 and its $O(\sqrt{TA \ln G_T})$ regret bound guarantee for sleeping bandits, where $G_T \leq K$ is the number of arms that were active at least once after T rounds. Its analysis relies on a new technique for bounding the growth of the potential function by decomposing the potential at round $t + 1$ based on the set of active arms in round t . In Section 3.2, the $O(\sqrt{T\sqrt{AK}})$ bound is obtained by the Follow-the-Active-and-Regularized-Leader (FTARL) algorithm, an adaptation of FTRL with Tsallis entropy to sleeping bandits. Section 3.3 considers the bandit-feedback version of the experts that report their confidences setting. Applying SB-EXP3 to this setting leads to new regret bounds which replace the dependence on T and A by the cumulative confidence over T rounds.
- Section 4 studies the *bandits with advice from sleeping experts* setting and presents SE-EXP4, a generalized version of EXP4 algorithm. The analysis developed for SB-EXP3 also works for SE-EXP4, leading to the same $O(\sqrt{TK \ln M})$ regret bound of EXP4 with M experts. For standard bandits, this implies an $O(\sqrt{TK \ln(KT)})$ bound on adaptive regret. This bound is the same as the one obtained by Luo et al. (2018), but with a different proof based on sleeping bandits instead of Fixed Share. This also implies both the $O(S\sqrt{KT \ln(KT)})$ and $O(\sqrt{SKT \ln(KT)})$ tracking regret bounds (Auer et al., 2002; Neu, 2015) for unknown and known number of arm-switches S , respectively, where the latter is obtained via restarting SE-EXP4 after every T/S rounds.
- Section 5 defines the per-action strongly adaptive regret bound as a bound that depends only on T_a for every action a , where T_a is the number of active rounds of arm a . Extending the construction of Daniely et al. (2015) for non-sleeping bandits to sleeping bandits, we show a linear $\Omega(T_a)$ per-action strongly adaptive lower bound. This implies that no algorithm can simultaneously guarantee an optimal per-action regret and sublinear $o(T_a)$ per-action regret for all arms.

¹ \tilde{O} hides terms in K , number of switches S and $\ln T$.

Table 1: A Summary of Bounds on Per-Action Regret. Hyphens indicate bounds that are either not comparable to a per-action regret bound or unavailable.

Algorithms	Fully Adversarial?	Pseudo-Regret	High-Probability
AUER (Kleinberg et al., 2010)	No (stochastic losses)	$\sqrt{TK \ln T}$	-
Sleeping-EXP3 (Saha et al., 2020)	No (stochastic \mathbb{A}_t)	-	-
SR_MAB (Blum and Mansour, 2007)	Yes	$K^2 \sqrt{TA \ln K}$	-
SI-EXP3 (Gaillard et al., 2023)	Yes	$K \sqrt{TA \ln K}$	-
SB-EXP3 (this work)	Yes	$\sqrt{TA \ln K}$	$\sqrt{TA \ln(K/\delta)}$
FTARL (this work)	Yes	$\sqrt{T \sqrt{AK}}$	$\sqrt{T \sqrt{AK}} + \sqrt{TA \ln\left(\frac{1}{\delta}\right)} + \ln\left(\frac{1}{\delta}\right)$

2 PRELIMINARIES

We consider the adversarial multi-armed bandit problem with K underlying arms, where K might be unknown. Let $[K] = \{1, 2, \dots, K\}$. In round $t = 1, 2, \dots, T$, a (possibly non-oblivious) adversary selects and reveals a set $\mathbb{A}_t \subseteq [K]$ of active arms to the learner. Let $I_{i,t} = 1$ (resp. $I_{i,t} = 0$) indicates that arm i is active (resp. inactive) in round t . Then, for each arm $i \in \mathbb{A}_t$, the adversary selects a (hidden) loss value $\ell_{i,t} \in [0, 1]$. The learner pulls one active arm $i_t \in \mathbb{A}_t$ and observes loss $\ell_{i_t,t}$.

The learner’s goal is to compete with the best arm in hindsight. For an arm $a \in [K]$, the regret of the learner with respect to arm a is the difference in the cumulative loss of the learner and that of arm a over its active rounds:

$$R(a) = \sum_{t=1}^T I_{a,t} (\ell_{i_t,t} - \ell_{a,t}). \quad (1)$$

We prove two types of regret bounds. The first is

$$\max_{a \in [K]} \mathbb{E}_{i_1, \dots, i_T} [R(a)] \leq \epsilon, \quad (2)$$

where the expectation is taken over the sequence of the learner’s selected arms. In standard non-sleeping bandits, this corresponds to the notion of *pseudo-regret* (Auer et al., 2002). If the adversary is oblivious, the pseudo-regret is equivalent to the expected regret. The second type of bound is

$$\Pr \left(\max_{a \in [K]} R(a) \leq \epsilon \right) \geq 1 - \delta, \quad (3)$$

where the probability is taken over the sequence of the learner’s selected arms.

Notations. Let $A_t = |\mathbb{A}_t|$ be the number of active arms in round t and $A = \max_{t \in [T]} A_t$ be the maximum

value of A_t over T rounds. Let $\mathbb{G}_t = \cup_{s=1, \dots, t} \mathbb{A}_s$ be the set of arms that are active at least once in the first t rounds. Let $G_t = |\mathbb{G}_t|$ be the size of \mathbb{G}_t . We write $\hat{\ell}_t = \ell_{i_t,t}$ for the learner’s loss in round t . Let $\Delta_n = \{p \in \mathbb{R}^n \mid p_i \geq 0, \sum_{i=1}^n p_i = 1\}$ be the n -dimensional probability simplex.

Remark 1. *The total number of underlying arms K is fixed before learning, and the adversary cannot change K . On the other hand, A_t and G_t are decided by the adversary and vary over time. As some arms may never be active, G_T can be strictly smaller than K .*

3 NEAR-OPTIMAL REGRET UPPER BOUNDS

In sleeping bandits, for any constant $A \in \{2, 3, \dots, K\}$, there exists an $\Omega(\sqrt{TA})$ minimax pseudo-regret lower bound. The construction follows that of the minimax lower bound for standard bandits (Auer et al., 2002) with A arms always active and $K - A$ arms always inactive over T rounds. In Section 3.1, we present SB-EXP3 (Algorithm 1) and its near-optimal $O(\sqrt{TA \ln G_T})$ pseudo-regret and high probability regret bounds. Note that SB-EXP3 does not require knowing K . In Section 3.2, we show that when K is known, an FTRL-based algorithm called FTARL (Algorithm 2) obtains an $O(\sqrt{TA \ln K})$ bound with negative Shannon entropy and an $O(\sqrt{T \sqrt{AK}})$ bound with Tsallis entropy as the regularization function.

3.1 Generalized EXP3 for Sleeping Bandits

In standard (non-sleeping) bandits, the EXP3 and EXP3-IX algorithms compute a distribution p_t over arms in round t based on the estimated regrets in previous rounds. Specifically, arms for which the estimated regrets are higher have larger probability of being sampled. In sleeping bandits, because arms can

Algorithm 1: SB-EXP3 for sleeping bandits

Input: $\eta > 0, \gamma \geq 0$

 Initialize $\tilde{q}_{i,1} = 1$ for $i = 1, 2, \dots, K$;

for each round $t = 1, \dots, T$ **do**

 The adversary selects and reveals \mathbb{A}_t ;

 Compute $W_t = \sum_{i \in \mathbb{A}_t} I_{i,t} \tilde{q}_{i,t}$;

 Compute $p_{i,t}$ by (6);

 Draw $i_t \sim p_t$ and observe $\hat{\ell}_t = \ell_{i_t,t}$;

 Compute loss estimate $\tilde{\ell}_{i,t}$ by (4);

 Update $\tilde{q}_{i,t+1}$ by (5).

have different and even non-overlapping sets of rounds, it is unclear what kind of statistics about the arms should be maintained in each round. In particular, in any given round, the estimated regret for an arm might be higher than for other arms due to it being active in more previous rounds, not because its losses are small in those rounds. Nevertheless, we will show that the estimated regret can be used effectively for selecting arms in each round.

Algorithm 1 illustrates Sleeping Bandits using EXP3 (SB-EXP3), an adaptation of EXP3 and EXP3-IX to sleeping bandits. In round t , SB-EXP3 computes a probability vector $p_t \in \Delta_{G_t}$ over \mathbb{G}_t . In the first round, p_1 is the uniform distribution. The learner samples $i_t \sim p_t$ and computes the loss estimates $\tilde{\ell}_t$ as follows:

$$\tilde{\ell}_{i,t} = \begin{cases} \frac{\ell_{i,t} \mathbf{1}\{i_t=i\}}{p_{i,t} + \gamma I_{i,t}} & \text{for } i \in \mathbb{A}_t, \\ 0 & \text{for } i \in \mathbb{G}_t \setminus \mathbb{A}_t, \end{cases} \quad (4)$$

where $\gamma \geq 0$ is a parameter of the loss estimator. When $\gamma = 0$, (4) is equivalent to the unbiased loss estimate in EXP3. When $\gamma > 0$, due to $I_{i,t} = 1$ for all $i \in \mathbb{A}_t$, (4) is the IX-loss estimator in EXP3-IX with exploration factor γ (Neu, 2015). More generally, having different exploration factors for different arms may be beneficial. We explore such a case in Section 3.3.

The weight $\tilde{q}_{i,t}$ of arm i at the beginning of round $t + 1$ is defined as follows:

$$\tilde{q}_{i,t+1} = \exp \left(\eta \sum_{s=1}^t I_{i,s} (\ell_{i,s} - \tilde{\ell}_{i,s} - \gamma \sum_{j \in \mathbb{A}_s} I_{j,s} \tilde{\ell}_{j,s}) \right), \quad (5)$$

where $\eta > 0$ is the learning rate. The sampling probability of arm i is proportional to its $\tilde{q}_{i,t}$, i.e.,

$$p_{i,t} = \begin{cases} \frac{\tilde{q}_{i,t}}{W_t} & \text{for } i \in \mathbb{A}_t, \\ 0 & \text{for } i \in \mathbb{G}_t \setminus \mathbb{A}_t, \end{cases} \quad (6)$$

where $W_t = \sum_{k \in \mathbb{A}_t} I_{k,t} \tilde{q}_{k,t}$ is the normalization factor.

Apart from setting zero sampling probability for inactive arms, similar to EXP3 and EXP3-IX, SB-EXP3 still follows the strategy of setting the sampling probability proportional to the exponential of the estimated per-action regret. The key difference is the added $-\gamma \sum_{j \in \mathbb{A}_s} I_{j,s} \tilde{\ell}_{j,s}$ in the exponent, which is crucial for obtaining near-optimal high-probability bounds (where $\gamma > 0$). Intuitively, this term significantly reduces the sampling probability of arms that were frequently active in previous rounds but not frequently chosen. Theorems 2 and 3 state the regret bounds of SB-EXP3.

Theorem 2. *With $\gamma = 0$, for any $\eta > 0$, Algorithm 1 guarantees*

$$\max_{a \in [K]} \mathbb{E}[R(a)] \leq \frac{\ln G_T}{\eta} + \frac{\eta}{2} \sum_{t=1}^T A_t.$$

Tuning η leads to an $O\left(\sqrt{\ln(G_T) \sum_{t=1}^T A_t}\right)$ bound.

Theorem 3. *For any $\gamma \geq \frac{\eta}{2} > 0$, Algorithm 1 guarantees*

$$\max_{a \in [K]} R(a) \leq \frac{\ln G_T}{\eta} + \frac{\ln(2G_T/\delta)}{\gamma} + \left(\frac{\eta}{2} + \gamma\right) \sum_{t=1}^T A_t$$

with probability at least $1 - \delta$. Tuning η and γ leads to an $O\left(\sqrt{\ln(G_T/\delta) \sum_{t=1}^T A_t}\right)$ bound.

Note that $\sum_{t=1}^T A_t \leq TA$. Since $A \leq G_T \leq K$, the bounds in Theorems 2 and 3 are generally smaller than the $O(\sqrt{TK \ln K})$ bounds of EXP3 and EXP3-IX. The difference is significant whenever $A \ll K$, which holds in many practical applications where the sets of active arms are sparse.

Analysis Sketch. The analyses of EXP3 and EXP3-IX (Auer et al., 2002; Neu, 2015) treat the normalization factor W_t as a potential function and bound the growth of $\frac{W_{t+1}}{W_t}$. This was possible in standard bandits because all K arms are always active, hence $p_{i,t+1}$ can always be related to $p_{i,t}$. However, in sleeping bandits, the sets \mathbb{A}_t and \mathbb{A}_{t+1} might be non-overlapping, hence there might be no relationship between W_{t+1} and W_t . Instead, we use

$$\tilde{Q}_t = \sum_{i \in \mathbb{G}_T} \tilde{q}_{i,t} \quad (7)$$

as the potential function. The following key technical lemma bounds the growth of $\frac{\tilde{Q}_{t+1}}{\tilde{Q}_t}$.

Lemma 4. *For any $t \geq 0$,*

$$\frac{\tilde{Q}_{t+1}}{\tilde{Q}_t} \leq \sum_{i \in \mathbb{A}_t} p_{i,t} \exp \left(\eta (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t}) \right).$$

The proof of Lemma 4 is based on decomposing \tilde{Q}_{t+1} and \tilde{Q}_t by \mathbb{A}_t and $\bar{\mathbb{A}}_t = \mathbb{G}_T \setminus \mathbb{A}_t$. If arm $i \in \mathbb{A}_t$, its weight $\tilde{q}_{i,t+1}$ at time $t+1$ can be related to $\tilde{q}_{i,t}$ via the update in (5). If arm $i \notin \mathbb{A}_t$, by construction $\tilde{q}_{i,t+1} = \tilde{q}_{i,t}$. These two observations control the growth of \tilde{Q}_{t+1} over \tilde{Q}_t .

Lemma 4 implies a bound on $\tilde{q}_{a,T+1}$ for any arm a as

$$\ln \tilde{q}_{a,T+1} \leq \ln \tilde{Q}_{T+1} = \sum_{t=1}^T \ln \frac{\tilde{Q}_{t+1}}{\tilde{Q}_t}.$$

Since $\tilde{q}_{a,T+1}$ grows with the estimated regret, this bound leads to an upper bound on the estimated regret, which in turns bounds the actual regret.

Observe that the dependency on K in Algorithm 1 can be removed completely: the initialization step of assigning \tilde{q} to 1 can be done implicitly. All other explicit computations only use the sets \mathbb{A}_t and \mathbb{G}_t . As a result, SB-EXP3 is independent of K . This property is similar to that of the AdaNormalHedge algorithm (Luo and Schapire, 2015), which obtains a low regret bound for sleeping experts when the total number of experts is unknown. Because the analysis of AdaNormalHedge relies on $[0, 1]$ -bounded loss vectors, it does not apply to sleeping bandits with the loss estimates in (4).

Remark 5. Lemma 4 enables the proofs of both pseudo-regret and high-probability bounds without significant modifications to the algorithm. This is a major advantage over existing works in sleeping bandits, which provided only pseudo-regret bounds.

In both Theorems 2 and 3, optimally tuning the learning rate requires knowing G_T and $\sum_{t=1}^T A_t$, which may not be available a priori. In Appendix F, we present Algorithm 5 which uses a two-level doubling trick to obtain a pseudo-regret bound of the same order without knowing these quantities beforehand.

Theorem 6. For any $T \geq 2$, Algorithm 5 (in Appendix F) guarantees that

$$\max_{a \in [K]} \mathbb{E}[R(a)] \leq \frac{4}{(\sqrt{2} - 1)^2} \sqrt{\ln(G_T) \sum_{t=1}^T A_t}.$$

The same technique can be used to obtain a high-probability regret bound; however, the resulting bound is slightly larger (in the logarithmic term) due to the union bound.

3.2 Generalized FTRL for Sleeping Bandits

In standard non-sleeping bandits, FTRL with Tsallis entropy obtains an $O(\sqrt{TK})$ minimax pseudo-regret bound and an $O(\sqrt{TK} \ln(1/\delta) + \ln(1/\delta))$ high

Algorithm 2: FTARL for sleeping bandits

Input: $\eta > 0, \gamma > 0, K \geq 2$

Initialize $\tilde{L}_{i,0} = 0$ for $i = 1, 2, \dots, K$;

for each round $t = 1, \dots, T$ **do**

The adversary selects and reveals \mathbb{A}_t ;

Compute $q_t = \arg \min_{q \in \Delta_K} \psi_t(q) + \langle q, \tilde{L}_{t-1} \rangle$;

Compute $W_t = \sum_{i=1}^K I_{i,t} q_{i,t}$;

Compute $p_{i,t}$ by (8);

Draw $i_t \sim p_t$ and observe $\hat{\ell}_t = \ell_{i_t,t}$;

for each arm $i \in [K]$ **do**

If $I_{i,t} = 1$, compute $\tilde{\ell}_{i,t}$ by (4);

If $I_{i,t} = 0$, compute $\tilde{\ell}_{i,t}$ by (9);

Update $\tilde{L}_{i,t} = \tilde{L}_{i,t-1} + \tilde{\ell}_{i,t}$;

probability bound (Audibert and Bubeck, 2009; Luo, 2017). In sleeping bandits, under the assumption that K is known, the Follow-the-Active-and-Regularized-Leader (FTARL) strategy in Algorithm 2 with $\frac{1}{2}$ -Tsallis entropy obtains $O(\sqrt{T\sqrt{AK}})$ pseudo-regret and $O(\sqrt{T\sqrt{AK}} + \sqrt{TA \ln(1/\delta)} + \ln 1/\delta)$ high-probability bounds. When $A = K$, the bounds of FTARL recover those of FTRL on standard bandits.

In round t , Algorithm 2 computes the weight vectors $q_t \in \Delta_K$ using FTRL, i.e.,

$$q_t = \arg \min_{q \in \Delta_K} \psi_t(q) + \langle q, \tilde{L}_{t-1} \rangle,$$

where ψ_t is the regularization function and $\tilde{L}_t \in \mathbb{R}^K$ is the cumulative (estimated) loss vector of K arms. In particular, $\tilde{L}_{i,t} = \sum_{s=1}^t \tilde{\ell}_{i,s}$. Note that this step is possible because K and the simplex Δ_K are known.

While q_t is a valid probability vector over K arms, it cannot be used directly for sampling because inactive arms might have non-zero elements in q_t . The sampling probability vector p_t is computed by taking elements in \mathbb{A}_t from q_t and normalizing:

$$p_{i,t} = \frac{I_{i,t} q_{i,t}}{\sum_{i \in \mathbb{A}_t} q_{i,t}}. \quad (8)$$

Similar to Algorithm 1, once an arm $i_t \sim p_t$ is sampled, the loss estimates of active arms are constructed by (4). The key difference in Algorithm 2 compared to Algorithm 1 is that the inactive arms have non-zero loss estimates. For an arm $i \notin \mathbb{A}_t$,

$$\tilde{\ell}_{i,t} = \hat{\ell}_t - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t}. \quad (9)$$

In Appendix G, we show that by having non-zero loss estimates for inactive arms as in (9), Algorithm 2 with negative Shannon entropy regularizer is equivalent to

Algorithm 1. In general, the motivation behind (9) is mostly technical: in every round, it ensures that the loss vector \tilde{L}_t contains only non-negative values. This facilitates using a local norm analysis of standard FTRL (e.g. Orabona, 2019). For an unbiased loss estimator where $\gamma = 0$, using (9) implies that the estimated regret is equivalent to

$$\sum_{t=1}^T I_{a,t}(\hat{\ell}_t - \tilde{\ell}_{a,t}) = \sum_{t=1}^T (\hat{\ell}_t - \tilde{\ell}_{a,t}),$$

which resembles the regret of a standard non-sleeping experts problem with input loss vector $\tilde{\ell}$. A similar technique is used by Chernov and Vovk (2009).

Using the Tsallis entropy $\psi_t(x) = \frac{1}{\eta} \left(\frac{1 - \sum_{i=1}^K x_i^\beta}{1 - \beta} \right)$ with parameter β , we obtain the regret bounds of Algorithm 2 in Theorems 7 and 8.

Theorem 7. *With $\gamma = 0$, for any $\beta \in (0, 1), \eta > 0$, Algorithm 2 with Tsallis entropy guarantees*

$$\max_{a \in [K]} \mathbb{E}[R(a)] \leq \frac{K^{1-\beta}}{\eta(1-\beta)} + \frac{\eta}{2\beta} TA^\beta.$$

Setting $\beta = \frac{1}{2}$ and tuning η leads to an $O(\sqrt{2T\sqrt{AK}})$ bound.

Theorem 8. *For any $\beta, \gamma, \eta \in (0, 1)$, Algorithm 2 with Tsallis entropy guarantees*

$$R(a) \leq \frac{K^{1-\beta}}{\eta(1-\beta)} + \frac{\eta A^\beta T}{\beta} + \gamma AT + \left(\frac{\eta + \beta}{2\beta\gamma} + \frac{1}{2} \right) \ln(3/\delta)$$

simultaneously for all $a \in [K]$ with probability at least $1 - \delta$. Letting $\beta = \frac{1}{2}$ and tuning η and γ leads to an $O(\sqrt{T\sqrt{AK}} + \sqrt{TA \ln(3/\delta)} + \ln(3/\delta))$ bound.

Remark 9. *If $G_T \leq K$ is known, then we can replace K by G_T in Algorithm 2 and in Theorems 7 and 8.*

Analysis Sketch. Similar to the analysis of Algorithm 1, the regret bound of Algorithm 2 is obtained by bounding the estimated regret $\sum_{t=1}^T \hat{\ell}_t - \tilde{\ell}_{a,t}$. By the local norm analysis of FTRL and properties of Tsallis entropy, we have

$$\sum_{t=1}^T \hat{\ell}_t - \sum_{t=1}^T \tilde{\ell}_{a,t} \leq \frac{K^{1-\beta}}{\eta(1-\beta)} + \frac{1}{2} \sum_{t=1}^T \frac{\eta}{\beta} \sum_{i=1}^K \tilde{\ell}_{i,t}^2 q_{i,t}^{2-\beta}.$$

In both cases $\gamma = 0$ and $\gamma > 0$, because inactive arms have non-zero elements in both $\tilde{\ell}_t$ and \tilde{q}_t , they contribute a non-zero positive amount in the last term on the right-hand side. Furthermore, as the computation of $\tilde{\ell}_{i,t}$ uses $p_{i,t}$, we wish to replace $q_{i,t}$ by $p_{i,t}$. The following technical lemma achieves this.

Lemma 10. *For any $t \geq 1$ and $\beta \in (0, 1)$,*

$$\sum_{i=1}^K \tilde{\ell}_{i,t}^2 q_{i,t}^{2-\beta} \leq \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t}^2 p_{i,t}^{2-\beta}.$$

Then, we bound $\sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t}^2 p_{i,t}^{2-\beta} \leq A_t^\beta$ using standard tools such as Holder's inequality. Finally, either taking the expectation (when $\gamma = 0$) or using concentration inequalities of the IX-loss estimator (Neu, 2015) leads to the $O(\sqrt{T\sqrt{AK}})$ bound.

Remark 11. *The technique of assigning the observed loss to sleeping arms in (9) is similar to the reduction from sleeping experts to standard prediction with expert advice (Chernov and Vovk, 2009; Gaillard et al., 2014). However, without Lemma 10, this reduction by itself does not immediately imply Theorems 7 and 8.*

Remark 12. *When K is large compared to A (i.e. sparse action sets), the $O(\sqrt{T\sqrt{AK}})$ bound of FTARL can be much larger than the $O(\sqrt{TA \ln G_T})$ bound of SB-EXP3. In contrast, when $A = \Theta(K)$ as in standard non-sleeping bandits, FTARL gives smaller bounds.*

3.3 Bounds on Confidence Regret

To the adversary, selecting the set \mathbb{A}_t in round t is equivalent to selecting K binary values $I_{i,t} \in \{0, 1\}$. We generalize the setting further by having the adversary select real-valued $I_{i,t} \in [0, 1]$. This new setting is the bandit feedback variant of the experts that report their confidences setting, in which $I_{i,t}$ is the confidence of expert i in round t (Blum and Mansour, 2007). In this case, $R(a)$ is the *confidence regret* with respect to arm a (Gaillard et al., 2014).

More concretely, at the beginning of round t the adversary selects and reveals K real-valued $I_{i,t} \in [0, 1]$. The set of active arms becomes $\mathbb{A}_t = \{i \in [K] : I_{i,t} > 0\}$. We apply Algorithm 1 to this problem without any modifications: the computations of $\tilde{\ell}_{i,t}, \tilde{q}_{i,t}$ and $p_{i,t}$ are the same as in Equations (4), (5) and (6). The full protocol and algorithm are given in Appendix C. We state the following high probability regret bound.

Theorem 13. *With optimally tuned η and γ , Algorithm 1 guarantees*

$$R(a) \leq O \left(\sqrt{\sum_{t=1}^T \sum_{i=1}^K I_{i,t} \ln(G_T/\delta)} \right)$$

simultaneously for all $a \in [K]$ with probability $1 - \delta$.

Observe that in this setting, because $I_{i,t}$ can be different for different arms in \mathbb{A}_t , SB-EXP3 uses a different implicit exploration factor $\gamma_{i,t} = \gamma I_{i,t}$ in (4). We

Algorithm 3: SE-EXP4 for bandits with advice from sleeping experts

Input: $\eta > 0, \gamma > 0$

Initialize $\tilde{q}_{m,1} = 1$ for $m = 1, 2, \dots, M$;

for each round $t = 1, \dots, T$ **do**

An adversary selects and reveals \mathbb{B}_t and E_t ;

Compute $z_{m,t}$ by (12) and $p_t = z_t E_t$;

Draw $i_t \sim p_t$ and observe $\hat{\ell}_t = \ell_{i_t,t}$;

Compute $\tilde{\ell}_{k,t} = \frac{\ell_{k,t} \mathbb{1}\{i_t=k\}}{p_{k,t} + \gamma}$ for arms $k \in [K]$;

for each awake expert $m \in \mathbb{B}_t$ **do**

Construct loss estimate $\tilde{x}_{m,t} = \langle E_{m,t}, \tilde{\ell}_t \rangle$;

Update $\tilde{q}_{m,t+1}$ by (11).

call this novel strategy based on the arm-dependent IX-loss estimator *confident implicit exploration*. The proof of Theorem 13 mostly follows that of Theorem 3 with one added key insight: the concentration inequalities of the original IX-loss estimator also hold for the arm-dependent IX-loss estimator.

Remark 14. An $O\left(\sqrt{\sum_{t=1}^T \sum_{i=1}^K I_{i,t} \ln G_T}\right)$ pseudo-regret bound is obtained via a similar analysis. When $I_{i,t}$ are binary, since $\sum_{i=1}^K I_{i,t} = A_t \leq A$, these bounds recover the bounds in Theorems 2 and 3.

Remark 15. Similar to Theorem 6, the two-level doubling trick presented in Appendix F can be used on $\sum_{t=1}^T \sum_{i=1}^K I_{i,t}$ and $\ln(G_T)$ to obtain pseudo-regret and high-probability bounds of the same order (up to a logarithmic factor).

4 BANDITS WITH ADVICE FROM SLEEPING EXPERTS

The bandits with expert advice framework (Auer et al., 2002) considers the non-sleeping MAB where M experts give advice to the learner in each round. We study a variant of this problem where a time-varying set $\mathbb{B}_t \subseteq [M]$ of awake experts gives advice to the learner in round t . The advice of expert $m \in \mathbb{B}_t$ is $E_{m,t} \in \Delta_K$. The learner aims to compete with each expert during their active rounds. More formally, let $I_{m,t} = 0$ and $I_{m,t} = 1$ denote whether expert m is sleeping or awake in round t , respectively. Let ℓ_t be the (hidden) loss vector of K arms in round t . The regret with respect to m after T rounds is

$$R(m) = \sum_{t=1}^T I_{m,t} (\ell_{i_t,t} - \langle E_{m,t}, \ell_t \rangle). \quad (10)$$

All theorems in this section are high probability bounds. The pseudo-regret bounds are in Appendix D.

4.1 Generalized EXP4

In the non-sleeping setting where all M experts are always active, EXP4 and EXP4-IX (Auer et al., 2002; Neu, 2015) obtain an $O(\sqrt{TK \ln M})$ pseudo-regret and $O(\sqrt{TK \ln(M/\delta)})$ high-probability bound, respectively. We will show that SE-EXP4 (Algorithm 3) obtains these two bounds as well in the bandits with advice from sleeping experts setting.

Let $B_t = |\mathbb{B}_t|$ be the number of awake experts at round t . Let E_t be the $B_t \times K$ matrix whose columns are $(E_{m,t})_{m \in \mathbb{B}_t}$. In round t , SE-EXP4 samples an expert m_t from a distribution $z_t \in \Delta_{B_t}$, then samples an arm $i_t \sim E_{m_t,t}$. This is equivalent to sampling $i_t \sim p_t$ directly, where $p_t = z_t E_t$ (Lattimore and Szepesvári, 2020). The main idea of SE-EXP4 is to take the sleeping experts as ‘‘augmented’’ sleeping arms and apply SB-EXP3. In particular, the weight of expert m is

$$\tilde{q}_{m,t} = \exp\left(\sum_{s=1}^{t-1} I_{m,s} \left(\hat{\ell}_s - \gamma \sum_{j=1}^K \tilde{\ell}_{j,s} - \tilde{x}_{m,s}\right)\right), \quad (11)$$

where $\tilde{x}_{m,s}$ is the estimated loss of expert m and $\tilde{\ell}_{j,s}$ is the estimated loss of arm j in round s . Initially, $\tilde{q}_{m,1} = 1$. For an expert $m \in \mathbb{B}_t$, its sampling probability $z_{m,t}$ is proportional to $\tilde{q}_{m,t}$, i.e.,

$$z_{m,t} = \frac{\tilde{q}_{m,t}}{\sum_{j \in \mathbb{B}_t} \tilde{q}_{j,t}}. \quad (12)$$

Note that $z_{m,t} = 0$ for $m \notin \mathbb{B}_t$. After $i_t \sim p_t$ is sampled, the loss estimates $\tilde{\ell}_t$ of K arms are computed as in (4). Then, the losses of the awake experts are estimated by the inner product of their advice and $\tilde{\ell}_t$:

$$\tilde{x}_{m,t} = \langle E_{m,t}, \tilde{\ell}_t \rangle. \quad (13)$$

Using the same analysis of SB-EXP3 implies:

Theorem 16. For any $\gamma, \eta \in (0, 1), \eta \leq 2\gamma$, SE-EXP4 guarantees

$$R(u) \leq \frac{\ln M}{\eta} + \frac{\ln(2M/\delta)}{2\gamma} + \left(\gamma + \frac{\eta}{2}\right)TK + \ln(2/\delta) \quad (14)$$

simultaneously for all experts $u \in [M]$ with probability at least $1 - \delta$, where the probability is taken over the sequence of the learner’s selected arms. Tuning η and γ leads to an $O(\sqrt{TK \ln(M/\delta)})$ bound.

4.2 Adaptive and Tracking Bounds for Standard Adversarial Bandits

Following Hazan and Seshadhri (2009), the adaptive regret of the learner on an interval $[t_1, t_2]$ with respect

to arm k in standard non-sleeping bandits is

$$R_{[t_1, t_2]}(k) = \sum_{t=t_1}^{t_2} (\ell_{i_t, t} - \ell_{k, t}),$$

where i_t is the arm chosen by the learner in round t . To obtain an adaptive regret bound using SE-EXP4, we follow the “virtual experts” strategy similar to [Adamskiy et al. \(2016\)](#): for each interval $[t_1, t_2]$ and arm k , we create a virtual expert indexed by (k, t_1, t_2) that is awake from round t_1 to round t_2 with advice $E_{(k, t), s} = e_k$ for any $s \geq t$, where e_k is the k^{th} vector in the standard basis of \mathbb{R}^K . There are $K \binom{T}{2} = \frac{KT(T+1)}{2}$ such experts. For any interval $[t_1, t_2]$, the regrets of experts $(1, t_1, t_2), (2, t_1, t_2), \dots, (K, t_1, t_2)$ are bounded by [Theorem 16](#). This implies the following result.

Theorem 17. *For any $\gamma, \eta \in (0, 1), \eta \leq 2\gamma$, SE-EXP4 with virtual experts guarantees that*

$$R_{[t_1, t_2]}(k) \leq \frac{2 \ln(KT)}{\eta} + \frac{\ln(KT/\delta)}{\gamma} + (\gamma + \frac{\eta}{2})TK + \ln(2/\delta)$$

simultaneously for all intervals $[t_1, t_2]$ and arms $k \in [K]$ with probability $1 - \delta$. Tuning η and γ leads to an $O(\sqrt{TK \ln(KT/\delta)})$ bound.

Next, we use [Theorem 17](#) to bound the tracking regret. The tracking regret of algorithm \mathcal{A} with respect to a sequence of arms $j_{1:T} = (j_1, j_2, \dots, j_T)$ is

$$R(j_{1:T}) = \sum_{t=1}^T (\hat{\ell}_t - \ell_{j_t, t}).$$

Let $S = \sum_{t=2}^T \mathbb{1}\{j_t \neq j_{t-1}\}$ be the (unknown) number of switches in the sequence $j_{1:T}$. Since there are $S+1$ intervals of different competing arms, [Theorem 17](#) immediately implies an $O\left(S\sqrt{TK \ln(KT)}\right)$ tracking bound. Furthermore, the following corollary shows that if S is known, then a tracking bound of order $O(\sqrt{STK \ln(\frac{KT}{\delta S})})$ is attainable as well. Thus, techniques based on sleeping bandits recover the tracking bounds in [Auer et al. \(2002\)](#) and [Neu \(2015\)](#).

Corollary 18. *Restarting SE-EXP4 with virtual experts after every T/S rounds guarantees that for any $j_{1:T}$ where $H(j_{1:T}) \leq S$, with probability at least $1 - \delta$*

$$R(j_{1:T}) \leq O\left(\sqrt{STK \ln\left(\frac{KT}{\delta S}\right)}\right).$$

Remark 19. [Luo et al. \(2018\)](#) obtained adaptive and tracking pseudo-regret bounds for contextual bandits which recover the same bounds (up to constants and logarithmic factors) in standard bandits. Our results hold for both pseudo-regret and high probability bounds.

5 A PER-ACTION STRONGLY ADAPTIVE LOWER BOUND

Let $T_a = |\{t \in [T] : I_{a,t} = 1\}|$ be the number of rounds in which arm a is active. Prior works in sleeping experts established per-action regret bounds of order $O(\sqrt{T_a})$, independent of T (e.g. [Chernov and Vovk, 2010](#); [Gaillard et al., 2014](#)). In this work, we call T -independent bounds that guarantee $R(a) = o(T_a)$ for all action $a \in [K]$ *per-action strongly adaptive bounds*. The bounds obtained by SB-EXP3 and FTARL in [Section 3](#) have $O(\sqrt{T})$ dependency for all arms a , thus they are not strongly adaptive. We study whether strongly adaptive bounds are achievable in sleeping bandits. The following theorem shows a linear per-action strongly adaptive lower bound for a large class of algorithms that contains SB-EXP3, FTARL in [Section 3](#) as well as any minimax optimal algorithm.

Theorem 20. *For any (possibly randomized) algorithm with guarantee*

$$\sup_{a \in [K]} E[R(a)] \leq O(T^\gamma A^\beta (\ln(T))^\mu), \quad (15)$$

where $\gamma \in (0, 1), \beta \geq 0, \mu \geq 0$ are constants, there exists a number of arms K and sequence of sets of active arms and their losses such that $A = 2$ and for at least one arm $a \in [K]$,

$$T_a \geq \Omega(T^{1-\gamma} (\ln(T))^{-\mu}) \text{ and } \mathbb{E}[R(a)] \geq \Omega(T_a).$$

Remark 21. *By setting $\gamma = \beta = \frac{1}{2}, \mu = 0$, (15) corresponds to the $O(\sqrt{TA})$ bound of any minimax optimal algorithms. This implies that no algorithms can simultaneously have an optimal per-action regret bound and a sublinear per-action strongly adaptive bound.*

Our proof extends that of [Daniely et al. \(2015\)](#) to sleeping bandits. Specifically, the setup contains arm 1 that is always active with small losses, while the other $K - 1$ arms are active in $K - 1$ non-overlapping intervals with large losses, one interval for each arm. To guarantee small regret against arm 1, with high probability, the learner must pull only arm 1 in the interval of some arm $j > 1$. Consider a slightly different setup where the losses of arm j are smaller than that of arm 1. Because the sequence of active arms and their losses in the two setups are identical from the first round until the start of the interval of arm j , the learner is unable to distinguish between the two setups and incurs linear regret against arm j in the second setup.

The limitation of this construction is that K has to grow with T . In particular, we require K of order $T^{1-\gamma} 2^{-\beta} (\ln(T))^{-\mu}$. As a result, algorithms with

bounds that are sublinear in T but have large dependency on K , for example $O(T^\gamma K)$, do not satisfy (15). Note that SB-EXP3 and FTARL always satisfy (15), since $\sqrt{TA \ln K} \leq \sqrt{TA \ln T}$ and $\sqrt{T\sqrt{AK}} \leq T^{3/4}A^{1/4}$ for any $K \leq T$.

6 CONCLUSION

We derived algorithms for sleeping bandits and proved their near-optimal per-action regret bounds. These algorithms and their regret bounds strictly generalize existing approaches and results for standard non-sleeping bandits. We showed that sleeping bandits-based approaches both imply new bounds and recover a number of existing order-optimal $\tilde{O}(\sqrt{T})$ bounds in related settings with fundamentally different proofs. Furthermore, the analysis can be used to show both pseudo-regret and high probability bounds by using either the unbiased or IX-loss estimators.

A direction for future work is to either prove an $\Omega(\sqrt{TA \ln G_T})$ lower bound, or show that an $O(\sqrt{TA})$ upper bound is possible, thereby obtaining minimax optimal bounds. For the former, such a lower bound must hold only in restricted conditions such as $A \ln G_T < K$, so there is no contradiction to the optimal $O(\sqrt{TK})$ bound in non-sleeping bandits. The latter likely requires new analysis techniques other than the potential-based analysis and FTRL.

Acknowledgements

This work was supported by the NSERC Discovery Grant RGPIN-2018-03942.

References

- Abernethy, J. D., Lee, C., and Tewari, A. (2015). Fighting bandits with a new kind of smoothness. In *Advances in Neural Information Processing Systems*, volume 28.
- Adamskiy, D., Koolen, W. M., Chernov, A., and Vovk, V. (2016). A closer look at adaptive regret. *Journal of Machine Learning Research*, 17(23):1–21.
- Audibert, J.-Y. and Bubeck, S. (2009). Minimax policies for adversarial and stochastic bandits. In *Proceedings of the 22nd Annual Conference on Learning Theory (COLT)*.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. (2002). The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77.
- Blum, A. (1997). Empirical support for winnow and weighted-majority algorithms: Results on a calendar scheduling domain. *Machine Learning*, 26(1):5–23.
- Blum, A. and Mansour, Y. (2007). From external to internal regret. *Journal of Machine Learning Research*, 8:1307–1324.
- Bouneffouf, D., Rish, I., and Aggarwal, C. (2020). Survey on applications of multi-armed and contextual bandits. In *2020 IEEE Congress on Evolutionary Computation (CEC)*, pages 1–8.
- Chernov, A. and Vovk, V. (2009). Prediction with expert evaluators’ advice. In *Algorithmic Learning Theory*, pages 8–22, Berlin, Heidelberg.
- Chernov, A. V. and Vovk, V. (2010). Prediction with advice of unknown number of experts. *CoRR*, abs/1006.0475.
- Daniely, A., Gonen, A., and Shalev-Shwartz, S. (2015). Strongly adaptive online learning. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, pages 1405–1411, Lille, France.
- Freund, Y. and Schapire, R. E. (1997). A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting. *Journal of Computer and System Sciences*, 55(1):119–139.
- Freund, Y., Schapire, R. E., Singer, Y., and Warmuth, M. K. (1997). Using and combining predictors that specialize. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing, STOC ’97*, page 334–343, New York, NY, USA.
- Gaillard, P., Saha, A., and Dan, S. (2023). One arrow, two kills: A unified framework for achieving optimal regret guarantees in sleeping bandits. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206, pages 7755–7773.
- Gaillard, P., Stoltz, G., and Erven, T. (2014). A second-order bound with excess losses. *Journal of Machine Learning Research*, 35.
- Hazan, E. and Seshadhri, C. (2009). Efficient learning algorithms for changing environments. In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML ’09*, page 393–400, New York, NY, USA.
- Herbster, M. and Warmuth, M. K. (1998). Tracking the best expert. *Machine Learning*, 32(2):151–178.
- Kanade, V. and Steinke, T. (2014). Learning hurdles for sleeping experts. *ACM Trans. Comput. Theory*, 6(3).
- Kleinberg, R., Niculescu-Mizil, A., and Sharma, Y. (2010). Regret bounds for sleeping experts and bandits. *Machine Learning*, 80(2–3):245–272.
- Lattimore, T. and Szepesvári, C. (2020). *Bandit Algorithms*. Cambridge University Press.

- Luo, H. (2017). Lecture 13, Introduction to Online Learning. <https://haipeng-luo.net/courses/CSCI699/lecture13.pdf>.
- Luo, H. and Schapire, R. E. (2015). Achieving all with no parameters: AdaNormalHedge. In *Annual Conference Computational Learning Theory*.
- Luo, H., Wei, C.-Y., Agarwal, A., and Langford, J. (2018). Efficient contextual bandits in non-stationary worlds. In Bubeck, S., Perchet, V., and Rigollet, P., editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 1739–1776. PMLR.
- Neu, G. (2015). Explore no more: Improved high-probability regret bounds for non-stochastic bandits. In *Advances in Neural Information Processing Systems*, volume 28.
- Neu, G. and Valko, M. (2014). Online combinatorial optimization with stochastic decision sets and adversarial losses. In *Advances in Neural Information Processing Systems*, volume 27.
- Orabona, F. (2019). A modern introduction to online learning. *CoRR*, abs/1912.13213.
- Saha, A., Gaillard, P., and Valko, M. (2020). Improved sleeping bandits with stochastic actions sets and adversarial rewards. In *Proceedings of the 37th International Conference on Machine Learning*, ICML'20.
- Slivkins, A. (2013). Dynamic ad allocation: Bandits with budgets. *CoRR*, abs/1306.0155.
- Slivkins, A. (2014). Contextual bandits with similarity information. *Journal of Machine Learning Research*, 15(1):2533–2568.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Not Applicable]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Not Applicable]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Not Applicable]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Not Applicable]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Not Applicable]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

A PROOFS FOR SECTION 3.1

Recall that $\tilde{Q}_t = \sum_{i \in \mathbb{G}_T} \tilde{q}_{i,t}$. For any set $S \subseteq \mathbb{G}_T$, let $\tilde{Q}_{S,t} = \sum_{i \in S} \tilde{q}_{i,t}$ be the projection of \tilde{Q}_t onto the set S . Let $\bar{S} = \mathbb{G}_T \setminus S$ be the complement of S . Note that the normalization factor W_t is equal to $\sum_{i \in \mathbb{A}_t} \tilde{q}_{i,t} = \tilde{Q}_{\mathbb{A}_t,t}$. First, we prove a technical lemma which holds for any $I_{i,t} \in [0, 1]$ for all $i \in [K]$ and $t \in [T]$.

Lemma 22. *For any $t \geq 1$,*

$$\mathbb{E}_{i \sim p_t}[\tilde{\ell}_{i,t}] = \hat{\ell}_t - \gamma \sum_{j \in \mathbb{A}_t} I_{j,t} \tilde{\ell}_{j,t}.$$

Proof. Using $\tilde{\ell}_{j,t} = 0$ for $j \neq i_t$, we obtain

$$\begin{aligned} \mathbb{E}_{i \sim p_t}[\tilde{\ell}_{i,t}] &= \sum_{i \in \mathbb{A}_t} p_{i,t} \tilde{\ell}_{i,t} \\ &= p_{i_t,t} \tilde{\ell}_{i_t,t} \\ &= p_{i_t,t} \frac{\hat{\ell}_t}{p_{i_t,t} + \gamma I_{i_t,t}} \\ &= \hat{\ell}_t - \gamma I_{i_t,t} \frac{\hat{\ell}_t}{p_{i_t,t} + \gamma I_{i_t,t}} \\ &= \hat{\ell}_t - \gamma I_{i_t,t} \tilde{\ell}_{i_t,t} \\ &= \hat{\ell}_t - \gamma \sum_{j \in \mathbb{A}_t} I_{j,t} \tilde{\ell}_{j,t}. \end{aligned}$$

□

A.1 Proof of Lemma 4

We make use of the following two facts that can be proved easily:

- Fact 1: the function $f(x) = e^{-\eta x}$ is convex for any $\eta \in \mathbb{R}$.
- Fact 2: For any $a, b > 0, c \geq 0$, if $a \geq b$ then

$$\frac{a}{b} \geq \frac{a+c}{b+c}.$$

Lemma 4. *For any $t \geq 0$,*

$$\frac{\tilde{Q}_{t+1}}{\tilde{Q}_t} \leq \sum_{i \in \mathbb{A}_t} p_{i,t} \exp \left(\eta (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t}) \right).$$

Proof. By Jensen's inequality and Fact 1, we have

$$\begin{aligned} \sum_{i \in \mathbb{A}_t} p_{i,t} \exp(-\eta \tilde{\ell}_{i,t}) &= \mathbb{E}_{i \sim p_t}[\exp(-\eta \tilde{\ell}_{i,t})] \\ &\geq \exp \left(-\eta \mathbb{E}_{i \sim p_t}[\tilde{\ell}_{i,t}] \right) \\ &= \exp \left(-\eta \left(\hat{\ell}_t - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t} \right) \right), \end{aligned}$$

where the last equality is due to Lemma 22. Multiplying by $\exp\left(\eta\left(\hat{\ell}_t - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t}\right)\right) \tilde{Q}_{\mathbb{A}_t,t} > 0$ on both sides, we obtain

$$\sum_{i \in \mathbb{A}_t} p_{i,t} \tilde{Q}_{\mathbb{A}_t,t} \exp\left(\eta\left(\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t}\right)\right) \geq \tilde{Q}_{\mathbb{A}_t,t}. \quad (16)$$

For all $i \in \mathbb{A}_t$, we have $I_{i,t} = 1$ and thus $p_{i,t} = \frac{\tilde{q}_{i,t}}{\tilde{Q}_{\mathbb{A}_t,t}}$. Equation (16) is equivalent to

$$\sum_{i \in \mathbb{A}_t} \tilde{q}_{i,t} \exp\left(\eta\left(\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t}\right)\right) \geq \tilde{Q}_{\mathbb{A}_t,t}.$$

By our update rule, $\tilde{q}_{i,t+1} = \tilde{q}_{i,t}$ for $i \notin \mathbb{A}_t$ and $\tilde{q}_{i,t+1} = \tilde{q}_{i,t} \exp\left(\eta\left(\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t}\right)\right)$ for $i \in \mathbb{A}_t$. Hence,

$$\sum_{i \in \mathbb{A}_t} \tilde{q}_{i,t+1} = \sum_{i \in \mathbb{A}_t} \tilde{q}_{i,t} \exp\left(\eta\left(\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t}\right)\right) \geq \tilde{Q}_{\mathbb{A}_t,t}.$$

Applying Fact 2 for $a = \sum_{i \in \mathbb{A}_t} \tilde{q}_{i,t+1}$, $b = \tilde{Q}_{\mathbb{A}_t,t}$ and $c = \tilde{Q}_{\bar{\mathbb{A}}_t,t}$, we obtain

$$\begin{aligned} \sum_{i \in \mathbb{A}_t} p_{i,t} \exp\left(\eta\left(\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t}\right)\right) &= \frac{\sum_{i \in \mathbb{A}_t} \tilde{q}_{i,t} \exp\left(\eta\left(\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t}\right)\right)}{\tilde{Q}_{\mathbb{A}_t,t}} \\ &= \frac{\sum_{i \in \mathbb{A}_t} \tilde{q}_{i,t+1}}{\tilde{Q}_{\mathbb{A}_t,t}} \\ &\geq \frac{\sum_{i \in \mathbb{A}_t} \tilde{q}_{i,t+1} + \tilde{Q}_{\bar{\mathbb{A}}_t,t}}{\tilde{Q}_{\mathbb{A}_t,t} + \tilde{Q}_{\bar{\mathbb{A}}_t,t}} \\ &= \frac{\tilde{Q}_{t+1}}{\tilde{Q}_t}, \end{aligned} \quad (17)$$

where the last equality is due to the fact that $\sum_{i \in \bar{\mathbb{A}}_t} \tilde{q}_{i,t+1} = \sum_{i \in \bar{\mathbb{A}}_t} \tilde{q}_{i,t} = \tilde{Q}_{\bar{\mathbb{A}}_t,t}$. \square

A.2 Bounding the Estimated Regret

Using Lemma 4, we bound the estimated regret as a function of the cumulative estimated losses of all active arms over T rounds.

Lemma 23. *For any $\gamma \geq 0$, any arm $a \in [K]$,*

$$\sum_{t=1}^T I_{a,t} \hat{\ell}_t \leq \frac{\ln G_T}{\eta} + \sum_{t=1}^T I_{a,t} \tilde{\ell}_{a,t} + \left(\frac{\eta}{2} + \gamma\right) \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t}.$$

Proof. For any arm $a \in \mathbb{G}_T$, we have

$$\ln \tilde{Q}_{T+1} = \ln \sum_{i \in \mathbb{G}_T} \tilde{q}_{i,T+1} \geq \ln \tilde{q}_{a,T+1} = \eta \sum_{t=1}^T I_{a,t} (\hat{\ell}_t - \tilde{\ell}_{a,t} - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t}). \quad (18)$$

On the other hand, we have

$$\begin{aligned}
 \ln \tilde{Q}_{T+1} &= \ln \tilde{Q}_1 + \sum_{t=1}^T \ln \frac{\tilde{Q}_{t+1}}{\tilde{Q}_t} \\
 &\leq \ln G_T + \sum_{t=1}^T \ln \left(\sum_{i \in \mathbb{A}_t} p_{i,t} \exp \left(\eta \left(\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t} \right) \right) \right) \\
 &= \ln G_T + \sum_{t=1}^T \ln \left(\exp \left(\eta \left(\hat{\ell}_t - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t} \right) \right) \sum_{i \in \mathbb{A}_t} p_{i,t} \exp(-\eta \tilde{\ell}_{i,t}) \right) \\
 &= \ln G_T + \sum_{t=1}^T \left(\eta \left(\hat{\ell}_t - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t} \right) + \ln \left(\sum_{i \in \mathbb{A}_t} p_{i,t} \exp(-\eta \tilde{\ell}_{i,t}) \right) \right) \\
 &\leq \ln G_T + \sum_{t=1}^T \left(\eta \left(\hat{\ell}_t - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t} \right) + \ln \left(\sum_{i \in \mathbb{A}_t} p_{i,t} \left(1 + \frac{\eta^2 \tilde{\ell}_{i,t}^2}{2} - \eta \tilde{\ell}_{i,t} \right) \right) \right) \\
 &= \ln G_T + \sum_{t=1}^T \left(\eta \left(\hat{\ell}_t - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t} \right) + \ln \left(1 + \eta^2 \sum_{i \in \mathbb{A}_t} \frac{p_{i,t} \tilde{\ell}_{i,t}^2}{2} - \eta \sum_{i \in \mathbb{A}_t} p_{i,t} \tilde{\ell}_{i,t} \right) \right) \\
 &\leq \ln G_T + \sum_{t=1}^T \left(\eta \left(\hat{\ell}_t - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t} \right) + \eta^2 \sum_{i \in \mathbb{A}_t} \frac{p_{i,t} \tilde{\ell}_{i,t}^2}{2} - \eta \sum_{i \in \mathbb{A}_t} p_{i,t} \tilde{\ell}_{i,t} \right) \\
 &= \ln G_T + \eta^2 \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \frac{p_{i,t} \tilde{\ell}_{i,t}^2}{2},
 \end{aligned}$$

where

- the first inequality is due to Lemma 4,
- the second inequality is $\exp(-x) \leq 1 + \frac{x^2}{2} - x$ for all $x \geq 0$,
- the third inequality is $\ln(1+x) \leq x$ for all $x \geq -1$,
- the last equality is due to Lemma 22.

Since the losses are bounded in $[0, 1]$ we also have $\sum_{i \in \mathbb{A}_t} p_{i,t} \tilde{\ell}_{i,t}^2 \leq \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t}$. This implies

$$\ln \tilde{Q}_{T+1} \leq \ln G_T + \frac{\eta^2}{2} \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t}. \tag{19}$$

From (18) and (19), we obtain

$$\sum_{t=1}^T I_{a,t} (\hat{\ell}_t - \tilde{\ell}_{a,t} - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t}) \leq \frac{\ln G_T}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t}.$$

Adding $\sum_{t=1}^T I_{a,t}(\tilde{\ell}_{a,t} + \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t})$ to both sides yields

$$\begin{aligned} \sum_{t=1}^T I_{a,t} \hat{\ell}_t &\leq \frac{\ln G_T}{\eta} + \sum_{t=1}^T I_{a,t} \tilde{\ell}_{a,t} + \frac{\eta}{2} \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t} + \gamma \sum_{t=1}^T I_{a,t} \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t} \\ &\leq \frac{\ln G_T}{\eta} + \sum_{t=1}^T I_{a,t} \tilde{\ell}_{a,t} + \left(\frac{\eta}{2} + \gamma\right) \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t}, \end{aligned} \quad (20)$$

where the second inequality is due to $I_{a,t} \leq 1$. \square

A.3 Proof of Theorem 2

Theorem 2. *With $\gamma = 0$, for any $\eta > 0$, Algorithm 1 guarantees*

$$\max_{a \in [K]} \mathbb{E}[R(a)] \leq \frac{\ln G_T}{\eta} + \frac{\eta}{2} \sum_{t=1}^T A_t.$$

Tuning η leads to an $O\left(\sqrt{\ln(G_T) \sum_{t=1}^T A_t}\right)$ bound.

Proof. If an arm a is not in \mathbb{G}_T then it has never been active in any round, thus $R(a) = 0$. For an arm $a \in \mathbb{G}_T$, taking the expectation of both sides of Lemma 23 and using

$$\begin{aligned} \mathbb{E}_{i_t \sim p_t}[\tilde{\ell}_{i,t}] &= p_{i,t} \frac{\ell_{i,t}}{p_{i,t}} \\ &= \ell_{i,t} \\ &\leq 1, \end{aligned}$$

we obtain

$$\mathbb{E}[R(a)] \leq \frac{\ln G_T}{\eta} + \frac{\eta}{2} \sum_{t=1}^T A_t. \quad \square$$

A.4 Proof of Theorem 3

Before proving Theorem 3, we state the following lemma and its corollary which provide high-probability guarantees that the sum of the loss estimators is a lower confidence bound for the sum of the true losses of all active arms. This lemma is adapted from Lemma 1 of Neu (2015).

Lemma 24 (Lemma 1 of Neu (2015)). *For all $i \in [K], t \in [T]$ and $I_{i,t} \in [0, 1]$, let $\alpha_{i,t}$ satisfy $0 \leq \alpha_{i,t} \leq 2\gamma I_{i,t}$. With probability $1 - \delta'$,*

$$\sum_{t=1}^T \sum_{i=1}^K \alpha_{i,t} \mathbb{1}\{I_{i,t} > 0\} (\tilde{\ell}_{i,t} - \ell_{i,t}) \leq \ln(1/\delta').$$

The proof of Lemma 24 is identical to that of Lemma 1 of Neu (2015) (with one additional straightforward step of handling $I_{i,t} = 0$) and thus is omitted here. For any fixed $j \in \mathbb{G}_T$, applying Lemma 24 with $\alpha_{i,t} = 2\gamma I_{i,t} \mathbb{1}\{i = j\}$ and taking a union bound over all $j \in \mathbb{G}_T$ leads to the following corollary.

Corollary 25. *With probability at least $1 - \delta'$,*

$$\sum_{t=1}^T I_{j,t} (\tilde{\ell}_{j,t} - \ell_{j,t}) \leq \frac{\ln(G_T/\delta')}{2\gamma}$$

holds simultaneously for all $j \in \mathbb{G}_T$.

Theorem 3. For any $\gamma \geq \frac{\eta}{2} > 0$, Algorithm 1 guarantees

$$\max_{a \in [K]} R(a) \leq \frac{\ln G_T}{\eta} + \frac{\ln(2G_T/\delta)}{\gamma} + \left(\frac{\eta}{2} + \gamma\right) \sum_{t=1}^T A_t$$

with probability at least $1 - \delta$. Tuning η and γ leads to an $O\left(\sqrt{\ln(G_T/\delta) \sum_{t=1}^T A_t}\right)$ bound.

Proof. Applying Lemma 24 with $\alpha_{i,t} = (\eta/2 + \gamma)I_{i,t}$, $\delta' = \delta/2$ and Corollary 25 with $\delta' = \delta/2$ gives

$$\left(\frac{\eta}{2} + \gamma\right) \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t} \leq \ln(2/\delta) + \left(\frac{\eta}{2} + \gamma\right) \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \ell_{i,t} \quad (21)$$

and

$$\sum_{t=1}^T I_{a,t} \tilde{\ell}_{a,t} \leq \frac{\ln(2G_T/\delta)}{2\gamma} + \sum_{t=1}^T I_{a,t} \ell_{a,t} \quad \text{for any } a \in [K]. \quad (22)$$

Plugging (21) and (22) into the right-hand side of Lemma 23 and taking a union bound, we have with probability at least $1 - \delta$, simultaneously for all $a \in [K]$,

$$\sum_{t=1}^T I_{a,t} \hat{\ell}_t \leq \frac{\ln G_T}{\eta} + \frac{\ln(2G_T/\delta)}{2\gamma} + \sum_{t=1}^T I_{a,t} \ell_{a,t} + \ln(2/\delta) + \left(\frac{\eta}{2} + \gamma\right) \sum_{t=1}^T \sum_{i=1}^K I_{i,t} \ell_{i,t}.$$

Subtracting $\sum_{t=1}^T I_{a,t} \ell_{a,t}$ on both sides and using $I_{i,t} \ell_{i,t} \leq 1$, we obtain

$$R(a) \leq \frac{\ln G_T}{\eta} + \frac{\ln(2G_T/\delta)}{2\gamma} + \ln(2/\delta) + \left(\frac{\eta}{2} + \gamma\right) \sum_{t=1}^T A_t \quad (23)$$

with probability at least $1 - \delta$. \square

B PROOFS FOR SECTION 3.2

Let $D_{\psi_t} : \mathbb{R}^K \times \mathbb{R}^K \rightarrow \mathbb{R}_+$ be the Bregman divergence generated by ψ_t . Let $\tilde{\ell}_t = \begin{bmatrix} \tilde{\ell}_{1,t} \\ \dots \\ \tilde{\ell}_{K,t} \end{bmatrix}$ be the estimated loss

vector in round t . We write $Q_t = \sum_{i=1}^K q_{i,t}$ for the sum of the weights of all arms at the beginning round t . Note that $Q_t = 1$. For a set $S \subseteq [K]$, we write $Q_{S,t} = \sum_{i \in S} q_{i,t}$ for the projection of Q_t on S . The complement of S is $\bar{S} = [K] \setminus S$.

B.1 Proof of Lemma 10

Lemma 10. For any $t \geq 1$ and $\beta \in (0, 1)$,

$$\sum_{i=1}^K \tilde{\ell}_{i,t}^2 q_{i,t}^{2-\beta} \leq \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t}^2 p_{i,t}^{2-\beta}.$$

Proof. Since $\tilde{\ell}_{i,t} = 0$ if $i \in \bar{\mathbb{A}}_t$ and $i \neq i_t$, the right-hand side can be reduced to

$$\sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t}^2 p_{i,t}^{2-\beta} = \tilde{\ell}_{i_t,t}^2 p_{i_t,t}^{2-\beta}.$$

For the left-hand side, the estimated losses of inactive arms are equal to

$$\begin{aligned}\hat{\ell}_t - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t} &= \hat{\ell}_t - \gamma \tilde{\ell}_{i_t,t} \\ &= \hat{\ell}_t - \gamma \frac{\hat{\ell}_t}{p_{i_t,t} + \gamma} \\ &= \frac{p_{i_t,t} \hat{\ell}_t}{p_{i_t,t} + \gamma}.\end{aligned}$$

Therefore,

$$\begin{aligned}\sum_{i=1}^K \tilde{\ell}_{i,t}^2 q_{i,t}^{2-\beta} &= \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t}^2 q_{i,t}^{2-\beta} + \sum_{i \in \bar{\mathbb{A}}_t} \tilde{\ell}_{i,t}^2 q_{i,t}^{2-\beta} \\ &= \tilde{\ell}_{i_t,t}^2 q_{i_t,t}^{2-\beta} + \frac{p_{i_t,t}^2 \hat{\ell}_t^2}{(p_{i_t,t} + \gamma)^2} \sum_{i \in \bar{\mathbb{A}}_t} q_{i,t}^{2-\beta} \\ &= \frac{\hat{\ell}_t^2}{(p_{i_t,t} + \gamma)^2} q_{i_t,t}^{2-\beta} + \frac{p_{i_t,t}^2 \hat{\ell}_t^2}{(p_{i_t,t} + \gamma)^2} \sum_{i \in \bar{\mathbb{A}}_t} q_{i,t}^{2-\beta} \\ &= \frac{\hat{\ell}_t^2}{(p_{i_t,t} + \gamma)^2} p_{i_t,t}^{2-\beta} (Q_{\bar{\mathbb{A}}_t}^{2-\beta} + p_{i_t,t}^\beta \sum_{i \in \bar{\mathbb{A}}_t} q_{i,t}^{2-\beta}) \quad \text{since } p_{i_t,t} = \frac{q_{i_t,t}}{Q_{\bar{\mathbb{A}}_t,t}} \\ &\leq \frac{\hat{\ell}_t^2}{(p_{i_t,t} + \gamma)^2} p_{i_t,t}^{2-\beta} (Q_{\bar{\mathbb{A}}_t}^{2-\beta} + \sum_{i \in \bar{\mathbb{A}}_t} q_{i,t}^{2-\beta}) \quad \text{since } p_{i_t,t}^\beta \leq 1 \\ &\leq \tilde{\ell}_{i_t,t}^2 p_{i_t,t}^{2-\beta} (Q_{\bar{\mathbb{A}}_t} + \sum_{i \in \bar{\mathbb{A}}_t} q_{i,t}) \\ &= \sum_{i \in \bar{\mathbb{A}}_t} \tilde{\ell}_{i,t}^2 p_{i,t}^{2-\beta},\end{aligned}$$

where the last two inequalities are due to applying $x^\alpha \leq x$ for all $x \in [0, 1]$, $\alpha > 1$ on $p_{i_t,t}$, $Q_{\bar{\mathbb{A}}_t,t}$ and each $q_{i,t}$ for $i \in \bar{\mathbb{A}}_t$ as well as the fact that $q_t \in \Delta_K$. \square

B.2 Bounding the Estimated Regret

Lemma 26. *For any arm $a \in [K]$,*

$$\sum_{t=1}^T I_{a,t} (\hat{\ell}_t - \tilde{\ell}_{a,t}) \leq \frac{K^{1-\beta}}{\eta(1-\beta)} + \gamma \sum_{t=1}^T \sum_{j \in \bar{\mathbb{A}}_t} \tilde{\ell}_{j,t} + \frac{\eta}{2\beta} \sum_{t=1}^T \sum_{i \in \bar{\mathbb{A}}_t} \tilde{\ell}_{i,t} p_{i,t}^{1-\beta},$$

where e_a is the a -th standard basis vector of \mathbb{R}^K .

Proof. From the regret bound of FTRL using local norms (e.g. [Orabona, 2019](#), Lemma 7.12) we have

$$\sum_{t=1}^T \langle \tilde{\ell}_t, q_t - e_a \rangle \leq \psi_{T+1}(e_a) - \min_{x \in \Delta_K} \psi_1(x) + \frac{1}{2} \sum_{t=1}^T \left\| \tilde{\ell}_t \right\|_{(\nabla^2 \psi_t(u_t))^{-1}}^2, \quad (24)$$

where u_t is a point between q_t and

$$\bar{q}_{t+1} = \arg \min_{x \in \mathbb{R}^K} \langle \tilde{\ell}_t, x \rangle + D_{\psi_t}(x; q_t). \quad (25)$$

First, we examine the left-hand side in (24). Obviously $\langle \tilde{\ell}_t, e_a \rangle = \tilde{\ell}_{a,t}$. Furthermore,

$$\begin{aligned}
 \langle \tilde{\ell}_t, q_t \rangle &= \sum_{i=1}^K \tilde{\ell}_{i,t} q_{i,t} \\
 &= \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t} q_{i,t} + \sum_{i \in \bar{\mathbb{A}}_t} \tilde{\ell}_{i,t} q_{i,t} \\
 &= \hat{\ell}_t \frac{q_{i_t,t}}{p_{i_t,t} + \gamma} + \frac{p_{i_t,t} \hat{\ell}_t}{p_{i_t,t} + \gamma} \sum_{i \in \bar{\mathbb{A}}_t} q_{i,t} \quad (\text{since } \tilde{\ell}_{i,t} = \frac{p_{i_t,t} \hat{\ell}_t}{p_{i_t,t} + \gamma} \text{ for } i \in \bar{\mathbb{A}}_t) \\
 &= \frac{p_{i_t,t} \hat{\ell}_t}{p_{i_t,t} + \gamma} \left(\frac{q_{i_t,t}}{p_{i_t,t}} + \sum_{i \in \bar{\mathbb{A}}_t} q_{i,t} \right) \\
 &= \frac{p_{i_t,t} \hat{\ell}_t}{p_{i_t,t} + \gamma} \left(\sum_{i \in \mathbb{A}_t} q_{i,t} + \sum_{i \in \bar{\mathbb{A}}_t} q_{i,t} \right) \quad (\text{since } p_{i,t} = \frac{q_{i,t}}{\sum_{j \in \mathbb{A}_j} q_{j,t}}) \\
 &= \frac{p_{i_t,t} \hat{\ell}_t}{p_{i_t,t} + \gamma} \quad (\text{since } q_t \in \Delta_K) \\
 &= \hat{\ell}_t - \gamma \sum_{j \in \bar{\mathbb{A}}_t} \tilde{\ell}_{j,t}.
 \end{aligned} \tag{26}$$

Therefore, $\langle \tilde{\ell}_t, q_t - e_a \rangle = \hat{\ell}_t - \gamma \sum_{j \in \bar{\mathbb{A}}_t} \tilde{\ell}_{j,t} - \tilde{\ell}_{a,t}$. By construction, if $I_{a,t} = 0$ then $\tilde{\ell}_{a,t} = \hat{\ell}_t - \gamma \sum_{j \in \bar{\mathbb{A}}_t} \tilde{\ell}_{j,t}$. Hence,

$$\sum_{t=1}^T \langle \tilde{\ell}_t, q_t - e_a \rangle = \sum_{t=1}^T I_{a,t} \left(\hat{\ell}_t - \gamma \sum_{j \in \bar{\mathbb{A}}_t} \tilde{\ell}_{j,t} - \tilde{\ell}_{a,t} \right).$$

Plugging this into (24) implies that

$$\sum_{t=1}^T I_{a,t} (\hat{\ell}_t - \tilde{\ell}_{a,t}) \leq \psi_{T+1}(e_a) - \min_{x \in \Delta_K} \psi_1(x) + \gamma \sum_{t=1}^T I_{a,t} \sum_{j \in \bar{\mathbb{A}}_t} \tilde{\ell}_{j,t} + \frac{1}{2} \sum_{t=1}^T \left\| \tilde{\ell}_t \right\|_{(\nabla^2 \psi_t(u_t))^{-1}}^2. \tag{27}$$

Next, we bound $\left\| \tilde{\ell}_t \right\|_{(\nabla^2 \psi_t(u_t))^{-1}}^2$ on the right-hand side. It can be shown (e.g. [Orabona, 2019](#), Section 10.1.2) that for the Tsallis entropy regularizer, the solution of the optimization problem (25) satisfies $\bar{q}_{i,t+1} \leq q_{t,i}$ whenever $\tilde{\ell}_{i,t} \geq 0$ for all $i \in [K]$. In our construction,

- $\tilde{\ell}_{i,t} = \frac{\ell_{i_t,t}}{p_{i_t,t} + \gamma} \geq 0$ if $i = i_t$,
- $\tilde{\ell}_{i,t} = 0$ if $i \in \mathbb{A}_t$ and $i \neq i_t$,
- $\tilde{\ell}_{i,t} = \ell_{i_t,t} - \gamma \sum_{j \in \bar{\mathbb{A}}_t} \tilde{\ell}_{j,t} = \frac{p_{i_t,t} \hat{\ell}_t}{p_{i_t,t} + \gamma} \geq 0$ if $i \notin \mathbb{A}_t$.

Hence, the condition $\tilde{\ell}_{i,t} \geq 0$ holds for all $i \in [K]$. It follows that $u_{i,t} \leq q_{i,t}$ for all $i \in [K]$.

It is well-known (e.g. [Abernethy et al., 2015](#)) that the Hessian of Tsallis entropy is diagonal and equal to

$$(\nabla^2 \psi_t(x))_{ii} = \frac{\beta}{\eta x_i^{2-\beta}}.$$

It follows that its inverse is a diagonal matrix with entries $\left(\eta \frac{x_i^{2-\beta}}{\beta}\right)_{i=1,2,\dots,K}$ on the main diagonal. Hence,

$$\begin{aligned} \left\| \tilde{\ell}_t \right\|_{(\nabla^2 \psi_t(u_t))^{-1}}^2 &= \frac{\eta}{\beta} \sum_{i=1}^K \tilde{\ell}_{i,t}^2 u_{i,t}^{2-\beta} \\ &\leq \frac{\eta}{\beta} \sum_{i=1}^K \tilde{\ell}_{i,t}^2 q_{i,t}^{2-\beta} \\ &\leq \frac{\eta}{\beta} \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t}^2 p_{i,t}^{2-\beta} \\ &\leq \frac{\eta}{\beta} \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t} p_{i,t}^{1-\beta}, \end{aligned}$$

where the first inequality is due to $u_{i,t} \leq q_{i,t}$, the second inequality is due to Lemma 10 and the last inequality is due to $\tilde{\ell}_{i,t} p_{i,t} \leq 1$ for all $i \in \mathbb{A}_t$. Plugging this into (27) and using $I_{a,t} \leq 1$ gives

$$\begin{aligned} \sum_{t=1}^T I_{a,t} (\hat{\ell}_t - \tilde{\ell}_{a,t}) &\leq \psi_{T+1}(e_a) - \min_{x \in \Delta_K} \psi_1(x) + \gamma \sum_{t=1}^T \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t} + \frac{\eta}{2\beta} \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t} p_{i,t}^{1-\beta} \\ &\leq \frac{K^{1-\beta}}{\eta(1-\beta)} + \gamma \sum_{t=1}^T \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t} + \frac{\eta}{2\beta} \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t} p_{i,t}^{1-\beta}, \end{aligned}$$

where, in the last inequality, where we used $\psi_{T+1}(e_a) - \min_{x \in \Delta_K} \psi_1(x) \leq \frac{K^{1-\beta}}{\eta(1-\beta)}$ as a standard property of the Tsallis entropy regularizer (e.g. Abernethy et al., 2015). \square

B.3 Proof of Theorem 7

Theorem 7. *With $\gamma = 0$, for any $\beta \in (0, 1)$, $\eta > 0$, Algorithm 2 with Tsallis entropy guarantees*

$$\max_{a \in [K]} \mathbb{E}[R(a)] \leq \frac{K^{1-\beta}}{\eta(1-\beta)} + \frac{\eta}{2\beta} T A^\beta.$$

Setting $\beta = \frac{1}{2}$ and tuning η leads to an $O\left(\sqrt{2T\sqrt{AK}}\right)$ bound.

Proof. With $\gamma = 0$, Lemma 26 implies that

$$\begin{aligned} \sum_{t=1}^T I_{a,t} (\hat{\ell}_t - \tilde{\ell}_{a,t}) &\leq \frac{K^{1-\beta}}{\eta(1-\beta)} + \frac{\eta}{2\beta} \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t} p_{i,t}^{1-\beta} \\ &= \frac{K^{1-\beta}}{\eta(1-\beta)} + \frac{\eta}{2\beta} \sum_{t=1}^T \ell_{i_t,t} p_{i_t,t}^{-\beta}, \end{aligned}$$

where the equality is due to the fact that for $i \in \mathbb{A}_t$, $\ell_{i,t} = 0$ if $i \neq i_t$ and $\tilde{\ell}_{i_t,t} = \frac{\ell_{i_t,t}}{p_{i_t,t}}$. Taking the expectation over $i_t \sim p_t$ on both sides and using

$$\begin{aligned} \mathbb{E}_{i_t \sim p_t} [\ell_{i_t,t} p_{i_t,t}^{-\beta}] &= \sum_{i \in \mathbb{A}_t} p_{i,t}^{1-\beta} \ell_{i,t} \\ &\leq \sum_{i \in \mathbb{A}_t} p_{i,t}^{1-\beta} \\ &\leq A_t^\beta, \end{aligned}$$

where the first inequality is due to $\ell_{i,t} \in [0, 1]$ and the second inequality is Holder's inequality, we obtain

$$\mathbb{E}[R(a)] \leq \frac{K^{1-\beta}}{\eta(1-\beta)} + \frac{\eta}{2\beta} T A^\beta.$$

\square

B.4 Proof of Theorem 8

Theorem 8. For any $\beta, \gamma, \eta \in (0, 1)$, Algorithm 2 with Tsallis entropy guarantees

$$R(a) \leq \frac{K^{1-\beta}}{\eta(1-\beta)} + \frac{\eta A^\beta T}{\beta} + \gamma AT + \left(\frac{\eta + \beta}{2\beta\gamma} + \frac{1}{2} \right) \ln(3/\delta)$$

simultaneously for all $a \in [K]$ with probability at least $1 - \delta$. Letting $\beta = \frac{1}{2}$ and tuning η and γ leads to an $O\left(\sqrt{T\sqrt{AK}} + \sqrt{TA \ln(3/\delta)} + \ln(3/\delta)\right)$ bound.

Proof. We apply Lemma 24 twice:

- the first time with $\delta' = \delta/3, \alpha_{i,t} = 2\gamma I_{i,t} p_{i,t}^{1-\beta}$ to obtain

$$\sum_{t=1}^T \sum_{i \in \mathbb{A}_t} p_{i,t}^{1-\beta} \tilde{\ell}_{i,t} \leq \frac{\ln(3/\delta)}{2\gamma} + \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} p_{i,t}^{1-\beta} \ell_{i,t} \quad (28)$$

with probability at least $1 - \delta/3$,

- the second time with $\delta' = \delta/3, \alpha_{i,t} = 2\gamma I_{i,t}$ to obtain

$$\sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \tilde{\ell}_{i,t} \leq \frac{\ln(3/\delta)}{2\gamma} + \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \ell_{i,t} \quad (29)$$

with probability at least $1 - \delta/3$.

We also apply Corollary 25 once with $\delta' = \frac{\delta}{3}$ to obtain

$$\sum_{t=1}^T I_{a,t} \tilde{\ell}_{a,t} \leq \frac{\ln(3G_T/\delta)}{2\gamma} + \sum_{t=1}^T I_{a,t} \ell_{a,t} \quad (30)$$

with probability at least $1 - \delta/3$. Plugging (28), (29), (30) into Lemma 26 and taking a union bound, we obtain with probability at least $1 - \delta$,

$$\sum_{t=1}^T I_{a,t} \hat{\ell}_t \leq \frac{K^{1-\beta}}{\eta(1-\beta)} + \frac{\ln(3G_T/\delta)}{2\gamma} + \sum_{t=1}^T I_{a,t} \ell_{a,t} + \frac{\ln(3/\delta)}{2} + \gamma \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \ell_{i,t} + \frac{\eta \ln(3/\delta)}{4\beta\gamma} + \frac{\eta}{2\beta} \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} p_{i,t}^{1-\beta} \ell_{i,t}.$$

Subtracting $\sum_{t=1}^T I_{a,t} \ell_{a,t}$ on both sides, we obtain

$$\begin{aligned} R(a) &\leq \frac{K^{1-\beta}}{\eta(1-\beta)} + \frac{\ln(3G_T/\delta)}{2\gamma} + \frac{\ln(3/\delta)}{2} + \gamma \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} \ell_{i,t} + \frac{\eta \ln(3/\delta)}{4\beta\gamma} + \frac{\eta}{2\beta} \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} p_{i,t}^{1-\beta} \ell_{i,t} \\ &\leq \frac{K^{1-\beta}}{\eta(1-\beta)} + \frac{\ln(3G_T/\delta)}{2\gamma} + \left(\frac{\eta}{4\beta\gamma} + \frac{1}{2} \right) \ln(3/\delta) + \gamma AT + \frac{\eta}{2\beta} \sum_{t=1}^T \sum_{i \in \mathbb{A}_t} p_{i,t}^{1-\beta} \\ &\leq \frac{K^{1-\beta}}{\eta(1-\beta)} + \frac{\eta}{2\beta} TA^\beta + \gamma AT + \left(\frac{\eta}{4\beta\gamma} + \frac{1}{2} \right) \ln(3/\delta) + \frac{\ln(3G_T/\delta)}{2\gamma}, \end{aligned}$$

where

- the second inequality is due to $\ell_{i,t} \in [0, 1]$ and $A_t \leq A$,
- the last inequality is Holder's inequality and $A_t \leq A$.

□

Algorithm 4: SB-EXP3 for experts that report their confidences with bandit feedback

Input: $\eta > 0, \gamma > 0, \eta \leq \gamma$

Initialize $\tilde{q}_{i,1} = 1$ for $i = 1, 2, \dots, K$;

for each round $t = 1, \dots, T$ **do**

An adversary selects and reveals K values $I_{i,t} \in [0, 1]$;

Compute $w_{i,t} = I_{i,t}\tilde{q}_{i,t}$ and $W_t = \sum_{i=1}^K w_{i,t}$;

Compute $p_{t,i} = \frac{w_{t,i}}{W_t}$;

Draw $i_t \sim p_t$ and observe $\hat{\ell}_t = \ell_{i_t,t}$;

Construct loss estimate $\tilde{\ell}_{i,t}$ by Equation (31) ;

Update $\tilde{q}_{i,t+1}$ by Equation (32).

end

C PROOFS FOR SECTION 3.3

First, we state a more general definition of the active sets of arms. An arm i is active in round t if $I_{i,t} > 0$ i.e. $\mathbb{A}_t = \{i \in [K] : I_{i,t} > 0\}$. Let $\mathbb{G}_t = \cup_{s=1}^t \mathbb{A}_s$ and $G_t = |\mathbb{G}_t|$. Let the potential function \tilde{Q}_t be

$$\tilde{Q}_t = \sum_{i \in \mathbb{G}_t} \tilde{q}_{i,t}.$$

The protocol is given in Algorithm 4. In round t , the loss estimator is

$$\tilde{\ell}_{i,t} = \begin{cases} \frac{\ell_{i,t} \mathbb{1}\{i_t=i\}}{p_{i,t} + \gamma I_{i,t}} & \text{if } p_{i,t} > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (31)$$

and the update rule for $\tilde{q}_{i,t+1}$ becomes

$$\tilde{q}_{i,t+1} = \tilde{q}_{i,t} \exp \left(\eta I_{i,t} \left(\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t} \right) \right). \quad (32)$$

Before proving Theorem 13, we present the following lemma which bounds the growth of $\frac{\tilde{Q}_{t+1}}{\tilde{Q}_t}$.

Lemma 27. *For any $t \geq 1$,*

$$\frac{\tilde{Q}_{t+1}}{\tilde{Q}_t} \leq 1 + \eta \frac{W_t}{\tilde{Q}_t} \sum_{i=1}^K p_{i,t} \left((\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t}) + \eta I_{i,t} (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t})^2 \right)$$

Proof. By the definition of \tilde{Q}_{t+1} , we have

$$\begin{aligned}
 \tilde{Q}_{t+1} - \tilde{Q}_t &= \sum_{i \in \mathbb{G}_T} \tilde{q}_{i,t} \exp \left(\eta I_{i,t} \left(\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t} \right) \right) - \sum_{i \in \mathbb{G}_T} \tilde{q}_{i,t} \\
 &= \sum_{i \in \mathbb{G}_T} \tilde{q}_{i,t} \left(\exp \left(\eta I_{i,t} (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t}) \right) - 1 \right) \\
 &\leq \sum_{i \in \mathbb{G}_T} \tilde{q}_{i,t} \left(\eta I_{i,t} (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t}) + \eta^2 I_{i,t}^2 (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t})^2 \right) \\
 &= \eta \sum_{i \in \mathbb{G}_T} I_{i,t} \tilde{q}_{i,t} \left((\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t}) + \eta I_{i,t} (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t})^2 \right) \\
 &= \eta W_t \sum_{i \in \mathbb{G}_T} p_{i,t} \left((\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t}) + \eta I_{i,t} (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t})^2 \right),
 \end{aligned}$$

where

- the inequality is obtained by applying $\exp(x) - 1 \leq x + x^2$ for any $x \leq 1$ on $x = \eta I_{i,t} (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t})$ and multiplying both sides by $\tilde{q}_{i,t} \geq 0$,
- the last equality is due to the computation of $p_{i,t} = \frac{I_{i,t} \tilde{q}_{i,t}}{W_t}$.

Dividing both sides by $\tilde{Q}_t > 0$, we obtain the desired expression. \square

Note that in each round t , Lemma 22 still holds for $I_{i,t} \in [0, 1]$. This implies the following corollary.

Corollary 28. For any $t \geq 1$,

$$\frac{\tilde{Q}_{t+1}}{\tilde{Q}_t} \leq 1 + \eta^2 \sum_{i=1}^K p_{i,t} I_{i,t} \left(\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t} \right)^2.$$

Proof. We write the second term in the sum on the right-hand side of Lemma 27 as follows:

$$\begin{aligned}
 \frac{\tilde{Q}_{t+1}}{\tilde{Q}_t} &\leq 1 + \eta \frac{W_t}{\tilde{Q}_t} \sum_{i=1}^K p_{i,t} \left((\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t}) + \eta I_{i,t} (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t})^2 \right) \\
 &= 1 + \underbrace{\eta \frac{W_t}{\tilde{Q}_t} \sum_{i=1}^K p_{i,t} (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t})}_{(a)} + \underbrace{\eta^2 \frac{W_t}{\tilde{Q}_t} \sum_{i=1}^K p_{i,t} I_{i,t} (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t})^2}_{(b)}
 \end{aligned}$$

We bound (a) and (b) separately. By Lemma 22 we have

$$\begin{aligned}
 \sum_{i=1}^K p_{i,t} (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t}) &= \hat{\ell}_t - \sum_{i=1}^K p_{i,t} \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t} \\
 &= 0,
 \end{aligned}$$

and thus the quantity (a) is equal to 0. To bound (b), we have

$$\begin{aligned} W_t &= \sum_{i=1}^K I_{i,t} \tilde{q}_{i,t} \\ &\leq \sum_{i=1}^K \tilde{q}_{i,t} \\ &= \tilde{Q}_t, \end{aligned}$$

where the inequality is due to $I_{i,t} \in [0, 1]$ and $\tilde{q}_{i,t} > 0$. This implies that $0 < \frac{W_t}{\tilde{Q}_t} \leq 1$. Multiplying both sides by $\eta^2 \sum_{i=1}^K p_{i,t} I_{i,t} (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t})^2 \geq 0$, we obtain

$$(b) \leq \eta^2 \sum_{i=1}^K p_{i,t} I_{i,t} \left(\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t} \right)^2,$$

which implies the desired statement. \square

C.1 Proof of Theorem 13

Theorem 13. *With optimally tuned η and γ , Algorithm 1 guarantees*

$$R(a) \leq O \left(\sqrt{\sum_{t=1}^T \sum_{i=1}^K I_{i,t} \ln(G_T/\delta)} \right)$$

simultaneously for all $a \in [K]$ with probability $1 - \delta$.

Proof. We still employ the standard strategy of lower and upper bounding Q_{T+1} . We have

$$\ln \tilde{Q}_{T+1} \geq \ln \tilde{q}_{a,T+1} = \eta \sum_{t=1}^T I_{a,t} \left(\hat{\ell}_t - \tilde{\ell}_{a,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t} \right). \quad (33)$$

On the other hand,

$$\begin{aligned} \ln \tilde{Q}_{T+1} &= \ln \tilde{Q}_1 + \sum_{t=1}^T \ln \frac{\tilde{Q}_{t+1}}{\tilde{Q}_t} \\ &\leq \ln G_T + \sum_{t=1}^T \ln \left(1 + \eta^2 \sum_{i=1}^K p_{i,t} I_{i,t} (\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t})^2 \right) \\ &\leq \ln G_T + \eta^2 \sum_{t=1}^T \sum_{i=1}^K p_{i,t} I_{i,t} \left(\hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t} \right)^2 \end{aligned} \quad (34)$$

where the first inequality is due to Corollary 28 and the second inequality is due to $\ln(1+x) \leq x$ for all $x \geq -1$. Let $c_t = \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t}$. For any $t \geq 1$, we have:

$$\begin{aligned} \hat{\ell}_t - c_t &= \hat{\ell}_t - \gamma I_{i_t,t} \tilde{\ell}_{i_t,t} \quad (\text{since } \tilde{\ell}_{j,t} = 0 \text{ if } j \neq i_t) \\ &= \hat{\ell}_t - \frac{\gamma I_{i_t,t} \hat{\ell}_t}{p_{i_t,t} + \gamma I_{i_t,t}} \\ &= \frac{p_{i_t,t} \hat{\ell}_t}{p_{i_t,t} + \gamma I_{i_t,t}} \\ &= p_{i_t,t} \tilde{\ell}_{i_t,t}. \end{aligned}$$

It follows that

$$\begin{aligned}
 \sum_{i=1}^K p_{i,t} I_{i,t} (\hat{\ell}_t - \tilde{\ell}_{i,t} - c_t)^2 &= \sum_{i=1}^K p_{i,t} I_{i,t} (\hat{\ell}_t - c_t)^2 + \sum_{i=1}^K p_{i,t} I_{i,t} \tilde{\ell}_{i,t}^2 - 2 \sum_{i=1}^K p_{i,t} I_{i,t} (\hat{\ell}_t - c_t) \tilde{\ell}_{i,t} \\
 &\leq (\hat{\ell}_t - c_t)^2 \sum_{i=1}^K p_{i,t} I_{i,t} + \sum_{i=1}^K p_{i,t} I_{i,t} \tilde{\ell}_{i,t}^2 \\
 &\leq \sum_{i=1}^K p_{i,t} I_{i,t} + \sum_{i=1}^K I_{i,t} \tilde{\ell}_{i,t},
 \end{aligned} \tag{35}$$

where

- the first inequality is due to $\hat{\ell}_t - c_t = p_{i,t} \tilde{\ell}_{i,t} \geq 0$,
- the second inequality is due to $\hat{\ell}_t - c_t = p_{i,t} \tilde{\ell}_{i,t} \leq 1$.

Combining (33), (34) and (35), we obtain

$$\sum_{t=1}^T I_{a,t} (\hat{\ell}_t - \tilde{\ell}_{a,t} - \gamma \sum_{j=1}^K I_{j,t} \tilde{\ell}_{j,t}) \leq \frac{\ln G_T}{\eta} + \eta \sum_{t=1}^T \left(\sum_{i=1}^K p_{i,t} I_{i,t} + \sum_{i=1}^K I_{i,t} \tilde{\ell}_{i,t} \right), \tag{36}$$

which implies that

$$\sum_{t=1}^T I_{a,t} \hat{\ell}_t \leq \sum_{t=1}^T I_{a,t} \tilde{\ell}_{a,t} + \frac{\ln G_T}{\eta} + \eta \sum_{t=1}^T \sum_{i=1}^K p_{i,t} I_{i,t} + \sum_{t=1}^T \sum_{i=1}^K (I_{a,t} \gamma + \eta) I_{i,t} \tilde{\ell}_{i,t} \tag{37}$$

$$\leq \sum_{t=1}^T I_{a,t} \tilde{\ell}_{a,t} + \frac{\ln G_T}{\gamma} + \gamma \sum_{t=1}^T \sum_{i=1}^K p_{i,t} I_{i,t} + 2\gamma \sum_{t=1}^T \sum_{i=1}^K I_{i,t} \tilde{\ell}_{i,t}, \tag{38}$$

where the last inequality is due to $I_{a,t} \leq 1$ and picking $\eta = \gamma$.

Note that Lemma 24 holds for $I_{i,t} \in [0, 1]$. Applying Lemma 24 with $\alpha_{i,t} = 2\gamma I_{i,t}$, $\delta' = \delta/2$, applying Corollary 25 with $\delta' = \delta/2$ and using $\sum_{i=1}^K p_{i,t} I_{i,t} \leq \sum_{i=1}^K I_{i,t}$, we obtain that with probability at least $1 - \delta$, simultaneously for all $a \in [K]$,

$$\sum_{t=1}^T I_{a,t} \hat{\ell}_t \leq \sum_{t=1}^T I_{a,t} \ell_{a,t} + \frac{\ln(2G_T/\delta)}{2\gamma} + \frac{\ln G_T}{\gamma} + \gamma \sum_{t=1}^T \sum_{i=1}^K I_{i,t} + \ln(2/\delta) + 2\gamma \sum_{t=1}^T \sum_{j=1}^K I_{j,t} \ell_{j,t}.$$

Subtracting $\sum_{t=1}^T I_{a,t} \ell_{a,t}$ on both sides and using $\ell_{j,t} \leq 1$ and $\gamma, \delta \in (0, 1)$, we obtain

$$R(a) \leq \frac{\ln(2G_T/\delta)}{\gamma} + \frac{\ln G_T}{\gamma} + \ln(2/\delta) + 3\gamma \sum_{t=1}^T \sum_{j=1}^K I_{j,t} \tag{39}$$

$$\leq \frac{3 \ln(G_T/\delta)}{\gamma} + 3\gamma \sum_{t=1}^T \sum_{i=1}^K I_{i,t}. \tag{40}$$

Setting $\gamma = \eta = \sqrt{\frac{\ln(G_T/\delta)}{\sum_{t=1}^T \sum_{i=1}^K I_{i,t}}}$ leads to the desired bound. \square

D PROOFS FOR SECTION 4

Recall that the sleeping experts are considered *sleeping augmented arms*. Note that $\mathbb{B}_t = \{m \in [M] : I_{m,t} = 1\}$ is the set of awake experts as defined in the main text. For an expert $u \in [M]$, the actual loss of expert u in round t is

$$x_{u,t} = \langle E_{u,t}, \ell_t \rangle.$$

First, we prove a technical lemma showing that the z_t -weighted average of these estimated losses of the augmented arms is equivalent to $\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t}$. This lemma is the counterpart of Lemma 22.

Lemma 29. For any $t \in [K]$ and $m \in \mathbb{B}_t$,

$$\mathbb{E}_{m \sim z_t}[\tilde{x}_{m,t}] = \hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t}.$$

Proof. Let $E_{m,t}^{(k)}$ be the value of the element at index k in $E_{m,t}$. We have

$$\begin{aligned} \mathbb{E}_{m \sim z_t}[\tilde{x}_{m,t}] &= \sum_{m \in \mathbb{B}_t} z_{m,t} \tilde{x}_{m,t} \\ &= \sum_{m \in \mathbb{B}_t} z_{m,t} \sum_{k=1}^K E_{m,t}^{(k)} \tilde{\ell}_{k,t} \\ &= \sum_{m \in \mathbb{B}_t} z_{m,t} E_{m,t}^{(i_t)} \tilde{\ell}_{i_t,t} \\ &= \frac{\hat{\ell}_t}{p_{i_t,t} + \gamma} \sum_{m \in \mathbb{B}_t} z_{m,t} E_{m,t}^{(i_t)} \\ &= \frac{\hat{\ell}_t p_{i_t,t}}{p_{i_t,t} + \gamma} \\ &= \hat{\ell}_t - \gamma \frac{\hat{\ell}_t}{p_{i_t,t} + \gamma} \\ &= \hat{\ell}_t - \gamma \tilde{\ell}_{i_t,t}, \end{aligned}$$

where

- the second equality is by Equation (13)
- the third equality is due to $\tilde{\ell}_{k,t} = 0$ whenever $k \neq i_t$
- the fourth equality is due to $p_{k,t} = \sum_{m \in \mathbb{B}_t} z_{m,t} E_{m,t}^{(k)}$ for all $k \in [K]$.

□

Let $\tilde{Q}_t = \sum_{m=1}^M \tilde{q}_{m,t}$. For a set $S \in [M]$, let $\tilde{Q}_{S,t} = \sum_{m \in S} \tilde{q}_{m,t}$ be the projection of \tilde{Q}_t on S . Let $\bar{S} = [M] \setminus S$ for any $S \subseteq [M]$. Lemma 29 leads to the following technical lemma that resembles Lemma 4.

Lemma 30. For any $t \geq 0$,

$$\frac{\tilde{Q}_{t+1}}{\tilde{Q}_t} \leq \sum_{m=1}^M z_{m,t} \exp \left(\eta \left(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t} - \tilde{x}_{m,t} \right) \right).$$

Proof. The proof makes use of the following two facts that can be proved easily:

- Fact 1: The function $f(x) = e^{-\eta x}$ is convex for any $\eta \in \mathbb{R}$.
- Fact 2: For any $a, b > 0, c \geq 0$, if $a \geq b$ then

$$\frac{a}{b} \geq \frac{a+c}{b+c}.$$

By Jensen's inequality and Fact 1, we have

$$\begin{aligned} \sum_{m=1}^M z_{m,t} \exp(-\eta \tilde{x}_{m,t}) &= \mathbb{E}_{m \sim z_t} [\exp(-\eta \tilde{x}_{m,t})] \\ &\geq \exp(-\eta \mathbb{E}_{m \sim z_t} [\tilde{x}_{m,t}]) \\ &= \exp\left(-\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t})\right), \end{aligned}$$

where the last equality is due to Lemma 29. Since $z_{m,t} = 0$ for $m \notin \mathbb{B}_t$, the expression above is equivalent to

$$\sum_{m \in \mathbb{B}_t} z_{m,t} \exp(-\eta \tilde{x}_{m,t}) \geq \exp\left(-\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t})\right).$$

Multiplying $\exp(\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t})) \tilde{Q}_{\mathbb{B}_t,t} > 0$ on both sides, we obtain

$$\sum_{m \in \mathbb{B}_t} z_{m,t} \tilde{Q}_{\mathbb{B}_t,t} \exp\left(\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t} - \tilde{x}_{m,t})\right) \geq \tilde{Q}_{\mathbb{B}_t,t}. \quad (41)$$

By definition, $z_{m,t} = \frac{\tilde{q}_{m,t}}{\tilde{Q}_{\mathbb{B}_t,t}}$. Hence, Equation (41) is equivalent to

$$\sum_{m \in \mathbb{B}_t} \tilde{q}_{m,t} \exp\left(\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t} - \tilde{x}_{m,t})\right) \geq \tilde{Q}_{\mathbb{B}_t,t}.$$

By our update rule, $\tilde{q}_{m,t+1} = \tilde{q}_{m,t}$ for $m \notin \mathbb{B}_t$ and $\tilde{q}_{m,t+1} = \tilde{q}_{m,t} \exp\left(\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t} - \tilde{x}_{m,t})\right)$ for $m \in \mathbb{B}_t$. Hence,

$$\sum_{m \in \mathbb{B}_t} \tilde{q}_{m,t+1} = \sum_{m \in \mathbb{B}_t} \tilde{q}_{m,t} \exp\left(\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t} - \tilde{x}_{m,t})\right) \geq \tilde{Q}_{\mathbb{B}_t,t}.$$

Applying Fact 2 for $a = \sum_{m \in \mathbb{B}_t} \tilde{q}_{m,t+1}$, $b = \tilde{Q}_{\mathbb{B}_t,t}$ and $c = \tilde{Q}_{\mathbb{B}_t,t}$, we obtain

$$\begin{aligned} \sum_{i=m}^M z_{m,t} \exp\left(\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t} - \tilde{x}_{m,t})\right) &= \sum_{m \in \mathbb{B}_t} z_{m,t} \exp\left(\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t} - \tilde{x}_{m,t})\right) \\ &= \frac{\sum_{m \in \mathbb{B}_t} \tilde{q}_{m,t} \exp\left(\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t} - \tilde{x}_{m,t})\right)}{\tilde{Q}_{\mathbb{B}_t,t}} \\ &= \frac{\sum_{m \in \mathbb{B}_t} \tilde{q}_{m,t+1}}{\tilde{Q}_{\mathbb{B}_t,t}} \\ &\geq \frac{\sum_{m \in \mathbb{B}_t} \tilde{q}_{m,t+1} + \tilde{Q}_{\mathbb{B}_t,t}}{\tilde{Q}_{\mathbb{B}_t,t} + \tilde{Q}_{\mathbb{B}_t,t}} \\ &= \frac{\tilde{Q}_{t+1}}{\tilde{Q}_t}, \end{aligned} \quad (42)$$

where the last equality is due to the fact that $\sum_{m \in \mathbb{B}_t} \tilde{q}_{m,t+1} = \sum_{m \in \mathbb{B}_t} \tilde{q}_{m,t} = \tilde{Q}_{\mathbb{B}_t,t}$. \square

D.1 Bounding the Estimated Regret

The following lemma bounds the estimated regret of each expert $u \in [M]$.

Lemma 31. For any $\gamma \geq 0$, for any $u \in [M]$, SE-EXP₄ guarantees that

$$\sum_{t=1}^T I_{u,t}(\hat{\ell}_t - \tilde{x}_{u,t}) \leq \frac{\ln M}{\eta} + \left(\gamma + \frac{\eta}{2}\right) \sum_{t=1}^T \sum_{j=1}^K \tilde{\ell}_{j,t}. \quad (43)$$

Proof. We have

$$\begin{aligned} \ln \tilde{Q}_{T+1} &= \ln \sum_{m=1}^K \tilde{q}_{m,T+1} \\ &\geq \ln \tilde{q}_{u,T+1} \\ &= \eta \sum_{t=1}^T I_{u,t}(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t} - \tilde{x}_{u,t}). \end{aligned} \quad (44)$$

On the other hand, we have

$$\begin{aligned} \ln \tilde{Q}_{T+1} &= \ln \tilde{Q}_1 + \sum_{t=1}^T \ln \frac{\tilde{Q}_{t+1}}{\tilde{Q}_t} \\ &\leq \ln M + \sum_{t=1}^T \ln \left(\sum_{m \in \mathbb{B}_t} z_{m,t} \exp \left(\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t} - \tilde{x}_{m,t}) \right) \right) \\ &= \ln M + \sum_{t=1}^T \ln \left(\exp(\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t})) \sum_{m \in \mathbb{B}_t} z_{m,t} \exp(-\eta \tilde{x}_{m,t}) \right) \\ &= \ln M + \sum_{t=1}^T \left(\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t}) + \ln \left(\sum_{m \in \mathbb{B}_t} z_{m,t} \exp(-\eta \tilde{x}_{m,t}) \right) \right) \\ &\leq \ln M + \sum_{t=1}^T \left(\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t}) + \ln \left(\sum_{m \in \mathbb{B}_t} z_{m,t} \left(1 + \frac{\eta^2 \tilde{x}_{m,t}^2}{2} - \eta \tilde{x}_{m,t} \right) \right) \right) \\ &= \ln M + \sum_{t=1}^T \left(\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t}) + \ln \left(1 + \eta^2 \sum_{m \in \mathbb{B}_t} \frac{z_{m,t} \tilde{x}_{m,t}^2}{2} - \eta \sum_{m \in \mathbb{B}_t} z_{m,t} \tilde{x}_{m,t} \right) \right) \\ &\leq \ln M + \sum_{t=1}^T \left(\eta(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t}) + \eta^2 \sum_{m \in \mathbb{B}_t} \frac{z_{m,t} \tilde{x}_{m,t}^2}{2} - \eta \sum_{m \in \mathbb{B}_t} z_{m,t} \tilde{x}_{m,t} \right) \\ &= \ln M + \eta^2 \sum_{t=1}^T \sum_{m \in \mathbb{B}_t} \frac{z_{m,t} \tilde{x}_{m,t}^2}{2}, \end{aligned}$$

where

- the first inequality is due to Lemma 30,
- the second inequality is $\exp(-x) \leq 1 + \frac{x^2}{2} - x$ for all $x \geq 0$,
- the third inequality is $\ln(1+x) \leq x$ for all $x \geq -1$,
- the last equality is due to Lemma 29.

We obtain

$$\sum_{t=1}^T I_{u,t}(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t} - \tilde{x}_{u,t}) \leq \frac{\ln M}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \sum_{m \in \mathbb{B}_t} z_{m,t} \tilde{x}_{m,t}^2. \quad (45)$$

Next, we proceed in the same way as in [Neu \(2015\)](#). We have

$$\sum_{m \in \mathbb{B}_t} z_{m,t} \tilde{x}_{m,t}^2 = \sum_{m \in \mathbb{B}_t} z_{m,t} \left(\sum_{k=1}^K E_{m,t}^{(k)} \tilde{\ell}_{k,t} \right)^2 \quad (46)$$

$$\leq \sum_{m \in \mathbb{B}_t} z_{m,t} \sum_{k=1}^K E_{m,t}^{(k)} (\tilde{\ell}_{k,t})^2 \quad (47)$$

$$= \sum_{k=1}^K (\tilde{\ell}_{k,t})^2 \sum_{m \in \mathbb{B}_t} z_{m,t} E_{m,t}^{(k)} \quad (48)$$

$$= \sum_{k=1}^K p_{k,t} (\tilde{\ell}_{k,t})^2 \quad (49)$$

$$\leq \sum_{k=1}^K \tilde{\ell}_{k,t}, \quad (50)$$

where the first inequality is Jensen's inequality. This implies

$$\sum_{t=1}^T I_{u,t}(\hat{\ell}_t - \gamma \sum_{j=1}^K \tilde{\ell}_{j,t} - \tilde{x}_{u,t}) \leq \frac{\ln M}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \sum_{k=1}^K \tilde{\ell}_{k,t}. \quad (51)$$

Moving $\gamma \sum_{t=1}^T I_{u,t} \sum_{j=1}^K \tilde{\ell}_{j,t}$ to the right-hand side and using $I_{u,t} \in [0, 1]$, we obtain the desired statement. \square

D.2 Proof of Theorem 16

Theorem 16. *For any $\gamma, \eta \in (0, 1), \eta \leq 2\gamma$, SE-EXP4 guarantees*

$$R(u) \leq \frac{\ln M}{\eta} + \frac{\ln(2M/\delta)}{2\gamma} + \left(\gamma + \frac{\eta}{2}\right)TK + \ln(2/\delta) \quad (14)$$

simultaneously for all experts $u \in [M]$ with probability at least $1 - \delta$, where the probability is taken over the sequence of the learner's selected arms. Tuning η and γ leads to an $O(\sqrt{TK \ln(M/\delta)})$ bound.

Proof. Lemma 31 implies that

$$\begin{aligned} \sum_{t=1}^T I_{u,t} \hat{\ell}_t &\leq \frac{\ln M}{\eta} + \sum_{t=1}^T I_{u,t} \tilde{x}_{u,t} + \sum_{t=1}^T \sum_{j=1}^K \left(\gamma + \frac{\eta}{2}\right) \tilde{\ell}_{j,t} \\ &= \frac{\ln M}{\eta} + \sum_{t=1}^T I_{u,t} \sum_{j=1}^K E_{u,t}^{(j)} \tilde{\ell}_{j,t} + \sum_{t=1}^T \sum_{j=1}^K \left(\gamma + \frac{\eta}{2}\right) \tilde{\ell}_{j,t}. \end{aligned} \quad (52)$$

Next, we apply Lemma 24 twice, where

- the first time with $\alpha_{i,t} = 2\gamma I_{u,t} E_{u,t}^{(i)}$, $\delta' = \frac{\delta}{2M}$ and a union bound over $[M]$ implies

$$\sum_{t=1}^T \sum_{j=1}^K I_{u,t} E_{u,t}^{(j)} (\tilde{\ell}_{j,t} - \ell_{j,t}) \leq \frac{\ln(2M/\delta)}{2\gamma} \quad (53)$$

with probability at least $1 - \delta/2$;

- the second time with $\alpha_{i,t} = \gamma + \eta/2$, $\delta' = \delta/2$ implies

$$\left(\gamma + \frac{\eta}{2}\right) \sum_{t=1}^T \sum_{j=1}^K (\tilde{\ell}_{j,t} - \ell_{j,t}) \leq \ln(2/\delta) \quad (54)$$

with probability at least $1 - \delta/2$.

Plugging (53) and (54) into (52) yields

$$\begin{aligned} \sum_{t=1}^T I_{u,t} \hat{\ell}_t &\leq \frac{\ln M}{\eta} + \sum_{t=1}^T I_{u,t} \sum_{j=1}^K E_{u,t}^{(j)} \ell_{j,t} + \frac{\ln(2M/\delta)}{2\gamma} + \left(\gamma + \frac{\eta}{2}\right) \sum_{t=1}^T \sum_{j=1}^K \ell_{j,t} + \ln(2/\delta) \\ &\leq \frac{\ln M}{\eta} + \sum_{t=1}^T I_{u,t} x_{u,t} + \frac{\ln(2M/\delta)}{2\gamma} + \left(\gamma + \frac{\eta}{2}\right) TK + \ln(2/\delta). \end{aligned}$$

Moving $\sum_{t=1}^T I_{u,t} x_{u,t}$ to the left-hand side, we obtain

$$R(u) \leq \frac{\ln M}{\eta} + \frac{\ln(2M/\delta)}{2\gamma} + \left(\gamma + \frac{\eta}{2}\right) TK + \ln(2/\delta). \quad (55)$$

Letting $\eta = 2\gamma$ and tuning η implies the $O(\sqrt{TK \ln(M/\delta)})$ bound. \square

D.3 A Pseudo-Regret Bound of SE-EXP4

We bound the pseudo-regret $\mathbb{E}[R(u)]$ for any expert $u \in [M]$. Note that $\gamma = 0$. Taking the expectation on both side of Lemma 31 and using

$$\mathbb{E}_{i_t \sim p_t} [\tilde{\ell}_{j,t}] = \ell_{j,t} \leq 1,$$

we obtain

$$\mathbb{E}\left[\sum_{t=1}^T I_{u,t} (\hat{\ell}_t - \tilde{x}_{u,t})\right] \leq \frac{\ln M}{\eta} + \frac{\eta}{2} TK. \quad (56)$$

On the other hand,

$$\begin{aligned} \mathbb{E}_{i_t \sim p_t} [\tilde{x}_{u,t}] &= \sum_{k=1}^K p_{k,t} \mathbb{E}[\tilde{x}_{u,t} \mid i_t = k] \\ &= \sum_{k=1}^K p_{k,t} \mathbb{E}\left[\sum_{j=1}^K E_{u,t}^{(j)} \tilde{\ell}_{j,t} \mid i_t = k\right] \\ &= \sum_{k=1}^K p_{k,t} E_{u,t}^{(k)} \frac{\ell_{k,t}}{p_{k,t}} \\ &= \sum_{k=1}^K E_{u,t}^{(k)} \ell_{k,t} \\ &= \langle E_{u,t}, \ell_t \rangle. \end{aligned}$$

We conclude that

$$\mathbb{E}[R(u)] \leq \frac{\ln M}{\eta} + \frac{\eta TK}{2}. \quad (57)$$

By setting $\eta = \sqrt{\frac{2 \ln M}{TK}}$ we obtain the bound

$$\mathbb{E}[R(u)] \leq \sqrt{2TK \ln M}. \quad (58)$$

D.4 Proof of Theorem 17

Theorem 17. For any $\gamma, \eta \in (0, 1), \eta \leq 2\gamma$, SE-EXP4 with virtual experts guarantees that

$$R_{[t_1, t_2]}(k) \leq \frac{2 \ln(KT)}{\eta} + \frac{\ln(KT/\delta)}{\gamma} + (\gamma + \frac{\eta}{2})TK + \ln(2/\delta)$$

simultaneously for all intervals $[t_1, t_2]$ and arms $k \in [K]$ with probability $1 - \delta$. Tuning η and γ leads to an $O(\sqrt{TK \ln(KT/\delta)})$ bound.

Proof. For any triple (k, t_1, t_2) we create a virtual expert that is active from round t_1 to round t_2 and give advice e_k . There are $M = K \binom{T}{2} = \frac{KT(T+1)}{2}$ such experts. Because $M \leq KT^2 \leq (KT)^2$, we have

$$\ln M \leq \ln((KT)^2) = 2 \ln(KT).$$

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

$$\ln(2M/\delta) \leq \ln(4M/\delta^2) \leq \ln((2KT/\delta)^2) = 2 \ln(2KT/\delta).$$

Let u_{k, t_1, t_2} denote the expert indexed by (k, t_1, t_2) . By Theorem 16, with probability at least $1 - \delta$,

$$\begin{aligned} R_{[t_1, t_2]}(k) &= R(u_{k, t_1, t_2}) \\ &\leq \frac{\ln M}{\eta} + \frac{\ln(2M/\delta)}{2\gamma} + (\gamma + \frac{\eta}{2})TK + \ln(2/\delta) \\ &\leq \frac{2 \ln(KT)}{\eta} + \frac{\ln(2KT/\delta)}{\gamma} + (\gamma + \frac{\eta}{2})TK + \ln(2/\delta) \end{aligned} \quad (59)$$

holds simultaneously for all u_{k, t_1, t_2} . □

D.5 Proof of Corollary 18

Corollary 18. Restarting SE-EXP4 with virtual experts after every T/S rounds guarantees that for any $j_{1:T}$ where $H(j_{1:T}) \leq S$, with probability at least $1 - \delta$

$$R(j_{1:T}) \leq O\left(\sqrt{STK \ln\left(\frac{KT}{\delta S}\right)}\right).$$

Proof. Without loss of generality, assume that T is divisible by S . We use the following algorithm called SE-EXP4-Restart, which runs in S episodes. In each episode, a new instance of SE-EXP4 with virtual experts is run for T/S rounds. Let $b = 1, 2, \dots, S$ be an index for the episodes, and $t_{(b)} = \frac{bT}{S}$ be the ending round of episode b . Note that each episode b starts from round $\frac{(b-1)T}{S} + 1$ to round $\frac{bT}{S}$.

We examine the regret of this SE-EXP4-Restart with respect to the competing arms $j_{1:T}$ in each episode b . Divide T rounds into $Z = H(j_{1:T}) + 1$ non-overlapping segments, where the competing arms are the same within each segment $z = 1, \dots, Z$. Let S_z, E_z be the first and last rounds of segment z . For every pair (b, z) of episode b and segment z , let

$$F_{b,z} = [S_z, E_z] \cap \left[\frac{(b-1)T}{S}, \frac{bT}{S}\right] \quad (60)$$

be the intersection between the rounds of episode b and segment z . Since the episodes are non-overlapping and the segments are non-overlapping, the intervals $F_{b,z}$'s are non-overlapping. In addition, their union is $\cup_{b,z} F_{b,z} = [T]$.

Fix an episode b and a segment z . There are two cases:

- $F_{b,z} = \emptyset$: Obviously, the regret of SE-EXP4-Restart on this empty interval is zero.

- $F_{b,z} \neq \emptyset$: in this case, because $F_{b,z}$ is an interval within episode b , the tracking regret of SE-EXP4-Restart on $F_{b,z}$ cannot exceed the adaptive regret of running (a new instance of) SE-EXP4 with virtual experts under horizon T/S . By Theorem 17, this is bounded by

$$R_{F_{b,s}} \leq \frac{2}{\eta} \ln \left(\frac{KT}{S\delta} \right) + \eta TK/S + \ln(2S/\delta) \quad (61)$$

with probability at least $1 - \delta/S$.

Taking a union bound over all S episodes and setting $\eta = 2\gamma$ implies that with probability at least $1 - \delta$,

$$R_{F_{b,s}} \leq \frac{2}{\eta} \ln \left(\frac{KT}{S\delta} \right) + \eta TK/S + \ln(2S/\delta)$$

simultaneously for all intervals $F_{b,s}$. Because $F_{b,z}$'s are non-overlapping and their union is $[T]$, the tracking regret of SE-EXP4-Restart is bounded by

$$R(j_{1:T}) = \sum_{b,z} R_{F_{b,z}} \quad (62)$$

$$= \left(\sum_{b,z} \mathbf{1}\{F_{b,z} \neq \emptyset\} \right) \left(\frac{2}{\eta} \ln \left(\frac{KT}{S\delta} \right) + \eta TK/S + \ln(2S/\delta) \right). \quad (63)$$

Next, we show that the count $C = \sum_{b,z} \mathbf{1}\{F_{b,z} \neq \emptyset\}$ is smaller than $2S$. Observe that the S episodes split the sequence of T rounds into S intervals with $S - 1$ splitting points (not counting the two ends at 0 and T). Similarly, the $S + 1$ segments of $j_{1:T}$ have S splitting points. In total, there are at most $2S - 1$ splitting points from the episodes and the segments of $j_{1:T}$. Each non-empty interval $F_{b,z}$ has an ending point that is either T or one of the $2S - 1$ splitting points. Therefore, there can be at most $2S$ such $F_{b,z}$. We conclude that $C \leq 2S$. As a result, with probability at least $1 - \delta$,

$$R(j_{1:T}) \leq \frac{4S}{\eta} \ln \left(\frac{KT}{S\delta} \right) + 2\eta TK + 2S \ln(2S/\delta).$$

Letting $\eta = \sqrt{\frac{2S \ln(\frac{KT}{S\delta})}{TK}}$ leads to the desired bound. \square

E PROOFS FOR SECTION 5

Theorem 20. *For any (possibly randomized) algorithm with guarantee*

$$\sup_{a \in [K]} E[R(a)] \leq O(T^\gamma A^\beta (\ln(T))^\mu), \quad (15)$$

where $\gamma \in (0, 1)$, $\beta \geq 0$, $\mu \geq 0$ are constants, there exists a number of arms K and sequence of sets of active arms and their losses such that $A = 2$ and for at least one arm $a \in [K]$,

$$T_a \geq \Omega(T^{1-\gamma} (\ln(T))^{-\mu}) \text{ and } \mathbb{E}[R(a)] \geq \Omega(T_a).$$

Proof. Let \mathcal{A} be any algorithm with the stated worst-case guarantee and $f(T, A) = O(T^\gamma A^\beta (\ln(T))^\mu)$ represent the worst-case regret bound of \mathcal{A} . We show a construction with $A_t = 2$ for all $t = 1, \dots, T$. Our construction is adapted from the lower bound construction of Daniely et al. (2015) for strongly adaptive regret in the standard adversarial MAB setting. Without loss of generality, assume $4f(T, A)$ divides T . Let $L = \frac{T}{4f(T, A)}$ and $K = 1 + 4f(T, A)$. Obviously $L \geq \Omega(T^{1-\gamma} A^{-\beta} (\ln(T))^{-\mu})$. Consider an environment V_0 defined as follows:

- Arm 1 is always active. Its losses are $\ell_{1,t} = 0.5$ for all $t \in [T]$.
- Arm $k = 2, 3, \dots, K$ are active only for the rounds in the interval $\mathcal{I}_k = \left[\frac{(k-2)T}{4f(T, A)} + 1, \frac{(k-1)T}{4f(T, A)} \right]$, respectively. The length of each interval is L . Their losses are $\ell_{k,t} = 1$ for the rounds t in which they are active.




Figure 1: Environment V_0 . All arms have loss equal 1 when they are active




Figure 2: Environment V_k . Except for arm k which has losses equal to 0, all arms have losses equal to 1 when they are active.

Figure 1 illustrates this environment V_0 . We also define $K - 1$ competing environments V_k for $k = 2, 3, \dots, K$, defined as follows:

- The number of arms and their active rounds are identical to that of V_0 . That is, arm 1 is always active for all rounds while each of the arms $k = 2, 3, \dots, K$ are active for rounds within \mathcal{I}_k , respectively.
- The losses of every arm are the same as in V_0 , except for that of arm k : all losses of arm k are 0. Figure 2 illustrates environment $V_{k,j}$.

Following the standard strategy of comparing the behavior of the algorithm on V_0 and competing environments (Auer et al., 2002; Daniely et al., 2015), we first consider the neutral environment V_0 . Let \mathbb{E}_0 and \Pr_0 indicate the expectation and probability taken in this environment over the randomness of the algorithm \mathcal{A} , respectively. Let $U = \{t : i_t \neq 1\}$ be the rounds in which the arm chosen by \mathcal{A} is not arm 1. On V_0 , since arm 1 is the best arm and the gaps between the losses of arm 1 and that of every other arm is 0.5, the inequality $\sup_a \mathbb{E}_0[R(a)] \leq f(T, A)$ implies

$$\mathbb{E}_0 [|U|] \leq 2f(T, A). \quad (64)$$

For any $k \in \{2, 3, \dots, K\}$, let

$$\mathcal{E}_k = \{U \cap \mathcal{I}_k = \emptyset\}$$

be the event that only arm 1 is chosen by \mathcal{A} on \mathcal{I}_k . Because the \mathcal{I}_k are non-overlapping and $\cup_{k=2, \dots, K} \mathcal{I}_k = [T]$, we can write

$$U = \cup_{k=2, \dots, K} (U \cap \mathcal{I}_k),$$

and

$$|U| = \sum_{k=2}^K |U \cap \mathcal{I}_k|.$$

Next, we show that for some segment \mathcal{I}_{k^*} of size $L = \frac{T}{4f(T, A)}$, we have $\mathbb{E}_0 [|U \cap \mathcal{I}_{k^*}|] \leq \frac{1}{2}$. Assume on the

contrary that $\mathbb{E}_0 [|U \cap \mathcal{I}_k|] > \frac{1}{2}$ for all $k = 2, \dots, K$. Then,

$$\begin{aligned} \mathbb{E}_0[U] &= \sum_{k=2}^K \mathbb{E}_0[U \cap \mathcal{I}_k] \\ &> \frac{K-1}{2} \\ &= 2f(T, A), \end{aligned}$$

which contradicts (64). Since $|U \cap \mathcal{I}_{k^*}|$ is a non-negative integer, the inequality $\mathbb{E}_0 [|U \cap \mathcal{I}_{k^*}|] \leq \frac{1}{2}$ implies $\Pr_0 [\mathcal{E}_{k^*}] \geq \frac{1}{2}$. Next, we consider the environment V_{k^*} . From round 1 up to (and including) round $t^* = \frac{(k^*-2)T}{4f(T, A)}$, the set of active arms and their losses on V_0 and V_{k^*} are identical. As a result, the distribution over the past chosen arms and observed losses induced by \mathcal{A} up to round t^* is the same on both environment. Moreover, once the algorithm enters \mathcal{I}_{k^*} at round $t^* + 1$ and chooses only arm 1 subsequently, it also observes the same sequence of chosen arms and losses on both V_0 and V_{k^*} until the end of \mathcal{I}_{k^*} . Hence,

$$\Pr_{k^*} [\mathcal{E}_{k^*}] = \Pr_0 [\mathcal{E}_{k^*}] \geq \frac{1}{2}, \quad (65)$$

where the subscript k^* indicates a probability measured in environment V_{k^*} . In other words, on V_{k^*} , with probability at least 0.5, arm 1 is always chosen on \mathcal{I}_{k^*} and arm k^* is never chosen. Under this event \mathcal{E}_{k^*} , the regret of \mathcal{A} with respect to arm k^* is $(0.5 - 0)L = 0.5L$. When \mathcal{E}_{k^*} does not hold, the regret with respect to k^* is non-negative because $\ell_{k^*, t} = 0$. Overall, on V_{k^*} the expected regret is at least

$$\mathbb{E}_{k^*} [R(k^*)] \geq \frac{L}{4}.$$

□

F PROOF OF THEOREM 6: DOUBLING-TRICK FOR ADAPTING TO $\sum_{t=1}^T A_t$ AND G_T

For SB-EXP3, computing the optimal learning rates require the fraction $\sqrt{\frac{\ln G_T}{\sum_{t=1}^T A_t}}$ of $\sqrt{\ln G_T}$ and the sum $\sqrt{\sum_{t=1}^T A_t}$. Both of these quantities are monotonically non-decreasing. Therefore, we can apply the doubling trick on these two quantities. First, we prove a simple lemma justifying doing this.

Lemma 32. *Let $a, b, c, d > 0$ be constants such that $a \leq c$ and $b \leq d$. Let*

$$f(x) = \frac{a}{x} + \frac{bx}{2}$$

be a function on \mathbb{R}_+ . Then,

$$f\left(\sqrt{\frac{2c}{d}}\right) \leq \sqrt{2cd}.$$

Proof. Due to $a \leq c$ and $b \leq d$, for any $x \geq 0$ we have

$$f(x) \leq \frac{c}{x} + \frac{dx}{2}.$$

Plugging $x = \sqrt{\frac{2c}{d}}$ into the right-hand side gives

$$\begin{aligned} f\left(\sqrt{\frac{2c}{d}}\right) &\leq c\sqrt{\frac{d}{2c}} + d\sqrt{\frac{c}{2d}} \\ &= \sqrt{2cd}. \end{aligned}$$

□

Algorithm 5: SB-EXP3-ATGT adapted to G_T and $\sum_{t=1}^T A_t$

Initialize $U = 0, C = 1, b = 1, \mathbb{V} = \emptyset$;

Initialize $L_i = 0$ for $i = 1, 2, \dots, K$;

for each round $t = 1, \dots$, **do**

An adversary selects and reveals \mathbb{A}_t ;

if $\ln(|\mathbb{V} \cup \mathbb{A}_t|) > 2^C$ **then**

$C = C + 1$;

while $\ln(A_t) > 2^C$ **do**

$C = C + 1$;

Set $\mathbb{V} = \emptyset$;

Set $U = 0, b = 1$;

for arm $i \in \mathbb{G}_t$ **do**

$L_i = 0$;

if $U + A_t > 2^b$ **then**

$b = b + 1$;

while $A_t > 2^b$ **do**

$b = b + 1$;

Set $U = 0$;

for arm $i \in \mathbb{G}_t$ **do**

$L_i = 0$;

Update $\mathbb{V} = \mathbb{V} \cup \mathbb{A}_t$;

Update $U = U + A_t$;

Compute $\eta = \sqrt{\frac{2^{C+1}}{2^b}}$;

for arm $i \in \mathbb{A}_t$ **do**

$\hat{q}_{i,t} = \exp(\eta L_i)$

Compute p_t by (6);

Sample $i_t \sim p_t$;

Compute $\tilde{\ell}_{i,t}$ by (4);

for arm $i \in \mathbb{A}_t$ **do**

$L_i = L_i + \hat{\ell}_t - \tilde{\ell}_{i,t} - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t}$;

Lemma 32 implies that for any horizon T , if we set $\eta_t = \sqrt{\frac{2^c}{d}}$ for c and d such that $c \geq \ln G_T$ and $d \geq \sum_{t=1}^T A_t$ then we obtain a regret of bound $\sqrt{2cd}$.

We proceed to perform the doubling trick on $\ln G_T$ and $\sum_{t=1}^T A_t$. The full procedure is given in Algorithm 5. The main idea is a two-level doubling trick which divides the learning process into episodes as follows:

- The first level: throughout the learning process, we maintain a set \mathbb{V} for the arms that have been active at least once in each episode and an upper bound 2^C for $\ln(|\mathbb{V}|)$. Initially, $\mathbb{V} = \emptyset$ and $C = 1$. At the beginning of round t , we check if $\ln(|\mathbb{V} \cup \mathbb{A}_t|)$ exceeds 2^C . If $\ln(|\mathbb{V} \cup \mathbb{A}_t|) \leq 2^C$ then we continue the learning process and update $\mathbb{V} = \mathbb{V} \cup \mathbb{A}_t$. Otherwise, we reset \mathbb{V} to \emptyset , increment C by at least one until $2^C \geq \ln(A_t)$ and start a new episode from round t .
- The second level: throughout the rounds of each episode, we maintain a cumulative sum U for the sum of A_t and an upper bound 2^b for U . Note that C is fixed within these rounds. Before the first round of an episode, we initialize $U = 0, b = 1$. As long as $U + A_t \leq 2^b$, we run SB-EXP3 with $\eta = \sqrt{\frac{2^{C+1}}{2^b}}$ and update $U = U + A_t$. Once U exceeds 2^b at some round t , we increment b by at least one until $A_t \leq 2^b$, reset $U = 0$ and run a new instance of SB-EXP3 onwards.

The pseudo-regret of Algorithm 5 is shown in the following theorem.

Theorem 6. For any $T \geq 2$, Algorithm 5 (in Appendix F) guarantees that

$$\max_{a \in [K]} \mathbb{E}[R(a)] \leq \frac{4}{(\sqrt{2} - 1)^2} \sqrt{\ln(G_T) \sum_{t=1}^T A_t}.$$

Proof. Let C_T be the last value of C after T rounds. Note that C_T is also the number of episodes. Let $c = 1, 2, \dots, C_T$ be the index of the episodes. We first bound the regret within each episode c , and then sum up this bound over C_T episodes to get the total regret bound.

Bounding The Regret Within Each Episode

Fix an episode c . Let \mathbb{T}_c be the rounds in this episode, and $T_c = |\mathbb{T}_c|$. Let \mathbb{V}_c be the set of arms that are active at least once during this episode, and $V_c = |\mathbb{V}_c|$. By construction, $b = 1$ at the beginning of this episode. Let B be the last value of b after T_c rounds starting from the first round of episode c . Divide the rounds in \mathbb{T}_c into B time intervals, where b does not change in each interval. For $b = 1, 2, \dots, B$, let F_b be the time interval of b . Let $U_b = \sum_{t \in F_b} A_t$ be the sum of A_t within F_b . Since $U_b \leq 2^b$ and $\ln(V_c) \leq 2^c$, by Lemma 32 and Theorem 2, the regret of the learner in this interval is bounded by

$$\max_{a \in [K]} \mathbb{E}[R_{F_b}(a)] \leq \sqrt{2^b 2^{c+1}}.$$

Let $R_c(a)$ be the regret incurred during episode c with respect to arm a . We have

$$\begin{aligned} \max_{a \in [K]} \mathbb{E}[R_c(a)] &\leq \sum_{b=1}^B \max_{a \in [K]} \mathbb{E}[R_{F_b}(a)] \\ &\leq \sum_{b=1}^B \sqrt{2^b 2^{c+1}} \\ &\leq \frac{\sqrt{2}}{\sqrt{2} - 1} \sqrt{2^B 2^{c+1}}. \end{aligned}$$

For $b = 1, \dots, B$, let t_b be the first round of F_b . Since b is increased from $B - 1$ to B at the beginning of round t_B , we have $U_{B-1} + A_{t_B} > 2^{B-1}$. It follows that $2^B \leq 2(U_{B-1} + A_{t_B}) \leq 2 \sum_{t \in \mathbb{T}_c} A_t$. Hence,

$$\max_{a \in [K]} \mathbb{E}[R_c(a)] \leq \frac{\sqrt{2}}{\sqrt{2} - 1} \sqrt{2^{c+1} \sum_{t \in \mathbb{T}_c} A_t}. \quad (66)$$

Bounding The Total Regret

Summing up (66) for $c = 1, \dots, C_T$, we obtain

$$\begin{aligned} \max_{a \in [K]} \mathbb{E}[R(a)] &\leq \sum_{c=1}^{C_T} \max_{a \in [K]} \mathbb{E}[R_c(a)] \\ &\leq \frac{\sqrt{2}}{\sqrt{2} - 1} \sum_{c=1}^{C_T} \sqrt{2^{c+1} \sum_{t \in \mathbb{T}_c} A_t} \\ &\leq \frac{\sqrt{2}}{\sqrt{2} - 1} \left(\sqrt{\sum_{t=1}^T A_t} \right) \sum_{c=1}^{C_T} \sqrt{2^{c+1}} \\ &\leq \frac{2\sqrt{2}}{(\sqrt{2} - 1)^2} \left(\sqrt{\sum_{t=1}^T A_t} \right) \sqrt{2^{C_T}}, \end{aligned}$$

where the third inequality is due to $\sum_{t \in \mathbb{T}_c} A_t \leq \sum_{t=1}^T A_t$ and the last inequality is due to $\sum_{c=1}^{C_T} \sqrt{2^c} = \frac{\sqrt{2}}{\sqrt{2}-1}(\sqrt{2^{C_T}} - 1)$.

Assume C was increased at least once i.e. $C_T > 1$, otherwise we immediately have an $O\left(\sqrt{\sum_{t=1}^T A_t}\right)$ total regret bound. Let τ be the first round of the last episode. Since C was increased from $C_T - 1$ to C_T at round τ , we have $2^{C_T-1} < \ln(|\mathbb{V}_{C_T-1} \cup \mathbb{A}_\tau|)$. On the other hand, since $(\mathbb{V}_{C_T-1} \cup \mathbb{A}_\tau) \subseteq \mathbb{G}_T$, we have $\ln(|\mathbb{V}_{C_T-1} \cup \mathbb{A}_\tau|) \leq \ln(G_T)$. This implies that $2^{C_T} \leq 2 \ln(G_T)$, hence the total regret is bounded by

$$\max_{a \in [K]} \mathbb{E}[R(a)] \leq \frac{4}{(\sqrt{2}-1)^2} \sqrt{\ln(G_T) \sum_{t=1}^T A_t}.$$

□

Remark 33. While the two-level doubling trick works, for practical purposes we can set a small constant (e.g. 16) to be an upper bound for $\ln G_T$ and perform a one-level doubling trick only on $\sum_{t=1}^T A_t$. This is because $\ln G_T$ increases exponentially slowly: if the upper bound for $\ln G_T$ is doubled in each increment starting from $2^0 = 1$, then to have k such increments G_T must be as large as $G_T \geq e^{2^k}$. For $k = 5$, this is approximately 8×10^{13} , which is exceedingly large for the number of arms. Overall, this implies that the doubling level on $\ln G_T$ would change at most 4 times in any practical scenario, and setting $\ln(G_T) \leq 16$ would contribute at most a multiplicative factor of $\sqrt{16} = 4$ on the total regret bound.

G FTARL WITH NEGATIVE ENTROPY IS EQUIVALENT TO SB-EXP3

In Algorithm 2, the loss estimate of non-active arms $a \notin \mathbb{A}_t$ is $\tilde{\ell}_{a,t} = \hat{\ell}_t - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t}$. Because $I_{a,t} = 1$ for $a \in \mathbb{A}_t$ and $I_{a,t} = 0$ for $a \notin \mathbb{A}_t$, it follows that in Algorithm 2, for all $a \in [K]$,

$$\sum_{t=1}^T \left(\hat{\ell}_t - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t} - \tilde{\ell}_{a,t} \right) = \sum_{t=1}^T I_{a,t} \left(\hat{\ell}_t - \gamma \sum_{j \in \mathbb{A}_t} \tilde{\ell}_{j,t} - \tilde{\ell}_{a,t} \right) \quad (67)$$

Next, we show that in each round, the sampling probability p_t of FTARL with negative entropy is the same as that of SB-EXP3. With $\psi_t(x) = \frac{1}{\eta} \sum_{i=1}^K x_i \ln x_i$ the negative Shannon entropy, in Algorithm 2 the weight vector q_t of FTARL is the solution of the optimization problem

$$q_t = \min_{x \in \Delta_K} \frac{1}{\eta} \sum_{i=1}^K x_i \ln x_i + \sum_{i=1}^K x_i \tilde{L}_{i,t-1}.$$

Solving for q_t , we obtain for any $i \in [K]$,

$$q_{i,t} = \frac{\exp(-\eta \tilde{L}_{i,t-1})}{\sum_{k=1}^K \exp(-\eta \tilde{L}_{k,t-1})}.$$

It follows that for an active arm $i \in \mathbb{A}_t$,

$$\begin{aligned} p_{i,t} &= \frac{q_{i,t}}{\sum_{k \in \mathbb{A}_t} q_{k,t}} \\ &= \frac{\exp(-\eta \tilde{L}_{i,t-1})}{\sum_{k \in \mathbb{A}_t} \exp(-\eta \tilde{L}_{k,t-1})} \\ &= \frac{\exp(\eta \sum_{s=1}^{t-1} \hat{\ell}_s - \gamma \sum_{j \in \mathbb{A}_s} \tilde{\ell}_{j,s} - \tilde{\ell}_{i,s})}{\sum_{k \in \mathbb{A}_t} \exp(\eta \sum_{s=1}^{t-1} \hat{\ell}_s - \gamma \sum_{j \in \mathbb{A}_s} \tilde{\ell}_{j,s} - \tilde{\ell}_{k,s})} \\ &= \frac{\exp(\eta \sum_{s=1}^{t-1} I_{i,s} (\hat{\ell}_s - \gamma \sum_{j \in \mathbb{A}_s} \tilde{\ell}_{j,s} - \tilde{\ell}_{i,s}))}{\sum_{k \in \mathbb{A}_t} \exp(\eta \sum_{s=1}^{t-1} I_{k,s} (\hat{\ell}_s - \gamma \sum_{j \in \mathbb{A}_s} \tilde{\ell}_{j,s} - \tilde{\ell}_{k,s}))}, \end{aligned} \quad (68)$$

where

- the second-to-last equality is by multiplying $\exp(\eta \sum_{s=1}^{t-1} \hat{\ell}_s - \gamma \sum_{j \in \mathbb{A}_s} \tilde{\ell}_{j,s})$ to both the denominator and numerator.
- the last equality is due to (67).

Observe that (68) is equal to the sampling probability of arm $i \in \mathbb{A}_t$ computed in (6) of Algorithm 3.1. Moreover, $p_{i,t} = 0$ for $i \notin \mathbb{A}_t$ in both Algorithms 1 and 2. This implies that FTARL with negative Shannon entropy is equivalent to SB-EXP3.