

Upper Confidence Bounds for Combining Stochastic Bandits

Ashok Cutkosky
 Boston University
 Boston, MA
 ashok@cutkosky.com

Abhimanyu Das
 Google Research
 Mountain View, CA
 abhidas@google.com

Manish Purohit
 Google Research
 Mountain View, CA
 mpurohit@google.com

Abstract

We provide a simple method to combine stochastic bandit algorithms. Our approach is based on a “meta-UCB” procedure that treats each of N individual bandit algorithms as arms in a higher-level N -armed bandit problem that we solve with a variant of the classic UCB algorithm. Our final regret depends only on the regret of the base algorithm with the best regret in hindsight. This approach provides an easy and intuitive alternative strategy to the CORRAL algorithm for adversarial bandits, without requiring the stability conditions imposed by CORRAL on the base algorithms. Our results match lower bounds in several settings, and we provide empirical validation of our algorithm on misspecified linear bandit and model selection problems.

1 Introduction

This paper studies the classic *contextual bandit problem* in a stochastic setting [1, 2], which is a generalization of the even more classical multi-armed bandit problem [3]. In each of T rounds indexed by $t = 1, \dots, T$, we observe an i.i.d. random *context* $c_t \in C$, which we use to select an *action* $a_t = x_t(c_t) \in A$ based on some *policy* $x_t : C \rightarrow A$. Then, we receive a noisy reward $\hat{r}_t \in [0, 1]$, whose expectation is a function of only c_t and a_t : $\mathbb{E}[r_t | c_t, a_t] = r(c_t, a_t)$. The goal is to perform nearly as well as the best policy in hindsight by minimizing the *regret*:

$$\text{Regret} = \sum_{t=1}^T r_{\star} - r(c_t, a_t)$$

where $r_{\star} = \inf_{x \in X} \mathbb{E}_c[r(c, x(c))]$, and X is some space of possible policies.

This problem and variants has been extensively studied under diverse assumptions about the space of policies X and distributions of the rewards and values for r_{\star} . (e.g. see [4, 1, 2, 5, 6, 7]). Many of these algorithms have different behaviors in different environments (e.g. one algorithm might do much better if the reward $r(c_t, a_t)$ is a linear function of the context, while another might do better if the reward is independent of the context). This plethora of prior algorithms necessitates a “meta-decision”: If the environment is not known in advance, which algorithm should be used for the task at hand? Even in hindsight, it may not be obvious which algorithm was most optimized for the experienced environment, and so this meta-decision can be quite difficult.

We model this meta-decision by assuming we have access to N *base bandit algorithms* $\mathcal{A}_1, \dots, \mathcal{A}_N$. We will attempt to design a meta-algorithm whose regret is comparable to the best regret experienced

by any base algorithm in hindsight for the current environment. Since we don't know in advance which base algorithm will be optimal for the current environment, we need to address this problem in an online fashion. On the t th round, we will choose some index i_t and play the action suggested by the algorithm \mathcal{A}_{i_t} . The primary difficulty is that some base algorithms might perform poorly at first, and then begin to perform well later. A naive strategy might discard such algorithms based on the poor early performance, so some enhanced form of exploration is necessary for success. A pioneering prior work on this setting has considered the *adversarial* rather than stochastic case [8], and utilizes a sampling based on mirror descent with a clever mirror-map. Somewhat simplifying these results, suppose each algorithm \mathcal{A}_i guarantees regret $C_i T^{\alpha_i}$ for some C_i and α_i ¹. Given a user-specified learning rate parameter $\eta \in \mathbb{R}$, [8, 9] guarantee:

$$\text{Regret} \leq \min_{j \leq N} C_j^{\frac{1}{\alpha_j}} T \eta^{\frac{1-\alpha_j}{\alpha_j}} + T \eta + \frac{N}{\eta} \quad (1)$$

The value of η is chosen a priori, but the values of C_i need not be known. To gain a little intuition for this expression, suppose all $\alpha_i = 1/2$, and set $\eta = \frac{\sqrt{N}}{\sqrt{T}}$. Then the regret is $\min_{i \leq N} C_i^2 \sqrt{NT}$.

In this paper, we leverage the stochastic environment to avoid requiring some technical stability conditions needed in [8]: our result is a true black-box meta-algorithm that does not require any modifications to the base algorithms. Moreover, our general technique is in our view both different and much simpler. Our regret bound also improves on (1) by virtue of being non-uniform over the base algorithms: given any parameters η_1, \dots, η_N , we obtain

$$\text{Regret} \leq \min_{j \leq N} C_j^{\frac{1}{\alpha_j}} T \eta_j^{\frac{1-\alpha_j}{\alpha_j}} + T \eta_j + \sum_{i \neq j} \frac{1}{\eta_i} \quad (2)$$

This recovers (1) when all η_i are equal. In general one can think of the values $\frac{1}{\eta_i}$ as specifying a kind of *prior* over the \mathcal{A}_i which allows us to develop more delicate trade-offs between their performances. For example, consider again the setting when all $\alpha_i = 1/2$. If we believe that \mathcal{A}_1 is more likely to perform well, then by setting $\eta_1 = \frac{1}{\sqrt{T}}$ and $\eta_j = \frac{N}{\sqrt{T}}$ for $j \neq 1$, we can obtain a regret of $\min(C_1^2 \sqrt{T}, \min_{2 \leq j \leq N} C_j^2 N \sqrt{T})$.

This type of bound is in some sense a continuation of a general trend in the bandit community towards finding algorithms that adapt to various properties of the environment that are unknown in foresight (e.g [10, 11, 12, 13]). However, instead of committing to some property (e.g. large value of r_* , or small variance of \hat{r}), we instead design an algorithm that is in some sense “future-proof”, as new algorithms can be easily incorporated as new base algorithms \mathcal{A}_i .

In a recent independent work, [9] has extended the techniques of [8] to our same stochastic setting. They use a clever smoothing technique to also dispense with the stability condition required by [8], and achieve the regret bound (1). In addition to achieving the non-uniform bound (2), we also improve upon their work in two other ways: our algorithm requires only $O(N)$ space in contrast to $O(TN)$ space, and we allow for our base algorithms to only guarantee in-expectation rather than high-probability bounds.

In the stochastic setting, it is frequently possible to obtain *logarithmic* regret subject to various forms of “gap” assumptions on the actions. However, the method of [8], even when considered in the stochastic setting as in [9], seems unable to obtain such bounds. Instead, [14] has recently provided

¹in many common settings, $\alpha_i = 1/2$

a method based on UCB that can achieve such results. Their algorithm is similar to ours, but we devise a somewhat more intricate method that is able to not only obtain the results outlined previously, but also match the logarithmic regret bounds provided by [14].

In the stochastic setting, [15] introduces a new model selection technique called Regret Balancing. However this approach requires knowledge of the exact regret bounds of the optimal base algorithm. Our approach not only avoids this requirement, but also results in stronger regret guarantees than those in that paper.

We also consider an extension of our techniques to the setting of linear contextual bandits with *adversarial* features. In this case we require some modifications to our combiner algorithm and assume that the base algorithms are in fact instances of linUCB [16] or similar confidence-ellipsoid based algorithms. However, subject to these restrictions we are able to recover essentially similar results as in the non-adversarial setting, which we present in Section G.

The rest of this paper is organized as follows. In section 2 we describe our formal setup and some assumptions. In section 3 we provide our algorithm and main regret bound in a high-probability setting. In section 5, we extend our analysis to in-expectation regret bounds as well as providing an automatic tuning of some parameters in the main algorithm. In section 6 we sketch some ways in which our algorithm can be applied, and show that it matches some prior lower bound frontiers. Finally, in section 7 we provide empirical validation of our algorithm.

2 Problem Setup

We let A be a space of actions, X a space of policies, and C a space of contexts. Each policy is a function $C \rightarrow A$. (random policies can be modeled by pushing the random bits into the context). In each round $t = 1, \dots, T$ we choose $x_t \in X$, then see an i.i.d. random context c_t , and then receive a random reward $\hat{r}_t \in [0, 1]$. Let H_t denote the sequence $x_1, c_1, \hat{r}_1, \dots, \hat{r}_{t-1}, x_t, c_t$. There is an unknown function $r : C \times A \rightarrow \mathbb{R}$ such that $\mathbb{E}[\hat{r}_t | H_t] = r(c_t, x_t(c_t))$ for any H_t . The distribution of \hat{r}_t is independent of all other values conditioned on x_t and c_t . We will also write $r(x) = \mathbb{E}_{c \sim \mathcal{D}_c}[r(c, x)]$ and $r_t = r(c_t, x_t(c_t))$. Let $x_\star \in \operatorname{argmax} r(x)$. Define $r_\star = \mathbb{E}_c r(c, x_\star(c))$. Then we define the regret (often called “pseudo-regret” instead) as:

$$\text{Regret} = \sum_{t=1}^T r_\star - r_t$$

Each base bandit algorithm \mathcal{A}_i can be viewed as a randomized procedure that takes any sequence $x_1, c_1, \hat{r}_1, \dots, x_{t-1}, c_{t-1}, \hat{r}_{t-1}$ and outputs some $x_t \in X$. At the t th round of the bandit game, our algorithm will choose some index $i_t \in \{1, \dots, N\}$. Then we obtain policy $x_{T(i_t, t)}^{i_t}$ from \mathcal{A}_{i_t} , and take action $x_t = x_{T(i_t, t)}^{i_t}(c_t)$. The policy $x_{T(i_t, t)}^{i_t}$ is the output of \mathcal{A}_{i_t} on the input sequence of policies, contexts, and rewards for all the prior rounds for which we have chosen this same index i_t . After receiving the reward \hat{r}_t , we send this reward as feedback to \mathcal{A}_{i_t} .

In order to formalize this analysis more cleanly, for all i , we define T independent random variables c_1^i, \dots, c_T^i , where each $c_t^i \sim \mathcal{D}_c$. Further, we define random variables x_1^i, \dots, x_T^i and $\hat{r}_1^i, \dots, \hat{r}_T^i$ such that x_t^i is the output of \mathcal{A}_i on the input sequence $x_1^i, c_1^i, \hat{r}_1^i, \dots, x_{t-1}^i, c_{t-1}^i, \hat{r}_{t-1}^i$ and \hat{r}_t^i is a reward obtained by choosing policy x_t^i with context c_t^i . We also define random variables $r_t^i = \mathbb{E}[\hat{r}_t^i | x_1^i, c_1^i, \hat{r}_1^i, \dots, \hat{r}_{t-1}^i, c_{t-1}^i]$.

Then we can rephrase the high-level action of our algorithm as follows: We choose some index $i_t \in \{1, \dots, N\}$. Then we play policy $x_t = x_{T(i_t, t)}^{i_t}$, see context $c_t = c_{T(i_t, t)}^{i_t}$, and obtain reward $\hat{r}_t = \hat{r}_{T(i_t, t)}^{i_t}$. We define $r_t = r_{T(i_t, t)}^{i_t}$. Note that the distribution of the observed reward $\sum_{t=1}^T \hat{r}_t$ as well as the expected reward $\sum_{t=1}^T r_t$ is maintained by this description of the random variables.

2.1 Assumptions

We will consider two settings for the base algorithms in our analysis. First, a high-probability setting in which we wish to provide a regret bound that holds with probability at least $1 - \delta$ for some given δ . Second, an in-expectation setting in which we simply wish to bound the expected value of the regret. In the first setting, we require high-probability bounds on the regret of the base algorithms, while in the second setting we do not. Our approach to the in-expectation setting will be a construction that uses our high-probability algorithm as a black-box. Thus, the majority of our analysis takes place in the high-probability setting (in Section 3), for which we now describe the assumptions.

In the high-probability setting, we assume there are (known) numbers C_1, \dots, C_N and $\alpha_1, \dots, \alpha_N \in [\frac{1}{2}, 1]$, a $\delta > 0$ and an (unknown) set $S \subset \{1, \dots, N\}$ such that with probability at least $1 - \delta$, there is some $J \in S$ such that for all $t \leq T$,

$$\sum_{\tau=1}^t r_{\star} - r_{\tau}^J \leq C_J t^{\alpha_J} \quad (3)$$

Intuitively, this is saying that each algorithm \mathcal{A}_i comes with a *putative* regret bound $C_i t^{\alpha_i}$, and with high probability there is *some* \mathcal{A}_i with $i \in S$ for which its claimed regret bound is in fact correct. The assumption is stronger as S becomes smaller, and our final results will depend on the size of S . In Section F, we provide some examples of algorithms that satisfy the requirement (3). Generally, it turns out that most algorithms based the optimism principle can be made to work in this setting for any desired δ . We will refer to an index J that satisfied (3) as being “well-specified”.

To gain intuition about our results, we recommend that the reader supposes that S is a singleton $S = \{J\}$ for some unknown index J , and $\alpha_i = 1/2$ for all i (note that S may be much smaller than the set of indices for which 3 holds). As a concrete example, suppose $|X| = K$ and $c_t = 0$ for all t , so that we are playing a classic K -armed bandits problem. Let \mathcal{A}_1 be an instance of UCB, but restricted to the first 10 arms, while \mathcal{A}_2 is instance of UCB restricted to the last $K - 10$ arms. Then, we can set $C_1 = O(\sqrt{10 \log(T^2/\delta)})$ and $C_2 = O(\sqrt{(K - 10) \log(T^2/\delta)})$. Now, depending on which is the optimal arm, we may have $S = \{1\}$ or $S = \{2\}$. Although $S = \{1, 2\}$ would also satisfy the assumptions, since our results improve when S is smaller and S is unknown to the algorithm, we are free to choose the smallest possible S .

These assumptions may seem stronger than prior work at first glance: not only do we require high-probability rather than in-expectation regret bounds, we require at least one base algorithm to be well-specified, and we require knowledge of the putative regret bounds through the coefficients C_i and α_i . Prior work in this setting (e.g. [8, 9]) dispenses with the last two requirements, and [8] also dispenses with the high-probability requirement. However, it turns out that through a simple doubling construction, we can easily incorporate base algorithms with unknown regret bounds that hold only in expectation into our framework. The ability to handle unknown regret bounds dispenses with the well-specified assumption. We describe this construction and the relevant assumptions in Section 5, and show that it only weakens our analysis by log factors.

3 UCB over Bandits

In this Section, we describe our meta-learner for the high-probability setting. The intuition is based on upper-confidence bounds: first, we observe that the unknown well-specified algorithm \mathcal{A}_J 's rewards \hat{r}_τ^J behave very similarly to independent bounded random variables with mean r_\star in that their average value $\hat{\mu}_t^J = \frac{1}{t} \sum_{\tau=1}^t \hat{r}_\tau^J$ concentrates about r_\star with radius $C_J t^{\alpha_J - 1}$. From this, one might imagine that for any index i , the value of $\hat{\mu}_t^i + C_i t^{\alpha_i - 1}$ gives some kind of upper-confidence bound for the final total reward of algorithm i . We could then feed these estimates into a UCB-style algorithm that views the N base algorithms as N arms. Unfortunately, such an approach is complicated by two issues. First, the putative regret bounds for each \mathcal{A}_i may not actually hold, which could damage the validity of our confidence estimates. Second, the confidence bounds for different algorithms may have very imbalanced behavior due to the different values of C_i and α_i , and we would like our final regret bound to depend only on C_J and α_J .

We address the first issue by keeping track of the statistic $\sum_{\tau=1}^t \hat{\mu}_\tau^i - \hat{r}_\tau^i$. If $\sum_{\tau=1}^t \hat{\mu}_\tau^i - \hat{r}_\tau^i \geq C_i t^{\alpha_i}$ at any time, then we can conclude that i is not the well-specified index J and so we simply discard \mathcal{A}_i . Moreover, it turns out that so long as $\sum_{\tau=1}^t \hat{\mu}_\tau^i - \hat{r}_\tau^i \leq C_i t^{\alpha_i}$, the rewards \hat{r}_τ^i are “well-behaved” enough that our meta-UCB algorithm can operate correctly.

We address the second issue by employing *shifted confidence intervals* in a construction analogous to that employed by [17], who designed a K -armed bandit algorithm with a regret bound that depends on the identity of the best arm. Essentially, each algorithm \mathcal{A}_i is associated with a *target regret bound*, R_i , and each confidence interval is decreased by $\frac{R_i}{T}$. Assuming R_i satisfies some technical conditions, this will guarantee that the regret is at most $O(R_J)$ for any J such that \mathcal{A}_J is well-specified.

Formally, our algorithm is provided in Algorithm 1, and its analysis is given in Theorem 1, proved in Appendix B. Note that little effort has been taken to improve the constants or log factors.

Algorithm 1 Bandit Combiner

Input: Bandit algorithms $\mathcal{A}_1, \dots, \mathcal{A}_N$, numbers C_1, \dots, C_N , $\alpha_1, \dots, \alpha_N$, R_1, \dots, R_N , T .

Set $T(i, 0) = 0$ for all i , set $\hat{\mu}_0^i = 0$ for all i , and set $I_1 = \{1, \dots, N\}$

for $t = 1 \dots T$ **do**

Set $U(i, t-1) = \hat{\mu}_{T(i,t-1)}^i + \min \left(1, \frac{C_i T(i,t-1)^{\alpha_i} + \sqrt{8 \log(T^3 N / \delta) T(i,t-1)}}{T(i,t-1)} \right) - \frac{R_i}{T}$ for all i .

Set $i_t = \operatorname{argmax}_{i \in I_t} U(i, t-1)$.

Update $T(i_t, t) = T(i_t, t-1) + 1$ and $T(j, t) = T(j, t-1)$ for $j \neq i_t$.

Get $T(i_t, t)$ th policy $x_t = x_{T(i_t, t)}^{i_t}$ from \mathcal{A}_{i_t} . See context c_t and play action $x_t(c_t)$.

Receive reward $\hat{r}_t = \hat{r}_{T(i_t, t)}^{i_t}$, provide reward \hat{r}_t and context c_t as feedback to \mathcal{A}_{i_t} .

Update $\hat{\mu}_{T(i_t, t)}^{i_t} = \frac{1}{T(i_t, t)} \sum_{\tau=1}^{T(i_t, t)} \hat{r}_\tau^{i_t}$.

if $\sum_{\tau=1}^{T(i_t, t)} \hat{\mu}_{\tau-1}^{i_t} - \hat{r}_\tau^{i_t} \geq C_{i_t} T(i_t, t)^{\alpha_{i_t}} + 3\sqrt{\log(T^3 N / \delta) T(i_t, t)}$ **then**

$I_t = I_{t-1} - \{i_t\}$.

else

$I_t = I_{t-1}$.

end if

end for

Theorem 1. Suppose there is a set $S \subset \{1, \dots, N\}$ such that with probability at least $1 - \delta$, there is some $J \in S$ such that

$$\sum_{\tau=1}^t r_{\star} - r_{\tau}^J \leq C_J t^{\alpha_J}$$

for all $t \leq T$. Further, suppose the C_i and α_i are known, and the R_i satisfy:

$$R_i \geq C_i T^{\alpha_i}$$

$$R_i \geq \sum_{k \neq i} \max \left[\frac{(1 - \alpha_k)(1 + \alpha_k)^{\frac{1}{1 - \alpha_k}} (2C_k)^{\frac{1}{1 - \alpha_k}} T^{\frac{\alpha_k}{1 - \alpha_k}}}{\alpha_k R_k^{\frac{\alpha_k}{1 - \alpha_k}}}, \frac{288 \log(T^3 N / \delta) T}{R_k} \right]$$

Let $r_t = \mathbb{E}[\hat{r}_{T(i_t, t)}^{i_t}]$ be the expected reward of Algorithm 1 at time t . Then, with probability at least $1 - 3\delta$, the regret satisfies:

$$\sum_{t=1}^T r_{\star} - r_t \leq 3 \sup_{j \in S} R_j$$

Note that the algorithm does not know the set S .

The conditions on R_i in this Theorem are somewhat opaque, so to unpack this a bit we provide the following corollary:

Corollary 2. Suppose there is a set $S \subset \{1, \dots, N\}$ such that with probability at least $1 - \delta$, there is some $J \in S$ such that $\sum_{\tau=1}^t r_{\star} - r_{\tau}^J \leq C_J t^{\alpha_J}$ for all $t \leq T$. Further, suppose we are given N positive real numbers η_1, \dots, η_N . Set R_i via:

$$R_i = C_i T^{\alpha_i} + \frac{(1 - \alpha_i)^{\frac{1 - \alpha_i}{\alpha_i}} (1 + \alpha_i)^{\frac{1}{\alpha_i}}}{\alpha_i^{\frac{1 - \alpha_i}{\alpha_i}}} C_i^{\frac{1}{\alpha_i}} T \eta_i^{\frac{1 - \alpha_i}{\alpha_i}}$$

$$+ 288 \log(T^3 N / \delta) T \eta_i + \sum_{k \neq i} \frac{1}{\eta_k}$$

Then, with probability at least $1 - 3\delta$, the regret of Algorithm 1 satisfies:

$$\text{Regret} \leq 3 \sup_{j \in S} R_j = \tilde{O} \left(\sup_{j \in S} C_j T^{\alpha_j} + C_j^{\frac{1}{\alpha_j}} T \eta_j^{\frac{1 - \alpha_j}{\alpha_j}} + T \eta_j + \sum_{k \neq j} \frac{1}{\eta_k} \right)$$

In most settings, $C_J \geq 1$ and $\eta \geq T^{-\alpha}$, so that $C_J T^{\alpha_J}$ is smaller than $C_J^{\frac{1}{\alpha_J}} T \eta_J^{\frac{1 - \alpha_J}{\alpha_J}}$.

Proof sketch of Theorem 1. While the full proof of Theorem 1 is deferred to the appendix, we sketch the main ideas here. For simplicity, we consider the case that $\alpha_i = \frac{1}{2}$ for all i , assume that $S = \{J\}$ is a singleton set, assume $C_i \geq 1$, and drop all log factors and constants. Then by some martingale concentration bounds combined with the high-probability regret bound on \mathcal{A}_J , we have that $r_\star \leq U(J, t-1) + \frac{R_J}{T}$ for all t with high probability. Furthermore, by martingale concentration again, we have that $\hat{\mu}_t^i \leq r_\star + \sqrt{t}$. Therefore, an algorithm is dropped from the set I_t only if $\sum_{\tau=1}^{T(i,t)} r_\star - r_\tau^i \geq C_i \sqrt{T(i,t)}$, which does not happen for algorithm \mathcal{A}_J with probability at least $1 - \delta$. Further, by definition of I_t and another martingale bound, all algorithms i satisfy $\sum_{\tau=1}^{T(i,T)} \hat{\mu}_{\tau-1}^i - r_\tau^i \leq C_i \sqrt{T(i,T)} + \sqrt{T(i,T)} \leq 2C_i \sqrt{T(i,T)}$ with high probability. Let us consider the instantaneous regret $r_\star - r_t$ on some round in which $i_t \neq J$. In this case, we must have $U(i_t, t-1) \geq U(J, t-1)$, so that we can write:

$$\begin{aligned} r_\star - r_t &= r_\star - r_{T(i_t,t)}^{i_t} \\ &= r_\star - U(J, t-1) + U(J, t-1) - U(i_t, t-1) \\ &\quad + U(i_t, t-1) - r_{T(i_t,t)}^{i_t} \\ &\leq \frac{R_J}{T} + \hat{\mu}_{T(i_t,t-1)}^{i_t} - r_{T(i_t,t)}^{i_t} - \frac{R_{i_t}}{T} \\ &= \frac{R_J}{T} + \hat{\mu}_{T(i_t,t)-1}^{i_t} - r_{T(i_t,t)}^{i_t} - \frac{R_{i_t}}{T} \end{aligned}$$

Summing over all timesteps for which algorithm i is chosen, we have

$$\begin{aligned} \sum_{i_t=i} r_\star - r_t &\leq \frac{R_J T(i, T)}{T} + 2C_i \sqrt{T(i, T)} - \frac{R_i T(i, T)}{T} \\ &\leq \frac{R_J T(i, T)}{T} + \sup_{Z \geq 0} 2C_i \sqrt{Z} - \frac{R_i Z}{T} \\ &\leq \frac{R_J T(i, T)}{T} + \frac{C_i^2 T}{R_i} \end{aligned}$$

Now summing over all indices $i \neq J$, we use the fact that $\sum_i T(i, T) \leq T$ and the assumption that $\sum_{i \neq J} \frac{C_i^2 T}{R_i} \leq O(R_J)$ to conclude that the regret over all rounds in which \mathcal{A}_J is not chosen is at most $O(R_J)$. For the rounds in which \mathcal{A}_J is chosen, we experience regret $C_J \sqrt{T(J, T)} \leq C_J \sqrt{T} \leq R_J$, which concludes the Theorem. \square

4 Gap-dependent regret bounds

In this section, we provide an analog of the standard ‘‘gap-dependent’’ bound for UCB. As a motivating example, consider the setting in which all \mathcal{A}_i for $i \neq J$ never play any policy x with $r(x) \geq r_\star - \Delta_i$ for some $\Delta_i \geq 0$. In this case, we might hope to perform much better, in the same way that standard UCB obtains logarithmic regret when the suboptimal arms have a non-negligible gap between their rewards and the optimal rewards. Specifically, we have the following result, whose proof is deferred to Section C:

Theorem 3. Suppose that there is some $J \in \{1, \dots, N\}$ such that with probability at least $1 - \delta$, we have:

$$\sum_{\tau=1}^t r_{\star} - r_{\tau}^J \leq C_J t^{\alpha_J}$$

Also, for all i , define $T_i = T(i, T)$ and Δ_i by:

$$\Delta_i = \frac{1}{T(i, T_i - 1)} \sum_{\tau=1}^{T(i, T_i - 1)} r_{\star} - r_{\tau}^i$$

And let $B \subset \{1, \dots, N\}$ with $J \notin B$ be the set of indices with $\Delta_i > \frac{2R_J}{T}$ for $i \in B$.

For $i \neq J$, let $C_i > 0$ and $\alpha_i \leq 1$ for $i \neq J$ be arbitrary. Then with probability at least $1 - 3\delta$, the regret of Algorithm 1 satisfies:

$$\begin{aligned} \sum_{t=1}^T r_{\star} - r_t \leq & \sum_{\tau=1}^{T(J, T)} r_{\star} - r_{\tau}^J + \sum_{i \in B} 1 + \frac{512 \log(T^3 N / \delta)}{\Delta_i} + \frac{4 \frac{1}{1-\alpha_i} C_i^{\frac{1}{1-\alpha_i}}}{\Delta_i^{\frac{\alpha_i}{1-\alpha_i}}} \\ & + \min \left[\sum_{k \neq J, k \notin B} 2R_J, R_J + \sum_{k \neq J, k \notin B} \max \left[\frac{(1-\alpha_k)(1+\alpha_k)^{\frac{1}{1-\alpha_k}} (2C_k)^{\frac{1}{1-\alpha_k}} T^{\frac{\alpha_k}{1-\alpha_k}}}{\alpha_k R_k^{\frac{\alpha_k}{1-\alpha_k}}}, \frac{288 \log(T^3 N / \delta) T}{R_k} \right] \right] \end{aligned}$$

Note that we have made no conditions on R_i in this expression. In particular, consider the case that each algorithm \mathcal{A}_i considers only a subset of the possible policies, and that \mathcal{A}_J is the only algorithm that is allowed to choose the optimal policy with reward r_{\star} . Then Δ_i is at least the gap between the reward of the best policy available to \mathcal{A}_i and r_{\star} . Thus for large enough T , B will be all indices except J , so that the overall regret provided in Theorem 3 is $\sum_{\tau=1}^{T(J, T)} r_{\star} - r_{\tau}^J + \tilde{O} \left(\sum_{i \neq J} \frac{C_i^{\frac{1}{1-\alpha_i}}}{\Delta_i^{\frac{\alpha_i}{1-\alpha_i}}} \right)$.

As another example of this Theorem in action, let us consider the setting studied by [14]. Specifically, each \mathcal{A}_i has a putative regret bound of $\sum_{\tau=1}^t r_{\star} - r_{\tau}^i \leq \sqrt{k_i \log(t)t}$ for all $t \leq T$ for some k_i , and \mathcal{A}_J in fact obtains its bound. However, for all $i \neq J$, \mathcal{A}_i also suffers $\sum_{\tau=1}^t r_{\star} - r_{\tau}^i \geq \Delta_i t$ for all t for some constant Δ_i . For example, this might occur if each \mathcal{A}_i is restricted to some subset of actions that does not include the best action. Now, recall that we made no restrictions of R_i in Theorem 3, so we are free to set $R_i = 0$ for all i . Then, we will have $B = \{1, \dots, J-1, J+1, N\}$ and obtain the following Corollary:

Corollary 4. Suppose k_1, \dots, k_N are such that for some J , \mathcal{A}_J guarantees $\sum_{\tau=1}^t r_{\star} - r_{\tau}^J \leq \sqrt{k_J \log(t)t}$ for all $t \leq T$. Further, suppose that for all $i \neq J$, $\sum_{\tau=1}^t r_{\star} - r_{\tau}^i \geq \Delta_i t$ for all t for some constant Δ_i . Then with $C_i = \sqrt{k_i \log(T)}$, $\alpha_i = \frac{1}{2}$ and $R_i = 0$, with probability at least $1 - 3\delta$, Algorithm 1 guarantees regret:

$$\sum_{t=1}^{T(J, T)} r_{\star} - r_t^J + \sum_{i \neq J} \frac{512 \log(T^3 N / \delta)}{\Delta_i} + \frac{16k_i \log(T)}{\Delta_i}$$

Notably, the first term is the *actual* regret of if \mathcal{A}_J rather than the regret *bound* $C_J \sqrt{T}$. Thus if \mathcal{A}_J outperforms this bound and obtains logarithmic regret, our combiner algorithm will also obtain logarithmic regret, which is not obviously possible using techniques based on the Corral algorithm [8]. Note that this result also appears to improve upon [14] (Theorem 4.2) by removing a $\log(T)$ factor, but this is because we have assumed knowledge of the time horizon T in order to set C_i .

5 Unknown and In-Expectation Bounds on Base Algorithms

In this section, we show how to remedy two surface-level issues with Algorithm 1. First, we require knowledge of the values C_i and α_i . Second, we require a high-probability regret bound for the well-specified base algorithm \mathcal{A}_J . Here, we show that a simple duplication and doubling-based technique suffices to address both issues.

First, let us gain some intuition for how to convert an in-expectation bound into a high-probability bound suitable for use in Theorem 1. Suppose we are given an algorithm that maintains expected regret $C_i T^{\alpha_i}$. We duplicate this algorithm $M = O(\log_2(1/\delta))$ times. Then by Markov inequality, each individual duplicate obtains regret at most $2C_i T^{\alpha_i}$ with probability at least $1/2$. Therefore, with probability at least $1 - \frac{1}{2^M} = 1 - \delta$, at least one of the duplicates obtains regret at most $2C_i T^{\alpha_i}$. Then in the terminology of Theorem 1, we let S be the set of duplicate algorithms and so we satisfy the hypothesis of the Theorem. This argument is slightly flawed as-is because we need an *anytime* regret bound for the base algorithms, but it turns out this is fixable by another use of Markov and union bound inequality. We then use a variant on the doubling trick to avoid requiring knowledge of C and α . The full construction is described below in Theorem 5, with proof in Appendix D.

Theorem 5. *Suppose that for some J , there is some unknown \bar{C}_J and $\bar{\alpha}_J$ such that \mathcal{A}_J ensures $\mathbb{E}[\sum_{\tau=1}^t r_\star - r_\tau^J] \leq \bar{C}_J t^{\bar{\alpha}_J}$ for all $t \leq T$. Further, suppose we are given positive numbers η_1, \dots, η_N . Let $\delta \in (0, 1)$ be some user-specified failure probability. For $M = \lceil \log_2(1/\delta) \rceil$ and $K = \lceil \log_2(T) \rceil$ and $L = \lceil \frac{\log_2(T)}{2} \rceil$, we duplicate each \mathcal{A}_i MKL times, specifying each duplicate $\mathcal{A}_{i,x,y,z}$ by a multi-index $(i, x, y, z) \in [N] \times [M] \times [K] \times [L]$. To each duplicate we associate $C_{i,x,y,z} = 2^y$ and $\alpha_{i,x,y,z} = \min\left(1, 1/2 + \frac{z}{\log(T)}\right)$. Let $\eta_{i,x,y,z} = \eta_i$. Specify $R_{i,x,y,z}$ as a function of $\eta_{i,x,y,z}$ as described in Theorem 2. Then with probability at least $1 - 3\delta$, Algorithm 1 guarantees regret:*

$$\begin{aligned} \text{Regret} \leq O & \left(\bar{C}_J T^{\bar{\alpha}_J} + \bar{C}_J^{\frac{1}{\bar{\alpha}_J}} T^{\bar{\alpha}_J} \eta_J^{\frac{1-\bar{\alpha}_J}{\bar{\alpha}_J}} \right. \\ & \left. + \log\left(\frac{T^3 N}{\delta}\right) T \eta_J + \sum_{i=1}^N \frac{\log(1/\delta) \log^2(T)}{\eta_i} \right) \end{aligned}$$

so that the expected regret is bounded by:

$$\begin{aligned} \mathbb{E}[\text{Regret}] \leq O & \left(\bar{C}_J T^{\bar{\alpha}_J} + \bar{C}_J^{\frac{1}{\bar{\alpha}_J}} T^{\bar{\alpha}_J} \eta_J^{\frac{1-\bar{\alpha}_J}{\bar{\alpha}_J}} \right. \\ & \left. + \log\left(\frac{T^3 N}{\delta}\right) T \eta_J + \sum_{i=1}^N \frac{\log(1/\delta) \log^2(T)}{\eta_i} + T\delta \right) \end{aligned}$$

6 Examples and Optimality

In this section, we provide some illustrative examples of how our approach can be used. We will also highlight a few examples in which our construction matches lower bounds. The proofs are straightforward applications of Theorems 1 and 2, and are deferred to Appendix E.

6.1 K -Armed Bandits

For our first example, suppose that the space A is a finite set of K arms and X consists of the K constant functions mapping all contexts to a single arm. This setup describes the classic K -armed bandit problem. Let $N = K$ and suppose each \mathcal{A}_i is a naive algorithm that simply pulls arm i on every round. We consider a \mathcal{A}_i to be well-specified if the i th arm is in fact the optimal arm, in which case it is clear we may set $C_i = 0$, $\alpha_i = \frac{1}{2}$ for all i . In the high-probability setting, we let S be the singleton set containing only the unknown optimal index. In this case, the conditions on R_i of Theorem 1 correspond almost exactly (up to constants and log factors) with the pareto frontier for regret bounds described in [17], showing that using our construction in this setting allows us to match this lower bound frontier.

6.2 Misspecified Linear Bandit

For our second example, suppose that the space A is a finite set of K arms, and that the context c_t is a constant $c_t = c$ and provides a feature $c(a) \in \mathbb{R}^d$ for each arm $a \in A$. The space of policies is the set of K constant functions again. In this case, it is possible the reward $r(c, a)$ is a fixed linear function $\langle \beta, c(a) \rangle$ for some $\beta \in \mathbb{R}^d$, in which case the linUCB algorithm [16] can obtain regret $\tilde{O}(\sqrt{d \log(K)t})$. On the other hand, in general the reward might be totally unrelated to the context, in which case one might wish to fall back on the UCB algorithm which obtains regret $\tilde{O}(\sqrt{Kt})$. By setting \mathcal{A}_1 to be linUCB and \mathcal{A}_2 to be ordinary UCB, we say that $S = \{1\}$ if the rewards are indeed linear, and $S = \{2\}$ otherwise. Further, we set $C_1 = \sqrt{d \log(K)}$, $C_2 = \sqrt{K}$ and $\alpha_1 = \alpha_2 = \frac{1}{2}$. Now let P and Q be any two numbers such that $PQ = KT$ and both P and Q are greater than $\sqrt{d \log(k)T}$. Then appropriate application of Corollary 2, yields regret $O(P)$ in the linear setting and regret $O(Q)$ in general. This again matches the frontier of regret bounds for this scenario described in Theorem 24.4 of [5] (see also Lemma 6.1 of [9]). Formally, we have the following Corollary:

Corollary 6. *Suppose $c_t : A \rightarrow \mathbb{R}^d$ for all t . Suppose $K \geq d \log(K)$. Let \mathcal{A}_1 be an instance of linUCB and \mathcal{A}_2 be an instance of the ordinary UCB algorithm. Let P and Q be any two numbers such that $PQ = KT$ and both are greater than $\sqrt{d \log(K)T}$. We consider two cases, either the reward is a linear function of c_t , or it is not. Then \mathcal{A}_1 guarantees regret $\tilde{O}(\sqrt{d \log(K)t})$ with probability $1 - \delta$ in the first case, while \mathcal{A}_2 guarantees regret $\tilde{O}(\sqrt{Kt})$ in the second case. Set $\eta_1 = \frac{P}{d \log(K)T} = \frac{K}{Qd \log(K)}$ and $\eta_2 = \frac{1}{P} = \frac{Q}{KT}$. Then using the R_i construction of Corollary 2, with probability at least $1 - 3\delta$, we guarantee regret $\tilde{O}(P)$ with linear rewards, and $\tilde{O}(Q)$ otherwise.*

Note that we leverage our ability to use non-uniform η values in this Corollary. It is not so obvious how to obtain this full frontier using the prior uniform bound (1), although it is of course conceivable that more detailed analysis of prior algorithms might allow for this same result.

6.3 Linear Model Selection

For our third example, we consider the case of model selection for linear bandits. In this setting, the context c_t again specifies features $c_t(a) \in \mathbb{R}^d$ for each arm $a \in A$, and we are guaranteed that the reward is a linear function of the context. The question now is whether the full d -dimensions are actually necessary. Specifically, if there is some d_\star such that the reward is in fact a linear function of the first d_\star coordinates of the context only, then we would like our regret to depend on d_\star rather than d . This setting has been studied before in the context of a finite set of actions in [18, 19]. These

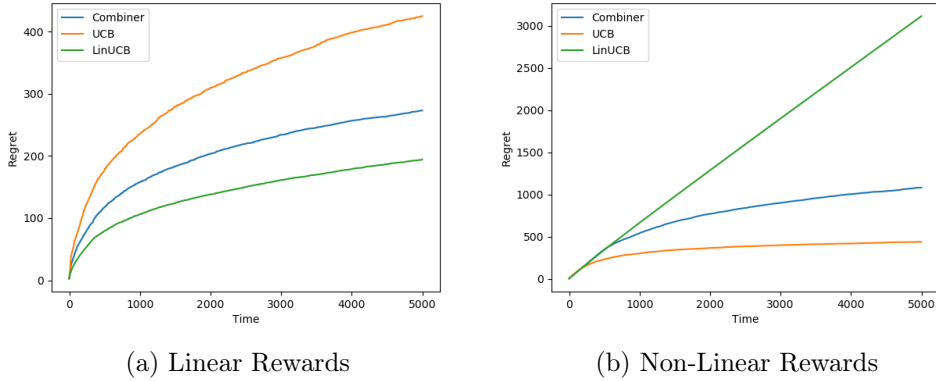


Figure 1: Misspecified Linear Bandit

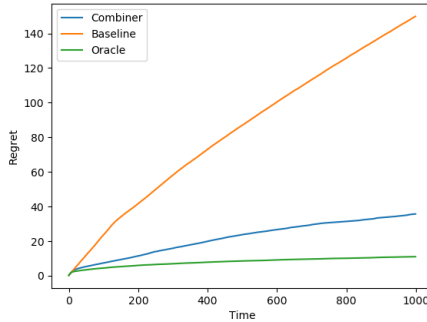


Figure 2: Model Selection Experiments

prior works impose some additional technical conditions on the distribution of rewards and contexts provided by the environment. Under their conditions, [19] obtains regret $\tilde{O}(\sqrt{d_*T})$ while under somewhat weaker conditions, [18] obtains regret $\tilde{O}(\sqrt{d_*T} + T^{3/4})$. In contrast, we require no extra conditions, and obtain regret $\tilde{O}(d_*\sqrt{T})$. The construction is detailed in the following Corollary:

Corollary 7. *Suppose $c_t \in \mathbb{R}^d$ for all t and the reward is always a linear function of c_t . Suppose that the reward is in fact purely a linear function of the first d_* coordinates of c_t . Suppose the action set A has finite cardinality K . Let \mathcal{A}_i be an instance of *linUCB* of [16] restricted to the first 2^i coordinates of the context. Set $C_i = \sqrt{2^i \log(K)}$, $\alpha_i = \frac{1}{2}$, and $\eta_i = \frac{1}{\sqrt{T}}$. Then using the instantiation of Algorithm 1 from Theorem 2, we obtain regret $\tilde{O}(d_* \log(K) \sqrt{T})$. If instead the set A is infinite, let \mathcal{A}_i be an instance of the *linUCB* algorithm for infinite arms [20, 21] restricted to the first 2^i coordinates. Set $C_i = 2^i$, $\alpha_i = \frac{1}{2}$, and $\eta_i = \frac{1}{\sqrt{T}}$. Then we obtain regret $\tilde{O}(d_*^2 \sqrt{T})$.*

7 Experimental Validation

We now demonstrate empirical validation of our results in two different application settings. For our first experiment, we consider the misspecified linear bandit setting. We use ordinary UCB and

linUCB as the two base algorithms. For simplicity we focus on the ordinary stochastic bandits framework (i.e. we assume the context remains fixed over time). Each arm $a \in [K]$ is associated with a feature vector $x_a \in \mathbb{R}^d$ that is chosen from the uniform distribution on the unit sphere. Let $\beta \in \mathbb{R}^d$ also chosen from the unit sphere be a fixed unknown parameter vector. Finally, for each arm a , we choose $\mu_a \in \mathbb{R}$ to be specified later. The reward for arm a at any time step t is set to be $\alpha\mu_a + (1 - \alpha)\sqrt{d} \cdot \langle \beta, x_a \rangle + \eta_t$ where $\eta_t \sim \mathcal{N}(0, \sigma)$ is independently sampled noise and $0 \leq \alpha \leq 1$ is a fixed constant. Let $a^* = \operatorname{argmin}_{a \in [K]} \langle \beta, x_a \rangle$ be the worst arm with respect to the linear component of the reward. For any arm $a \neq a^*$, we set $\mu_a = 0.25\sqrt{d} \cdot \langle \beta, x_a \rangle$, whereas we set $\mu_{a^*} = 1$. We consider two different settings of α . When $\alpha = 0$, the reward is simply a linear function of the arms and we expect the linUCB algorithm to outperform the ordinary UCB algorithm. On the other hand, when $\alpha = 1$, the rewards are constructed so that the linUCB algorithm essentially never chooses the arm a^* and thus incurs linear regret whereas the ordinary UCB algorithm still guarantees $O(\sqrt{T})$ regret. Figures 1a and 1b show the performance on these two settings respectively. In both settings, the combiner uses UCB and linUCB as the base algorithms and the putative regret bounds for UCB and linUCB are computed empirically on independent instances of the non-linear and linear reward settings respectively.

For our second experiment, we consider the model selection problem in linear bandits. As earlier, each arm $a \in K$ is associated with a feature vector x_a sampled independently from the unit sphere. The parameter vector $\beta \in \mathbb{R}^d$ is chosen so that first each $\beta_j \sim \mathcal{N}(0, 1)$, $\forall j \leq d^*$ and $\beta_j = 0$, otherwise, and then normalized to be of unit length. The reward for arm a at any time step t is set to be $\langle \beta, x_a \rangle + \eta_t$ where as earlier $\eta_t \sim \mathcal{N}(0, \sigma)$ is independently sampled noise. Figure 2 demonstrates the performance of three different algorithms in this setting with $d = 128$, $d^* = 8$ and $K = 1000$. The *Baseline* algorithm is a vanilla LinUCB algorithm that works on the ambient dimension d , while the *Oracle* is a LinUCB algorithm that works on the true dimensionality of the reward parameter d^* . The *Combiner* algorithm is an implementation of Algorithm 1 with $\log d$ base algorithms where each base algorithm A_i is an instance of LinUCB restricted to the first 2^i dimensions of the features. The putative regret bound C_i for each base algorithm A_i is computed empirically on independent instances where the corresponding algorithm is well-specified. We set the target regret bound $R_i \leftarrow (C_i^2 + N)\sqrt{T}$ where $N = \log d$ is the number of base algorithms. The experimental results validate our theoretical findings and show that the combiner algorithm is able to adapt to the true dimensionality of the rewards.

8 Conclusion

We have introduced a new method for combining stochastic bandit algorithms in such a way that our final regret may depend only on the regret of the best base algorithm in hindsight. Our method is based on upper-confidence techniques, and provides a contrast to prior work based on mirror descent [8]. We verify empirically that our technique can be used to solve some model selection and misspecification problems. In the future, we hope to see advancements in this area. For example, can we maintain logarithmic regret in benign settings? Further, in the full-information setting, one can often combine algorithms with minimal overhead. Are there any bandit settings in which such ideal behavior is possible?

References

- [1] John Langford and Tong Zhang. The epoch-greedy algorithm for multi-armed bandits with side information. In *Advances in neural information processing systems*, pages 817–824, 2008.
- [2] Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandit algorithms with supervised learning guarantees. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 19–26, 2011.
- [3] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- [4] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- [5] Tor Lattimore and Csaba Szepesvári. Bandit algorithms. *preprint*, page 28, 2018.
- [6] Ambuj Tewari and Susan A Murphy. From ads to interventions: Contextual bandits in mobile health. In *Mobile Health*, pages 495–517. Springer, 2017.
- [7] Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *International Conference on Machine Learning*, pages 1638–1646, 2014.
- [8] Alekh Agarwal, Haipeng Luo, Behnam Neyshabur, and Robert E Schapire. Corraling a band of bandit algorithms. In *Conference on Learning Theory*, pages 12–38, 2017.
- [9] Aldo Pacchiano, My Phan, Yasin Abbasi-Yadkori, Anup Rao, Julian Zimmert, Tor Lattimore, and Csaba Szepesvari. Model selection in contextual stochastic bandit problems. *arXiv preprint arXiv:2003.01704*, 2020.
- [10] Sébastien Bubeck and Aleksandrs Slivkins. The best of both worlds: Stochastic and adversarial bandits. In *Conference on Learning Theory*, pages 42–1, 2012.
- [11] Chen-Yu Wei and Haipeng Luo. More adaptive algorithms for adversarial bandits. *Proceedings of Machine Learning Research*, 75, 2018.
- [12] Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):1876–1902, 2009.
- [13] Avishek Ghosh, Abishek Sankararaman, and Kannan Ramchandran. Problem-complexity adaptive model selection for stochastic linear bandits. *arXiv preprint arXiv:2006.02612*, 2020.
- [14] Raman Arora, Teodor V Marinov, and Mehryar Mohri. Corraling stochastic bandit algorithms. *arXiv preprint arXiv:2006.09255*, 2020.
- [15] Yasin Abbasi-Yadkori, Aldo Pacchiano, and My Phan. Regret balancing for bandit and rl model selection. *arXiv preprint arXiv:2006.05491*, 2020.

- [16] Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 208–214, 2011.
- [17] Tor Lattimore. The pareto regret frontier for bandits. In *Advances in Neural Information Processing Systems*, pages 208–216, 2015.
- [18] Dylan J Foster, Akshay Krishnamurthy, and Haipeng Luo. Model selection for contextual bandits. In *Advances in Neural Information Processing Systems*, pages 14714–14725, 2019.
- [19] Niladri S Chatterji, Vidya Muthukumar, and Peter L Bartlett. Osom: A simultaneously optimal algorithm for multi-armed and linear contextual bandits. *arXiv preprint arXiv:1905.10040*, 2019.
- [20] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320, 2011.
- [21] Varsha Dani, Thomas P. Hayes, and Sham M. Kakade. Stochastic linear optimization under bandit feedback. In *COLT*, 2008.
- [22] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- [23] Aleksandrs Slivkins. Introduction to multi-armed bandits. *Foundations and Trends in Machine Learning*, 12(1-2):1–286, 2019.
- [24] Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Online-to-confidence-set conversions and application to sparse stochastic bandits. In *Artificial Intelligence and Statistics*, pages 1–9, 2012.
- [25] H Brendan McMahan. A survey of algorithms and analysis for adaptive online learning. *The Journal of Machine Learning Research*, 18(1):3117–3166, 2017.
- [26] Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.

A Lemmas for analysis of Algorithm 1

Before proving Theorem 1, we need a few Lemmas. For the most part, these are straightforward verification of intuitive concentration bounds.

Lemma 8. *With probability at least $1 - \delta/T \geq 1 - \delta$, for all i , for all t , we have:*

$$\sum_{\tau=1}^{T(i,t)} r_{\tau}^i - \hat{r}_{\tau}^i \leq \sqrt{8 \log(T^3 N/\delta) T(i,t)}$$

Similarly, with probability at least $1 - \delta$, for all i , for all t , we have:

$$\sum_{\tau=1}^{T(i,t)} \hat{r}_{\tau}^i - r_{\tau}^i \leq \sqrt{8 \log(T^3 N/\delta) T(i,t)}$$

Proof. We prove only the first statement. The proof of the second is entirely symmetrical.

Fix a given i . For any $\tau \leq T$ define $z_{\tau} = r_{\tau}^i - \hat{r}_{\tau}^i$, we recall that r_{τ}^i is the expected value of the τ th action of A_i given the past history - if t is the first time that $T(i,t) = \tau$, we have $r_{\tau}^i = \mathbb{E}[\hat{r}_{\tau}^i | \hat{r}_1, \dots, \hat{r}_{t-1}]$. Notice that $z_{\tau} \in [-1, 1]$ is a martingale difference sequence.

Now for any ϵ we apply Azuma's inequality to obtain for any t :

$$\begin{aligned} \text{Prob} \left[\sum_{\tau=1}^t r_{\tau}^i - \hat{r}_{\tau}^i \geq \epsilon \right] &= \text{Prob} \left[\sum_{\tau=1}^t z_{\tau} \geq \epsilon \right] \\ &\leq \exp \left(\frac{-\epsilon^2}{8t} \right) \end{aligned}$$

In particular, setting $\epsilon = \sqrt{8t \log(T^3 N/\delta)}$, we have that

$$\text{Prob} \left[\sum_{\tau=1}^t r_{\tau}^i - \hat{r}_{\tau}^i \geq \sqrt{8 \log(T^3 N/\delta) t} \right] \leq \frac{\delta}{T^3 N}$$

Therefore:

$$\begin{aligned} \text{Prob} \left[\sum_{\tau=1}^{T(i,t)} r_{\tau}^i - \hat{r}_{\tau}^i \geq \sqrt{8 \log(T^3 N/\delta) T(i,t)} \right] &= \sum_{t'=1}^T \text{Prob} \left[\sum_{\tau=1}^{t'} r_{\tau}^i - \hat{r}_{\tau}^i \geq \sqrt{8 \log(T^3 N/\delta) t'} \text{ and } T(i,t) = t' \right] \\ &\leq \sum_{t'=1}^T \text{Prob} \left[\sum_{\tau=1}^{t'} r_{\tau}^i - \hat{r}_{\tau}^i \geq \sqrt{8 \log(T^3 N/\delta) t'} \right] \\ &\leq \frac{\delta}{NT^2} \end{aligned}$$

Then union bound over the N indices i and T values of t completes the Lemma. \square

Similarly, we have the following result:

Lemma 9. *With probability at least $1 - \delta$, for all i , for all t , we have:*

$$\sum_{\tau=1}^{T(i,t)} \hat{\mu}_\tau^i - r_\star \leq 2\sqrt{8 \log(T^3 N/\delta) T(i,t)}$$

Proof. Following the same argument as in Lemma 8, we have that for any i and t ,

$$\text{Prob} \left[\sum_{\tau=1}^t \hat{r}_\tau^i - r_\tau^i \geq \sqrt{8 \log(T^3 N/\delta) t} \right] \leq \frac{\delta}{T^3 N}$$

Further, note that since $r_\tau^i \leq r_\star$ for all i and t , $\sum_{\tau=1}^t \hat{r}_\tau^i - r_\tau^i \leq \sqrt{8 \log(T^3 N/\delta) t}$ implies $\hat{\mu}_t^i - r_\star \leq \frac{\sqrt{8 \log(T^3 N/\delta)}}{\sqrt{t}}$. Therefore, for any i and t ,

$$\text{Prob} \left[\hat{\mu}_t^i - r_\star \geq \frac{\sqrt{8 \log(T^3 N/\delta)}}{\sqrt{t}} \right] \leq \frac{\delta}{T^3 N}$$

Further, observe that if $\hat{\mu}_t^i - r_\star \leq \frac{\sqrt{8 \log(T^3 N/\delta)}}{\sqrt{t}}$ for all t , then $\sum_{\tau=1}^t \hat{\mu}_\tau^i - r_\star \leq 2\sqrt{8 \log(T^3 N/\delta) t}$. Therefore for any given t and i :

$$\begin{aligned} \text{Prob} \left[\sum_{\tau=1}^{T(i,t)} \hat{\mu}_\tau^i - r_\star \geq 2\sqrt{8 \log(T^3 N/\delta)} \sqrt{T(i,t)} \right] &= \sum_{t'=1}^T \text{Prob} \left[\sum_{\tau=1}^{t'} \hat{\mu}_\tau^i - r_\star \geq 2\sqrt{8 \log(T^3 N/\delta)} \sqrt{t'} \text{ and } T(i,t) = t' \right] \\ &\leq \sum_{t'=1}^T \text{Prob} \left[\sum_{\tau=1}^{t'} \hat{\mu}_\tau^i - r_\star \geq 2\sqrt{8 \log(T^3 N/\delta)} \sqrt{t'} \right] \\ &\leq \sum_{t'=1}^T \text{Prob} \left[\hat{\mu}_\tau^i - r_\star \leq \frac{\sqrt{8 \log(T^3 N/\delta)}}{\sqrt{\tau}} \text{ for all } \tau \leq t' \right] \\ &\leq \sum_{t'=1}^T \sum_{\tau=1}^{t'} \text{Prob} \left[\hat{\mu}_\tau^i - r_\star \leq \frac{\sqrt{8 \log(T^3 N/\delta)}}{\sqrt{\tau}} \right] \\ &\leq \frac{\delta}{TN} \end{aligned}$$

Now a union bound over the T values of t and N values of i completes the statement. \square

Lemma 10. *With probability at least $1 - 3\delta$, there is some $J \in S$ such that for all t we have simultaneously:*

$$\begin{aligned} \sum_{\tau=1}^{T(J,t)} r_\star - \hat{r}_\tau^J &\leq C_J T(J,t)^{\alpha_J} \\ r_\star - \hat{\mu}_{T(i,t-1)}^J &\leq \min \left(1, C_J T(J,t-1)^{\alpha_J-1} + \frac{\sqrt{8 \log(T^3 N/\delta)}}{\sqrt{T(J,t-1)}} \right) \\ \sum_{\tau=1}^{T(J,t-1)} \hat{\mu}_{\tau-1}^J - \hat{r}_\tau^J &\leq C_J T(J,t-1)^{\alpha_J} + 3\sqrt{8 \log(TN/\delta) T(J,t-1)} \end{aligned}$$

Proof. By assumption, with probability $1 - \delta$ we have that there is some $J \in S$ such that for all t ,

$$\sum_{t=1}^{T(J,t-1)} r_{\star} - r_{\tau}^J \leq C_J T(J, t-1)^{\alpha_J} \quad (4)$$

Next, by Lemma 8 we have with probability at least $1 - \delta$, for all t ,

$$\sum_{\tau=1}^{T(J,t-1)} r_{\tau}^J - \hat{r}_{\tau}^J \leq \sqrt{8 \log(T^3 N / \delta) T(J, t-1)} \quad (5)$$

and by Lemma 9, with probability $1 - \delta$ we have for all t :

$$\sum_{\tau=1}^{T(J,t)} \hat{\mu}_{\tau}^J - r_{\star} \leq 2\sqrt{8 \log(T^3 N / \delta) T(J, t)} \quad (6)$$

All three equations then hold with probability at least $1 - 3\delta$. Let us condition on this probability $1 - 3\delta$ event. Then the first equation in the Lemma is immediately satisfied.

For the second equation, observe that since $\hat{\mu}_{T(J,t-1)}^J$ and r_{\star} are both in $[0, 1]$, the statement is true so long as the minimum binds to 1. Further, for $T(J, t-1) \neq 0$, we have

$$\begin{aligned} r_{\star} - \hat{\mu}_{T(J,t-1)}^J &= \frac{1}{T(J, t-1)} \sum_{\tau=1}^{T(J,t-1)} r_{\star} - \hat{r}_{\tau}^J \\ &= \frac{1}{T(J, t-1)} \sum_{\tau=1}^{T(J,t-1)} r_{\star} - r_{\tau}^J + \frac{1}{T(J, t-1)} \sum_{\tau=1}^{T(J,t-1)} r_{\tau}^J - \hat{r}_{\tau}^J \\ &\leq C_J T(J, t-1)^{\alpha_J-1} + \frac{\sqrt{8 \log(T^3 N / \delta)}}{\sqrt{T(J, t-1)}} \end{aligned}$$

Now for the third equation, observe $\hat{\mu}_0^J = 0$, so that

$$\begin{aligned} \sum_{\tau=1}^{T(J,t-1)} \hat{\mu}_{\tau-1}^J - r_{\star} &\leq \sum_{\tau=1}^{T(J,t-1)} \hat{\mu}_{\tau}^J - r_{\star} \\ &\leq \sum_{\tau=1}^{T(J,t-1)} \frac{\sqrt{8 \log(T^3 N / \delta)}}{\sqrt{\tau}} \\ &\leq 2\sqrt{8 \log(T^3 N / \delta) T(J, t-1)} \end{aligned}$$

Combining these equations yields:

$$\begin{aligned} \sum_{\tau=1}^{T(J,t-1)} \hat{\mu}_{\tau-1}^J - \hat{r}_{\tau}^J &= \sum_{\tau=1}^{T(J,t-1)} \hat{\mu}_{\tau-1}^J - r_{\star} + r_{\star} - r_{\tau}^J + r_{\tau}^J - \hat{r}_{\tau}^J \\ &\leq 2\sqrt{8 \log(T^3 N / \delta) T(J, t-1)} + C_J T(J, t-1)^{\alpha_J} + \sqrt{8 \log(T^3 N / \delta) T(J, t-1)} \\ &\leq C_J T(J, t-1)^{\alpha_J} + 3\sqrt{8 \log(T^3 N / \delta) T(J, t-1)} \end{aligned}$$

□

We also need a bound on the performance of the misspecified algorithms:

Lemma 11. *With probability at least $1 - \delta$, all indices i satisfy:*

$$\sum_{\tau=1}^{T(i,T)} \hat{\mu}_{\tau-1}^i - r_{\tau}^i \leq C_i T(i, T)^{\alpha_i} + 4\sqrt{8 \log(TN/\delta) T(i, T)} + 1$$

Proof. Let t be the smallest time such that $T(i, t) = T(i, T)$. Then $i_t = i$, so we must have that $i \in I_t$. Therefore:

$$\sum_{\tau=1}^{T(i,t-1)} \hat{\mu}_{\tau-1}^i - \hat{r}_{\tau}^i \leq C_i T(i, t-1)^{\alpha_i} + 3\sqrt{8 \log(T^3 N/\delta) (T(i, t-1))}$$

Next, by Lemma 8, we have that with probability at least $1 - \delta$,

$$\sum_{\tau=1}^{T(i,t-1)} \hat{r}_{\tau}^i - r_{\tau}^i \leq \sqrt{8 \log(T^3 N/\delta) T(i, t-1)}$$

Then observe that since all rewards are in $[0, 1]$, we must have $\hat{\mu}_{T(i,t)-1}^i - r_{T(i,t)}^i \leq 1$ and the statement follows. \square

One more technical fact:

Lemma 12. *Suppose $\alpha \in [1/2, 1)$. Then for any constants $A, B \geq 0$,*

$$\sup_{Z \geq 0} AZ^{\alpha} - BZ = \alpha^{\frac{\alpha}{1-\alpha}} (1 - \alpha) \frac{A^{\frac{1}{1-\alpha}}}{B^{\frac{\alpha}{1-\alpha}}}$$

Proof. We differentiate with respect to Z to obtain:

$$\begin{aligned} \alpha AZ^{\alpha-1} &= B \\ Z &= \left(\frac{\alpha A}{B} \right)^{\frac{1}{1-\alpha}} \end{aligned}$$

Now substitute this value back into the original expression:

$$\begin{aligned} AX^{\alpha} - BX &= \left(\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}} \right) \frac{A^{\frac{1}{1-\alpha}}}{B^{\frac{\alpha}{1-\alpha}}} \\ &= \alpha^{\frac{\alpha}{1-\alpha}} (1 - \alpha) \frac{A^{\frac{1}{1-\alpha}}}{B^{\frac{\alpha}{1-\alpha}}} \end{aligned}$$

\square

B Proof of Theorem 1

Now we are in a position to prove Theorem 1, which we restate below for reference:

Theorem 1. *Suppose there is a set $S \subset \{1, \dots, N\}$ such that with probability at least $1 - \delta$, there is some $J \in S$ such that*

$$\sum_{\tau=1}^t r_{\star} - r_{\tau}^J \leq C_J t^{\alpha_J}$$

for all $t \leq T$. Further, suppose the C_i and α_i are known, and the R_i satisfy:

$$R_i \geq C_i T^{\alpha_i}$$

$$R_i \geq \sum_{k \neq i} \max \left[\frac{(1 - \alpha_k)(1 + \alpha_k)^{\frac{1}{1 - \alpha_k}} (2C_k)^{\frac{1}{1 - \alpha_k}} T^{\frac{\alpha_k}{1 - \alpha_k}}}{\alpha_k R_k^{\frac{1}{1 - \alpha_k}}}, \frac{288 \log(T^3 N / \delta) T}{R_k} \right]$$

Let $r_t = \mathbb{E}[\hat{r}_{T(i_t, t)}^i]$ be the expected reward of Algorithm 1 at time t . Then, with probability at least $1 - 3\delta$, the regret satisfies:

$$\sum_{t=1}^T r_{\star} - r_t \leq 3 \sup_{j \in S} R_j$$

Note that the algorithm does not know the set S .

Proof. To begin, observe that by Lemma 10 and Lemma 11, with probability $1 - 4\delta$, there is some $J \in S$ such that for all t and all i , we have

$$\sum_{\tau=1}^{T(J, T)} r_{\star} - r_{\tau}^J \leq C_J T(J, T)^{\alpha_J} \tag{7}$$

$$r_{\star} - \hat{\mu}_{T(J, t-1)}^J \leq \min \left(1, C_J T(J, t-1)^{\alpha_J - 1} + \frac{\sqrt{8 \log(T^3 N / \delta)}}{\sqrt{T(J, t-1)}} \right) \tag{8}$$

$$J \in I_t$$

$$\sum_{\tau=1}^{T(i, T)} \hat{\mu}_{\tau-1}^i - r_{\tau}^i \leq C_i T(i, T)^{\alpha_i} + 4\sqrt{8 \log(TN / \delta) T(i, T)} + 1 \tag{9}$$

All of our analysis is conditioned on this probability $1 - 4\delta$ event.

First, notice that by (8), we have that we have that:

$$r_{\star} \leq U(J, t-1) + \frac{R_J}{T}$$

for all t .

Now we write the regret:

$$\begin{aligned} \sum_{t=1}^T r_{\star} - r_t &= \sum_{t=1}^T r_{\star} - r_{T(i_t, t)}^{i_t} \\ &= \sum_{i_t=J} r_{\star} - r_{T(J, t)}^J + \sum_{i_t \neq J} r_{\star} - r_{T(i_t, t)}^{i_t} \end{aligned}$$

First let's consider the indices for which $i_t = J$:

$$\sum_{i_t=J} r_{\star} - r_{T(J, t)}^J \leq C_J T(J, T)^{\alpha_J} \leq R_J$$

Next, if $i_t \neq J$, we must have $U(i_t, t-1) \geq U(J, t-1)$ since $J \in I_t$ for all t . Then:

$$\begin{aligned} r_{\star} &= r_{\star} - U(J, t-1) + U(J, t-1) - U(i_t, t-1) + U(i_t, t-1) \\ &\leq \frac{R_J}{T} + U(i_t, t-1) \end{aligned}$$

Therefore:

$$\begin{aligned} \sum_{i_t \neq J} r_{\star} - r_{T(i_t, t)}^{i_t} &\leq \sum_{i_t \neq J} \frac{R_J}{T} + U(i_t, t-1) - r_{T(i_t, t)}^{i_t} \\ &\leq R_J + \sum_{i_t \neq J} U(i_t, t-1) - r_{T(i_t, t)}^{i_t} \end{aligned}$$

Let us focus on the sum $\sum_{i_t \neq J} U(i_t, t-1) - r_{T(i_t, t)}^{i_t}$. As a first step, we observe the following identity:

$$\begin{aligned} \sum_{t \mid i_t=i} U(i, t-1) - \hat{\mu}_{T(i, t-1)}^i &= \sum_{\tau=1}^{T(i, T)} \min \left(1, C_i (\tau-1)^{\alpha_i-1} + \frac{\sqrt{8 \log(T^3 N / \delta)}}{\sqrt{\tau-1}} \right) - \frac{R_i}{T} \\ &\leq 1 - \frac{T(i, T) R_i}{T} + \sum_{\tau=1}^{T(i, T)} C_i \tau^{\alpha_i-1} + \frac{\sqrt{8 \log(T^3 N / \delta)}}{\sqrt{\tau}} \\ &\leq 1 - \frac{T(i, T) R_i}{T} + \frac{C_i T(i, T)^{\alpha_i}}{\alpha_i} + 2\sqrt{8 \log(T^3 N / \delta) T(i, T)} \end{aligned}$$

Now we use this observation as follows:

$$\begin{aligned}
\sum_{i_t \neq J} U(i_t, t-1) - \hat{r}_{T(i_t, t)}^{i_t} &= \sum_{i_t \neq J} \hat{\mu}_{T(i_t, t-1)}^{i_t} - r_{T(i_t, t)}^{i_t} + U(i_t, t-1) - \hat{\mu}_{T(i_t, t-1)}^{i_t} \\
&\leq \sum_{i \neq J} \left[\sum_{t \mid i_t = i} \hat{\mu}_{T(i, t-1)}^i - r_{T(i, t)}^i + U(i, t-1) - \hat{\mu}_{T(i, t-1)}^i \right] \\
&\leq \sum_{i \neq J} 1 - \frac{T(i, T)R_i}{T} + \frac{C_i T(i, T)^{\alpha_i}}{\alpha_i} + 2\sqrt{8 \log(T^3 N/\delta) T(i, T)} + \sum_{\tau=1}^{T(i, T)} \hat{\mu}_{\tau-1}^i - r_\tau^i \\
&\leq \sum_{i \neq J} -\frac{T(i, T)R_i}{T} + \frac{C_i T(i, T)^{\alpha_i}}{\alpha_i} + 3\sqrt{8 \log(T^3 N/\delta) T(i, T)} + \sum_{\tau=1}^{T(i, T)} \hat{\mu}_{\tau-1}^i - r_\tau^i \\
&\leq \sum_{i \neq J} \frac{(1 + \alpha_i) C_i T(i, T)^{\alpha_i}}{\alpha_i} + 6\sqrt{8 \log(T^3 N/\delta) T(i, T)} - \frac{T(i, T)R_i}{T}
\end{aligned}$$

Now we use a simple trick that allows us to apply Lemma 12:

$$\begin{aligned}
&\frac{(1 + \alpha_i) C_i T(i, T)^{\alpha_i}}{\alpha_i} + 6\sqrt{8 \log(T^3 N/\delta) T(i, T)} - \frac{T(i, T)R_i}{T} \\
&\leq \sup_Z \frac{(1 + \alpha_i) C_i Z^{\alpha_i}}{\alpha_i} + 6\sqrt{8 \log(T^3 N/\delta) Z} - \frac{Z R_i}{T} \\
&\leq \max \left[\sup_Z \frac{(1 + \alpha_i) 2 C_i Z^{\alpha_i}}{\alpha_i} - \frac{Z R_i}{T}, \sup_Z 12\sqrt{8 \log(T^3 N/\delta) Z} - \frac{Z R_i}{T} \right] \\
&\leq \max \left[\frac{(1 - \alpha_i)(1 + \alpha_i)^{\frac{1}{1-\alpha_i}} (2 C_i)^{\frac{1}{1-\alpha_i}} T^{\frac{\alpha_i}{1-\alpha_i}}}{\alpha_i R^{\frac{\alpha_i}{1-\alpha_i}}}, \frac{288 \log(T^3 N/\delta) T}{R_i} \right]
\end{aligned}$$

Therefore by our conditions on the R_i , we have

$$\begin{aligned}
&\sum_{i \neq J} \frac{(1 + \alpha_i) C_i T(i, T)^{\alpha_i}}{\alpha_i} + 6\sqrt{8 \log(T^3 N/\delta) T(i, T)} - \frac{T(i, T)R_i}{T} \\
&\leq \sum_{i \neq J} \max \left[\frac{(1 - \alpha_i)(1 + \alpha_i)^{\frac{1}{1-\alpha_i}} (2 C_i)^{\frac{1}{1-\alpha_i}} T^{\frac{\alpha_i}{1-\alpha_i}}}{\alpha_i R^{\frac{\alpha_i}{1-\alpha_i}}}, \frac{288 \log(T^3 N/\delta) T}{R_i} \right] \\
&\leq R_J
\end{aligned}$$

And so putting all this together we have with probability at least $1 - 3\delta$:

$$\begin{aligned}
\sum_{t=1}^T r_\star - \hat{r}_t &= \sum_{i_t=J} r_\star - \hat{r}_{T(J, t)}^J + \sum_{i_t \neq J} r_\star - \hat{r}_{T(i_t, t)}^{i_t} \\
&\leq R_J + R_J + R_J \leq 3R_J
\end{aligned}$$

for some $J \in S$. Therefore with probability at least $1 - 3\delta$, the regret is bounded by $3 \sup_{j \in S} R_j$. \square

B.1 Proof of Corollary 2

Now we prove Corollary 2:

Corollary 2. *Suppose there is a set $S \subset \{1, \dots, N\}$ such that with probability at least $1 - \delta$, there is some $J \in S$ such that $\sum_{\tau=1}^t r_\star - r_\tau^J \leq C_J t^{\alpha_J}$ for all $t \leq T$. Further, suppose we are given N positive real numbers η_1, \dots, η_N . Set R_i via:*

$$R_i = C_i T^{\alpha_i} + \frac{(1 - \alpha_i)^{\frac{1-\alpha_i}{\alpha_i}} (1 + \alpha_i)^{\frac{1}{\alpha_i}}}{\alpha_i^{\frac{1-\alpha_i}{\alpha_i}}} C_i^{\frac{1}{\alpha_i}} T \eta_i^{\frac{1-\alpha_i}{\alpha_i}} \\ + 288 \log(T^3 N / \delta) T \eta_i + \sum_{k \neq i} \frac{1}{\eta_k}$$

Then, with probability at least $1 - 3\delta$, the regret of Algorithm 1 satisfies:

$$\text{Regret} \leq 3 \sup_{j \in S} R_j = \tilde{O} \left(\sup_{j \in S} C_j T^{\alpha_j} + C_j^{\frac{1}{\alpha_j}} T \eta_j^{\frac{1-\alpha_j}{\alpha_j}} \right. \\ \left. + T \eta_j + \sum_{k \neq j} \frac{1}{\eta_k} \right)$$

Proof. Observe that by definition of R_i we have:

$$R_i \geq C_i T^{\alpha_i}$$

Further, due to the second and third terms in the definition of R_i respectively, we have the following:

$$\frac{1}{\eta_i} \geq \frac{(1 - \alpha_i)(1 + \alpha_i)^{\frac{1}{1-\alpha_i}} (2C_i)^{\frac{1}{1-\alpha_i}} T^{\frac{\alpha_i}{1-\alpha_i}}}{\alpha_i R_i^{\frac{\alpha_i}{1-\alpha_i}}} \\ \frac{1}{\eta_i} \geq \frac{288 \log(T^3 N / \delta) T}{R_i}$$

Finally, due to the last term in the definition of R_i , we have:

$$R_i \geq \sum_{k \neq i} \frac{1}{\eta_k} \\ \geq \sum_{k \neq i} \max \left\{ \frac{(1 - \alpha_k)(1 + \alpha_k)^{\frac{1}{1-\alpha_k}} (2C_k)^{\frac{1}{1-\alpha_k}} T^{\frac{\alpha_k}{1-\alpha_k}}}{\alpha_k R_k^{\frac{\alpha_k}{1-\alpha_k}}}, \frac{288 \log(T^3 N / \delta) T}{R_k} \right\}$$

Therefore these R_i satisfy the conditions of Theorem 1, and so we are done. \square

C Proof of Theorem 3

Theorem 3. *Suppose that there is some $J \in \{1, \dots, N\}$ such that with probability at least $1 - \delta$, we have:*

$$\sum_{\tau=1}^t r_{\star} - r_{\tau}^J \leq C_J t^{\alpha_J}$$

Also, for all i , define $T_i = T(i, T)$ and Δ_i by:

$$\Delta_i = \frac{1}{T(i, T_i - 1)} \sum_{\tau=1}^{T(i, T_i - 1)} r_{\star} - r_{\tau}^i$$

And let $B \subset \{1, \dots, N\}$ with $J \notin B$ be the set of indices with $\Delta_i > \frac{2R_J}{T}$ for $i \in B$.

For $i \neq J$, let $C_i > 0$ and $\alpha_i \leq 1$ for $i \neq J$ be arbitrary. Then with probability at least $1 - 3\delta$, the regret of Algorithm 1 satisfies:

$$\begin{aligned} \sum_{t=1}^T r_{\star} - r_t &\leq \sum_{\tau=1}^{T(J, T)} r_{\star} - r_{\tau}^J + \sum_{i \in B} 1 + \frac{512 \log(T^3 N / \delta)}{\Delta_i} + \frac{4^{\frac{1}{1-\alpha_i}} C_i^{\frac{1}{1-\alpha_i}}}{\Delta_i^{\frac{\alpha_i}{1-\alpha_i}}} \\ &+ \min \left[\sum_{k \neq J, k \notin B} 2R_J, R_J + \sum_{k \neq J, k \notin B} \max \left[\frac{(1 - \alpha_k)(1 + \alpha_k)^{\frac{1}{1-\alpha_k}} (2C_k)^{\frac{1}{1-\alpha_k}} T^{\frac{\alpha_k}{1-\alpha_k}}}{\alpha_k R_k^{\frac{\alpha_k}{1-\alpha_k}}}, \frac{288 \log(T^3 N / \delta) T}{R_k} \right] \right] \end{aligned}$$

Proof. We begin with the regret decomposition:

$$\sum_{t=1}^T r_{\star} - r_t = \sum_{\tau=1}^{T(J, T)} r_{\star} - r_{\tau}^J + \sum_{i \notin B, i \neq J} \sum_{\tau=1}^{T(i, T)} r_{\star} - r_{\tau}^i + \sum_{i \in B} \sum_{\tau=1}^{T(i, T)} r_{\star} - r_{\tau}^i$$

The first sum in the above is already in the form we desire. The second sum could be bounded using the same argument as in the proof of Theorem 1 by:

$$R_J + \sum_{k \neq J, k \notin B} \max \left[\frac{(1 - \alpha_k)(1 + \alpha_k)^{\frac{1}{1-\alpha_k}} (2C_k)^{\frac{1}{1-\alpha_k}} T^{\frac{\alpha_k}{1-\alpha_k}}}{\alpha_k R_k^{\frac{\alpha_k}{1-\alpha_k}}}, \frac{288 \log(T^3 N / \delta) T}{R_k} \right]$$

However, we can also bound it another way. By definition of B , for any $i \notin B$, we have:

$$\sum_{\tau=1}^{T(i, T)} r_{\star} - r_{\tau}^i \leq T \Delta_i \leq 2R_J$$

so that the second sum is also bounded by

$$\sum_{k \neq J, k \notin B} 2R_J$$

This leads to the expression involving a minimum in the Theorem statement.

For the last term, our proof proceeds by counting how many times we choose an algorithm \mathcal{A}_i for $i \in B$. We only choose such an algorithm when $U(i, t-1) \geq U(J, t-1)$. Further, note from the proof of Theorem 1, that with probability at least $1 - \delta$, we have:

$$r_\star \leq U(J, t-1) + \frac{R_J}{T}$$

for all t . For any i , define T_i as the last time at which \mathcal{A}_i is chosen by Algorithm 1. Our goal is to bound T_i for all $i \in B$. By definition, we must have $U(i, T_i - 1) \geq U(J, T_i - 1)$. Then, observe that from Lemma 8, we have that with probability at least $1 - \delta$, for all $i \in B$:

$$\hat{\mu}_{T(i, T_i - 1)}^i - r_\star + \Delta_i = \frac{1}{T(i, T_i - 1)} \sum_{\tau=1}^{T(i, T_i - 1)} \hat{r}_\tau^i - r_\tau^i \leq \frac{2\sqrt{8 \log(T^3 N / \delta)}}{\sqrt{T(i, T_i - 1)}}$$

Define $Z(i, t) = \min\left(1, \frac{C_i T(i, t)^\alpha + \sqrt{8 \log(T^3 N / \delta) T(i, t)}}{T(i, t)}\right)$. Then by the previous line and by definition of $U(i, T_i - 1)$, we have with probability at least $1 - 2\delta$,

$$\begin{aligned} U(i, T_i - 1) &\leq \frac{2\sqrt{8 \log(T^3 N / \delta)}}{\sqrt{T(i, T_i - 1)}} + r_\star - \Delta_i + Z(i, T_i - 1) - \frac{R_i}{T} \\ &\leq \frac{2\sqrt{8 \log(T^3 N / \delta)}}{\sqrt{T(i, T_i - 1)}} + U(J, T_i - 1) + \frac{R_J}{T} - \Delta_i + Z(i, T_i - 1) - \frac{R_i}{T} \end{aligned}$$

Therefore, if $U(i, T_i - 1) \geq U(J, T_i - 1)$, we must have:

$$\begin{aligned} \Delta_i + \frac{R_i}{T} - \frac{R_J}{T} &\leq \frac{2\sqrt{8 \log(T^3 N / \delta)}}{\sqrt{T(i, T_i - 1)}} + \min\left(1, \frac{C_i T(i, T_i - 1)^\alpha + \sqrt{8 \log(T^3 N / \delta) T(i, T_i - 1)}}{T(i, T_i - 1)}\right) \\ &\quad (10) \\ \frac{\Delta_i}{2} &\leq \frac{2\sqrt{8 \log(T^3 N / \delta)}}{\sqrt{T(i, T_i - 1)}} + \min\left(1, \frac{C_i T(i, T_i - 1)^\alpha + \sqrt{8 \log(T^3 N / \delta) T(i, T_i - 1)}}{T(i, T_i - 1)}\right) \\ T(i, T_i - 1) &\leq \frac{512 \log(T^3 N / \delta)}{\Delta_i^2} + \frac{4^{\frac{1}{1-\alpha_i}} C_i^{\frac{1}{1-\alpha_i}}}{\Delta_i^{\frac{1}{1-\alpha_i}}} \end{aligned}$$

Therefore, observing that by definition we have $T(i, T_i - 1) + 1 = T_i = T(i, T)$, we have:

$$\begin{aligned} \sum_{\tau=1}^{T(i, T)} r_\star - r_\tau^i &\leq 1 + T(i, T_i - 1) \Delta_i \\ &\leq 1 + \frac{512 \log(T^3 N / \delta)}{\Delta_i} + \frac{4^{\frac{1}{1-\alpha_i}} C_i^{\frac{1}{1-\alpha_i}}}{\Delta_i^{\frac{1}{1-\alpha_i}}} \end{aligned}$$

□

D Proof of Theorem 5

In this section, we provide a proof of Theorem 5, restated below:

Theorem 5. *Suppose that for some J , there is some unknown \bar{C}_J and $\bar{\alpha}_J$ such that \mathcal{A}_J ensures $\mathbb{E}[\sum_{\tau=1}^t r_\star - r_\tau^J] \leq \bar{C}_J t^{\bar{\alpha}_J}$ for all $t \leq T$. Further, suppose we are given positive numbers η_1, \dots, η_N . Let $\delta \in (0, 1)$ be some user-specified failure probability. For $M = \lceil \log_2(1/\delta) \rceil$ and $K = \lceil \log_2(T) \rceil$ and $L = \lceil \frac{\log_2(T)}{2} \rceil$, we duplicate each \mathcal{A}_i MKL times, specifying each duplicate $\mathcal{A}_{i,x,y,z}$ by a multi-index $(i, x, y, z) \in [N] \times [M] \times [K] \times [L]$. To each duplicate we associate $C_{i,x,y,z} = 2^y$ and $\alpha_{i,x,y,z} = \min\left(1, 1/2 + \frac{z}{\log(T)}\right)$. Let $\eta_{i,x,y,z} = \eta_i$. Specify $R_{i,x,y,z}$ as a function of $\eta_{i,x,y,z}$ as described in Theorem 2. Then with probability at least $1 - 3\delta$, Algorithm 1 guarantees regret:*

$$\begin{aligned} \text{Regret} &\leq O\left(\bar{C}_J T^{\bar{\alpha}_J} + \bar{C}_J^{\frac{1}{\bar{\alpha}_J}} T^{\bar{\alpha}_J} \eta_J^{\frac{1-\bar{\alpha}_J}{\bar{\alpha}_J}}\right. \\ &\quad \left. + \log\left(\frac{T^3 N}{\delta}\right) T \eta_J + \sum_{i=1}^N \frac{\log(1/\delta) \log^2(T)}{\eta_i}\right) \end{aligned}$$

so that the expected regret is bounded by:

$$\begin{aligned} \mathbb{E}[\text{Regret}] &\leq O\left(\bar{C}_J T^{\bar{\alpha}_J} + \bar{C}_J^{\frac{1}{\bar{\alpha}_J}} T^{\bar{\alpha}_J} \eta_J^{\frac{1-\bar{\alpha}_J}{\bar{\alpha}_J}}\right. \\ &\quad \left. + \log\left(\frac{T^3 N}{\delta}\right) T \eta_J + \sum_{i=1}^N \frac{\log(1/\delta) \log^2(T)}{\eta_i} + T\delta\right) \end{aligned}$$

Proof. Let $\bar{y} = \lceil \log_2(\bar{C}_J) \rceil$, and let $\bar{z} = \lceil \frac{\bar{\alpha}_J - \frac{1}{2}}{\log_2(T)} \rceil$. Then observe since $T^{1/\log_2(T)} = 2$, we have that for all x and $t \leq T$,

$$\bar{C}_J t^{\bar{\alpha}_J} \leq C_{J,x,\bar{y},\bar{z}} t^{\alpha_{J,x,\bar{y},\bar{z}}} \leq 4\bar{C}_J t^{\bar{\alpha}_J}$$

Next, let $\bar{S} = \{(J, x, \bar{y}, \bar{z}) : x \in [M]\}$. Then we claim that with probability at least $1 - \delta$, there exists some $\bar{J} = (J, \bar{x}, \bar{u}, \bar{z}) \in \bar{S}$ such that $\mathcal{A}_{\bar{J}}$ obtains regret $\sum_{\tau=1}^t r_\star - r_\tau^{\bar{J}} \leq C_{J,x,\bar{y},\bar{z}} t^{\alpha_{J,x,\bar{y},\bar{z}}}$ for all $t \leq T$. The conclusion of the Theorem will then follow from Theorem 2, observing that the total number of duplicates for each algorithm is $MKL = O(\log(1/\delta) \log^2(T))$.

Let us prove the claim. Define $V = \lceil \log_2(T) \rceil$ and define the sequence t_1, \dots, t_V by $t_i = 2^i$ for $i < V$ and $t_V = T$. Then for any $\epsilon > 0$ and any i and any j , by Markov inequality, with probability at least $1 - \epsilon$ we have

$$\sum_{\tau=1}^{t_i} r_\star - r_\tau^{J,j,\bar{y},\bar{z}} \leq \frac{C_{J,j,\bar{y},\bar{z}} t_i^{\alpha_{J,j,\bar{y},\bar{z}}}}{\epsilon}$$

With $\epsilon = \frac{1}{2(\log_2(T)+1)}$, by union bound over the $V \leq \log_2(T) + 1$ values of i , we have that for any given j , with probability at least $1/2$,

$$\sum_{\tau=1}^{t_i} r_\star - r_\tau^{J,j,\bar{y},\bar{z}} \leq 2(\log_2(T) + 1) C_{J,j,\bar{y},\bar{z}} t_i^{\alpha_{J,j,\bar{y},\bar{z}}} \quad (11)$$

for all i .

There are M multi-index \bar{J} of the form (J, j, \bar{y}, \bar{z}) . Thus the probability that *none* of them satisfies (11) for all i is at most $\frac{1}{2^M} \leq \delta$. Therefore, with probability at least $1 - \delta$, we have that (11) holds for all i *some* j and multi-index $\bar{J} = (J, j, \bar{y}, \bar{z})$.

Next, observe that $r_\star - r_t^{\bar{J}} \geq 0$ for all t with probability 1. Therefore, with probability $1 - \delta$, for any t we set t_i such that $t \leq t_i \leq 2t$ and have

$$\begin{aligned} \sum_{\tau=1}^t r_\star - r_\tau^{\bar{J}} &\leq \sum_{\tau=1}^{t_i} r_\star - r_\tau^{\bar{J}} \\ &\leq 2C(\log_2(T) + 1)\sqrt{t_i} \\ &\leq 4C(\log_2(T) + 1)\sqrt{t} \end{aligned}$$

□

E Proofs for Section 6

Corollary 6. *Suppose $c_t : A \rightarrow \mathbb{R}^d$ for all t . Suppose $K \geq d \log(K)$. Let \mathcal{A}_1 be an instance of *linUCB* and \mathcal{A}_2 be an instance of the ordinary *UCB* algorithm. Let P and Q be any two numbers such that $PQ = KT$ and both are greater than $\sqrt{d \log(K)T}$. We consider two cases, either the reward is a linear function of c_t , or it is not. Then \mathcal{A}_1 guarantees regret $\tilde{O}(\sqrt{d \log(K)t})$ with probability $1 - \delta$ in the first case, while \mathcal{A}_2 guarantees regret $\tilde{O}(\sqrt{Kt})$ in the second case. Set $\eta_1 = \frac{P}{d \log(K)T} = \frac{K}{Qd \log(K)}$ and $\eta_2 = \frac{1}{P} = \frac{Q}{KT}$. Then using the R_i construction of Corollary 2, with probability at least $1 - 3\delta$, we guarantee regret $\tilde{O}(P)$ with linear rewards, and $\tilde{O}(Q)$ otherwise.*

Proof. The regret bounds on \mathcal{A}_1 and \mathcal{A}_2 are standard from analysis of the respective algorithms. We set $\alpha_1 = \alpha_2 = \frac{1}{2}$ and set $C_1 = \tilde{O}(\sqrt{d \log(K)})$ and $C_2 = \tilde{O}(\sqrt{K})$, where the constants and log factors in the \tilde{O} notations come from the regret bounds on \mathcal{A}_1 and \mathcal{A}_2 . Then if the rewards are linear, then applying Theorem 2 yields regret $\tilde{O}\left(d \log(K)T\eta_1 + \frac{1}{\eta_2}\right) = \tilde{O}(X)$ (with $S = \{1\}$), while otherwise we have $\tilde{O}\left(KT\eta_2 + \frac{1}{\eta_1}\right) = \tilde{O}\left(Y + \frac{d \log(K)Y}{K}\right) = \tilde{O}(Y)$ (with $S = \{2\}$). □

Corollary 7. *Suppose $c_t \in \mathbb{R}^d$ for all t and the reward is always a linear function of c_t . Suppose that the reward is in fact purely a linear function of the first d_\star coordinates of c_t . Suppose the action set A has finite cardinality K . Let \mathcal{A}_i be an instance of *linUCB* of [16] restricted to the first 2^i coordinates of the context. Set $C_i = \sqrt{2^i \log(K)}$, $\alpha_i = \frac{1}{2}$, and $\eta_i = \frac{1}{\sqrt{T}}$. Then using the instantiation of Algorithm 1 from Theorem 2, we obtain regret $\tilde{O}\left(d_\star \log(K)\sqrt{T}\right)$. If instead the set A is infinite, let \mathcal{A}_i be an instance of the *linUCB* algorithm for infinite arms [20, 21] restricted to the first 2^i coordinates. Set $C_i = 2^i$, $\alpha_i = \frac{1}{2}$, and $\eta_i = \frac{1}{\sqrt{T}}$. Then we obtain regret $\tilde{O}\left(d_\star^2 \sqrt{T}\right)$.*

Proof. For the finite action setting, [16] proves, supposing $d_\star \leq \hat{d}$, then an instance of *linUCB* restricted to the first \hat{d} dimensions will obtain regret at most $\tilde{O}(\sqrt{\hat{d} \log(K)T})$. In the infinite action setting, [20] shows that the OFUL algorithm obtains regret $\tilde{O}(\hat{d}\sqrt{T})$. The corollary is then a consequence of Theorem 2. □

F High-Probability Regret and the Optimism Principle

In this section we sketch a proof that UCB-like algorithms tend to satisfy the high-probability conditions required by Theorem 1, without resorting to the duplication technique needed in Theorem 5. The analysis in this section is standard and more detailed discussion can be found in various texts (e.g. [5, 22, 23]).

UCB and related algorithms make use of the *optimism principle*. The intuitive idea is to maintain for each policy x a bound on the best possible value for the expected reward of x given the data seen so far. At each round, the policy with the largest such bound is chosen. In general, the bound for a policy is some mixture of its past observed performance and some uncertainty due to lack of data, so that the upper bound is larger both if a policy has done well in the past, or if it has not been explored in the past. More formally, we consider an algorithm UCB-like if for each policy x and round index t it produces a confidence set $C(x, t, \delta) \subset \mathbb{R}$ such that with probability at least $1 - \delta$, the expected reward r_x satisfies $r_x \in C(x, t, \delta)$ for all x for all t . At each time step, the algorithm plays the arm $x_t = \operatorname{argmax}_x U(x, t, \delta)$, where $U(x, t, \delta) = \sup C(x, t, \delta)$. Conditioned on this probability $1 - \delta$ event, the analysis proceeds as follows:

$$\begin{aligned} r_\star - r_{x_t} &\leq r_\star - U(\star, t, \delta) + U(\star, t, \delta) - U(x_t, t, \delta) + U(x_t, t, \delta) - r_{x_t} \\ &\leq U(x_t, t, \delta) - r_{x_t} \\ &\leq U(x_t, t, \delta) - L(x_t, t, \delta) \end{aligned}$$

where $L(x, t, \delta) = \inf C(x, t, \delta)$. Thus, with probability at least $1 - \delta$, the total regret at any time $\tau \leq T$ is

$$\sum_{t=1}^{\tau} r_\star - r_t \leq \sum_{t=1}^{\tau} U(x_t, t, \delta) - L(x_t, t, \delta)$$

Many UCB-like bounds operate by providing a bound on $\sum_{t=1}^T U(x_t, t, \delta) - L(x_t, t, \delta)$ that holds for *every possible sequence* of x_t - not just those ones that are chosen by the algorithm. Algorithms with regret bounds that can be derived in this manner include the standard UCB algorithm for K -armed bandits [3], as well as the linear UCB algorithm for potentially infinitely armed linear bandits [21, 20]. In the case of ordinary UCB, $C(x, t, \delta) = [\mu_x - \frac{C(\delta)}{\sqrt{T(x, t-1)}}, \mu_x + \frac{C(\delta)}{\sqrt{T(x, t-1)}}]$ where C is some constant that is $O(\sqrt{\log(T/\delta)})$ and μ_x is the empirical average of all prior rewards observed from arm x . Here we observe that $U(x_t, t, \delta) - L(x_t, t, \delta)$, conditioned on the sequence of arm pulls, is independent of the actual observed rewards. It can be shown that

$$\sum_{t=1}^{\tau} U(x_t, t, \delta) - L(x_t, t, \delta) \leq O(C(\delta)\sqrt{K\tau})$$

for any sequence x_1, \dots, x_τ . Therefore so long as the confidence intervals are correct for all times t , which happens with probability at least $1 - \delta$, we have a bound on the regret of $O(C(\delta)\sqrt{K\tau})$

In the case of linUCB, we have the more complicated set:

$$C(x, t, \delta) = [\langle \hat{\theta}, x \rangle - C(\delta)\sqrt{x^\top V_{t-1} x}, \langle \hat{\theta}, x \rangle + C(\delta)\sqrt{x^\top V_{t-1} x}]$$

where $\hat{\theta}$ is the regularized least-squares regression solution for the linear parameter θ , $V_{t-1} = \lambda I + \sum_{\tau=1}^{t-1} x_{i_\tau} x_{i_\tau}^\top$, and $C(\delta)$ again is constant that is $O(\sqrt{d \log(T/\delta)})$. Again, we observe that

$U(x_t, t, \delta) - L(x_t, t, \delta)$, conditioned on the sequence of arm pulls, is independent of the actual observed rewards. It can be shown that

$$\sum_{t=1}^{\tau} U(x_t, t, \delta) - L(x_t, t, \delta) \leq O(C(\delta)\sqrt{d\tau \log(\tau)})$$

for any sequence x_1, \dots, x_τ . Thus with probability at least $1 - \delta$, the confidence intervals are all correct, and so we have regret at most $O(C(\delta)\sqrt{d \log(\tau)\tau}) = O(d\sqrt{\log(T/\delta)\tau})$.

G Combining Linear Bandits with Adversarial Contexts

In this section we consider linear contextual bandits with *adversarial* rather than stochastic contexts and show that a variant of our combiner Algorithm 1 can combine instances of standard confidence-ellipsoid based bandit algorithms. Formally, let us allow each learner \mathcal{A}_i to receive a context mapping $\psi_t^i : A \rightarrow \mathbb{R}^{d_i}$ that maps actions in A to feature vectors in \mathbb{R}^{d_i} . The learner will use this mapping to minimize regret. We suppose there is some J such that for all t and a , the reward for action a at time t is $\langle \psi_t^J(a), \theta_\star \rangle$ for some $\theta_\star \in \mathbb{R}^{d_J}$. For simplicity of exposition, we consider the case in which $\|\theta_\star\| \leq 1$, $\|\psi_t^i(a)\| \leq 1$ for all i, t , and a , and the observed rewards satisfy $|\hat{r}_t^i| \leq 1$ with probability 1. We consider the regret with respect to the optimal action at time t . That is, define

$$r_{\star, t} = \max_{a \in A} \langle \psi_t^J(a), \theta_\star \rangle$$

and the regret as:

$$\sum_{t=1}^T r_{\star, t} - r_t$$

where r_t is the expected loss of the chosen action. The standard base learner for this problem is LinUCB (e.g see [16, 5], which would obtain a base regret of $\tilde{O}(d_J\sqrt{T})$). We would like to combine many base learners to obtain a regret bound that depends on T and d_J , but not d_i for $i \neq J$. Specifically, we will obtain a result analogous to Theorem 1 that will enable regret $\tilde{O}(d_J^2\sqrt{T})$.

By mild abuse of notation, we will write $a_{T(i_t, t)}^i$ to indicate the feature under the map ψ_t^i of the $T(i_t, t)$ th action taken by algorithm \mathcal{A}_{i_t} . Define the regularized empirical covariance matrix M_t^i as:

$$M_t^i = \lambda I + \sum_{\tau=1}^t (a_\tau^i)^\top a_\tau^i$$

and define the linear regression estimate for θ_\star :

$$\hat{\mu}_t^i = \operatorname{argmin}_{\theta} \lambda \|\theta\|^2 + \sum_{\tau=1}^t (\langle a_\tau^i, \theta \rangle - r_\tau^i)^2$$

Now, we redefine our UCB combiner with confidence intervals. First, set

$$\beta = 4160 \log \left(\frac{T \log_2(\sqrt{T/\log(T/\delta)}) + 2}{\delta} \right) + 6\lambda + 16d \log(1 + T/\lambda)$$

and then define:

$$U(i, t-1) = \max_{a \in A} \langle \psi_t^i(a), \hat{\mu}_{T(i,t)-1}^i \rangle + \beta \sqrt{\psi_t^i(a)^\top (M_{T(i,t)-1}^i)^{-1} \psi_t^i(a)} - \frac{R_i}{T}$$

We set \mathcal{A}_i to be the linUCB algorithm (e.g. [16]). At time t , it suggests the action:

$$\operatorname{argmax}_{a \in A} \langle \psi_t^i(a), \hat{\mu}_{T(i,t)-1}^i \rangle + \beta \sqrt{\psi_t^i(a)^\top (M_{T(i,t)-1}^i)^{-1} \psi_t^i(a)} \quad (12)$$

We note that it is more typical to use a time-varying value for β , but in order to simplify the use of indices somewhat we present this fixed scaling. The arguments are essentially identical using standard time-varying schedules for β .

Next, we change our criterion for eliminating algorithms from the active set I_t . We define the quantity:

$$z_t^i = \langle a_t^i, \hat{\mu}_{t-1}^i \rangle - \hat{r}_t^i - \beta \sqrt{(a_t^i)^\top (M_{t-1}^i)^{-1} a_t^i}$$

and we decide to eliminate an algorithm if it ever does not satisfy:

$$\sum_{\tau=1}^t z_\tau^i \leq 2\sqrt{t \log(T/\delta)} \quad (13)$$

or if it does not satisfy

$$2 \sum_{\tau=1}^t \beta \sqrt{(a_\tau^i)^\top (M_{\tau-1}^i)^{-1} a_\tau^i} \leq C_i t^{\alpha_i} \quad (14)$$

Formally, the algorithm is stated in Algorithm 2.

We will use a proof similar to that of Theorem 1 to show the following:

Theorem 13. *Suppose there is some $J \in \{1, \dots, N\}$ such that with probability at least $1 - \delta$:*

$$\begin{aligned} \sum_{\tau=1}^t r_\star - r_\tau^J &\leq 2 \sum_{\tau=1}^t \beta \sqrt{(a_\tau^J)^\top (M_{\tau-1}^J)^{-1} a_\tau^J} \\ &\leq C_J t^{\alpha_J} \end{aligned} \quad (15)$$

and also for all $a \in A$, the expected reward $r_t(a) = \langle \psi_t^J(a), \theta_\star \rangle$ for some θ_\star with $\|\theta_\star\| \leq 1$ for all $t \leq T$. Further, suppose all the base learners \mathcal{A}_i are linUCB algorithms using prediction (12), C_i and α_i are known, and the R_i satisfy:

$$\begin{aligned} R_i &\geq C_i T^{\alpha_i} \\ R_i &\geq \sum_{k \neq i} \max \left[\frac{24T}{R_k}, \frac{(1 - \alpha_k) \alpha_k^{\frac{\alpha_k}{1-\alpha_k}} (2C_k)^{\frac{1}{1-\alpha_k}} T^{\frac{\alpha_k}{1-\alpha_k}}}{R_k^{\frac{\alpha_k}{1-\alpha_k}}} \right] \end{aligned}$$

Let $r_t = \mathbb{E}[\hat{r}_{T(i,t)}^i]$ be the expected reward of Algorithm 2 at time t . Then, with probability at least $1 - 3\delta$, the regret satisfies:

$$\sum_{t=1}^T r_\star - r_t \leq 3R_J$$

Algorithm 2 Adversarial Linear Bandit Combiner

Input: Numbers $\{C_i\}$, $\{\alpha_i\}$, $\{R_i\}$, T .

Set $T(i, 0) = 0$ for all i , set $\hat{\mu}_0^i = 0$ for all i , and set $I_1 = \{1, \dots, N\}$

Set $\beta = 4160 \log \left(\frac{T \log_2(\sqrt{T/\log(T/\delta)} + 2)}{\delta} \right) + 6\lambda + 16d \log(1 + T/\lambda)$

for $t = 1 \dots T$ **do**

Set $M_t^i = \lambda I + \sum_{\tau=1}^t (a_\tau^i)^\top a_\tau^i$

Set $U(i, t-1) = \max_{a \in A} \langle \psi_t^i(a), \hat{\mu}_{T(i,t)-1}^i \rangle + \beta \sqrt{\psi_t^i(a)^\top (M_{T(i,t)-1}^i)^{-1} \psi_t^i(a)} - \frac{R_i}{T}$.

Set $i_t = \operatorname{argmax}_{i \in I_t} U(i, t-1)$.

Update $T(i_t, t) = T(i_t, t-1) + 1$ and $T(j, t) = T(j, t-1)$ for $j \neq i_t$.

Get $T(i_t, t)$ th action $a_t = a_{T(i_t, t)}^{i_t}$ from \mathcal{A}_{i_t} .

Receive reward $\hat{r}_t = \hat{r}_{T(i_t, t)}^{i_t}$, provide reward \hat{r}_t as feedback to \mathcal{A}_{i_t} .

Set $\hat{\mu}_{T(i_t, t)}^{i_t} = \operatorname{argmin}_\theta \lambda \|\theta\|^2 + \sum_{\tau=1}^t (\langle a_\tau^{i_t}, \theta \rangle - \hat{r}_\tau^{i_t})^2$.

Set $z_t^{i_t} = \langle a_{T(i_t, t)}^{i_t}, \hat{\mu}_{T(i_t, t)-1}^{i_t} \rangle - \hat{r}_t - \beta \sqrt{(a_t)^{\top} (M_{T(i_t, t)-1}^{i_t})^{-1} a_t}$

if $\sum_{\tau=1}^{T(i_t, t)} z_\tau^{i_t} > 2\sqrt{T(i_t, t) \log(T/\delta)}$ or $2 \sum_{\tau=1}^{T(i_t, t)} \beta \sqrt{(a_\tau^{i_t})^\top (M_{\tau-1}^{i_t})^{-1} a_\tau^{i_t}} > C_{i_t} T(i_t, t)^{\alpha_{i_t}}$ **then**
 $I_t = I_{t-1} - \{i_t\}$.

else

$I_t = I_{t-1}$.

end if

end for

Note that standard matrix analysis (e.g. see Lemma 17) shows that

$$\sum_{\tau=1}^T \beta \sqrt{(a_\tau^J)^\top (M_{\tau-1}^J)^{-1} a_\tau^J} \leq \tilde{O}(d_J \sqrt{T})$$

so that by setting $C_i = \tilde{O}(d_i^2)$ and $\alpha_i = 1/2$ for all i , Theorem 13 yields a regret bound of $\tilde{O}(d_J^2 \sqrt{T})$, just as we encountered in the case of stochastic contexts.

Proof. First, recall that by our linUCB update for \mathcal{A}_i (12), we have

$$U(i_t, t-1) = \langle a_t, \hat{\mu}_{T(i_t, t)-1}^{i_t} \rangle + \beta \sqrt{a_t^\top (M_{T(i_t, t)-1}^{i_t})^{-1} a_t} - \frac{R_{i_t}}{T}$$

Now, we write the regret:

$$\begin{aligned} \sum_{t=1}^T r_{\star, t} - r_t &= \sum_{t=1}^T r_{\star, t} - r_{T(i_t, t)}^{i_t} \\ &= \sum_{i_t=J} r_{\star, t} - r_{T(J, t)}^J + \sum_{i_t \neq J} r_{\star, t} - r_{T(i_t, t)}^{i_t} \end{aligned}$$

First let's consider the indices for which $i_t = J$. Notice that by Corollary 15 and Lemma 17, a well-specified learner \mathcal{A}_J is never eliminated with probability at least $1 - 2\delta$, and we further have by

Theorem 16:

$$\sum_{i_t=J} r_{\star,t} - r_{T(J,t)}^J \leq C_J T(J, T)^{\alpha_J} \leq R_J$$

Next, if $i_t \neq J$, we must have $U(i_t, t-1) \geq U(J, t-1)$ since $J \in I_t$ for all t . Then:

$$\begin{aligned} r_{\star,t} - r_{T(i_t,t)}^{i_t} &\leq r_{\star,t} - U(J, t-1) + U(i_t, t-1) - r_{T(i_t,t)}^{i_t} \\ &= r_{\star,t} - U(J, t-1) + \hat{\mu}_{T(i_t,t-1)}^{i_t} a_t + \beta \sqrt{a_t^\top (M_{T(i_t,t-1)}^{i_t})^{-1} a_t} - \frac{R_{i_t}}{T} - \hat{r}_{T(i_t,t)}^{i_t} + \hat{r}_{T(i_t,t)}^{i_t} - r_{T(i_t,t)}^{i_t} \\ &= r_{\star,t} - U(J, t-1) + z_{T(i_t,t)}^{i_t} + 2\beta \sqrt{a_t^\top (M_{T(i_t,t-1)}^{i_t})^{-1} a_t} - \frac{R_{i_t}}{T} + \hat{r}_{T(i_t,t)}^{i_t} - r_{T(i_t,t)}^{i_t} \end{aligned}$$

where the last inequality follows from the definition of z_t^i .

Now, $r_{\star,t} - U(J, t-1) - \frac{R_J}{T} = r_{\star,t} - \max_{a \in A} \left[\hat{\mu}_{T(J,t-1)}^J a + \beta \sqrt{a^\top (M_{T(J,t-1)}^J)^{-1} a} \right] \leq 0$, with probability $1 - \delta$ by Corollary 15. Therefore, conditioned on this $1 - \delta$ event, we have

$$\begin{aligned} \sum_{i_t \neq J} r_{\star,t} - r_{T(i_t,t)}^{i_t} &= \sum_{i \neq J} \sum_{t|i_t=i} r_{\star,t} - r_{T(i_t,t)}^{i_t} \\ &\leq \sum_{i \neq J} \sum_{t|i_t=i} \frac{R_J}{T} + z_{T(i_t,t)}^{i_t} + 2\beta \sqrt{a_t^\top (M_{T(i_t,t-1)}^i)^{-1} a_t} - \frac{R_i}{T} + \hat{r}_{T(i_t,t)}^i - r_{T(i_t,t)}^i \\ &\leq R_J + \sum_{i \neq J} \left[-\frac{T(i, T)R_i}{T} + 2\sqrt{T(i, T) \log(T/\delta)} + \sum_{t|i_t=i} \hat{r}_{T(i_t,t)}^i - r_{T(i_t,t)}^i \right. \\ &\quad \left. + \sum_{t|i_t=i} 2\beta \sqrt{a_t^\top (M_{T(i_t,t-1)}^i)^{-1} a_t} \right] \\ &\leq R_J + \sum_{i \neq J} \left[-\frac{T(i, T)R_i}{T} + 2\sqrt{T(i, T) \log(T/\delta)} + \sqrt{8 \log(T^3 N/\delta) T(i, T)} \right. \\ &\quad \left. + \sum_{t|i_t=i} 2\beta \sqrt{a_t^\top (M_{T(i_t,t-1)}^i)^{-1} a_t} \right] \\ &\leq R_J + \sum_{i \neq J} \left[-\frac{T(i, T)R_i}{T} + \sqrt{24 \log(T^3 N/\delta) T(i, T)} \right. \\ &\quad \left. + \sum_{t|i_t=i} 2\beta \sqrt{a_t^\top (M_{T(i_t,t-1)}^i)^{-1} a_t} \right] \end{aligned}$$

where the third inequality is from equation (13), and the last inequality follows from Lemma 8

Now, $\sum_{t|i_t=i} 2\beta \sqrt{a_t^\top (M_{T(i_t,t-1)}^i)^{-1} a_t} \leq C_i T(i, T)^{\alpha_i}$, where the inequality follows from the second condition for being in I_t (14).

Therefore,

$$\begin{aligned} \sum_{i_t \neq J} r_{\star, t} - r_{T(i_t, t)}^{i_t} &\leq R_J + \sum_{i \neq J} \left[-\frac{T(i, T)R_i}{T} + \sqrt{24 \log(T^3 N/\delta) T(i, T)} + C_i T(i, T)^{\alpha_i} \right] \\ &\leq R_J + \sum_{i \neq J} \max \left\{ \sup_Z \left[-\frac{ZR_i}{T} + 2\sqrt{24 \log(T^3 N/\delta) Z} \right], \sup_Z \left[-\frac{ZR_i}{T} + 2C_i T(i, T)^{\alpha_i} \right] \right\} \end{aligned}$$

By Lemma 12:

$$\begin{aligned} &\leq R_J + \sum_{i \neq J} \max \left[\frac{24T}{R_i}, \frac{(1 - \alpha_i) \alpha_i^{\frac{1}{1-\alpha_i}} (2C_i)^{\frac{1}{1-\alpha_i}} T^{\frac{\alpha_i}{1-\alpha_i}}}{R_i^{\frac{\alpha_i}{1-\alpha_i}}} \right] \\ &\leq 2R_J \end{aligned}$$

This component of the regret due to $i_t \neq J$ is now in the same form as in the proof of the Theorem 1 and can be bounded by $O(R_J)$ by an almost identical analysis.

And so, putting all this together we have with probability at least $1 - 3\delta$, $\sum_{t=1}^T r_{\star, t} - r_t \leq 3R_J$. Therefore with probability at least $1 - 3\delta$, the regret is bounded by $3R_j$. \square

H Adversarial linear contextual bandits analysis

In this section we analyze the standard confidence ellipsoid technique for linear contextual bandits in the case of adversarial contexts. We do not claim these results are novel (essentially similar and tighter analysis can be found in [16, 20]). We include this section for completeness and to complement the analysis in Section G. Our proof technique is however slightly different from these works - it is a variant on the online-to-confidence set conversion idea of [24] that might have some independent interest.

In the linear contextual bandit problem, in each round the adversary reveals a mapping ψ_t from the set of actions A to \mathbb{R}^d . There is an unknown fixed $\theta_\star \in \mathbb{R}^d$ such that the expected reward of action a is always $\langle \psi_t(a), \theta_\star \rangle$. We assume $\|\psi_t(a)\| \leq 1$ for all a and t and $\|\theta_\star\| \leq 1$. Let a_1, \dots, a_T be the actions chosen by the algorithm over T rounds. Then the regret is defined as:

$$\sum_{t=1}^T \max_{a \in A} \langle \psi_t(a) - \psi_t(a_t), \theta_\star \rangle$$

By abuse of notation, we will identify actions with their associated vectors in order to write a_t^\star and a_t instead of $\arg\max_a \psi_t(a)$ and $\psi_t(a_t)$. Thus, the regret is

$$\sum_{t=1}^T \langle a_t^\star - a_t, \theta_\star \rangle$$

After playing action a_t , we observe $y_t \in [-1, 1]^2$ with $\mathbb{E}[y_t | a_t] = \langle a_t, \theta_\star \rangle$.

²it is also possible to consider subgaussian y_t , but we stick with bounded values here for simplicity

Define $A_t = \lambda I + \sum_{\tau=1}^t a_\tau^\top a_\tau$, where $\lambda > 0$ is some scalar parameter, and I is the $d \times d$ identity matrix. Define

$$\hat{\mu}_t = \operatorname{argmin}_{\|\theta\| \leq 1} \frac{\lambda}{2} \|\theta\|^2 + \frac{1}{2} \sum_{\tau=1}^{t-1} (\langle a_\tau, \theta \rangle - y_\tau)^2$$

In round t , we choose

$$a_t = \operatorname{argmax}_a \langle a, \mu_{t-1} \rangle + \beta_{t-1} \sqrt{a^\top A_{t-1}^{-1} a} \quad (16)$$

for some appropriate scaling factor β_t we will choose shortly. We claim the following:

Lemma 14. *With probability at least $1 - \delta$,*

$$(\theta_\star - \hat{\mu}_t)^\top A_t (\theta_\star - \hat{\mu}_t) \leq 4160 \log \left(\frac{\log_2(\sqrt{T/\log(1/\delta)}) + 2}{\delta} \right) + 6\lambda + 16d \log(1 + T/\lambda)$$

Proof. Observe that $\hat{\mu}_t$ is the output of the follow-the-regularized-leader algorithm (e.g. [25]) on the losses $\ell_t(\mu) = \frac{1}{2} (\langle a_t, \mu \rangle - y_t)^2$ with constant regularizer $\frac{\lambda}{2} \|\mu\|^2$ restricted to vectors of norm at most 1. Thus, by [25] Theorem 2, we have

$$\begin{aligned} \sum_{t=1}^T \frac{1}{2} (\langle a_t, \mu_t \rangle - y_t)^2 - \frac{1}{2} (\langle a_t, \theta_\star \rangle - y_t)^2 &\leq \frac{\lambda}{2} \|\theta_\star\|^2 + \sum_{t=1}^T (\langle a_t, \mu_t \rangle - y_t)^2 \frac{a_t^\top A_t^{-1} a_t}{2} \\ &\leq \frac{\lambda}{2} + 4 \sum_{t=1}^T \frac{a_t^\top A_t^{-1} a_t}{2} \\ &\leq \frac{\lambda}{2} + 4d \log(1 + T/\lambda) \end{aligned}$$

where the last line is by [26] Lemma 11.

Now, let $z_t = y_t - \langle a_t, \theta_\star \rangle$. Observe that $\mathbb{E}[z_t] = 0$ and z_t is 2-subgaussian. Therefore:

$$\sum_{t=1}^T \frac{1}{2} (\langle a_t, \hat{\mu}_t \rangle - y_t)^2 - \frac{1}{2} (\langle a_t, \theta_\star \rangle - y_t)^2 = \sum_{t=1}^T \frac{1}{2} \langle a_t, \hat{\mu}_t - \theta_\star \rangle^2 - z_t \langle a_t, \hat{\mu}_t - \theta_\star \rangle$$

Now by Lemma 18, with probability at least $1 - \delta$ we have:

$$\begin{aligned} &\geq \sum_{t=1}^T \frac{1}{2} \langle a_t, \hat{\mu}_t - \theta_\star \rangle^2 - 16 \sqrt{\log \left(\frac{\log_2(\sqrt{T/\log(1/\delta)}) + 2}{\delta} \right) \sum_{t=1}^T \langle a_t, \hat{\mu}_t - \theta_\star \rangle^2} \\ &\quad - 16 \log \left(\frac{\log_2(\sqrt{T/\log(1/\delta)}) + 2}{\delta} \right) \end{aligned}$$

Next, use the identity $A - B\sqrt{A} \geq \frac{A}{2} - 2B^2$:

$$\geq \frac{1}{4} \sum_{t=1}^T \langle a_t, \hat{\mu}_t - \theta_\star \rangle^2 - 1040 \log \left(\frac{\log_2(\sqrt{T/\log(1/\delta)}) + 2}{\delta} \right)$$

Thus with probability at least $1 - \delta$:

$$\begin{aligned} \sum_{t=1}^T \frac{1}{4} \langle a_t, \hat{\mu}_t - \theta_\star \rangle^2 &\leq 1040 \log \left(\frac{\log_2(\sqrt{T/\log(1/\delta)}) + 2}{\delta} \right) + \frac{\lambda}{2} + 4d \log(1 + T/\lambda) \\ \lambda \|\hat{\mu}_t - \theta_\star\|^2 + \sum_{t=1}^T \langle a_t, \hat{\mu}_t - \theta_\star \rangle^2 &\leq 4160 \log \left(\frac{\log_2(\sqrt{T/\log(1/\delta)}) + 2}{\delta} \right) + 6\lambda + 16d \log(1 + T/\lambda) \\ (\hat{\mu}_t - \theta_\star)^\top A_t (\hat{\mu}_t - \theta_\star) &\leq 4160 \log \left(\frac{\log_2(\sqrt{T/\log(1/\delta)}) + 2}{\delta} \right) + 6\lambda + 16d \log(1 + T/\lambda) \end{aligned}$$

□

This Lemma has an important corollary:

Corollary 15. *Define*

$$\beta_t^2 = 4160 \log \left(\frac{T \log_2(\sqrt{T/\log(T/\delta)}) + 2}{\delta} \right) + 6\lambda + 16d \log(1 + T/\lambda)$$

and

$$z_t = \langle a_t, \hat{\mu}_t \rangle - y_t - \beta_t \sqrt{a_t^\top A_t a_t}$$

Then

$$\mathbb{E}[z_t] \leq 0$$

and with probability at least $1 - 2\delta$

$$\sum_{t=1}^{\tau} z_t \leq 2\sqrt{\tau \log(T/\delta)}$$

and also

$$\max_a \langle a, \theta_\star \rangle - \max_a \langle a, \hat{\mu}_{t-1} \rangle + \beta_t \sqrt{a^\top A_{t-1}^{-1} a} \leq 0$$

for all $\tau \leq T$

Proof. First, we claim that with probability at least $1 - \delta$, for any a :

$$\langle a, \hat{\mu}_{t-1} - \theta_\star \rangle - \beta_t \sqrt{a^\top A_{t-1} a} \leq 0$$

for all t . To see this, notice that by Lemma 14, we have with probability $1 - \delta$ for all t .

$$(\theta_\star - \hat{\mu}_t)^\top A_t (\theta_\star - \hat{\mu}_t) \leq \beta_t^2$$

Therefore

$$\begin{aligned} \langle a, \theta_\star \rangle &= \langle a, \hat{\mu}_{t-1} \rangle + \langle a, \theta_\star - \hat{\mu}_{t-1} \rangle \\ &\leq \langle a, \hat{\mu}_{t-1} \rangle + \beta_t \sqrt{a^\top A_{t-1}^{-1} a} \end{aligned}$$

This shows the last claim of the Corollary. Now let $x_t = \langle a_t, \theta_\star \rangle - y_t$. Notice that $\mathbb{E}[x_t] = 0$ and $x_t \in [-2, 2]$. Thus we have with probability at least $1 - \delta$,

$$\sum_{t=1}^{\tau} x_t \leq 2\sqrt{\tau \log(T/\delta)}$$

for all τ . Combining all together proves the Corollary. \square

Next, we need to bound the actual regret of this algorithm:

Theorem 16. *Set*

$$\beta_t^2 = 4160 \log \left(\frac{T \log_2(\sqrt{T/\log(T/\delta)}) + 2}{\delta} \right) + 6\lambda + 16d \log(1 + T/\lambda)$$

Then, if $\lambda \geq 2$, with probability at least $1 - \delta$:

$$\begin{aligned} \sum_{t=1}^T \langle a_t^\star - a_t, \theta_\star \rangle &\leq 1 \sum_{t=1}^T \beta_t \sqrt{a_t^\top A_{t-1}^{-1} a_t} \\ &\leq \beta_T \sqrt{dT \log(1 + 2T/\lambda)} \\ &= \tilde{O}(d\sqrt{T}) \end{aligned}$$

Proof. By Lemma 14, we have

$$(\theta_\star - \hat{\mu}_t)^\top A_t (\theta_\star - \hat{\mu}_t) \leq \beta_t^2$$

for all t with probability at least $1 - \delta$. Therefore conditioned on this $1 - \delta$ probability event, for any a :

$$\begin{aligned} \langle a, \theta_\star \rangle &= \langle a \hat{\mu}_{t-1} \rangle + \langle a, \theta_\star - \hat{\mu}_{t-1} \rangle \\ &\leq \langle a, \hat{\mu}_{t-1} \rangle + \beta_{t-1} \sqrt{a^\top A_{t-1}^{-1} a} \end{aligned}$$

so by definition of a_t :

$$\leq \langle a_t, \hat{\mu}_{t-1} \rangle + \beta_{t-1} \sqrt{a_t^\top A_{t-1}^{-1} a_t}$$

Similarly, we have

$$\begin{aligned} \langle a_t, \theta_\star \rangle &= \langle a_t, \hat{\mu}_{t-1} \rangle + \langle a_t, \theta_\star - \hat{\mu}_{t-1} \rangle \\ &\geq \langle a_t, \hat{\mu}_{t-1} \rangle - \beta_t \sqrt{a_t^\top A_{t-1}^{-1} a_t} \end{aligned}$$

Therefore:

$$\langle a_t^\star - a_t, \theta_\star \rangle \leq 2\beta_t \sqrt{a_t^\top A_{t-1}^{-1} a_t}$$

Thus by Cauchy-Schwarz and the monotonicity of β_t :

$$\sum_{t=1}^T \langle a_t^* - a_t, \theta_* \rangle \leq 2\beta_T \sqrt{T \sum_{t=1}^T a_t^\top A_{t-1}^{-1} a_t}$$

Now define $\hat{A}_t = \frac{\lambda}{2}I + \sum_{\tau=1}^t a_\tau^\top a_\tau$. Since $\|a_t\| \leq 1$, if we set $\lambda \geq 2$, we must have

$$a_t^\top A_{t-1}^{-1} a_t \leq a_t^\top \hat{A}_t^{-1} a_t$$

Now, again by [26] Lemma 11, we have:

$$\sum_{t=1}^T a_t^\top \hat{A}_t^{-1} a_t \leq d \log(1 + 2T/\lambda)$$

which concludes the Theorem. □

Finally, one more check:

Lemma 17. *Suppose $\lambda \geq 2$. Then for all t ,*

$$\sum_{\tau=1}^t \beta_\tau \sqrt{a_\tau^\top A_{\tau-1}^{-1} a_\tau} \leq \beta_t \sqrt{dt \log(1 + 2t/\lambda)}$$

Proof. Since $\|a_t\| \leq 1$, we have

$$a_t^\top A_{t-1}^{-1} a_t \leq a_t^\top \hat{A}_{t-1}^{-1} a_t$$

where

$$\hat{A}_t = \frac{\lambda}{2}I + \sum_{\tau=1}^t a_\tau^\top a_\tau$$

Then by Cauchy-Schwarz and the monotonicity of β_t :

$$\begin{aligned} \sum_{\tau=1}^t \beta_\tau \sqrt{a_\tau^\top A_{\tau-1}^{-1} a_\tau} &\leq \beta_t \sqrt{t} \sqrt{\sum_{\tau=1}^t a_\tau^\top A_{\tau-1}^{-1} a_\tau} \\ &\leq \beta_t \sqrt{t} \sqrt{\sum_{\tau=1}^t a_\tau^\top \hat{A}_{\tau-1}^{-1} a_\tau} \end{aligned}$$

now use [26] Lemma 11 once again:

$$\leq \beta_t \sqrt{dt \log(1 + 2t/\lambda)}$$

□

Lemma 18. Suppose x_1, \dots, X_T are arbitrary random variables with $|x_t| \leq 1$ almost surely. Suppose z_1, \dots, z_T are such that $\mathbb{E}[z_t | x_1, \dots, x_t] = 0$ and z_t is 1-subgaussian given x_1, \dots, x_t . Then with probability at least $1 - \delta$:

$$\sum_{t=1}^T z_t x_t \geq -4 \sqrt{\log \left(\frac{\log_2(\sqrt{T/\log(1/\delta)}) + 2}{\delta} \right)} \sum_{t=1}^T x_t^2 + 4 \log \left(\frac{\log_2(\sqrt{T/\log(1/\delta)}) + 2}{\delta} \right)$$

Proof. The proof should follow from standard inequalities. Here we just check that x_t being random does not cause a significant problem... Define $S_\tau = \sum_{t=1}^\tau z_t x_t$. Note that S_1, S_2, \dots is a martingale. For any η , we have:

$$\begin{aligned} P[S_T \geq \epsilon] &\leq \exp^{-\eta\epsilon} \mathbb{E}[\exp(\eta S_T)] \\ &= \exp^{-\eta\epsilon} \mathbb{E}[\exp(\eta S_{T-1}) \exp(\eta z_T x_T)] \\ &= \exp^{-\eta\epsilon} \mathbb{E}[\exp(\eta S_{T-1}) \mathbb{E}[\exp(\eta z_T x_T) | x_1, \dots, x_T]] \\ &\leq \exp(-\eta\epsilon) \mathbb{E}[\exp(\eta S_{T-1}) \exp(\eta^2 x_T^2 / 2)] \end{aligned}$$

repeating for T steps:

$$\leq \exp(-\eta\epsilon + \eta^2 \sum_{t=1}^T x_t^2 / 2)$$

Therefore with probability at least $1 - \delta$,

$$S_T \leq \frac{\log(1/\delta)}{\eta} + \eta \sum_{t=1}^T x_t^2$$

Then, considering $\eta = 2^k$ for $-\lceil \log_2(\sqrt{T/\log(1/\delta)}) \rceil \leq k \leq 0$, we have that with probability at least $1 - \delta$,

$$\begin{aligned} S_T &\leq \min_{-\lceil \log_2(\sqrt{T/\log(1/\delta)}) \rceil \leq k \leq 0} \frac{\log \left(\frac{\log_2(\sqrt{T/\log(1/\delta)}) + 2}{\delta} \right)}{2^k} + 2^k \sum_{t=1}^T x_t^2 \\ &\leq 2 \inf_{\eta \in [\sqrt{\log(1/\delta)/T}, 1]} \frac{\log \left(\frac{\log_2(\sqrt{T/\log(1/\delta)}) + 2}{\delta} \right)}{\eta} + \eta \sum_{t=1}^T x_t^2 \\ &\leq 4 \sqrt{\log \left(\frac{\log_2(\sqrt{T/\log(1/\delta)}) + 2}{\delta} \right)} \sum_{t=1}^T x_t^2 + 4 \log \left(\frac{\log_2(\sqrt{T/\log(1/\delta)}) + 2}{\delta} \right) \end{aligned}$$

□