

---

# Oracle-Efficient Algorithms for Online Linear Optimization with Bandit Feedback\*

---

**Shinji Ito<sup>†</sup>**

NEC Corporation, The University of Tokyo  
i-shinji@nec.com

**Daisuke Hatano**

RIKEN AIP  
daisuke.hatano@riken.jp

**Hanna Sumita**

Tokyo Metropolitan University  
sumita@tmu.ac.jp

**Kei Takemura**

NEC Corporation  
kei\_takemura@nec.com

**Takuro Fukunaga<sup>‡</sup>**

Chuo University, RIKEN AIP, JST PRESTO  
fukunaga.07s@g.chuo-u.ac.jp

**Naonori Kakimura<sup>§</sup>**

Keio University  
kakimura@math.keio.ac.jp

**Ken-ichi Kawarabayashi<sup>§</sup>**

National Institute of Informatics  
k-keniti@nii.ac.jp

## Abstract

We propose computationally efficient algorithms for *online linear optimization with bandit feedback*, in which a player chooses an *action vector* from a given (possibly infinite) set  $\mathcal{A} \subseteq \mathbb{R}^d$ , and then suffers a loss that can be expressed as a linear function in action vectors. Although existing algorithms achieve an optimal regret bound of  $\tilde{O}(\sqrt{T})$  for  $T$  rounds (ignoring factors of  $\text{poly}(d, \log T)$ ), computationally efficient ways of implementing them have not yet been specified, in particular when  $|\mathcal{A}|$  is not bounded by a polynomial size in  $d$ . A standard way to pursue computational efficiency is to assume that we have an efficient algorithm referred to as *oracle* that solves (offline) linear optimization problems over  $\mathcal{A}$ . Under this assumption, the computational efficiency of a bandit algorithm can then be measured in terms of *oracle complexity*, i.e., the number of oracle calls. Our contribution is to propose algorithms that offer optimal regret bounds of  $\tilde{O}(\sqrt{T})$  as well as low oracle complexity for both *non-stochastic settings* and *stochastic settings*. Our algorithm for non-stochastic settings has an oracle complexity of  $\tilde{O}(T)$  and is the first algorithm that achieves both a regret bound of  $\tilde{O}(\sqrt{T})$  and an oracle complexity of  $\tilde{O}(\text{poly}(T))$ , given only linear optimization oracles. Our algorithm for stochastic settings calls the oracle only  $O(\text{poly}(d, \log T))$  times, which is smaller than the current best oracle complexity of  $O(T)$  if  $T$  is sufficiently large.

---

\*This work was supported by JST, ERATO, Grant Number JPMJER1201, Japan.

<sup>†</sup>This work was supported by JST, ACT-I, Grant Number JPMJPR18U5, Japan.

<sup>‡</sup>This work was supported by JST, PRESTO, Grant Number JPMJPR1759, Japan.

<sup>§</sup>This work was supported by JSPS, KAKENHI, Grant Number JP18H05291, Japan.

# 1 Introduction

*Online linear optimization with bandit feedback*, or *bandit linear optimization*, is an important problem that has a wide range of applications. In it, a player is given  $\mathcal{A} \subseteq \mathbb{R}^d$ , referred to as a set of *action vectors*, and  $T$ , the number of *rounds* of decision-making. In each round  $t \in [T] := \{1, 2, \dots, T\}$ , the player chooses an action  $a_t \in \mathcal{A}$ , and then observes loss  $\ell_t^\top a_t$ , where  $\ell_t \in \mathbb{R}^d$  is an unknown *loss vector* that can change over rounds. The bandit linear optimization includes a variety of important online decision-making problems as special cases. For example, given a graph  $G = (V, E)$  and  $s, t \in V$ , by setting  $\mathcal{A} \subseteq \mathbb{R}^{|E|}$  to be the set of all characteristic vectors of  $s$ - $t$  paths, we can take into account *bandit shortest path* or *adaptive routing* [9]. In this setting,  $\ell_t \in \mathbb{R}^{|E|}$  corresponds to (unknown) lengths of the edges, and bandit feedback  $\ell_t^\top a_t$  represents the length of a chosen  $s$ - $t$  path  $a_t$ . In addition to this application, bandit linear optimization includes bandit versions of combinatorial optimization problems such as minimum spanning tree, minimum cut, and knapsack problem, as well as continuous optimization problems such as linear programming and semidefinite programming.

The performance of the player is evaluated in terms of *regret*  $R_T(a^*)$ , defined as  $R_T(a^*) = \sum_{t=1}^T \ell_t^\top a_t - \sum_{t=1}^T \ell_t^\top a^*$  for  $a^* \in \mathcal{A}$ , which represents the difference between the cumulative loss for decision  $\{a_t\}$  of the player and that for an arbitrarily fixed strategy  $a^*$ . The primary goal in bandit linear optimization is to achieve small regret for arbitrary  $a^* \in \mathcal{A}$ . Some existing algorithms achieve regret bounds of  $\tilde{O}(\sqrt{T})$ ,<sup>1</sup> as shown in Tables 1 and 2. In contrast, papers [6; 8; 12; 21] showed that any algorithm will suffer at least  $\Omega(\sqrt{T})$  regret in the worst case. Thus, algorithms with  $\tilde{O}(\sqrt{T})$ -regret bounds achieve *optimal* performance w.r.t. dependence on  $T$ .

Algorithms achieving an optimal  $\tilde{O}(\sqrt{T})$ -regret, however, have computational issues, especially when the action set  $\mathcal{A}$  is exponentially large or is an infinite set. For example, well-known LinUCB methods [1; 16; 29] need to solve quadratic programming over  $\mathcal{A}$ , which has time complexity of  $\Omega(|\mathcal{A}|)$  if there are no additional assumptions. The ComBand algorithm [11] runs efficiently if there is an efficient sampling algorithm for  $\mathcal{A}$  (such as  $k$ -sets, spanning trees, or bipartite perfect matchings), but such sampling algorithms are open for many important examples, including  $s$ - $t$  paths. For the special case in which the convex hull of  $\mathcal{A}$  can be expressed by  $c$  linear inequalities, CombExp [13] runs in  $O(\text{poly}(c, d)T)$ -time. However,  $c$  (the size of the linear inequality expression) can be exponentially large for many examples.

In this study, we aim to develop computationally efficient algorithms that achieve an  $\tilde{O}(\sqrt{T})$  regret bound, under the assumption that we can call a *linear optimization oracle*. The oracle solves offline linear optimization problems over  $\mathcal{A}$ , i.e, given a loss vector  $\ell \in \mathbb{R}^d$ , the oracle outputs  $a^* \in \arg \min_{a \in \mathcal{A}} \ell^\top a$ . This assumption is standard in the context of online optimization [15; 23]. Under it, the computational efficiency of online optimization algorithms is evaluated in terms of *oracle complexity*: the number of oracle calls for the linear optimization oracle.

For online linear optimization with *full information*, in which a player can observe all entries of  $\ell_t \in \mathbb{R}^d$  after choosing  $a_t$ , Kalai and Vempala [23] have proposed algorithms with an  $\tilde{O}(\sqrt{T})$ -regret bound and an oracle complexity of  $O(T)$ . Using this algorithm, McMahan and Blum [26] and Dani and Hayes [15] showed that one can achieve  $\tilde{O}(T^{2/3})$ -regret and  $O(T^{1/2})$ -oracle complexity for bandit linear optimization. However, it has been an open question as to whether or not we can achieve  $\tilde{O}(\sqrt{T})$ -regret and  $\tilde{O}(\text{poly}(T))$ -oracle complexity for bandit linear optimization, with only linear optimization oracles. In this study, we solve this open problem by proposing an algorithm that achieves  $\tilde{O}(\sqrt{T})$ -regret as well as  $\tilde{O}(T)$ -oracle complexity.

Here, we separately consider here two different settings for bandit linear optimization: a *non-stochastic setting* and a *stochastic setting*. In the non-stochastic setting, we do not assume any generative models, but  $\ell_t$  may be chosen in an adversarial manner, depending on previous actions  $a_1, \dots, a_{t-1}$ . The performance of an algorithm is measured in terms of the expectation of regret  $R_T(a^*)$  w.r.t. the randomness of the algorithm's internal randomness and  $\ell_t$ . In the stochastic setting, by way of contrast, the loss vectors  $\ell_t$  are assumed to follow a probability distribution  $D$  over  $\mathbb{R}^d$ , i.i.d. for  $t = 1, \dots, T$ .

<sup>1</sup> In  $\tilde{O}(\cdot)$  notation, we ignore factors of polynomials in  $d$  and  $\log(T)$ .

## 1.1 Our Contribution

In this paper, we present computationally efficient algorithms that achieve  $O(\text{poly}(d)\sqrt{T})$ -regret. Specifically, we present algorithms with a small oracle complexity, i.e., algorithms that call the oracle as infrequently as possible. Our contribution is summarized in Tables 1 and 2.

For the non-stochastic setting, we propose an algorithm (Algorithm 1) that achieves  $O(\sqrt{d^3 T \log T})$ -regret in expectation and has  $O(\text{poly}(d, \log T)T)$ -oracle complexity.

**Theorem 1.** *For the non-stochastic setting, Algorithm 1 satisfies the following conditions:*

- *The output of the algorithm satisfies  $\mathbf{E}[R_T(a^*)] = O(\sqrt{d^3 T \log T})$  for all  $a^* \in \mathcal{A}$ .*
- *The algorithm calls the linear optimization oracle  $O(\text{poly}(d, \log T)T)$  times.*
- *The computational time, except for that of the oracle, is of  $O(\text{poly}(d, T))$ .*

As shown in Table 1, our Algorithm 1 achieves the smallest oracle complexity among algorithms with  $\tilde{O}(\sqrt{T})$ -regret. Noting that GeometricHedge assumes  $\mathcal{A}$  to be a convex body, we can see that Algorithm 1 is the first algorithm that is applicable to discrete  $\mathcal{A}$  and that achieves  $\tilde{O}(\sqrt{T})$ -regret and  $\tilde{O}(\text{poly}(T))$ -oracle complexity.

Although the first algorithm in Table 1 with  $\tilde{O}(T^{2/3})$ -regret and  $O(T^{2/3})$ -oracle complexity might look incomparable to our results, algorithms with such bounds can be easily constructed from our Algorithm 1. In fact, by dividing  $T$  rounds into  $T/B$  blocks of size  $B > 1$  and regarding each block as an individual round, we obtain the following statement:

**Proposition 1.** *Suppose there exists an algorithm with  $O(f(T))$ -regret and  $O(g(T))$ -oracle complexity. Then, for arbitrary positive integer  $B$ , there exists an algorithm with  $O(B \cdot f(T/B))$ -regret and  $\tilde{O}(g(T/B))$ -oracle complexity.*

By setting the block size to be  $B = \Theta(T^{1/3})$  and applying Algorithm 1, we can achieve  $O(B\sqrt{T/B}) = \tilde{O}(T^{2/3})$ -regret and  $O(T/B) = \tilde{O}(T^{2/3})$ -oracle complexity, which is equivalent to the uppermost result in Table 1. Note that Proposition 1 does *not* lead to an  $\tilde{O}(\sqrt{T})$ -regret algorithm given an  $\tilde{O}(T^{2/3})$ -regret algorithm, conversely, since the block size  $B$  must be at least 1.

For the stochastic setting, we propose an algorithm (Algorithm 2) that achieves  $O(\sqrt{d^3 T \log(d \log T / \delta)})$ -regret with probability  $1 - \delta$  and has  $O(\text{poly}(d, \log T))$ -oracle complexity, where  $\delta \in (0, 1)$  is an arbitrary parameter.

**Theorem 2.** *Suppose  $\ell_t$  follows a distribution over  $\mathbb{R}^d$ , i.i.d. for  $t = 1, 2, \dots, T$ . Algorithm 2 then satisfies the following conditions:*

- *The output of the algorithm satisfies  $R_T(a^*) = O(\sqrt{d^3 T \log(d \log T / \delta)})$  for all  $a^* \in \mathcal{A}$ , with probability  $1 - \delta$ .*
- *The algorithm calls the linear optimization oracle  $O(\text{poly}(d, \log T))$  times.*
- *The computational time, except for that of the oracle, is of  $O(\text{poly}(d, T))$ .*

A complete description of Algorithm 2 and a proof of this theorem are given in Appendix B. As shown in Table 2, all existing algorithms that achieve  $\tilde{O}(\sqrt{T})$ -regret require at least  $\Omega(T)$  oracle complexity, and our Algorithm 2 is the first with an  $\tilde{O}(\sqrt{T})$ -regret bound and a sublinear oracle complexity in  $T$ .

In both Algorithms 1 and 2, we use the well-known techniques [30] of reduction among linear optimization, separation, and decomposition over a given convex body. Definitions of these three problems are given in Section 4. The reduction algorithms enable us to solve separation and decomposition problems by calling the linear optimization oracle  $O(\text{poly}(d))$  times. Using these

<sup>2</sup> In this algorithm,  $\mathcal{A}$  is assumed to be a convex body, and a membership oracle for  $\mathcal{A}$  is assumed. Because we can construct a membership oracle from a linear optimization oracle and vice versa by a polynomial-time reduction [30], the assumption regarding the oracle is equivalent to ours, modulo polynomial-time reduction.

Table 1: Regret Bound and Oracle Complexity of Non-Stochastic Bandit Linear Optimization

Algorithm	Regret Bound	Oracle Complexity
MV algorithm [15; 26] with FPL [23]	$\tilde{O}(T^{2/3})$	$O(T^{2/3})$
ComBand [11], GeometricHedge [17], Exp2 [6]	$\tilde{O}(T^{1/2})$	–
GeometricHedge with Volumetric Spanners <sup>2</sup> [19]	$\tilde{O}(T^{1/2})$	$\tilde{O}(T^7)$
<b>Algorithm 1 [This paper]</b>	$\tilde{O}(T^{1/2})$	$\tilde{O}(T)$

Table 2: Regret Bound and Oracle Complexity for Stochastic Bandit Linear Optimization

Algorithm	Regret Bound	Oracle Complexity
LinRel [7], LinUCB with $\ell_2$ -ball [1; 16; 29]	$\tilde{O}(T^{1/2})$	–
LinUCB with $\ell_1$ -ball [16],	$\tilde{O}(T^{1/2})$	$O(T)$
Linear Thompson sampling [2; 4]	$\tilde{O}(T^{1/2})$	$O(T)$
<b>Algorithm 2 [This paper]</b>	$\tilde{O}(T^{1/2})$	$O(\text{poly}(d, \log T))$

reduction techniques, Algorithms 1 and 2 maintain, respectively, supersets and subsets of the convex hull of  $\mathcal{A}$  ( $=: \text{Conv}(\mathcal{A})$ ).

To construct Algorithm 1 for the non-stochastic setting, we extend a *cutting-plane approach* to our bandit-feedback setting. The cutting-plane approach, a way of reducing oracle complexity, has been applied only to full-information settings [20], not a bandit-feedback setting. A major difference between bandit-feedback and full-information settings is that the former needs *exploration*, i.e., chosen actions should be randomized with sufficiently large variance, whereas the latter does not need it and chooses actions deterministically. In full-information settings, hence, it suffices to focus on a deterministically chosen action alone. In contrast to this, to deal with the bandit-feedback setting, the difficulty lies in constructing a distribution of actions with sufficiently large variance, for which cutting planes can be efficiently computed and the number of them can be bounded.

To this end, we design relevant probability distributions so that the cutting-plane approach works, which successfully reduces oracle complexity. Specifically, the cutting-plane approach maintains convex bodies  $\mathcal{K}_t$  that include and approximate  $\text{Conv}(\mathcal{A})$ , from which we choose *candidates* for actions, employing the support of the probability distribution of actions to choose. It is only when some candidates are invalid, i.e., when some are outside of  $\text{Conv}(\mathcal{A})$ , that  $\mathcal{K}_t$  is updated with a cutting plane excluding such an invalid candidate. To bound the number of oracle calls, we design candidates for actions that satisfy two conditions: the set of candidates has a bounded cardinality, and each candidate is sufficiently close to the weighted center of  $\mathcal{K}_t$ . Thanks to the first condition, we can efficiently decide if invalid candidates exist. The second condition is essential for bounding the number of oracle calls in each update of  $\mathcal{K}_t$ .

Our Algorithm 2 for the stochastic setting is based on the framework of *phased elimination of actions*, in which  $T$  rounds are divided into *phases* that are segments of consequent rounds, and, in each phase, actions are eliminated so that only promising ones are left. Existing works employing this framework [7; 24; 31] are computationally inefficient, mainly for the following two reasons: (i) We need to maintain a set of promising actions that may be an exponentially large combinatorial set, and, (ii) when choosing actions, we need to solve hard optimization problems, e.g., G-optimal design [24] or quadratic programming [7].

Our idea for resolving the first computational issue is to maintain the set of promising actions as a convex set instead of a subset of actions. The convex set here can be represented with only  $O(\text{poly}(d) \log T)$  linear inequalities, which implies that operations over it can be conducted efficiently. We resolve the second computational issue by combining *barycentric spanners* [9] and the decomposition technique over convex bodies [30], both of which are efficiently computed with  $O(\text{poly}(d))$  oracle calls. We show that, thanks to these techniques, we can estimate the loss vector with enough accuracy to achieve an  $\tilde{O}(\sqrt{T})$ -regret bound. The oracle complexity is bounded as follows: In our algorithm, all  $a_t$  chosen in each phase are determined at the beginning of the phase, which means that the oracle complexity depends not on the number of rounds, but on the number of phases. The number of phases is of  $O(\text{poly}(d) \log T)$  and that of oracle calls in each phase is of  $O(\text{poly}(d, \log T))$ , which results in overall  $O(\text{poly}(d, \log T))$ -oracle complexity.

## 2 Related Work

For the *full-information setting* in which a player can observe  $\ell_t \in \mathbb{R}^d$  rather than  $\ell_t^\top a_t$ , Follow the Perturbed Leader (FPL) by Kalai and Vempala [23] achieves  $O(\sqrt{T})$ -regret and  $O(T)$ -oracle complexity. This algorithm is used as a subroutine in MV algorithm [15; 26] (see Table 2).

For a more general problem referred to as *online improper learning*, in which only an approximate linear optimization oracle is given, Kakade et al. [22] have proposed the first efficient algorithms that achieve approximate regret of  $O(\sqrt{T})$  for the full-feedback setting, and  $O(T^{2/3})$  in the bandit-feedback setting. Recent papers by Garber [18] and Hazan et al. [20] have improved oracle complexity. Algorithms in [20] achieve oracle complexity of  $\tilde{O}(T)$  in the full-feedback setting and  $\tilde{O}(T^{2/3})$  in the bandit feedback setting with the same regret bound as in Kakade et al. [22]. For online improper learning with bandit feedback, however, constructing an efficient algorithm achieving  $\tilde{O}(\sqrt{T})$  poses difficulties that have yet to be overcome.

In addition to the studies listed in Tables 1 and 2, there exist efficient algorithms for bandit linear optimization that work under different assumptions. Abernethy et al. [3] proposed a computationally efficient algorithm achieving  $O(\sqrt{T})$ -regret under the assumption that  $\mathcal{A}$  is a convex body and that a self-concordant barrier [27] for  $\mathcal{A}$  is given. However, constructing self-concordant barriers is not always possible with a linear optimization oracle alone, and, hence, this algorithm does not always work under our assumptions of linear optimization oracle and Assumption 1 given in the next section.

## 3 Problem Setting

The *bandit linear optimization problem* is a repeated game described as follows: Before the game starts, a player is given the number  $T$  of rounds and the dimensionality  $d$  of the *action set*  $\mathcal{A} \subseteq \mathbb{R}^d$ . In each round  $t = 1, 2, \dots, T$ , the player chooses  $a_t \in \mathcal{A}$  while an environment chooses a loss vector  $\ell_t \in \mathbb{R}^d$ , and then the player observes a loss  $\ell_t^\top a_t$ . The goal of the player is to achieve a small *regret*  $R_T(a)$ , which is defined for an arbitrary  $a \in \mathcal{A}$  as  $R_T(a) := \sum_{t=1}^T \ell_t^\top a_t - \sum_{t=1}^T \ell_t^\top a$ .

We assume the action set  $\mathcal{A}$  to be a compact set. Suppose that we have an algorithm for linear optimization over  $\mathcal{A}$  for any vector  $w \in \mathbb{R}^d$ , which we call *linear optimization oracle*  $\mathcal{O}_{\mathcal{A}} : \mathbb{R}^d \rightarrow \mathcal{A}$  that receives an input  $w \in \mathbb{R}^d$  and returns a point  $\mathcal{O}_{\mathcal{A}}(w) \in \mathcal{A}$  satisfying  $w^\top \mathcal{O}_{\mathcal{A}}(w) = \min_{a \in \mathcal{A}} w^\top a$ .

*Assumption 1.* We assume that there exist positive real numbers  $L$  and  $R$  such that (a)  $\|\ell_t\|_2 \leq L$  for all  $t \in [T]$ , and (b)  $\|a\|_2 \leq R$  for all  $a \in \mathcal{A}$ . In addition, we assume that (c)  $\mathcal{K} := \text{Conv}(\mathcal{A})$  has a positive volume, i.e.,  $\text{Vol}(\mathcal{K}) := \int_{\mathcal{K}} 1 dx > 0$ .

The first two assumptions (a) and (b) are standard in bandit linear optimization. If we are given a linear optimization oracle over  $\mathcal{A}$ , we can assume (c) without loss of generality. In fact, if  $\mathcal{A}$  is included in a subspace with a smaller dimension than  $d$ , we can then detect it by calling the linear optimization oracle polynomial times (see, e.g., Corollary 14.1g in [30]), and we can make  $\mathcal{K}$  full-dimensional by ignoring redundant dimensions.

## 4 Preliminaries

### 4.1 Linear Optimization, Separation, and Decomposition

We define a *linear optimization problem* (LP), *separation problem* (SP), and *decomposition problem* (DP) for a compact convex body  $\mathcal{P} \subseteq \mathbb{R}^d$  as follows:

*Problem 1* (linear optimization problem, LP). Given a vector  $w \in \mathbb{R}^d$ , find a vector  $x^* \in \mathcal{P}$  such that  $w^\top x^* = \min_{x \in \mathcal{P}} w^\top x$ .

*Problem 2* (separation problem, SP). Given a vector  $y \in \mathbb{R}^d$ , decide whether  $y$  belong to  $\mathcal{P}$  or not, and, in the latter case, find a vector  $w \in \mathbb{R}^d$  such that  $w^\top y < \min_{x \in \mathcal{P}} w^\top x$ .

*Problem 3* (decomposition problem, DP). Given a vector  $x \in \mathcal{P}$ , find vertices  $x_0, \dots, x_d$  of  $\mathcal{P}$  and  $\lambda_0, \dots, \lambda_d \geq 0$  such that  $x = \lambda_0 x_0 + \dots + \lambda_d x_d$ .

Ellipsoid methods provide reductions among these problems, which imply that

$$\text{LP: solvable} \iff \text{SP: solvable} \implies \text{DP: solvable}.$$

**Theorem 3 (Corollaries 14.1a, 14.1b and 14.1g in [30]).** Suppose that  $\mathcal{P} \subseteq \mathbb{R}^d$  is a polytope of which each vertex can be expressed by rationals with bit-lengths of at most  $\varphi$ , and that each entry of  $x, y, w \in \mathbb{Q}^d$  is also a rational, with bit-length of at most  $\varphi$ . Then, the following holds:

- (a) If there is an algorithm SEP that solves the separation problem, we can solve the linear optimization problem for  $w \in \mathbb{Q}^d$  by calling SEP at most  $\text{poly}(d, \varphi)$  times.
- (b) If there is an algorithm OPT that solves the linear optimization problem, we can solve the separation problem for  $y \in \mathbb{Q}^d$  by calling OPT at most  $\text{poly}(d, \varphi)$  times.
- (c) If there is an algorithm OPT that solves the linear optimization problem, we can solve the decomposition problem for  $x \in \mathcal{P}$  by calling OPT at most  $\text{poly}(d, \varphi)$  times.

*Remark 1.* For any  $\varepsilon > 0$  and any real number  $x \in [-1, 1]$ , we can approximate  $x$  by a rational  $\hat{x} \in \mathbb{Q}$  with a bit-length of at most  $O(\log(1/\varepsilon))$  so that  $|x - \hat{x}| \leq \varepsilon$ . Hence, we can assume that  $\varphi$  in Theorem 3 is bounded as  $\varphi = O(\log T)$  by ignoring  $O(1/T)$  errors. This implies that the above reductions can be computed in  $O(\text{poly}(d, \log T))$  time.

## 4.2 Algorithms for Logconcave Distributions

If a probability distribution over convex body  $\mathcal{P} \subseteq \mathbb{R}^d$  has a probability density function (PDF)  $p : \mathcal{P} \rightarrow \mathbb{R}_{>0}$  such that  $\log p$  is a concave function, we refer to it as a *logconcave distribution*. The following theorem means that, given the value oracle of a convex function  $f : \mathcal{P} \rightarrow \mathbb{R}$ , we can approximately sample a vector in  $\mathcal{P}$  from a logconcave distribution  $p(x) \propto \exp(-f(x))$ .

**Theorem 4** (Theorems 2.1 and 2.2 in [25], Lemma 10 in [19]). Let  $\mathcal{P} \subseteq \mathbb{R}^d$  be a convex body with non-zero Lebesgue measure, and let  $f : \mathcal{P} \rightarrow \mathbb{R}$  be a convex function and let  $p$  be a logconcave distribution proportional to  $\exp(-f(x))$ . Suppose  $\varepsilon > 0$  and  $\delta \in (0, 1)$ . Then, given access to the membership oracle of  $\mathcal{P}$  and the value oracle of  $f$ , there is an algorithm that samples approximately from  $p$  such that (i) the total variation distance between the produced distribution and  $p$  is at most  $\varepsilon$ , and (ii) after preprocessing for a time of  $O(d^5 (\log d)^{O(1)})$ , each sample can be produced in a time of  $O(d^4 / \varepsilon^4 \cdot (\log(d/\varepsilon))^{O(1)})$ .

As an implication of this theorem, we can efficiently approximate mean  $\mu(p) \in \mathbb{R}^d$  and covariance matrix  $\text{Cov}(p) \in \mathbb{R}^{d \times d}$  of distribution  $p$ . In fact, from Corollary 5.52 in [32] and standard concentration of logconcave distribution (see, e.g., Lemma 5.17 in [25]), it takes  $(n \log(1/\delta)/\varepsilon)^{O(1)}$  samples to get a matrix  $\hat{\Sigma}$  such that  $(1 - \varepsilon)\text{Cov}(p) \preceq \hat{\Sigma} \preceq (1 + \varepsilon)\text{Cov}(p)$  with probability of at least  $1 - \delta$ .<sup>3</sup> Similarly, we can get  $\hat{\mu} \in \mathbb{R}^d$  such that  $\|\hat{\mu} - \mu(p)\|_{\text{Cov}(p)^{-1}} \leq \varepsilon$  from  $(n \log(1/\delta)/\varepsilon)^{O(1)}$  samples.<sup>4</sup> Accordingly, we obtain the following corollary:

**Corollary 1.** Suppose the same assumption as in Theorem 4. There is an algorithm that outputs a vector  $\hat{\mu} \in \mathbb{R}^d$  and a symmetric matrix  $\hat{\Sigma} \in \mathbb{R}^{d \times d}$  satisfying  $\frac{1}{2}\text{Cov}(p) \preceq \hat{\Sigma} \preceq 2\text{Cov}(p)$  and  $\|\hat{\mu} - \mu(p)\|_{\text{Cov}(p)^{-1}} \leq \varepsilon$  with a probability of at least  $1 - \delta$ . The computational time, the number of calls for the membership oracle of  $\mathcal{P}$ , and the value oracle of  $f$  are bounded by  $\text{poly}(d, \frac{1}{\varepsilon}, \log \frac{1}{\delta})$ .

## 5 Algorithm for Non-stochastic Bandit Linear Optimization

Our algorithm uses the framework of a continuous multiplicative weight update (CMWU) [5; 14; 33]. A straightforward way of applying CMWU is to maintain probability distributions over  $\mathcal{K} := \text{Conv}(\mathcal{A})$ , which, however, requires a large number of oracle calls. In fact, the algorithm by Hazan and Karnin [19] for bandit linear optimization over convex bodies calls an oracle  $\tilde{O}(T^7)$  times. This inefficiency is due to that we need to sample from  $\mathcal{K}$ ; the sampling algorithm in Theorem 4 requires  $O(d^4/\varepsilon^4)$ -oracle complexity.

We reduce oracle complexity by means of a *cutting-plane approach* [20]. In this approach, we maintain convex bodies  $\mathcal{K}_t^{(j)}$  that include and approximate  $\mathcal{K}$ , and we update a distribution over  $\mathcal{K}_t^{(j)}$  instead of  $\mathcal{K}$ . The advantage of this approach is that we can sample from  $\mathcal{K}_t^{(j)}$  without calling

<sup>3</sup>A similar argument can be found in Section 6.3 in [10].

<sup>4</sup> For a vector  $x \in \mathbb{R}^d$  and a positive semidefinite matrix  $A \in \mathbb{R}^{d \times d}$ , denote  $\|x\|_A := \sqrt{x^\top A x}$ .

an oracle. On the other hand, updating  $\mathcal{K}_t^{(j)}$  requires oracle calls; therefore, we need to bound the number of the updates as well as the number of oracle calls in each update. We design a strategy achieving these as follows: We set *candidates of actions*  $\mathcal{E}_t^{(j)} \subseteq \mathcal{K}_t^{(j)}$ , from which we choose action. When some actions among the candidates are invalid, i.e., outside of  $\mathcal{K}$ , we then reduce  $\mathcal{K}_t^{(j)}$  by a cutting plane excluding such an invalid candidate. With this strategy, we need oracle calls to check if invalid candidates exist. Our algorithm bounds the oracle complexity here by setting  $\mathcal{E}_t^{(j)}$  to have  $O(d)$  elements. Further, we design  $\mathcal{E}_t^{(j)}$  so that its elements are sufficiently close to the weighted center of  $\mathcal{K}_t^{(j)}$ . This plays an important role in bounding the number of updates of  $\mathcal{K}_t^{(j)}$ . Indeed, when a convex body is updated by a cutting plane that excludes a point close to its center, its volume then decreases by a constant factor less than 1 (see, e.g., [25]). On the other hand,  $\mathcal{K}_t^{(j)}$  always includes  $\mathcal{K}$  with a positive volume; hence, the volume of  $\mathcal{K}_t^{(j)}$  cannot be smaller than that of  $\mathcal{K}$ , which implies that the number of updates is bounded.

## 5.1 Algorithm

Our algorithm maintains a convex body  $\mathcal{K}_t^{(j)} \subseteq \mathbb{R}^d$  such that  $\mathcal{K}_1^{(0)} \supseteq \mathcal{K}_1^{(1)} \supseteq \dots \supseteq \mathcal{K}_1^{(s_1)} = \mathcal{K}_2^{(0)} \supseteq \mathcal{K}_2^{(1)} \supseteq \dots \supseteq \mathcal{K}_2^{(s_2)} \supseteq \dots \supseteq \mathcal{K}_T^{(s_T)} \supseteq \mathcal{K} = \text{Conv}(\mathcal{A})$ , where  $t$  corresponds to the round,  $j \in \{0, 1, \dots\}$  is an index, and  $s_t \in \{0, 1, \dots, T\}$  will be defined later. It also updates a logconcave function  $z_t : \mathbb{R}^d \rightarrow \mathbb{R}_{>0}$  in each round  $t$  based on the multiplicative weight update [5; 14; 33]. Before the first round,  $z_t$  is initialized to be a constant function  $z_1(x) = 1$ . Let  $q_t^{(j)}$  denote the PDF of a distribution over  $\mathcal{K}_t^{(j)}$  that is proportional to the function  $z_t$ , i.e.,

$$Z_t^{(j)} = \int_{\mathcal{K}_t^{(j)}} z_t(x) dx, \quad q_t^{(j)}(x) = \begin{cases} \frac{z_t(x)}{Z_t^{(j)}} & \text{if } x \in \mathcal{K}_t^{(j)}, \\ 0 & \text{if } x \in \mathbb{R}^d \setminus \mathcal{K}_t^{(j)}. \end{cases} \quad (1)$$

Let us denote the mean and the covariance matrix for distribution of  $q_t^{(j)}$  by  $\mu_t^{(j)} \in \mathbb{R}^d$  and  $\Sigma_t^{(j)} \in \mathbb{R}^{d \times d}$ , respectively:  $\mu_t^{(j)} = \mathbf{E}_{a \sim q_t^{(j)}}[a]$ ,  $\Sigma_t^{(j)} = \mathbf{E}_{a \sim q_t^{(j)}}[(a - \mu_t^{(j)})(a - \mu_t^{(j)})^\top]$ . From Corollary 1, we can compute estimators  $\hat{\mu}_t^{(j)}$  and  $\hat{\Sigma}_t^{(j)}$  of  $\mu_t^{(j)}$  and  $\Sigma_t^{(j)}$ , respectively, such that

$$\frac{1}{2} \Sigma_t^{(j)} \preceq \hat{\Sigma}_t^{(j)} \preceq 2 \Sigma_t^{(j)}, \quad \|\hat{\mu}_t^{(j)} - \mu_t^{(j)}\|_{(\Sigma_t^{(j)})^{-1}} \leq \varepsilon \quad (2)$$

with probability of at least  $1 - \delta_t^{(j)}$ , where  $\varepsilon > 0$  and  $\delta_t^{(j)} \in (0, 1)$ , which will be defined later. Let  $B_t^{(j)} = (b_{t1}^{(j)}, \dots, b_{td}^{(j)}) \in \mathbb{R}^{d \times d}$  be a matrix such that  $B_t^{(j)} B_t^{(j)\top} = \hat{\Sigma}_t^{(j)}$ . Define  $\mathcal{E}_t^{(j)} \subseteq \mathbb{R}^d$  as

$$\mathcal{E}_t^{(j)} = \left\{ \hat{\mu}_t^{(j)} + \frac{1}{4e} b_{ti}^{(j)} \mid i \in [d] \right\} \cup \left\{ \hat{\mu}_t^{(j)} - \frac{1}{4e} b_{ti}^{(j)} \mid i \in [d] \right\}. \quad (3)$$

In each round  $t$ , our algorithm checks if  $\mathcal{E}_t^{(j)}$  is included in  $\mathcal{K}$ , and if not, it updates  $\mathcal{K}_t^{(j)}$ , as described in Step 7 of Algorithm 1, to exclude an element in  $\mathcal{E}_t^{(j)} \setminus \mathcal{K}$ . Set  $\mathcal{E}_t^{(j)}$  is designed so that the following four conditions are satisfied:

1. The cardinality of  $\mathcal{E}_t^{(j)}$  is bounded as  $|\mathcal{E}_t^{(j)}| = O(d)$ . Hence, we can decide if  $\mathcal{E}_t^{(j)} \subseteq \mathcal{K}$  by  $O(\text{poly}(d))$  oracle calls.
2. Each  $y \in \mathcal{E}_t^{(j)}$  is sufficiently close to  $\mu_t^{(j)}$ , i.e., it satisfies  $\|y - \mu_t^{(j)}\|_{(\Sigma_t^{(j)})^{-1}} \leq 1/(2e)$ . This is important to bound the number of oracle calls.
3. The mean of  $\mathcal{E}_t^{(j)}$  is equal to  $\hat{\mu}_t^{(j)}$ . This implies that if  $y$  follows a uniform distribution over  $\mathcal{E}_t^{(j)}$ , we then have  $\mathbf{E}[\ell_t^\top y] = \ell_t^\top \hat{\mu}_t^{(j)} \approx \ell_t^\top \mu_t^{(j)} = \mathbf{E}[\ell_t^\top x]$  for  $x \sim q_t^{(j)}$ .
4. The covariance matrix  $\Sigma$  of a uniform distribution over  $\mathcal{E}_t^{(j)}$  satisfies  $\Sigma \succeq O(1/d^2) \cdot \Sigma_t^{(j)}$ .

Thanks to this, empirical estimates of  $\ell_t$  based on  $\mathcal{E}_t^{(j)}$  will have a sufficiently small variance.

The conditions 1 and 2 are used to bound the oracle complexity, and 3 and 4 are necessary to bound the regret. Once  $\mathcal{E}_t^{(j)}$  is included in  $\mathcal{K}$ , our algorithm escapes the loop of updating  $\mathcal{K}_t^{(j)}$ . An integer  $s_t$

---

**Algorithm 1** An oracle efficient algorithm for non-stochastic bandit linear optimization

---

**Require:** Learning rate  $\eta > 0$ , error bound  $\varepsilon > 0$ , time horizon  $T \in \mathbb{N}$ ,  $R > 0$  satisfying Assumption 1.

```
1: Set  $\mathcal{K}_1^{(0)} = B_\infty(0, R) = \{x \in \mathbb{R}^d \mid \|x\|_\infty \leq R\}$  and define  $z_1 : \mathbb{R}^d \rightarrow \mathbb{R}_{>0}$  by  $z_1(x) = 1$ .
2: for  $t = 1, 2, \dots, T$  do
3:   for  $j = 0, 1, 2, \dots$  do
4:     Compute  $\mathcal{E}_t^{(j)}$  on the basis of (1)  $\sim$  (3).
5:     Solve SP for  $\mathcal{P} = \mathcal{K}$  and for each  $y \in \mathcal{E}_t^{(j)}$ .
6:     if There is a hyperplane  $w \in \mathbb{R}^d$  s.t.  $w^\top y < \min_{x \in \mathcal{K}} w^\top x$  for some  $y \in \mathcal{E}_t^{(j)}$  then
7:       Update  $\mathcal{K}_t^{(j)}$  by  $\mathcal{K}_t^{(j+1)} = \mathcal{K}_t^{(j)} \cap \{x \in \mathbb{R}^d \mid w^\top x \geq w^\top y\}$ .
8:     else
9:       Set  $s_t = j$  and  $\mathcal{K}_t = \mathcal{K}_t^{(s_t)}$ . Break the for loop w.r.t. the index  $j$ .
10:    end if
11:  end for
12:  Let  $\hat{\mu}_t = \hat{\mu}_t^{(s_t)}$  and  $b_{ti} = b_{ti}^{(s_t)}$  for  $i \in [n]$ , which are defined in (1) – (3).
13:  Choose  $i_t \in [d]$  and  $\sigma_t \in \{1, -1\}$  uniformly at random.
14:  Solve DP for  $\mathcal{P} = \mathcal{K}$  and  $x = \hat{\mu}_t + \frac{\sigma_t}{4e} b_{ti_t}$  to get a decomposition  $x_{t0}, \dots, x_{td} \in \mathcal{A}$  and  $\lambda_{t0}, \dots, \lambda_{td}$  such that  $\hat{\mu}_t + \frac{\sigma_t}{4e} b_{ti_t} = \lambda_{t0}x_{t0} + \dots + \lambda_{td}x_{td}$ .
15:  Play  $a_t = x_{ts}$  with probability  $\lambda_{ts}$  ( $s = 0, \dots, d$ ), and receive loss  $\ell_t^\top a_t$ .
16:  Set  $\hat{\ell}_t$  by (4) and update  $z_t$  by  $z_{t+1}(x) = z_t(x) \exp(-\eta \hat{\ell}_t^\top (x - \hat{\mu}_t))$ .
17:  Set  $\mathcal{K}_{t+1}^{(0)} = \mathcal{K}_t^{(s_t)}$ .
18: end for
```

---

denotes the number of the updates in the round  $t$ . We denote  $\mathcal{E}_t = \mathcal{E}_t^{(s_t)}$ ,  $\hat{\Sigma}_t = \hat{\Sigma}_t^{(s_t)}$ ,  $\hat{\mu}_t = \hat{\mu}_t^{(s_t)}$ , and  $B_t = B_t^{(s_t)}$ . We randomly choose  $x$  from  $\mathcal{E}_t$  as follows: choose  $\sigma_t \in \{-1, 1\}$  and  $i_t \in [d]$  uniformly at random, and define  $x = \hat{\mu}_t + \frac{\sigma_t}{4e} b_{ti_t}$ . If we can play this  $x$ , then we can construct a good estimate of  $\ell_t$  from the above condition 4, which leads to a small degree of regret. However,  $x \in \mathcal{E}_t$  does not always belong to  $\mathcal{A}$ , particularly when  $\mathcal{A}$  is discrete. To address this issue, we solve DP for this  $x$  and  $\mathcal{P} = \mathcal{K}$  to derive a decomposition of  $x$ , i.e., compute  $x_{t0}, \dots, x_{td} \in \mathcal{K}$  and  $\lambda_{t0}, \dots, \lambda_{td} \geq 0$  as in Step 14. Then, the algorithm plays  $a_t = x_{ti}$  with probability  $\lambda_{ti}$ , and obtains feedback of  $\ell_t^\top a_t$ . Based on this feedback, we compute an estimator  $\hat{\ell}_t$  of the loss vector  $\ell_t$  as

$$\hat{\ell}_t = 4e d \sigma_t \ell_t^\top a_t \hat{\Sigma}_t^{-1} b_{ti_t}. \quad (4)$$

This is an unbiased estimator of  $\ell_t$ , i.e., we have  $\mathbf{E}[\hat{\ell}_t] = \ell_t$ . The existence of  $\hat{\Sigma}_t^{-1}$  follows from the definition of  $\hat{\Sigma}$  and Assumption 1. In fact, from  $\mathcal{A} \subseteq \mathcal{K}_t^{(j)}$  and Assumption 1,  $\mathcal{K}_t^{(j)}$  has a positive volume and  $q_t^{(j)}$  has a positive density over  $\mathcal{K}_t^{(j)}$ , which implies that the covariance matrix  $\Sigma_t^{(j)}$  of  $q_t^{(j)}$  is positive definite. From this and (2),  $\hat{\Sigma}_t^{(j)}$  is positive definite for all  $t$  and  $j$ . The function  $z_t$  is updated by  $z_{t+1}(x) = z_t(x) \exp(-\eta \hat{\ell}_t^\top (x - \hat{\mu}_t))$ , where  $\eta > 0$  is an input parameter standing for the learning rate, which will be optimized later. Let

$$\delta_t^{(j)} = 1/(T(j+2 + \sum_{i=1}^{t-1} (s_i+1))(j+3 + \sum_{i=1}^{t-1} (s_i+1))). \quad (5)$$

To compute  $\hat{\Sigma}_t^{(j)}$  and  $\hat{\mu}_t^{(j)}$  satisfying (2) with probability at least  $1 - \delta_t^{(j)}$ , we use the algorithm in Corollary 1.

Let  $S_T = \sum_{t=1}^T s_t$  denote the number of updates of  $\mathcal{K}_t^{(j)}$ . We show the following regret bound.

**Theorem 5.** Define  $\psi = \frac{1}{d} \log \frac{\text{Vol}(B_\infty(0, R))}{\text{Vol}(\mathcal{K})}$ . Suppose  $a_t$  is given by Algorithm 1 with parameters  $\varepsilon = \frac{1}{12eT}$  and  $\eta = \frac{1}{2eLR} \min\{\sqrt{\frac{1+\psi+\log T}{dT}}, \frac{1}{2^4 d^{3/2} (1+\psi+\log T)}\}$ . Then, for all  $a^* \in \mathcal{A}$ , we have

$$\mathbf{E}[R_T(a^*)] \leq 2^7 eLRd^{3/2} \max\{\sqrt{T(1+\psi+\log T)}, d(1+\psi+\log T)^2\} (1 - S_T/2^{10}). \quad (6)$$

We note that  $\psi$  in the above theorem satisfies  $\psi \leq \log \frac{R}{r}$  if  $\mathcal{K}$  includes an  $\ell_\infty$ -ball of radius  $r > 0$ . The proof of this theorem is given in Appendix A.



## 5.2 Oracle Complexity Analysis

Here, we show that Algorithm 1 calls the linear optimization oracles only  $O(\text{poly}(d)T)$  times.

To implement Algorithm 1, the linear optimization oracle is required only in Steps 5 and 14. In Step 5, we need to solve SP to decide if there exists  $x \in \mathcal{E}_t^{(j)}$  such that  $x \notin \mathcal{K}$ . From the definition (3) of  $\mathcal{E}_t^{(j)}$ , the number of elements in  $\mathcal{E}_t^{(j)}$  is equal to  $2d$  for each  $t$  and  $j$ , and, accordingly, the total number of solutions of SP is  $\sum_{t=1}^T \sum_{j=0}^{s_t} |\mathcal{E}_t^{(j)}| = 2d \sum_{t=1}^T (s_t + 1) = 2d(T + S_T)$ . The number  $S_T$  can be bounded as  $S_T = O(T)$ . Indeed, from Theorem 5, if  $S_T > 2^{10}(1+T)$  then  $\mathbf{E}[R_T(a^*)] < -2^7 e LRT$ , which contradicts to  $R_T(a^*) = \sum_{t=1}^T \ell_t^\top (a_t - a^*) \geq -2LRT$ . Consequently, the total number of solutions of SP is  $O(dT)$ . In Step 14, we solve DP in each round  $t$ ; hence the total number of solutions of DP is equal to  $T$ . Because we can solve SP and DP by calling the linear optimization oracle  $\text{poly}(d, \log T)$  times from Theorem 4 and Remark 1, we can implement Algorithm 1 so that it calls the linear optimization oracle  $O(\text{poly}(d, \log T)T)$  times.

## References

- [1] Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320, 2011.
- [2] M. Abeille, A. Lazaric, et al. Linear thompson sampling revisited. *Electronic Journal of Statistics*, 11(2):5165–5197, 2017.
- [3] J. D. Abernethy, E. Hazan, and A. Rakhlin. An efficient algorithm for bandit linear optimization. In *Conference on Learning Theory*, 2008.
- [4] S. Agrawal and N. Goyal. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, pages 127–135, 2013.
- [5] S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
- [6] J.-Y. Audibert and S. Bubeck. Minimax policies for adversarial and stochastic bandits. In *Conference on Learning Theory*, pages 217–226, 2009.
- [7] P. Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- [8] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.
- [9] B. Awerbuch and R. D. Kleinberg. Adaptive routing with end-to-end feedback: Distributed learning and geometric approaches. In *Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing*, pages 45–53. ACM, 2004.
- [10] S. Bubeck, Y. T. Lee, and R. Eldan. Kernel-based methods for bandit convex optimization. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 72–85. ACM, 2017.
- [11] N. Cesa-Bianchi and G. Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422, 2012.
- [12] A. Cohen, T. Hazan, and T. Koren. Tight bounds for bandit combinatorial optimization. In *Conference on Learning Theory*, pages 629–642, 2017.
- [13] R. Combes, M. S. T. M. Shahi, A. Proutiere, and M. Lelarge. Combinatorial bandits revisited. In *Advances in Neural Information Processing Systems*, pages 2116–2124, 2015.
- [14] T. M. Cover. Universal portfolios. In *The Kelly Capital Growth Investment Criterion: Theory and Practice*, pages 181–209. World Scientific, 2011.

- [15] V. Dani and T. P. Hayes. Robbing the bandit: Less regret in online geometric optimization against an adaptive adversary. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm*, pages 937–943, 2006.
- [16] V. Dani, T. P. Hayes, and S. M. Kakade. Stochastic linear optimization under bandit feedback. In *Conference on Learning Theory*, pages 355–366, 2008.
- [17] V. Dani, S. M. Kakade, and T. P. Hayes. The price of bandit information for online optimization. In *Advances in Neural Information Processing Systems*, pages 345–352, 2008.
- [18] D. Garber. Efficient online linear optimization with approximation algorithms. In *Advances in Neural Information Processing Systems*, pages 627–635, 2017.
- [19] E. Hazan and Z. Karnin. Volumetric spanners: an efficient exploration basis for learning. *The Journal of Machine Learning Research*, 17(1):4062–4095, 2016.
- [20] E. Hazan, W. Hu, Y. Li, and z. li. Online improper learning with an approximation oracle. In *Advances in Neural Information Processing Systems 31*, pages 5652–5660. 2018.
- [21] S. Ito, D. Hatano, H. Sumita, K. Takemura, T. Fukunaga, N. Kakimura, and K. Kawarabayashi. Improved regret bounds for bandit combinatorial optimization. In *Advances in Neural Information Processing Systems*, 2019, to appear.
- [22] S. M. Kakade, A. T. Kalai, and K. Ligett. Playing games with approximation algorithms. *SIAM Journal on Computing*, 39(3):1088–1106, 2009.
- [23] A. Kalai and S. Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- [24] T. Lattimore and C. Szepesvári. *Bandit Algorithms*. preprint, Revision: 1699, 2019.
- [25] L. Lovász and S. Vempala. The geometry of logconcave functions and sampling algorithms. *Random Structures & Algorithms*, 30(3):307–358, 2007.
- [26] H. B. McMahan and A. Blum. Online geometric optimization in the bandit setting against an adaptive adversary. In *International Conference on Computational Learning Theory*, pages 109–123, 2004.
- [27] Y. Nesterov and A. Nemirovskii. *Interior-Point Polynomial Algorithms in Convex Programming*, volume 13. SIAM, 1994.
- [28] A. Prékopa. Logarithmic concave measures with application to stochastic programming. *Acta Scientiarum Mathematicarum*, 32:301–316, 1971.
- [29] P. Rusmevichientong and J. N. Tsitsiklis. Linearly parameterized bandits. *Mathematics of Operations Research*, 35(2):395–411, 2010.
- [30] A. Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, 1998.
- [31] M. Valko, R. Munos, B. Kveton, and T. Kocák. Spectral bandits for smooth graph functions. In *International Conference on Machine Learning*, pages 46–54, 2014.
- [32] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- [33] V. G. Vovk. Aggregating strategies. *Proc. of Computational Learning Theory, 1990*, 1990.

## A Regret Analysis of Algorithm 1

We use the following two lemmas to prove Theorem 5

**Lemma 1.** *The conditional expectation of  $\hat{\ell}_t$  defined by (4), given  $\ell_t$  and  $\mathcal{E}_t$ , satisfies*

$$\mathbf{E}[\hat{\ell}_t] = \ell_t.$$

*Proof.* Because  $\sum_{s=0}^d \lambda_{ts} x_{ts} = \hat{\mu}_t + \frac{\sigma_t}{4e} b_{ti_t}$ , the expectation of  $a_t$  given  $\sigma_t, i_t$  satisfies

$$\mathbf{E}[a_t] = \hat{\mu}_t + \frac{\sigma_t}{4e} b_{ti_t}.$$

Hence, we have

$$\mathbf{E}[\sigma_t b_{ti_t} a_t^\top] = \mathbf{E} \left[ \sigma_t b_{ti_t} \left( \hat{\mu}_t + \frac{\sigma_t}{4e} b_{ti_t} \right)^\top \right] = \frac{1}{4e} \mathbf{E}[b_{ti_t} b_{ti_t}^\top] = \frac{1}{4ed} \sum_{i=1}^d b_{ti} b_{ti}^\top = \frac{1}{4ed} \hat{\Sigma}_t,$$

where the second equality comes from  $\mathbf{E}[\sigma_t] = 0$  and  $\sigma_t^2 = 1$ , the third equality holds because  $i_t$  follows a uniform distribution over  $[d]$ , and the last equality follows from the definition of  $\{b_{ti}\}_{i=1}^d$ . From the above equation and definition (4) of  $\hat{\ell}_t$ , we have

$$\mathbf{E}[\hat{\ell}_t] = 4ed \hat{\Sigma}_t^{-1} \mathbf{E}[\sigma_t b_{ti_t} a_t^\top] \ell_t = \ell_t.$$

□

**Lemma 2.** *Suppose that a random variable  $X$  follows a logconcave distribution and that  $\mathbf{E}[X^2] \leq 1/\alpha^2$  holds for given  $\alpha \geq 1$ . Then, we have*

$$\mathbf{E}[\exp(X) \mathbf{1}_{X>1}] \leq \frac{\exp(3-\alpha)}{1-\exp(1-\alpha)}.$$

*Proof.* From Lemma 5.7 in [25], we have

$$\text{Prob}[|X| \geq i] \leq \exp(-\alpha i + 1)$$

for all  $i \geq 1$ . Hence, we have

$$\begin{aligned} \mathbf{E}[\exp(X) \mathbf{1}_{X>1}] &= \sum_{i=1}^{\infty} \mathbf{E}[\exp(X) \mathbf{1}_{i \leq X \leq i+1}] \leq \sum_{i=1}^{\infty} \text{Prob}[i \leq X \leq i+1] \exp(i+1) \\ &\leq \sum_{i=1}^{\infty} \text{Prob}[|X| \geq i] \exp(i+1) \leq \sum_{i=1}^{\infty} \exp((1-\alpha)i + 2) = \frac{\exp(3-\alpha)}{1-\exp(1-\alpha)}. \end{aligned}$$

□

To prove Theorem 5, we introduce some notations. In the following, we denote

$$f_t(x) = \ell_t^\top (x - \hat{\mu}_t), \quad \hat{f}_t(x) = \hat{\ell}_t^\top (x - \hat{\mu}_t). \quad (7)$$

Then, we can express  $z_t$  as

$$z_t(x) = \exp \left( -\eta \sum_{i=1}^{t-1} \hat{f}_i(x) \right). \quad (8)$$

Then, the regret can be bounded by means of  $Z_{T+1}^{(0)}$  defined in (1), as follows:

**Lemma 3.** *For all  $a^* \in \mathcal{A}$  and  $\gamma \in (0, 1)$ , we have*

$$\mathbf{E}[R_T(a^*)] \leq \frac{1}{\eta(1-\gamma)} \left( \mathbf{E} \log Z_{T+1}^{(0)} - \log \text{Vol}(\mathcal{K}) + \eta \gamma LRT - d \log \gamma \right). \quad (9)$$

*Proof.* From  $\mathbf{E}[a_t|\hat{\mu}_t] = \hat{\mu}_t$  and  $\mathbf{E}[\hat{\ell}_t|\hat{\mu}_t] = \ell_t$ , we have

$$\mathbf{E} R_T(a^*) = \mathbf{E} \sum_{t=1}^T \ell_t^\top (a_t - a^*) = \mathbf{E} \sum_{t=1}^T \ell_t^\top (\hat{\mu}_t - a^*) = \mathbf{E} \sum_{t=1}^T \hat{\ell}_t^\top (\hat{\mu}_t - a^*) = -\mathbf{E} \sum_{t=1}^T \hat{f}_t(a^*). \quad (10)$$

We consider evaluating the rightmost-hand side, by using  $Z_{T+1}^{(0)}$ . Define a convex body  $\bar{\mathcal{K}}$  by  $\bar{\mathcal{K}} := (1-\gamma)a^* + \gamma\mathcal{K}$ . Because  $\mathcal{K}$  is convex and  $a^* \in \mathcal{K}$ , we have  $\bar{\mathcal{K}} \subseteq \mathcal{K}$ . Hence, we have  $\bar{\mathcal{K}} \subseteq \mathcal{K}_t^{(j)}$  for all  $t$  and  $j$ . Then, we have

$$\begin{aligned} \mathbf{E} \log Z_{T+1}^{(0)} &= \mathbf{E} \log \int_{\mathcal{K}_T} z_{T+1}(x) dx \geq \mathbf{E} \log \int_{\bar{\mathcal{K}}} z_{T+1}(x) dx \\ &= \mathbf{E} \log \int_{\mathcal{K}} \gamma^d z_{T+1}((1-\gamma)a^* + \gamma x) dx \\ &= d \log \gamma + \mathbf{E} \log \int_{\mathcal{K}} \exp \left( -\eta \sum_{t=1}^T ((1-\gamma)\hat{f}_t(a^*) + \gamma\hat{f}_t(x)) \right) dx \\ &= d \log \gamma - \eta(1-\gamma) \mathbf{E} \sum_{t=1}^T \hat{f}_t(a^*) + \mathbf{E} \log \int_{\mathcal{K}} \exp(-\eta\gamma \sum_{t=1}^T \hat{f}_t(x)) dx. \end{aligned} \quad (11)$$

The factor  $\gamma^d$  in the second equality comes from the change of variables  $x \leftarrow (1-\gamma)a^* + \gamma x$ . The last term in (11) can be bounded from below by means of  $\bar{x} := \frac{\int_{\mathcal{K}} x dx}{\int_{\mathcal{K}} 1 dx} \in \mathcal{K}$ . Indeed, because  $\hat{f}_t$  is an affine function and  $\exp$  is a convex function, from Jensen's inequality, we have

$$\begin{aligned} \mathbf{E} \log \int_{\mathcal{K}} \exp(-\eta\gamma \sum_{t=1}^T \hat{f}_t(x)) dx &\geq \mathbf{E} \log \int_{\mathcal{K}} \exp(-\eta\gamma \sum_{t=1}^T \hat{f}_t(\bar{x})) dx \\ &= -\eta\gamma \mathbf{E} \sum_{t=1}^T \hat{f}_t(\bar{x}) + \log \int_{\mathcal{K}} 1 dx = -\eta\gamma \mathbf{E} \sum_{t=1}^T f_t(\bar{x}) + \log \int_{\mathcal{K}} 1 dx \geq -\eta\gamma LRT + \log \text{Vol}(\mathcal{K}). \end{aligned}$$

By combining this and (11), we obtain

$$-\eta(1-\gamma) \mathbf{E} \sum_{t=1}^T \hat{f}_t(a^*) \leq \mathbf{E} \log Z_{T+1}^{(0)} - \log \text{Vol}(\mathcal{K}) + \eta\gamma LRT - d \log \gamma.$$

From this and (10), we have (9).  $\square$

The value  $\mathbf{E} \log Z_{T+1}^{(0)}$  can be expressed as

$$\mathbf{E} \log Z_{T+1}^{(0)} = \log Z_1^{(0)} + \sum_{t=1}^T \mathbf{E} \log \frac{Z_{t+1}^{(0)}}{Z_t^{(0)}} \leq \log Z_1^{(0)} + \sum_{t=1}^T \left( \mathbf{E} \log \frac{Z_{t+1}^{(0)}}{Z_t^{(s_t)}} + \mathbf{E} \log \frac{Z_t^{(s_t)}}{Z_t^{(0)}} \right). \quad (12)$$

We can evaluate  $\mathbf{E} \log \frac{Z_{t+1}^{(0)}}{Z_t^{(s_t)}}$  and  $\mathbf{E} \log \frac{Z_t^{(s_t)}}{Z_t^{(0)}}$  as in Lemmas 4 and 5, respectively.

**Lemma 4.** *Under the assumption that (2) and  $\eta \leq \frac{1}{2^3 e L R d^{3/2} (4 + \log T)}$  holds for all  $t \in [T]$ , we have*

$$\mathbf{E} \log \frac{Z_{t+1}^{(0)}}{Z_t^{(s_t)}} \leq L R \varepsilon \eta + 2^5 (e d L R \eta)^2 + \frac{1}{T}. \quad (13)$$

*Proof.* We denote  $q_t := q_t^{(s_t)}$ , which is defined in (1). Let  $\mu_t \in \mathbb{R}^d$  and  $\Sigma_t \in \mathbb{R}^{d \times d}$  denote the mean and the covariance matrix of  $q_t$ , respectively. Define  $\tilde{f}_t : \mathbb{R}^d \rightarrow \mathbb{R}$  by  $\tilde{f}_t(x) = \hat{\ell}_t^\top (x - \mu_t)$ . From

definition (7) of  $\hat{f}_t$ , we have  $\hat{f}_t(x) = \tilde{f}_t(x) + \hat{\ell}_t^\top(\mu_t - \hat{\mu}_t)$ . From definitions (8) and (1) of  $z_t$  and  $Z_t^j$ , we can express  $\frac{Z_{t+1}^{(0)}}{Z_t^{(s_t)}}$  as follows:

$$\begin{aligned} \frac{Z_{t+1}^{(0)}}{Z_t^{(s_t)}} &= \int_{\mathcal{K}_t} \frac{z_{t+1}(x)}{Z_t^{(s_t)}} dx = \int_{\mathcal{K}_t} \frac{z_t(x)}{Z_t^{(s_t)}} \exp(-\eta \hat{f}_t(x)) dx \\ &= \mathbf{E}_{x \sim q_t} \exp(-\eta \hat{f}_t(x)) = \exp(-\eta \hat{\ell}_t^\top(\mu_t - \hat{\mu}_t)) \cdot \mathbf{E}_{x \sim q_t} \exp(-\eta \tilde{f}_t(x)). \end{aligned} \quad (14)$$

Because  $\exp(x) \leq 1 + x + x^2$  holds for  $x \leq 1$ ,  $\mathbf{E}_{x \sim q_t} \exp(-\eta \tilde{f}_t(x))$  can be bounded as

$$\begin{aligned} \mathbf{E}_{x \sim q_t} \exp(-\eta \tilde{f}_t(x)) &\leq \mathbf{E}_{x \sim q_t} [(1 - \eta \tilde{f}_t(x) + \eta^2 \tilde{f}_t(x)^2) \mathbf{1}_{-\eta \tilde{f}_t(x) \leq 1}] + \mathbf{E}_{x \sim q_t} [\exp(-\eta \tilde{f}_t(x)) \mathbf{1}_{-\eta \tilde{f}_t(x) > 1}] \\ &\leq \mathbf{E}_{x \sim q_t} [(1 - \eta \tilde{f}_t(x) + \eta^2 \tilde{f}_t(x)^2)] + \mathbf{E}_{x \sim q_t} [\exp(-\eta \tilde{f}_t(x)) \mathbf{1}_{-\eta \tilde{f}_t(x) > 1}]. \end{aligned} \quad (15)$$

The first term in (15) can be, from definition of  $\mu_t$ ,  $\Sigma_t$ , and  $\tilde{f}_t$ , expressed as follows:

$$\begin{aligned} \mathbf{E}_{x \sim q_t} [(1 - \eta \tilde{f}_t(x) + \eta^2 \tilde{f}_t(x)^2)] &= 1 - \eta \mathbf{E}_{x \sim q_t} [\hat{\ell}_t^\top(x - \mu_t)] + \eta^2 \mathbf{E}_{x \sim q_t} [\hat{\ell}_t^\top(x - \mu_t)(x - \mu_t) \hat{\ell}_t^\top] \\ &= 1 + \eta^2 \hat{\ell}_t^\top \Sigma_t \hat{\ell}_t. \end{aligned} \quad (16)$$

To bound the second term in (15), we use Lemma 2 with  $X = -\eta \tilde{f}_t(x)$ : If  $x$  follows  $q_t$ , a log-concave distribution, and if  $\hat{\ell}_t$  is fixed, i.e.,  $\ell_t$ ,  $\sigma_t$ ,  $i_t$  and  $a_t$  are fixed, then  $\tilde{f}_t(x) = \hat{\ell}_t^\top(x - \mu_t) = 4ed\sigma_t \ell_t^\top a_t b_{ti_t} \hat{\Sigma}_t^{-1}(x - \mu_t)$  follows a logconcave distribution, because logconcavity is preserved under linear transformations (see, e.g., [28]). Furthermore, we have<sup>5</sup>

$$\begin{aligned} \mathbf{E}_{x \sim q_t} [(\eta \tilde{f}_t(x))^2] &= \eta^2 \mathbf{E}_{x \sim q_t} [\hat{\ell}_t^\top(x - \mu_t)(x - \mu_t) \hat{\ell}_t] = \eta^2 \hat{\ell}_t^\top \Sigma_t \hat{\ell}_t = (4ed\eta \ell_t^\top a_t)^2 b_{ti_t}^\top \hat{\Sigma}_t^{-1} \Sigma_t \hat{\Sigma}_t^{-1} b_{ti_t} \\ &\leq 2(4ed\eta LR)^2 b_{ti_t}^\top \hat{\Sigma}_t^{-1} b_{ti_t} = 2(4ed\eta LR)^2 \hat{\Sigma}_t^{-1} \bullet (b_{ti_t} b_{ti_t}^\top) \leq 2^5 (ed\eta LR)^2 \hat{\Sigma}_t^{-1} \bullet \hat{\Sigma}_t = 2^5 (e\eta LR)^2 d^3, \end{aligned}$$

where the first, second, and third equalities come from the definitions of  $\tilde{f}_t$ ,  $\Sigma_t$ , and  $\hat{\ell}_t$ , respectively, the first inequality follows from the condition  $\Sigma_t \preceq 2\hat{\Sigma}_t$  given in (2), and the second inequality follows from that  $\hat{\Sigma}_t = \sum_{i=1}^d b_{ti} b_{ti}^\top$ . Hence, under the assumption that  $2^5 (e\eta LR)^2 d^3 (4 + \log T)^2 \leq 1$  holds, it follows from Lemma 2 that

$$\mathbf{E}_{x \sim q_t} [\exp(-\eta \tilde{f}_t(x)) \mathbf{1}_{-\eta \tilde{f}_t(x) > 1}] \leq \frac{\exp(-1 - \log T)}{1 - \exp(-3 - \log T)} \leq \frac{1}{T}.$$

Combining this, (15), and (16), we obtain

$$\mathbf{E}_{x \sim q_t} \exp(-\eta \tilde{f}_t(x)) \leq 1 + \eta^2 \hat{\ell}_t^\top \Sigma_t \hat{\ell}_t + \frac{1}{T}.$$

From this, (14), Lemma 1, and the fact that  $\log(1 + x) \leq x$  holds for  $x \geq 0$ , we have

$$\begin{aligned} \mathbf{E} \log \frac{Z_{t+1}^{(0)}}{Z_t^{(s_t)}} &\leq -\eta \mathbf{E} \hat{\ell}_t^\top(\mu_t - \hat{\mu}_t) + \mathbf{E} \log \left( 1 + \eta^2 \hat{\ell}_t^\top \Sigma_t \hat{\ell}_t + \frac{1}{T} \right) \\ &\leq -\eta \mathbf{E} \ell_t^\top(\mu_t - \hat{\mu}_t) + \mathbf{E} \left[ \eta^2 \hat{\ell}_t^\top \Sigma_t \hat{\ell}_t \right] + \frac{1}{T}. \end{aligned}$$

From (2), we have

$$|\hat{\ell}_t^\top(\mu_t - \hat{\mu}_t)| \leq \|\ell_t\|_2 \|\mu_t - \hat{\mu}_t\|_2 \leq L \|\mu_t - \hat{\mu}_t\|_{\Sigma_t^{-1}} \|\Sigma_t\|_2 \leq LR\varepsilon,$$

where  $\|\Sigma_t\|_2$  stands for the  $\ell_2$  operator norm, i.e., the largest eigenvalue of  $\Sigma_t$ , and the last inequality holds because  $\Sigma_t$  is the covariance matrix of distribution over a region included in  $B_\infty(0, R)$ . Furthermore, we have

$$\begin{aligned} \mathbf{E} [\hat{\ell}_t^\top \Sigma_t \hat{\ell}_t] &\leq (4edLR)^2 \Sigma_t^{-1} \bullet \mathbf{E} [b_{ti_t} b_{ti_t}^\top] = (4edLR)^2 \Sigma_t^{-1} \bullet \left( \frac{1}{d} \sum_{i=1}^d b_{ti} b_{ti}^\top \right) \\ &= (4edLR)^2 \Sigma_t^{-1} \bullet \left( \frac{1}{d} \hat{\Sigma}_t \right) \leq 2(4edLR)^2. \end{aligned}$$

By combining the above three inequalities, we obtain (13).  $\square$

<sup>5</sup>  $A \bullet B$  means the Frobenius inner product of matrices  $A$  and  $B$ , i.e., its value is defined to be the trace of  $A^\top B$ .

**Lemma 5.** Suppose  $\varepsilon \leq 1/(12e)$ . For all  $t \in [T]$  and  $j \in \{0, 1, \dots, s_t - 1\}$ , under the assumption of (2), we have

$$\frac{Z_t^{(j+1)}}{Z_t^{(j)}} \leq \left(1 - \frac{1}{2e}\right). \quad (17)$$

*Proof.* Let  $(w, b)$  denote the hyperplane that the algorithm chooses for updating  $\mathcal{K}_t^{(j)}$ , which means that  $\mathcal{K}_t^{(j+1)} = \mathcal{K}_t^{(j)} \cap \{x \in \mathbb{R}^d \mid w^\top x \geq b\}$ . Then, we have

$$1 - \frac{Z_t^{(j+1)}}{Z_t^{(j)}} = \frac{\int_{\mathcal{K}_t^{(j)} \setminus \mathcal{K}_t^{(j+1)}} z_t(a) da}{\int_{\mathcal{K}_t^{(j)}} z_t(a) da} = \text{Prob}_{x \sim q_t^{(j)}}[w^\top x < b].$$

Because  $(w, b)$  satisfies  $w^\top x \leq b$  for some  $x \in \mathcal{E}_t^{(j)}$ , there exists  $i \in [d]$  and  $\sigma \in \{1, -1\}$  such that

$$w^\top \left( \hat{\mu}_t^{(j)} + \frac{\sigma}{4e} b_{ti}^{(j)} \right) \leq b.$$

Combining the above equality and inequality, we obtain

$$\begin{aligned} 1 - \frac{Z_t^{(j+1)}}{Z_t^{(j)}} &\geq \text{Prob}_{x \sim q_t^{(j)}} \left[ w^\top x < w^\top \left( \hat{\mu}_t^{(j)} + \frac{\sigma}{4e} b_{ti}^{(j)} \right) \right] \\ &= \text{Prob}_{x \sim q_t^{(j)}} \left[ w^\top \left( x - \mu_t^{(j)} \right) < w^\top \left( \hat{\mu}_t^{(j)} - \mu_t^{(j)} + \frac{\sigma}{4e} b_{ti}^{(j)} \right) \right]. \end{aligned}$$

The value  $w^\top \left( \hat{\mu}_t^{(j)} - \mu_t^{(j)} + \frac{\sigma}{4e} b_{ti}^{(j)} \right)$  can be bounded as

$$\begin{aligned} \left| w^\top \left( \hat{\mu}_t^{(j)} - \mu_t^{(j)} + \frac{\sigma}{4e} b_{ti}^{(j)} \right) \right| &\leq \|w\|_{\hat{\Sigma}_t^{(j)}} \left( \|\hat{\mu}_t^{(j)} - \mu_t^{(j)}\|_{\hat{\Sigma}_t^{(j)-1}} + \frac{1}{4e} \|b_{ti}^{(j)}\|_{\hat{\Sigma}_t^{(j)-1}} \right) \\ &\leq \|w\|_{\hat{\Sigma}_t^{(j)}} \left( \varepsilon + \frac{1}{4e} \right) \leq \sqrt{2} \|w\|_{\Sigma_t^{(j)}} \left( \varepsilon + \frac{1}{4e} \right) \leq \frac{\|w\|_{\Sigma_t^{(j)}}}{2e}, \end{aligned}$$

where the first inequality comes from the Cauchy-Schwarz inequality, the second inequality follows from (2) and that  $\Sigma_t^{(j)} = \sum_{i=1}^d b_{ti}^{(j)} b_{ti}^{(j)\top}$ , the third inequality follows from  $\hat{\Sigma}_t^{(j)} \succeq 2\Sigma_t^{(j)}$  in (2), and the last inequality comes from the assumption of  $\varepsilon < 1/(12e)$ . From the above two inequations, we have

$$1 - \frac{Z_t^{(j+1)}}{Z_t^{(j)}} \geq \text{Prob}_{x \sim q_t^{(j)}} \left[ \frac{w^\top (x - \mu_t^{(j)})}{\|w\|_{\Sigma_t^{(j)}}} < -\frac{1}{2e} \right].$$

When  $x$  follows  $q_t^{(j)}$ ,  $y := \frac{w^\top (x - \mu_t^{(j)})}{\|w\|_{\Sigma_t^{(j)}}}$  follows a distribution with mean 0 and variance 1, because

$w^\top x$  has mean  $w^\top \mu_t^{(j)}$  and variance  $w^\top \Sigma_t^{(j)} w$ . Moreover, the PDF of  $y$  is a logconcave function, because logconcavity is preserved under linear transformations [28]. Because we have  $\text{Prob}[y \leq 0] \geq 1/e$  from Lemma 5.4 in [25] and  $\text{Prob}[-1/(2e) \leq y \leq 0] \leq 1/(2e)$  from Lemma 5.5 in [25], we have  $\text{Prob}[y < -1/(2e)] \geq 1/(2e)$ .  $\square$

Combining (12) and Lemmas 4 and 5, under the assumption that (2) holds for all  $t \in [T]$  and  $j \in \{0, 1, \dots, s_t\}$ , we have

$$\begin{aligned} \mathbf{E} \log Z_{T+1}^{(0)} &\leq \log Z_1^{(0)} + T \left( LR\varepsilon\eta + 2^5(edLR\eta)^2 + \frac{1}{T} \right) + \sum_{t=1}^T s_t \log \left( 1 - \frac{1}{2e} \right) \\ &\leq \log Z_1^{(0)} + LRT\varepsilon\eta + 2^5T(edLR\eta)^2 + 1 - \frac{S_T}{5}, \end{aligned} \quad (18)$$

where we denote  $S_T = \sum_{t=1}^T s_t$  and the second inequality follows from  $\log(1 - \frac{1}{2e}) \leq \frac{1}{5}$ . Define  $\delta := \text{Prob}[\text{there exists } t \in [T] \text{ and } j \leq s_t \text{ such that (2) does not hold}]$ . From definition (5) of  $\delta_t^{(j)}$ ,

we have  $\delta \leq \sum_{t=1}^T \sum_{j=0}^{s_t} \leq \sum_{k=2}^{\infty} \frac{1}{T k(k+1)} = \frac{1}{2T}$ . Because (18) holds with probability at least  $1 - \delta$ , and it always holds that  $R_T(a^*) \leq 2LRT$ , from Lemma 3, we have

$$\begin{aligned} \mathbf{E}[R_T(a^*)] &\leq \frac{1-\delta}{1-\gamma} \left( 2^5 T (\text{ed}LR)^2 \eta + \frac{1}{\eta} \left( \log \frac{Z_1^{(0)}}{\text{Vol}(\mathcal{K})} + 1 - d \log \gamma - \frac{S_T}{5} \right) + LRT(\varepsilon + \gamma) \right) + 2\delta LRT \\ &\leq \frac{1}{1-\gamma} \left( 2^5 T (\text{ed}LR)^2 \eta + \frac{1}{\eta} \left( \log \frac{Z_1^{(0)}}{\text{Vol}(\mathcal{K})} + 1 - d \log \gamma - \frac{S_T}{10} \right) + LRT(\varepsilon + \gamma) \right) + LR, \end{aligned}$$

where the second inequality comes from  $0 \leq \delta \leq \frac{1}{2T}$ . We denote  $\psi := \frac{1}{d} \log \frac{Z_1^{(0)}}{\text{Vol}(\mathcal{K})} = \frac{1}{d} \log \frac{\text{Vol}(B_{\infty}(0, R))}{\text{Vol}(\mathcal{K})}$ . By setting  $\varepsilon = \gamma = \frac{1}{12eT}$ , we obtain

$$\mathbf{E}[R_T(a^*)] \leq 2 \left( 2^5 T (\text{ed}LR)^2 \eta + \frac{1}{\eta} \left( d(5 + \psi + \log T) - \frac{S_T}{10} \right) + 2LR \right).$$

In addition, by setting  $\eta = \frac{1}{2eLR} \min \left\{ \sqrt{\frac{1+\psi+\log T}{dT}}, \frac{1}{2^4 d^{3/2} (1+\psi+\log T)} \right\}$ , we obtain

$$\mathbf{E}[R_T(a^*)] \leq 2^7 eLRd^{3/2} \max \left\{ \sqrt{T(1+\psi+\log T)}, d(1+\psi+\log T)^2 \right\} \left( 1 - \frac{S_T}{2^{10}} \right),$$

which means that (6) holds.  $\square$

## B Algorithm for Stochastic Bandit Linear Optimization

In this section, we present an algorithm for stochastic bandit linear optimization, where we assume that  $\ell_t$  follows a distribution  $D$  over  $\mathbb{R}^d$ , i.i.d. for  $t = 1, 2, \dots, T$ . We denote  $\ell^* = \mathbf{E}_{\ell \sim D}[\ell] \in \mathbb{R}^d$  and  $\xi_t = \ell_t - \ell^*$ .

### B.1 Preliminary: Barycentric Spanner

*Definition 1.* Let  $S \in \mathbb{R}^d$  be a subset whose linear span is  $\mathbb{R}^d$ , and let  $C > 1$ . A set  $X = \{x_1, \dots, x_d\} \subseteq S$  is a  $C$ -barycentric spanner for  $S$  if every  $x \in S$  may be expressed as a linear combination of elements of  $X$  using coefficients in  $[-C, C]$ .

**Theorem 6 (Proposition 2.4. in [9]).** Suppose  $\mathcal{P} \subseteq \mathbb{R}^d$  is a compact set not contained in any proper linear subspace. Given an algorithm  $OPT$  for LP, for any  $C > 1$  we may compute a  $C$ -barycentric spanner for  $\mathcal{P}$  in polynomial time, using  $O(d^2 \log_C(d))$  calls to  $OPT$ . Its span is equal to  $\mathbb{R}^d$ .

### B.2 Algorithm

Our algorithm is summarized in Algorithm 2. In the algorithm, a parameter  $\delta > 0$  controls the probability of achieving a small regret. The rounds are divided into  $K = O(\log T)$  phases, so that the  $k$ -th phase consists of  $\Theta(2^k)$  rounds for each  $k \in \{1, \dots, K\}$ .

When the  $k$ -th phase begins, the algorithm maintains an action set  $\mathcal{P}_k$ . This action set is initialized by  $\mathcal{P}_1 = \text{Conv}(\mathcal{A})$  and is defined recursively by (20) (thus  $\mathcal{P}_1 \supseteq \mathcal{P}_2 \supseteq \dots \supseteq \mathcal{P}_K$ ).  $\mathcal{P}_k$  is designed so that the value of  $\ell^{*\top} x$  is smaller for all  $x \in \mathcal{P}_k$  as  $k$  gets larger (see Lemma 7). At the beginning of the  $k$ -th phase, the algorithm computes a 2-barycentric spanner  $X_k = \{x_{k1}, \dots, x_{kd}\}$  of  $\mathcal{P}_k$ . We can construct a good estimate of  $\ell^*$  if each element of  $X_k$  can be chosen as an action. However, elements in  $X_k$  do not always belong to  $\mathcal{A}$ , especially when  $\mathcal{A}$  is discrete. To address this issue, our algorithm decomposes each  $x_{ki}$  into the points  $x_{ki0}, \dots, x_{kid} \in \mathcal{A}$  with weight  $\lambda_{ki0}, \dots, \lambda_{kid} \geq 0$  so that  $\lambda_{ki0} + \dots + \lambda_{kid} = 1$  and  $\lambda_{ki0}x_{ki0} + \dots + \lambda_{kid}x_{kid} = x_{ki}$ . Then, it plays  $x_{kij}$ ,  $T_{kij} \propto \lambda_{kij}$  times, for each  $j = 0, 1, \dots, d$ . We denote the action played at the  $t$ -th round by  $a_t$ . The algorithm computes an empirical estimate  $\hat{\ell}_k$  of  $\ell^*$ , based on the feedback obtained in  $k$ -th phase, as defined in (19).

We note that  $\text{Vol}(\mathcal{P}_k) > 0$  holds for all  $k$ , which implies that we can apply the algorithm in Theorem 6 to  $\mathcal{P}_k$  and that  $V_k$  is invertible. In fact, we have  $\text{Vol}(\mathcal{P}_1) > 0$  from Assumption 1 and we can show

---

**Algorithm 2** An oracle efficient algorithm for stochastic bandit linear optimization

---

**Require:** Action set  $\mathcal{A} \subseteq \mathbb{R}^d$ , positive real numbers  $L$  and  $R$  satisfying Assumption 1,  $\delta \in (0, 1)$ .

- 1: Set  $\mathcal{P}_1 = \text{Conv}(\mathcal{A})$  and  $t_{1,1,0} = 0$ .
- 2: **for**  $k = 1, 2, \dots, K$  **do**
- 3:   Let  $X_k = \{x_{k1}, \dots, x_{kd}\} \subseteq \mathcal{P}_k$  be a 2-barycentric spanner for  $\mathcal{P}_k$ .
- 4:   Set  $\zeta_k = 2^{2k+9} d^2 \log \left( \frac{2dk(k+1)}{\delta} \right)$ .
- 5:   **for**  $i = 1, \dots, d$  **do**
- 6:     Solve DP for  $\mathcal{P} = \text{Conv}(\mathcal{A})$  and  $x = x_{ki}$  to get a decomposition  $x_{ki0}, \dots, x_{kidi} \in \mathcal{A}$  and  $\lambda_{ki0}, \dots, \lambda_{kidi}$  such that  $x_{ki} = \lambda_{ki0}x_{ki0} + \dots + \lambda_{kidi}x_{kidi}$ .
- 7:     **for**  $j = 0, \dots, d$  **do**
- 8:       Set  $T_{kij} = \lceil \zeta_k \lambda_{kij} \rceil$ ,  $t_{ki,j+1} = t_{kij} + T_{kij}$ .
- 9:       Choose action  $a_t = x_{kij}$  exactly  $T_{kij}$  times, from the  $(t_{kij} + 1)$ -th round to the  $(t_{ki,j+1})$ -th round.
- 10:    **end for**
- 11:    Set  $t_{k,i+1,0} = t_{ki,d+1}$ .
- 12:   **end for**
- 13:   Calculate empirical estimate  $\hat{\ell}_k$  of  $\ell^*$  by

$$V_k = \sum_{i=1}^d \sum_{j=0}^d T_{kij} x_{kij} x_{kij}^\top, \quad \hat{\ell}_k = V_k^{-1} \sum_{i=1}^d \sum_{j=0}^d \sum_{t=t_{kij}+1}^{t_{ki,j+1}} (\ell_t^\top x_{kij}) x_{kij}. \quad (19)$$

- 14:   Solve LP for  $\mathcal{P} = \mathcal{P}_k$  and  $w = \hat{\ell}_k$  to find a vector  $x_k^* \in \arg \min_{x \in \mathcal{P}_k} \hat{\ell}_k^\top x$ .
- 15:   Update  $\mathcal{P}_k$  by

$$\mathcal{P}_{k+1} = \{x \in \mathcal{P}_k \mid \hat{\ell}_k^\top (x - x_k^*) \leq LR2^{-k}\}. \quad (20)$$

- 16: **end for**

---

$\text{Vol}(\mathcal{P}_k) > 0$  by induction in  $k$ , from the definition (20). The linear span of a barycentric spanner  $X_k = \{x_{k1}, \dots, x_{kd}\}$  coincides with that of  $\mathcal{P}_k$  (see, e.g., [9]), which is equal to  $\mathbb{R}^d$ , because  $\text{Vol}(\mathcal{P}_k) > 0$ . Hence, we have  $V_k \succeq \zeta_k \sum_{i=1}^d \sum_{j=0}^d \lambda_{kij} x_{kij} x_{kij}^\top \succeq \zeta_k \sum_{i=1}^d x_{ki} x_{ki}^\top \succ O$ , which means that  $V_k$  is nonsingular.

Algorithm 2 satisfies the following regret bound:

**Theorem 7.** Suppose that  $\ell_t$  follows an i.i.d. distribution for  $t = 1, \dots, T$  with  $T \geq 2$ , and that  $\{a_t\}_{t=1}^T$  is given by Algorithm 2. With probability at least  $1 - \delta$ , the regret is bounded as follows:

$$\max_{a \in \mathcal{A}} R_T(a) \leq 2^{12} LR \sqrt{d^3 T \log \left( \frac{d \log T}{\delta} \right)}. \quad (21)$$

The proof of this theorem is given in Appendix C.

### B.3 Oracle Complexity Analysis

Step 3 in Algorithm 2 requires constructing a 2-barycentric spanner. From Theorem 6, we can construct it by calling an algorithm that solves LP for  $\mathcal{P} = \mathcal{P}_k$ ,  $O(\text{poly}(d))$  times. Theorem 3 (a) and Remark 1 imply that we can solve LP by solving SP  $O(\text{poly}(d, \log T))$  times. SP for  $\mathcal{P} = \mathcal{P}_k$  can be solved by the following procedure:

1. Decide if  $y \in \mathcal{P}_1$  or not, and, in the latter case, output a vector  $w \in \mathbb{R}^d$  such that  $w^\top y < \min_{x \in \mathcal{P}_1} w^\top x$ . From Theorem 3 (b), this can be done by calling the LP oracle for  $\mathcal{A}$   $O(\text{poly}(d))$  times. In the former case, i.e., if  $y \in \mathcal{P}_1$ , go to the next step.
2. For  $j = 1, \dots, k-1$ , if  $\hat{\ell}_j^\top (y - x_k^*) > LR2^{-j}$ , output  $\hat{\ell}_j$ . If  $\hat{\ell}_j^\top (y - x_k^*) \leq LR2^{-j}$  for all  $j = 1, \dots, k-1$ , it means that  $y \in \mathcal{P}_k$ .



This procedure calls the LP oracle for  $O(\text{poly}(d, \log T))$  times and runs in  $O(\text{poly}(d, K)) = O(\text{poly}(d, \log T))$  times. Hence, we have an efficient implementation of Step 3 that calls the LP oracle for  $\mathcal{A}$   $O(\text{poly}(d, \log T))$  times. Similarly, Steps 6 and 14 in Algorithm 2 can be executed by calling the LP oracle for  $\mathcal{A}$   $O(\text{poly}(d, \log T))$  times. The other steps are free from access to the oracle and can be efficiently implemented. Because the number  $K$  of iterations w.r.t.  $k$  is bounded as in (32), the number of oracle calls for solving LP over  $\mathcal{A}$  is of  $O(\text{poly}(d, \log T)K) = O(\text{poly}(d, \log T))$ .

## C Regret Analysis of Algorithm 2

In the proof of Theorem 7, we may assume that

$$T > 2d\zeta_1 = 2^{12}d^3 \log\left(\frac{4d}{\delta}\right). \quad (22)$$

Indeed, if  $T \leq 2d\zeta_1$ , then we see  $R_T(a) \leq 2LRT \leq 2LR\sqrt{2d\zeta_1 T}$ , which means that (21) holds.

To prove the above theorem, we start with analyzing the error of the estimators  $\hat{\ell}_k$  defined by (19):

**Lemma 6.** *With probability at least  $1 - \delta$ , for all  $k \in \{1, 2, \dots\}$  and  $x \in \mathcal{P}_k$ , we have*

$$|(\hat{\ell}_k - \ell^*)^\top x| \leq 2^{-1-k}LR. \quad (23)$$

*Proof.* Because  $X_k = \{x_{k1}, \dots, x_{kd}\}$  is a 2-barycentric spanner for  $\mathcal{P}_k$ , for all  $x \in \mathcal{P}_k$ , there exists a vector  $w = (w_1, \dots, w_d)^\top \in [-2, 2]^d$  such that  $x = w_1x_{k1} + \dots + w_dx_{kd}$ . By means of this  $w$ ,  $(\hat{\ell}_k - \ell^*)^\top x$  can be expressed as

$$\begin{aligned} (\hat{\ell}_k - \ell^*)^\top x &= \left( \sum_{i=1}^d \sum_{j=0}^d \sum_{t=t_{ki,j}+1}^{t_{ki,j+1}} \ell_t^\top x_{kij} x_{kij}^\top V_k^{-1} - \ell^{*\top} \right) \sum_{s=1}^d w_s x_{ks} \\ &= \left( \sum_{i=1}^d \sum_{j=0}^d \sum_{t=t_{ki,j}+1}^{t_{ki,j+1}} (\ell_t - \ell^*)^\top x_{kij} x_{kij}^\top V_k^{-1} \right) \sum_{s=1}^d w_s x_{ks} \\ &= \sum_{s=1}^d w_s \sum_{i=1}^d \sum_{j=0}^d \sum_{t=t_{ki,j}+1}^{t_{ki,j+1}} (\xi_t^\top x_{kij}) x_{kij}^\top V_k^{-1} x_{ks}, \end{aligned} \quad (24)$$

where the first equality comes from the definition (19) of  $\hat{\ell}_k$ , the second equality comes from  $\sum_{i=1}^d \sum_{j=0}^d \sum_{t=t_{ki,j}+1}^{t_{ki,j+1}} x_{kij} x_{kij}^\top V_k^{-1} = \sum_{i=1}^d \sum_{j=0}^d T_{kij} x_{kij} x_{kij}^\top V_k^{-1} = I$ , and the last equality comes from  $\xi_t = \ell_t - \ell^*$ . We give a uniform bound for this value by the following claim: with probability at least  $1 - \delta$ , it holds for all  $k = 1, 2, \dots$  and  $s = 1, \dots, d$  that

$$\left| \sum_{i=1}^d \sum_{j=0}^d \sum_{t=t_{ki,j}+1}^{t_{ki,j+1}} (\xi_t^\top x_{kij}) x_{kij}^\top V_k^{-1} x_{ks} \right| \leq \frac{LR}{d2^{k+2}}. \quad (25)$$

Because the expectation of  $\xi_t = \ell_t - \ell^*$  is equal to 0, and because  $\|\xi_t\|_2 \leq \|\ell_t\|_2 + \|\ell^*\|_2 \leq 2L$  and  $\|x_{kij}\|_2 \leq R$  hold from Assumption 1,  $\{\xi_t^\top x_{kij}\}$  are independent random variables with mean 0 and absolute value at most  $2LR$ . Hence, from Hoeffding's inequality, it holds with probability at least

$1 - \frac{\delta}{dk(k+1)}$  that

$$\begin{aligned}
& \left| \sum_{i=1}^d \sum_{j=0}^d \sum_{t=t_{kij}+1}^{t_{ki,j}+1} (\xi_t^\top x_{kij}) x_{kij}^\top V_k^{-1} x_{ks} \right| \\
& \leq 2LR \sqrt{8 \log \left( \frac{2dk(k+1)}{\delta} \right) \sum_{i=1}^d \sum_{j=0}^d \sum_{t=t_{kij}+1}^{t_{ki,j}+1} (x_{kij}^\top V_k^{-1} x_{ks})^2} \\
& = 2LR \sqrt{8 \log \left( \frac{2dk(k+1)}{\delta} \right) x_{ks}^\top V_k^{-1} \left( \sum_{i=1}^d \sum_{j=0}^d \sum_{t=t_{kij}+1}^{t_{ki,j}+1} x_{kij} x_{kij}^\top \right) V_k^{-1} x_{ks}} \\
& = 2LR \sqrt{8 \log \left( \frac{2dk(k+1)}{\delta} \right) x_{ks}^\top V_k^{-1} x_{ks}}. \tag{26}
\end{aligned}$$

The value  $x_{ks}^\top V_k^{-1} x_{ks}$  is bounded as  $x_{ks}^\top V_k^{-1} x_{ks} \leq \zeta_k^{-1}$ . Indeed, we have

$$\begin{aligned}
V_k & \succeq \sum_{j=0}^d T_{ksj} x_{ksj} x_{ksj}^\top \succeq \zeta_k \sum_{j=0}^d \lambda_{ksj} x_{ksj} x_{ksj}^\top \\
& = \zeta_k \left( x_{ks} x_{ks}^\top + \sum_{j=0}^d \lambda_{ksj} (x_{ksj} - x_{ks})(x_{ksj} - x_{ks})^\top \right) \succeq \zeta_k x_{ks} x_{ks}^\top,
\end{aligned}$$

where the first inequality comes from the definition (19) of  $V_k$ , the second inequality comes from  $T_{ksj} = \lceil \zeta_k \lambda_{ksj} \rceil \geq \zeta_k \lambda_{ksj}$  (Step 8 of Algorithm 2), and the equality comes from that  $\lambda_{ks0} + \dots + \lambda_{ksd} = 1$  and that  $\lambda_{ks0} x_{ks0} + \dots + \lambda_{ksd} x_{ksd} = x_{ks}$ . This inequality indicates

$$0 \leq (V_k^{-1} x_{ks})^\top (V_k - \zeta_k x_{ks} x_{ks}^\top) (V_k^{-1} x_{ks}) = x_{ks}^\top V_k^{-1} x_{ks} (1 - \zeta_k x_{ks}^\top V_k^{-1} x_{ks}),$$

from which we have  $x_{ks}^\top V_k^{-1} x_{ks} \leq \zeta_k^{-1}$ .

Plugging this bound on  $x_{ks}^\top V_k^{-1} x_{ks}$  into (26) and the definition of  $\zeta_k$  (Step 4 of Algorithm 2) show that the inequality (25) holds with probability at least  $1 - \frac{\delta}{dk(k+1)}$  for each  $k$  and  $s$ . Hence, we have

$$\begin{aligned}
\text{Prob}[(25) \text{ does not hold for some } k \text{ and } s] & \leq \sum_{k=1}^K \sum_{s=1}^d \text{Prob}[(25) \text{ does not hold for } k \text{ and } s] \\
& \leq \sum_{k=1}^K \sum_{s=1}^d \frac{\delta}{dk(k+1)} = \delta,
\end{aligned}$$

which means that, with probability at least  $1 - \delta$ , (25) holds for all  $k$  and  $s$ . Combining this, (24), and  $|w_s| \leq 2$  for all  $s \in [d]$ , we obtain (23).  $\square$

**Lemma 7.** Fix  $a^* \in \arg \min_{a \in \mathcal{A}} \ell^{*\top} a$ . With probability at least  $1 - \delta$ , for all  $k$ , we have

$$a^* \in \mathcal{P}_k, \text{ and } \ell^{*\top} x - \ell^{*\top} a^* \leq 2^{2-k} LR \text{ for all } x \in \mathcal{P}_k. \tag{27}$$

*Proof.* From Lemma 6, we can assume that (23) holds for all  $k$  and  $x \in \mathcal{P}_k$ . Under this assumption, we show (27) by induction in  $k$ . We can confirm that (27) holds for  $k = 1$ . Indeed,  $a^* \in \mathcal{P}_1$  follows from  $\mathcal{A} \subseteq \mathcal{P}_1$ , and  $\ell^{*\top} x - \ell^{*\top} a^* \leq \|\ell^*\|_2 \|x - a^*\|_2 \leq 2LR$  follows from  $\|\ell^*\|_2 \leq L$  and  $\|x\|_2 \leq R$  for  $x \in \mathcal{P}_1$ . Suppose that (27) holds for  $k = s$ . Then,  $\mathcal{P}_{s+1}$ , defined by (20), contains  $a^*$  because

$$\hat{\ell}_s^\top a^* \leq \ell^{*\top} a^* + 2^{-1-s} LR \leq \ell^{*\top} x_s^* + 2^{-1-s} LR \leq \hat{\ell}_s^\top x_s^* + 2^{-s} LR,$$

where the first and the third inequalities come from (23) and the second inequality comes from  $\ell^{*\top} a^* = \min_{a \in \mathcal{A}} \{\ell^{*\top} a\} = \min_{x \in \mathcal{P}_1} \{\ell^{*\top} x\} \leq \ell^{*\top} x_s^*$ . Furthermore, for all  $x \in \mathcal{P}_{s+1}$ , we have

$$\ell^{*\top} x \leq \hat{\ell}_s^\top x + 2^{-s-1} LR \leq \hat{\ell}_s^\top x_s^* + 3 \cdot 2^{-s-1} LR \leq \hat{\ell}_s^\top a^* + 3 \cdot 2^{-s-1} LR \leq \ell^{*\top} a^* + 2^{-s+1} LR,$$

where the first and the fourth inequalities come from (23), the second inequality comes from the definition (20) of  $\mathcal{P}_{k+1}$ , and the third inequality comes from  $a^* \in \mathcal{P}_s$  and  $x_s^* \in \arg \min_{x \in \mathcal{P}_s} \hat{\ell}_s^\top x$ . Hence, (27) holds for  $k = s + 1$ . By induction in  $k$ , (27) is proven to hold for all positive integers  $k$ .  $\square$

From this lemma and the definition of  $a_t$  in Algorithm 2, we can bound  $\sum_{t=t_{k,1,0}+1}^{t_{k+1,1,0}} \ell^{*\top}(a_t - a^*)$  as follows:

$$\begin{aligned} \sum_{t=t_{k,1,0}+1}^{t_{k+1,1,0}} \ell_t^{*\top}(a_t - a^*) &= \sum_{i=1}^d \sum_{j=0}^d T_{kij} \ell^{*\top}(x_{kij} - a^*) \leq \sum_{i=1}^d \sum_{j=0}^d (\zeta_k \lambda_{kij} \ell^{*\top}(x_{kij} - a^*) + 2LR) \\ &= \sum_{i=1}^d (\zeta_k \ell^{*\top}(x_{ki} - a^*) + 2LR(d+1)) \leq dLR(2^{2-k} \zeta_k + 2(d+1)), \end{aligned} \quad (28)$$

where the first inequality follows from  $T_{kij} = \lceil \zeta_k \lambda_{kij} \rceil$  (Step 8 of Algorithm 2) and  $|\ell^{*\top}(x_{kij} - a^*)| \leq 2LR$ , the second inequality follows from  $\sum_{j=0}^d \lambda_{kij} x_{kij} = x_{ki}$  and  $\sum_{j=0}^d \lambda_{kij} = 1$ , and the last inequality follows from Lemma 7.

Let  $T_k$  denote the number of rounds in the  $k$ -th phase, i.e.,

$$T_k = \sum_{i=1}^d \sum_{j=0}^d T_{kij} = t_{k+1,1,0} - t_{k,1,0}. \quad (29)$$

From the definition of  $T_{kij} = \lceil \zeta_k \lambda_{kij} \rceil$  (Step 8 of Algorithm 2), we have  $\zeta_k \lambda_{kij} \leq T_{kij} < \zeta_k \lambda_{kij} + 1$ . Combining this and the condition that  $\sum_{j=0}^d \lambda_{kij} = 1$ , we obtain

$$d\zeta_k \leq T_k \leq d\zeta_k + d(d+1). \quad (30)$$

Let  $K$  be the index of phases such that the  $K$ -th phase includes  $T$ -th round, i.e.,  $K$  is the number such that  $\sum_{k=1}^{K-1} T_k < T \leq \sum_{k=1}^K T_k$ . Note that  $K \geq 2$  follows from the assumption 22. From (30) and the definition of  $\zeta_k$  (Step 4 in Algorithm 2), we have

$$T > T_{K-1} \geq d\zeta_{K-1} = 2^{2K+7} d^3 \log \left( \frac{2dK(K-1)}{\delta} \right) \geq (2^{K+3})^2 d^3 \log \left( \frac{2dK(K+1)}{\delta} \right),$$

where the last inequality follows from  $2 \log \left( \frac{2dK(K-1)}{\delta} \right) = \log \left( \frac{2dK(K-1)}{\delta} \right)^2 \geq \log \left( \frac{2dK(K+1)}{\delta} \right)$ . This inequality implies the bound on  $2^K$  and  $K$ , as follows:

$$2^K \leq \frac{1}{2^3} \left( \frac{T}{d^3} \right)^{\frac{1}{2}} \left( \log \left( \frac{2dK(K+1)}{\delta} \right) \right)^{-\frac{1}{2}}, \quad (31)$$

$$K \leq \frac{1}{2} \log_2 T - 3. \quad (32)$$

By means of these inequalities, we can bound the value  $\sum_{t=1}^T \ell^{*\top}(a_t - a^*)$ , for the output  $a_t$  of Algorithm 2, and  $a^* \in \arg \min_{a \in \mathcal{A}} \ell^{*\top} a$ :

$$\begin{aligned}
\sum_{t=1}^T \ell^{*\top}(a_t - a^*) &\leq \sum_{k=1}^K \sum_{t=t_{k,1,0}+1}^{t_{k+1,1,0}} \ell^{*\top}(a_t - a^*) \\
&\leq dLR \sum_{k=1}^K (2^{2-k} \zeta_k + 2(d+1)) \quad (\text{from (28)}) \\
&= dLR \sum_{k=1}^K \left( 2^{k+11} d^2 \log \left( \frac{2dk(k+1)}{\delta} \right) + 2K(d+1) \right) \quad (\text{Step 4 in Algo. 2}) \\
&\leq dLR \left( 2^{K+12} d^2 \log \left( \frac{2dK(K+1)}{\delta} \right) + 2K(d+1) \right) \\
&\leq dLR \left( 2^9 d^2 \left( \frac{T}{d^3} \log \left( \frac{2dK(K+1)}{\delta} \right) \right)^{\frac{1}{2}} + 2K(d+1) \right) \quad (\text{from (31)}) \\
&\leq dLR \left( 2^9 \sqrt{dT \log \left( \frac{d(\log_2 T)^2}{\delta} \right)} + (d+1) \log_2 T \right) \quad (\text{from (32)}) \\
&\leq dLR \left( 2^{10} \sqrt{dT \log \left( \frac{d \log T}{\delta} \right)} + (d+1) \log_2 T \right). \quad (33)
\end{aligned}$$

By combining this and the following lemma, we obtain an upper bound on the regret  $R_T(a) = \sum_{t=1}^T \ell_t^\top(a_t - a)$ .

**Lemma 8.** *Let  $a^* \in \arg \min_{a \in \mathcal{A}} \ell^{*\top} a$ . With probability at least  $1 - \delta$ , it holds for all  $a \in \mathcal{A}$  that*

$$R_T(a) \leq \sum_{t=1}^T \ell^{*\top}(a_t - a^*) + 8LR \sqrt{dT \log \left( \frac{2d}{\delta} \right)}. \quad (34)$$

*Proof.* We show (34) by proving the following two inequalities:

$$\sum_{t=1}^T \ell_t^\top a_t - \sum_{t=1}^T \ell^{*\top} a_t \leq LR \sqrt{8T \log \left( \frac{2}{\delta} \right)}, \quad (35)$$

$$\sum_{t=1}^T \ell^{*\top} a^* - \sum_{t=1}^T \ell_t^\top a \leq LR \sqrt{8dT \log \left( \frac{2d}{\delta} \right)}. \quad (36)$$

Denote  $X_\tau := \sum_{t=1}^\tau (\ell_t - \ell^*)^\top a_t$ . Because  $\{X_\tau\}_{\tau=0}^T$  is a martingale such that  $|X_{\tau+1} - X_\tau| \leq 2LR$ , from Azuma's inequality, with probability  $1 - \frac{\delta}{2}$ , we have  $X_T \leq LR \sqrt{8T \log \left( \frac{2}{\delta} \right)}$ , which means that (35) holds. Similarly, from Hoeffding's inequality, we have

$$\left\| \sum_{t=1}^T (\ell_t - \ell^*) \right\|_2 \leq L \sqrt{dT \log \left( \frac{2d}{\delta} \right)} \quad (37)$$

with probability at least  $1 - \frac{\delta}{2}$ . Under this condition, we have

$$-\sum_{t=1}^T \ell_t^\top a \leq -\sum_{t=1}^T \ell^{*\top} a + LR \sqrt{dT \log \left( \frac{2d}{\delta} \right)} \leq -\sum_{t=1}^T \ell^{*\top} a^* + LR \sqrt{8dT \log \left( \frac{2d}{\delta} \right)},$$

where the first inequality follows from (37) and  $\|a\|_2 \leq R$ , and the second inequality follows from  $a^* \in \arg \min_{a \in \mathcal{A}} \ell^{*\top} a$ . Hence, we have (36) for all  $a \in \mathcal{A}$  with probability at least  $1 - \frac{\delta}{2}$ . Because each

of (35) and (36) holds with probability  $1 - \frac{\delta}{2}$ , both (35) and (36) hold with probability  $1 - \delta$ . Then, by taking the sum of each side of (35) and (36), we obtain (34).  $\square$

By combining (33), Lemma 8 and (22), we obtain

$$R_T(a) \leq 2^{11} L R \sqrt{d^3 T \log \left( \frac{d \log T}{\delta} \right)}$$

with probability at least  $1 - 2\delta$ . Replacing  $\delta$  with  $\delta/2$ , we obtain (21).