# **Regret Bounds for Online Portfolio Selection** with a Cardinality Constraint

Shinji Ito NEC Corporation	<b>Daisuke H</b> RIKEN	Hanna Tokyo Metropo	 Akihiro Yabe NEC Corporation
<b>Takuro Fuku</b> RIKEN AIP. JST	0	<b>ori Kakimura</b> o University	<b>awarabayashi</b> ute of Informatics

#### Abstract

Online portfolio selection is a sequential decision-making problem in which a learner repetitively selects a portfolio over a set of assets, aiming to maximize long-term return. In this paper, we study the problem with the cardinality constraint that the number of assets in a portfolio is restricted to be at most k, and consider two scenarios: (i) in the *full-feedback setting*, the learner can observe price relatives (rates of return to cost) for all assets, and (ii) in the *bandit-feedback setting*, the learner can observe price relatives only for invested assets. We propose efficient algorithms for these scenarios, which achieve sublinear regrets. We also provide regret (statistical) lower bounds for both scenarios which nearly match the upper bounds when k is a constant. In addition, we give a computational lower bound, which implies that no algorithm maintains both computational efficiency, as well as a small regret upper bound.

# 1 Introduction

Online portfolio selection [10, 22] is a fundamental problem in financial engineering, in which a learner sequentially selects a portfolio over a set of assets, aiming to maximize cumulative wealth. For this problem, principled algorithms (e.g., the universal portfolio algorithm [10]) have been proposed, which behave as if one knew the empirical distribution of future market performance. On the other hand, these algorithms work only under the strong assumption that we can hold portfolios of arbitrary combinations of assets, and that we can observe price relatives, the multiplicative factors by which prices change, for all assets. Due to these limitations, this framework does not directly apply to such real-world applications as investment in advertising or R&D, where the available combination of assets that have been invested in.

In order to overcome such issues, we consider the following problem setting: Suppose that there are T rounds and a market has d assets, represented by  $[d] := \{1, \ldots, d\}$ . In each round t, we design a portfolio, that represents the proportion of the current wealth invested in each of the d assets. That is, a *portfolio* can be expressed as a vector  $\mathbf{x}_t = [x_{t1}, \ldots, x_{td}]^{\top}$  such that  $x_{ti} \ge 0$  for all  $i \in [d]$  and  $\sum_{i=1}^{d} x_{ti} \le 1$ . The combination of assets is restricted with a set of *available combinations*  $S \subseteq 2^{[d]}$ , that is, a portfolio  $\mathbf{x}_t$  must satisfy  $\operatorname{supp}(\mathbf{x}_t) = \{i \in [d] \mid x_{ti} \ne 0\} \in S$ . Thus, in each period t, we choose  $S_t$  from S and determine a portfolio  $\mathbf{x}_t$  only from assets in  $S_t$ . A typical example of S can be given by cardinality constraints, i.e.,  $S_k := \{S \subseteq [d] \mid |S| = k\}$  for some  $k \le d$ . We denote by  $\mathbf{r}_t = [r_{t1}, \ldots, r_{td}]^{\top}$  a price relative vector, where  $1 + r_{ti}$  is the price relative for the *i*-th asset in the *t*-th period. Then the wealth  $A_T$  resulting from the sequentially rebalanced portfolios  $\mathbf{x}_1, \ldots, \mathbf{x}_t$  is given by  $A_T = \prod_{t=1}^T (1 + \mathbf{r}_t^\top \mathbf{x}_t)$ . The best constant portfolio strategy earns the wealth

32nd Conference on Neural Information Processing Systems (NeurIPS 2018), Montréal, Canada.

Table 1: Regret bounds for the full-feedback setting.

Constraints	Upper bound by Algorithm 1	Lower bound
Single asset ( $S = S_1$ )	$R_T = O(\sqrt{T \log d})$	$R_T = \Omega(\sqrt{T\log d})$
Combination ( $S = S_k$ )	$R_T = O\left(\sqrt{Tk\log\frac{d}{k}}\right)$	$R_T = \Omega\left(\sqrt{T\log\frac{d}{k}}\right) \text{ for } d \ge 17k$
	(run in $T\binom{d}{k}$ poly(k)-time)	and no $\operatorname{poly}(d, k, T)$ -time algorithm achieves $R_T \leq T^{1-\delta} \operatorname{poly}(d, k)$

Table 2: Regret bounds for the bandit-feedback setting.

Constraint	Upper bound by Algorithm 2	Lower bound
Single asset ( $S = S_1$ )	$R_T = O(\sqrt{dT \log T})$	$R_T = \Omega(\sqrt{dT})$
Combination ( $S = S_k$ )	$R_T = O\left(\sqrt{Tk\binom{d}{k}\log T}\right)$	$R_T = \Omega\left(\sqrt{T\left(\frac{d}{Ck^3}\right)^k}\right) \text{ for } d > k$
	(run in $T$ poly $(d, k)$ -time)	and no $poly(d, k, T)$ -time algorithm achieves $R_T \leq T^{1-\delta} poly(d, k)$

 $A_T^* := \max_{\mathbf{x}} \prod_{t=1}^T (1 + \mathbf{r}_t^\top \mathbf{x})$  subject to the constraint that  $\mathbf{x}$  is a portfolio satisfying  $\operatorname{supp}(\mathbf{x}) \in S$ . The performance of our portfolio selection is measured by  $R_T = \log A_T^* - \log A_T$ , which we call *regret*. The reason that we use  $\log A_T$  rather than  $A_T$  comes from capital growth theory [16, 21].<sup>1</sup> In terms of the observable information, we consider two different settings: (i) in the *full-feedback setting*, we can observe all the price relatives  $r_{ti}$  for  $i = 1, \ldots, d$ , and (ii) in the *bandit-feedback setting*, we can observe the price relatives  $r_{ti}$  only for  $i \in S_t$ . Note that in each round t a portfolio  $\mathbf{x}_t$  has to be determined before knowing  $r_{ti}$  in either of the settings. Note also that we do not make any statistical assumption about the behavior of  $r_{ti}$ , but we assume that  $r_{ti}$  is bounded in a closed interval  $[C_1, C_2]$ , where  $C_1$  and  $C_2$  are constants satisfying  $-1 < C_1 \leq C_2$ .

Our problem is a generalization of the standard online portfolio selection problem. In fact, if portfolios combining all assets are available, i.e., if  $S = 2^{[d]}$ , then our problem coincides with the standard online portfolio selection problem. For this special case, it has been shown that some online convex optimization (OCO) methods [18, 17, 27] (e.g., the online Newton step method) achieve regret of  $O(d \log T)$ , and that any algorithm will suffer from regret of  $\Omega(d \log T)$  in the worst case [26].

Our contribution is twofold; algorithms with sublinear regret upper bounds, and analyses proving regret lower bounds. First, we propose the following two algorithms:

- Algorithm 1 for the full-feedback setting, achieving regret of  $O(\sqrt{T \log |\mathcal{S}|})$ .
- Algorithm 2 for the bandit-feedback setting, achieving regret of  $O(\sqrt{Tk|S|\log T})$ , where k denotes the largest cardinality among elements in S, i.e.,  $k = \max_{S \in S} |S|$ .

Tables 1 and 2 summarize the regret bounds for the special case in which the cardinality of assets is restricted to be at most 1 or at most k. As shown in Table 1, Algorithm 1 can achieve regret of  $O(\sqrt{T}\text{poly}(d))$  even if  $k = \Omega(d)$  when S has an exponentially large size with respect to d. In such a case, however, Algorithm 1 requires exponentially large computational time. For the bandit-feedback setting, the regret upper bound can be exponential w.r.t. d if  $k = \Omega(d)$ , but it is still sublinear in T. One main idea behind our algorithms is to combine the multiplicative weight update method (MWU) [3, 14] (in the full-feedback setting) / multi-armed bandit algorithms (MAB) [5, 6] (in the bandit-feedback setting) with OCO. Specifically, for choosing the combination  $S_t$  of assets, we employ MWU/MAB, which are online decision making methods over a finite set of actions. For maintaining the proportion  $\mathbf{x}_t$  of portfolios, we use OCO, that is, online decision making methods for convex objectives over a convex set of actions.

Second, we show regret lower bounds for both the full-feedback setting and the bandit-feedback setting where  $S = S_k$ , which give insight into the tightness of regret upper bounds achieved with our algorithms. As shown in Table 1, the proven lower bounds for the full-feedback setting are tight up to the  $O(\sqrt{k})$  term. For the bandit-feedback setting, the lower bounds are also tight up to the  $O(\sqrt{\log T})$  term, if k = O(1). Note that, if k = d then the problem coincides with the standard online portfolio

<sup>&</sup>lt;sup>1</sup> For more details, see Appendix A in the supplementary material.

selection problem, and hence, there exist algorithms achieving  $R_T = O(\sqrt{T \log d})$ . This implies that the assumption of  $d = \Omega(k)$  is essential for proving the lower bounds of  $\Omega(\sqrt{T})$ . We also note that these *statistical* lower bounds are valid for arbitrary learners, including exponential-time algorithms. Besides statistical ones, we also show *computational* lower bounds suggesting that there is no polynomial-time algorithm achieving a regret bound with a sublinear term in T and a polynomial term w.r.t. d and k, unless **NP**  $\subseteq$  **BPP**. This means that we cannot improve the computational efficiency of Algorithm 1 to O(poly(d, k, T))-time while preserving its regret upper bound.

To prove the regret lower bounds, we use three different techniques: for the statistical lower bound for the full-feedback setting, we consider a completely random market and evaluate how well the "best" strategy worked after observing the market behavior, in a similar way to that for the lower bound for MWU [3]; for the bandit-feedback setting, we construct a "good" combination  $S^* \in S$ of assets so that it is hard to distinguish it from the others, and bound the number of choosing this "good" combination via a technique similar to that used in the proof of the regret lower bound for MAB [5]; to prove the computational lower bound, we reduce the 3-dimensional matching problem (3DM), one of Karp's 21 NP-complete problems [20], to our problem.

# 2 Related work

Online portfolio selection has been studied in many research areas, including finance, statistics, machine learning, and optimization [1, 10, 19, 22, 23] since Cover [10] formulated the problem setting and proposed a *universal portfolio algorithm* that achieves regret of  $O(d \log T)$  with exponential computation cost. This regret upper bound was shown to be optimal by Ordentlich and Cover [26]. The computation cost was reduced by the celebrated work on the online gradient method of Zinkervich [29] for solving online convex optimization (OCO) [17, 27], a general framework including online portfolio selection, but the regret bound is  $O(d\sqrt{T})$  and suboptimal for online portfolio selection. A breakthrough w.r.t. this suboptimality came with the online Newton step and the follow-the-approximation-leader method of Hazan et al. [18], which are computationally efficient and achieve regret of  $O(d \log T)$  for a special case of OCO, including online portfolio selection. Among studies on online portfolio selection, the work by Das et al. [12] has a motivation similar to ours: the aim of selecting portfolios with a group-sparse structure. However, their problem setting differs from ours in that they did not put constraints about sparsity but, rather, defined regret containing regularizer inducing group sparsity, and that they supposed that a learner can observe price relatives for all assets after determining portfolios. In contrast to this, our work deals with the sparsity constraint on portfolios, and our methods work even for the bandit-feedback setting, in which feedbacks are observed only on assets that have been invested in.

Another closely related topic is the multi-armed bandit problem (MAB) [4, 5, 6]. For nonstochastic MAB problems, a nearly optimal regret bound is achieved by the Exp3 algorithm [5], which our algorithm strongly relies on. For combinatorial bandit problems [7, 8, 9] in which each arm corresponds to a subset, the work by Chen et al. [8] gives solutions to a wide range of problems. However, this work does not directly apply to our setting, because we need to maintain not only subsets  $S_t$  but also continuous variables  $\mathbf{x}_t$ , and both of them affect regret.

# 3 Upper bounds

#### 3.1 Notation and preliminary consideration

Let us introduce some notations. For  $S \subseteq [d]$ , denote by  $\Delta^S$  the set of portfolios whose supports are included in S, i.e.,  $\Delta^S = \{ \mathbf{x} \mid x_i \ge 0 \ (i \in [d]), \sum_{i=1}^d x_i \le 1, \operatorname{supp}(\mathbf{x}) \subseteq S \}$ . Let  $(S^*, \mathbf{x}^*)$  denote the optimal fixed strategy for T rounds, i.e.,  $(S^*, \mathbf{x}^*) \in \underset{S \in S, \mathbf{x} \in \Delta^S}{\operatorname{arg max}} \sum_{t=1}^T \log(1 + \mathbf{r}_t^\top \mathbf{x})$ . Let  $\mathbf{x}_t$  denote the output of an algorithm for the t-th round. Then the regret  $R_T$  of the algorithm can be expressed as

$$R_T = \max_{S \in \mathcal{S}, \mathbf{x} \in \Delta^S} \sum_{t=1}^T \log(1 + \mathbf{r}_t^\top \mathbf{x}) - \sum_{t=1}^T \log(1 + \mathbf{r}_t^\top \mathbf{x}_t) = \sum_{t=1}^T \log(1 + \mathbf{r}_t^\top \mathbf{x}^*) - \sum_{t=1}^T \log(1 + \mathbf{r}_t^\top \mathbf{x}_t).$$

Algorithm 1 An algorithm for the full-feedback setting.

- **Input:** The number T of rounds. The number d of assets. The set of available subsets  $S \subseteq 2^{[d]}$ . Parameters  $\eta > 0$  and  $\beta > 0$ . 1: Set  $\mathbf{w}_1 = (w_1^S)_{S \in S} \in \mathbb{R}^S$  and  $(\mathbf{x}_1^S)_{S \in S}$  by  $w_1^S = 1$  and  $\mathbf{x}_1^S = \mathbf{0}$ , respectively, for  $S \in S$ . 2: for  $t = 1, \ldots, T$  do

- Set  $S_t$  by randomly choosing  $S \in S$  with a probability proportional to  $w_t^S$ , i.e., choose S with probability  $w_t^S / \|\mathbf{w}_t\|_1$ .
- 4:
- Output  $S_t$  and  $\mathbf{x}_t = \mathbf{x}_t^{S_t}$  and observe  $r_{ti}$  for all  $i \in [d]$ . Update  $\mathbf{w}_t$ ; set  $\mathbf{w}_{t+1}$  by  $w_{t+1}^S = w_t^S (1 + \mathbf{r}_t^\top \mathbf{x}_t^S)^\eta$  for  $S \in S$ . 5:
- Update  $\mathbf{x}_t^S$ ; set  $\mathbf{x}_{t+1}^S$  by equation (3) for  $S \in \mathcal{S}$ . 6:
- 7: end for

The algorithms presented in this section maintain vectors  $\mathbf{x}_t^S \in \Delta^S$  for all  $S \in \mathcal{S}$  at the beginning of the t-th round. They then choose  $S_t$  from S, and output  $(S_t, \mathbf{x}_t^{S_t})$ . Although other vectors  $\mathbf{x}_t^S$  $(S \neq S_t)$  do not appear in the output, they are used to compute outputs in subsequent rounds.

In the computation of  $\mathbf{x}_{t+1}^S$ , we refer to the following vectors  $\mathbf{g}_t$  and matrices  $H_t^S$ :

$$\mathbf{g}_{t}^{S} = \frac{\mathbf{r}_{t}|_{S}}{1 + \mathbf{r}_{t}^{\top}\mathbf{x}_{t}^{S}}, \quad H_{t}^{S} = \frac{(1 + C_{1})^{2}}{(1 + C_{2})^{2}} \mathbf{g}_{t}^{S} \mathbf{g}_{t}^{S\top} = C_{3} \mathbf{g}_{t}^{S} \mathbf{g}_{t}^{S\top}, \tag{1}$$

where  $\mathbf{r}_t|_S = [r'_{t1}, \ldots, r'_{td}]^\top$  is defined by  $r'_{ti} = r_{ti}$  for  $i \in S$  and  $r'_{ti} = 0$  for  $i \in [d] \setminus S$ . These  $\mathbf{g}_t^S$  and  $H_t^S$  have the following property which plays an important role in our analysis: **Lemma 1.** For any  $\mathbf{x} \in \Delta^S$ , it holds that

$$\log(1 + \mathbf{r}_t^{\mathsf{T}} \mathbf{x}) - \log(1 + \mathbf{r}_t^{\mathsf{T}} \mathbf{x}_t^S) \le \mathbf{g}_t^{S\mathsf{T}} (\mathbf{x} - \mathbf{x}_t^S) - \frac{1}{2} (\mathbf{x} - \mathbf{x}_t^S)^{\mathsf{T}} H_t^S (\mathbf{x} - \mathbf{x}_t^S).$$
(2)

For the proof, see Appendix B in the supplementary material.

#### 3.2 Algorithm for the full-feedback setting

We propose an algorithm for the full-feedback setting, created by combining the multiplicative weight update method (MWU) [3] and the follow-the-approximate-leader method (FTAL) [18]. More specifically, our proposed algorithm updates the probability of choosing a subset  $S \in S$  by MWU and updates the portfolio vector  $\mathbf{x}_t^S$  by FTAL. The entire algorithm is summarized in Algorithm 1.

Our algorithm maintains weight  $w_t^S \ge 0$  and a portfolio vector  $\mathbf{x}_t^S$  for each subset  $S \in S$  at the begining of the *t*-th round, where  $w_1^S$  and  $\mathbf{x}_1^S$  are initialized by  $w_1^S = 1$  and  $\mathbf{x}_1^S = \mathbf{0}$  for all  $S \in S$ . In each round *t*, a subset  $S_t$  is chosen with a probability proportional to  $w_t^S$ . Given the feedback  $\mathbf{r}_t$ , the algorithm computes  $w_{t+1}^S$  and  $\mathbf{x}_{t+1}^S$ . The weight  $w_{t+1}^S$  is obtained from  $w_t^S$  by multiplying  $(1 + \mathbf{r}_t^{\top} \mathbf{x}_t^S)^{\eta}$ , where  $\eta > 0$  is a parameter we optimize later. The portfolio vector  $\mathbf{x}_{t+1}^S$  is computed by FTAL as follows:

$$\mathbf{x}_{t+1}^{S} \in \operatorname*{arg\,max}_{\mathbf{x}\in\Delta^{S}} \left\{ \sum_{j=1}^{t} \left( \mathbf{g}_{j}^{S\top} (\mathbf{x} - \mathbf{x}_{j}^{S}) - \frac{1}{2} (\mathbf{x} - \mathbf{x}_{j}^{S})^{\top} H_{j}^{S} (\mathbf{x} - \mathbf{x}_{j}^{S}) \right) - \frac{\beta}{2} \|\mathbf{x}\|_{2}^{2} \right\}, \quad (3)$$

where  $\beta$  is a regularization parameter optimized later, and  $\|\cdot\|$  stands for the  $\ell_2$  norm:  $||[x_1, \ldots, x_d]^\top||_2^2 = \sum_{i=1}^d x_i^2$ . Since (3) is a convex quadratic programming problem with lin-ear constraints,  $\mathbf{x}_{t+1}^S$  can be computed efficiently by, e.g., interior point methods [24]. Recently, Ye et al. [28] have proposed a more efficient algorithm for solving (3). For the special case of the single asset selection setting, i.e., if  $S = S_1 = \{\{i\} \mid i \in [d]\}$ , then  $\mathbf{x}_{t+1}^{\{i\}} = (0, \dots, 0, x_{t+1,i}, 0, \dots, 0)$ has a closed-form expression:  $x_{t+1,i} = \pi_{[0,1]} \left(\frac{\sum_{j=1}^t g_{ji}}{\beta + C_3 \sum_{j=1}^t g_{ji}^2}\right)$ , where  $g_{ji} := \frac{r_{ji}}{1 + r_{ji}x_{ji}}$  and  $\pi_{[0,1]}(\cdot)$ stands for a projection onto [0, 1] defined by  $\pi_{[0,1]}(y) = 0$  for  $y < 0, \pi_{[0,1]}(y) = y$  for  $0 \le y \le 1$ , and  $\pi_{[0,1]}(y) = 1$  for y > 1.

Our algorithm achieves the regret described below for arbitrary inputs, where constants  $C_3$ ,  $C_4$ ,  $C_5$  are given by  $C_3 = \frac{(1+C_1)^2}{(1+C_2)^2}$ ,  $C_4 = \log \frac{1+C_2}{1+C_1}$ , and  $C_5 = \frac{\max\{C_1^2, C_2^2\}}{(1+C_1)^2}$ .

Algorithm 2 An algorithm for the bandit-feedback setting.

**Input:** The number T of rounds. The number d of assets. The set of available subsets  $S \subseteq 2^{[d]}$ . Parameters  $\eta > 0$ ,  $\gamma \in (0, 1)$  and  $\beta > 0$ . 1: Set  $\mathbf{w}_1 = (w_1^S)_{S \in S} \in \mathbb{R}^S$  and  $(\mathbf{x}_1^S)_{S \in S}$  by  $w_1^S = 1$  and  $\mathbf{x}_1^S = \mathbf{0}$ , respectively, for  $S \in S$ . 2: for  $t = 1, \ldots, T$  do

- Set the probability vector  $\mathbf{p}_t = (p_t^S)_{S \in S} \in [0,1]^S$  by  $p_t^S = \frac{\gamma}{|S|} + (1-\gamma) \frac{w_t^S}{\|\mathbf{w}_t\|_1}$ . Randomly choose  $S_t \in S$  on the basis of the probability vector  $\mathbf{p}_t$ . Output  $S_t$  and  $\mathbf{x}_t = \mathbf{x}_t^{S_t}$ , and observe  $r_{ti}$  for  $i \in S_t$ . 3:
- 4:
- 5:
- Update  $\mathbf{w}_t$ ; set  $\mathbf{w}_{t+1}^S$  by  $w_{t+1}^{S_t} = w_{ti_t} \left(\frac{1+\mathbf{r}_t^\top \mathbf{x}_t}{1+C_1}\right)^{\eta/p_{ti_t}}$  and  $w_{t+1}^S = w_t^S$  for  $S \in \mathcal{S} \setminus \{S_t\}$ . 6:
- Update  $\mathbf{x}_{t}^{S}$ ; set  $\mathbf{x}_{t+1}^{S}$  by equation (7). 7:
- 8: end for

**Theorem 2.** Algorithm 1 achieves the following regret upper bound if  $\eta \leq 1/C_4$ :

$$\mathbf{E}[R_T] \le \frac{\log|\mathcal{S}|}{\eta} + C_4^2 \eta T + \frac{1}{2}\beta + \frac{k}{C_3} \log\left(1 + \frac{C_3 C_5 T}{\beta}\right). \tag{4}$$

(5)

In particular, setting  $\eta = \frac{1}{C_4} \min \left\{ 1, \sqrt{\frac{\log |\mathcal{S}|}{T}} \right\}$  and  $\beta = 1$ , we obtain  $\mathbf{E}[R_T] = O\left(\sqrt{T \log |\mathcal{S}|} + k \log T + \log |\mathcal{S}|\right).$ 

**Running time** If (3) can be computed in p(k)-time, Algorithm 1 runs in  $O(|\mathcal{S}|p(k))$ -time per round. If S is an exponentially large set, e.g., if  $S = \{S \subseteq [d] \mid |S| = k\}$  and  $k = \Theta(d)$ , the computational time for O(|S|p(k)) will be exponentially large w.r.t. d. This computational complexity is shown to be inevitable in Section 4.1. For the special case of the single asset selection setting, i.e., if  $S = S_1 = \{\{i\} \mid i \in [d]\}, \text{ Algorithm 1 runs in } O(d) \text{-time per round since each } \mathbf{x}_t^{\{i\}} \text{ can be updated}$ in constant time.

#### 3.3 Algorithm for the bandit-feedback setting

We construct an algorithm for the bandit-feedback setting by combining the Exp3 algorithm [5] for the multi-armed bandit problem and FTAL. Similarly to the process used in Algorithm 1, the algorithm updates the probability of choosing  $S_t \in S$  by the Exp3 algorithm (in place of MWU) and updates portfolios  $\mathbf{x}_{k}^{F}$  by FTAL. The main difficulty comes from the fact that the learner cannot observe all the entries of  $(r_{ti})_{i=1}^d$ . Due to this limitation, we cannot always update  $\mathbf{x}_t^S$  for all  $S \in \mathcal{S}$ . In order to deal with this problem, we construct unbiased estimators of  $\mathbf{g}_t^S$  and  $H_t^S$  for each  $S \in \mathcal{S}$ by

$$\hat{\mathbf{g}}_{t}^{S_{t}} = \frac{\mathbf{g}_{t}^{S_{t}}}{p_{t}^{S_{t}}}, \quad \hat{H}_{t}^{S_{t}} = \frac{H_{t}^{S_{t}}}{p_{t}^{S_{t}}}, \qquad \qquad \hat{\mathbf{g}}_{t}^{S} = \mathbf{0}, \quad \hat{H}_{t}^{S} = O \quad (S \in \mathcal{S} \setminus \{S_{t}\}), \qquad \qquad (6)$$

where  $p_t^S$  is the probability of choosing S in round t, which is computed by a procedure similar to that used in the Exp3 algorithm. Note that  $\hat{\mathbf{g}}_t^S$  and  $\hat{H}_t^S$  can be calculated from the observed information alone. Using these unbiased estimators, we compute the portfolio vectors  $\mathbf{x}_{t+1}^S$  by FTAL as follows:

$$\mathbf{x}_{t+1}^{S} \in \operatorname*{arg\,max}_{\mathbf{x} \in \Delta^{S}} \left\{ \sum_{j=1}^{t} \left( \hat{\mathbf{g}}_{j}^{S\top} (\mathbf{x} - \mathbf{x}_{j}^{S}) - \frac{1}{2} (\mathbf{x} - \mathbf{x}_{j}^{S})^{\top} \hat{H}_{j}^{S} (\mathbf{x} - \mathbf{x}_{j}^{S}) \right) - \frac{1}{2} \beta \|\mathbf{x}\|_{2}^{2} \right\}.$$
(7)

Note that  $\mathbf{x}_{t+1}^S = \mathbf{x}_t^S$  for each  $S \in S \setminus \{S_t\}$  since  $\hat{\mathbf{g}}_t^S = \mathbf{0}$  and  $\hat{H}_t^S = O$ . Hence the convex quadratic programming problem (7) is solved only once in each round. The entire algorithm is summarized in Algorithm 2.

**Theorem 3.** Algorithm 2 achieves the following regret upper bound if  $\eta \leq \frac{\gamma}{C_4|S|}$ :

$$\mathbf{E}[R_T] \le \frac{\log|\mathcal{S}|}{\eta} + (C_4^2\eta|\mathcal{S}| + C_4\gamma)T + \frac{1}{2}\beta + \frac{k|\mathcal{S}|}{C_3\gamma}\log\left(1 + \frac{C_3C_5T}{\beta}\right).$$
(8)

Setting 
$$\gamma = \min\left\{1, \sqrt{\frac{k|\mathcal{S}|\log(1+T)}{T}}\right\}, \eta = \frac{\gamma}{C_4|\mathcal{S}|}\min\left\{1, \sqrt{\frac{\log|\mathcal{S}|}{k\log(1+T)}}\right\}$$
 and  $\beta = C_3C_5$ , we obtain  $\mathbf{E}[R_T] = O\left(\sqrt{T|\mathcal{S}|k\log T} + |\mathcal{S}|\sqrt{k\log|\mathcal{S}|\log T} + |\mathcal{S}|k\right).$ 

**Running time** Algorithm 2 runs in  $O(p(k) + \log^2(|\mathcal{S}|))$ -time per round, assuming that (7) can be computed in p(k)-time. In fact, from the definition (6) of  $\hat{\mathbf{g}}_t^S$  and  $\hat{H}_t^S$ , the update of  $\mathbf{x}_t^S$  given by (7) is needed only for  $S = S_t$ . Furthermore, for  $\mathcal{S} = \{S_1, S_2, \dots, S_{|\mathcal{S}|}\}$ , both updating  $w_t^S$  for some  $S \in \mathcal{S}$  and computing the prefix sum  $\sum_{j=1}^i w_t^{S_j}$  for some  $i \in [|\mathcal{S}|]$  can be performed in  $O(\log |\mathcal{S}|)$ -time by using a Fenwick tree [13]. This implies that sampling  $S_t$  w.r.t.  $p_t^S = \frac{\gamma}{|\mathcal{S}|} + \frac{w_t^S}{||\mathbf{w}_t||^S}$  can be performed in  $O(\log^2 |\mathcal{S}|)$ -time.

#### 4 Lower bounds

In this section, we present lower bounds on regrets achievable by algorithms for the online portfolio selection problem. We focus on the case of  $S = S_k = \{S \subseteq [d] \mid |S| = k\}$  throughout this section.

#### 4.1 Computational complexity

We show that, unless the complexity class **BPP** includes **NP**, there exists no algorithm for the online problem with a cardinality constraint such that its running time is polynomial both on d and T and its regret is bounded by a polynomial in d and sublinear in T. This fact is shown by presenting a reduction from the 3-dimensional matching problem (**3DM**). An instance U of **3DM** consists of 3-tuples  $(x_1, y_1, z_1), \ldots, (x_d, y_d, z_d) \in [k] \times [k] \times [k]$ . Two tuples,  $(x_i, y_i, z_i)$  and  $(x_j, y_j, z_j)$ , are called disjoint if  $x_i \neq x_j$ ,  $y_i \neq y_j$ , and  $z_i \neq z_j$ . The task of **3DM** is to determine whether or not there exist k pairwise-disjoint tuples; if they do exist, we write  $U \in$ **3DM**.

From a **3DM** instance  $U = \{(x_j, y_j, z_j)\}_{j=1}^d$ , we construct an input sequence  $(\mathbf{r}_t)_{t=1,...,T}$  of the online portfolio selection problem as follows. Let  $A = (a_{ij}) \in \{0, 1\}^{3k \times d}$  be a matrix such that  $a_{ij} = 1$  if  $i = x_j$  or  $i = k + y_j$  or  $i = 2k + z_j$ , and  $a_{ij} = 0$  otherwise. From A, we construct  $B \in \mathbb{R}^{3k \times (d+1)}$  by  $B = \frac{1}{3k} [A, -\mathbf{1}_{3k}]$ , where  $\mathbf{1}_{3k}$  is the all-one vector of dimension 3k. Let  $T \ge \max\{(4 \cdot 5184k^4)^2, (5184k^4 \cdot p_2(d))^{\frac{1}{\delta}}\}$  for an arbitrary polynomial  $p_2$  and an arbitrary positive parameter  $\delta$ . For each  $t \in [T]$ , take  $\mathbf{z}_t$  from the uniform random distribution on  $\{-1, 1\}^{3k}$ , independently. Then,  $\mathbf{r}_t$  can be defined by  $\mathbf{r}_t = \mathbf{1}_{d+1} + B^{\top} \mathbf{z}_t$  for each  $t \in [T]$ . Note that  $\mathbf{r}_t \in [0, 2]^{(d+1)}$  holds for each  $t \in [T]$ .

We give the sequence  $(\mathbf{r}_t)_{t=1,...,T}$  to an algorithm  $\mathcal{A}$ . Let  $(\mathbf{x}_t)_{t=1,...,T}$  denote the sequence output by  $\mathcal{A}$ . We determine that  $U \in \mathbf{3DM}$  if  $\sum_{t=1}^{T} \log(1 + \mathbf{r}_t^{\top} \mathbf{x}_t) \geq T(\log 2 - \frac{1}{5184k^4})$  holds, while otherwise we determine  $U \notin \mathbf{3DM}$  to hold. We can prove that this determination is correct with a probability of at least 2/3. For the proof, see Appendix E in the supplementary material.

**Theorem 4.** Let  $\delta$  be an arbitrary positive number, and  $p_1$  and  $p_2$  be arbitrary polynomials. Assume that there exists a  $p_1(d, T)$ -time algorithm  $\mathcal{A}$  for the full-feedback online portfolio selection problem with  $\mathcal{S} = \mathcal{S}_{k+1}$  that achieves regret  $R_T \leq p_2(d)T^{1-\delta}$  with a probability of at least 2/3. Then, given a **3DM** instance  $U \subseteq [k] \times [k] \times [k]$ , one can decide if  $U \in$ **3DM** with a probability of at least 2/3 in  $p_1(|U|, \max\{k^8, (k^4p_2(|U|))^{\frac{1}{\delta}}\})$ -time.

**Corollary 5.** Under the assumption of NP  $\not\subseteq$  BPP, if an algorithm achieves  $O(p(d, k)T^{1-\delta})$  regret for arbitrary d and arbitrary k, the algorithm will not run in polynomial time, i.e., the running time will be larger than any polynomial for some d and some k.

Note that the computational lower bounds described in Theorem 4 and Corollary 5 are also valid for the bandit-feedback setting, since algorithms for the bandit-feedback settings can be used for the full-feedback setting.

#### 4.2 Regret lower bound for the full-feedback setting

We show here that, for the full-feedback setting of the online portfolio selection problem with  $S = S_k$ , every algorithm (including exponential-time algorithms) suffers from regret of  $\Omega\left(\sqrt{T\log\frac{d}{k}}\right)$  in the

worst case. We can show this by analyzing the behavior of an algorithm for a certain random input. In the analysis, we use the fact that the following two inequalities hold when  $\mathbf{r}_t$  follows the discrete uniform distribution on  $\{0, 1\}^d$  independently:

$$\begin{split} \mathbf{E}_{\mathbf{r}_{t},\mathbf{x}_{t}}\left[\sum_{t=1}^{T}\log(1+\mathbf{r}_{t}^{\top}\mathbf{x}_{t})\right] &\leq T \mathbf{E}_{X}\left[\log\left(1+\frac{1}{k}X\right)\right],\\ \mathbf{E}_{\mathbf{r}_{t},\mathbf{x}_{t}}\left[\max_{S\in\mathcal{S}_{k},\mathbf{x}\in\Delta^{S}}\sum_{t=1}^{T}\log(1+\mathbf{r}_{t}^{\top}\mathbf{x})\right] &\geq T \cdot \mathbf{E}_{X}\left[\log\left(1+\frac{1}{k}X\right)\right] + \Omega\left(\sqrt{T\log\frac{d}{k}}\right), \end{split}$$

where X is a binomial random variable following B(k, 1/2). See Appendix F for details regarding the proof.

**Theorem 6.** Let  $d \ge 17k$ , and consider the online portfolio selection problem with d assets and available combinations  $S = S_k$ . There is a probability distribution of input sequences  $\{\mathbf{r}_t\}_{t=1}^T$  such that the regret of any algorithm for the full-feedback setting is bounded as  $\mathbf{E}[R_T] = \Omega\left(\sqrt{T\log\frac{d}{k}}\right)$ , where the expectation is with respect to the randomness of both  $\mathbf{r}$  and the algorithm.

#### 4.3 Regret lower bound for the bandit-feedback setting

In this subsection, we consider the bandit-feedback setting of the online portfolio selection problem with  $S = S_k$ . We show that every algorithm (including exponential-time algorithms) for this setting suffers from regret of  $\Omega\left(\sqrt{T(\frac{d}{Ck^3})^k}\right)$  when the input sequence is defined as follows. Let  $S^* \in S_k$ . We define a random distribution  $D_{S^*}$  on  $\{-1,1\}^d$  so that a random vector  $\mathbf{z} = [z_1, \ldots, z_d]^\top$ following this distribution satisfies

$$\prod_{i \in S^*} z_i = \begin{cases} 1 & \text{w.p. } 1/2 - \epsilon \\ -1 & \text{w.p. } 1/2 + \epsilon \end{cases}, \quad \prod_{i \in S} z_i = \begin{cases} 1 & \text{w.p. } 1/2 \\ -1 & \text{w.p. } 1/2 \end{cases} \quad (S \in 2^{[d]} \setminus \{\emptyset, S^*\}).$$

Such a distribution can be constructed as follows: fix an index  $i^* \in S^*$ , let  $z_i = \begin{cases} 1 & \text{w.p. } 1/2 \\ -1 & \text{w.p. } 1/2 \end{cases}$ 

for each  $i \in [d] \setminus \{i^*\}$ , and let  $z_0 = \begin{cases} 1 & \text{w.p. } 1/2 - \epsilon \\ -1 & \text{w.p. } 1/2 + \epsilon \end{cases}$  independently. Define  $z_{i^*} = z_0 \prod_{i \in S^* \setminus \{i^*\}} z_i$ . Then  $\mathbf{z} = [z_1, \ldots, z_d]^\top \sim D_{S^*}$ . The price relative vector  $\mathbf{r}_t$  in the *t*-th round can be defined by  $\mathbf{r}_t = \mathbf{1}_d - \mathbf{z}_t$ , where  $\mathbf{z}_t \sim D_S^*$  independently for  $t \in [T]$ . We can show that  $\mathbf{r}_t|_S$  follows a uniform distribution for any  $S \in \mathcal{S}_k \setminus \{S^*\}$  and only  $\mathbf{r}_t|_{S^*}$  follows a slightly different distribution. Because of this, it is difficult for algorithms to distinguish  $S^*$  from others, which makes their regrets large. For more details, see Appendix G.

**Theorem 7.** Let  $d \ge k - 1$ , and consider the online portfolio selection problem with d assets and available combinations  $S = S_k$ . There is a probability distribution of input sequences  $\{\mathbf{r}_t\}_{t=1}^T$  such that the regret of any algorithm for the bandit-feedback setting is bounded as  $\mathbf{E}[R_T] = \Omega\left(\min\left\{\frac{T}{k(Ck)^k}, \sqrt{T(\frac{d}{Ck^3})^k}\right\}\right)$ , where the expectation is with respect to the randomness of both  $\mathbf{r}$  and the algorithm, and C is a constant depending on  $C_1$  and  $C_2$ .

## 5 Experimental evaluation

We show the empirical performance of our algorithms through experiments over synthetic and realworld data. In this section, we consider the online portfolio selection problem with  $S = S_1$ . A problem instance is parameterized by a tuple  $(d, T, \{\mathbf{r}_t\}_{t=1}^T)$ . A synthetic instance is generated as follows: given parameters  $d, T, C_1$ , and  $C_2$ , we randomly choose an asset  $i^*$  from [d], and generate  $r_{ti^*} \sim U((C_2 + C_1)/2, C_2)$  and  $r_{ti} \sim U(C_1, C_2)$  for  $i \in [d] \setminus \{i^*\}$ .

We also conduct our experiments for two real-world instances. The first is based on crypto coin historical data<sup>2</sup>, including dates and price data for 19 crypto coins. From this data, we select 7 crypto coins, each having 929 prices, and obtain price relatives  $r_{ti}$  of coin *i* at time *t* by  $(p_{ti}/p_{t-1,i}) - 1$ , where  $p_{ti}$  indicates the price of coin *i* at time *t*. Thus, d = 7 and T = 928 in this instance. The other

<sup>&</sup>lt;sup>2</sup>https://www.kaggle.com/sudalairajkumar/cryptocurrencypricehistory

instance is based on S&P 500 stock data<sup>3</sup>, including dates and price data for 505 companies. From this data, we choose d = 470 companies, each having 1259 stock prices, and compute T = 1258 price relatives for each company in the same way.

For purposes of comparison, we prepare three baseline algorithms: Exp3\_cont, Exp3\_disc, and MWU\_disc. MWU\_disc (based on MWU [3]) works in the full-feedback setting and is compared with Algorithm 1. Exp3\_cont and Exp3\_disc (based on Exp3 [5]) work in the bandit-feedback setting and are compared with Algorithm 2. These baseline algorithms have different ways of updating  $\mathbf{x}_t^S$  from those of Algorithms 1 and 2. Note that since  $S = S_1 = \{\{i\} \mid i \in [d]\}, \mathbf{x}_t^S$  can be expressed as  $\mathbf{x}_t^S = \mathbf{x}_t^{\{i\}} = [0, \dots, 0, x_{ti}, 0, \dots, 0]^\top$ . Below, we offer a brief explanation of the comparisons.

- **MWU\_disc** Set  $x_{ti} = 1$  if  $\sum_{j=1}^{t-1} r_{ji} \ge 0$  and  $x_{ti} = 0$  otherwise. For each  $t \in [T]$ , select  $i_t$  by MWU, where rewards in the *t*-th round are given by  $[\log(1 + r_{ti}x_{ti})]_{i=1}^d$ , and output  $i_t, \mathbf{x}_t^{\{i_t\}}$ .
- **Exp3\_disc** Set  $x_{ti} = 1$  if  $\sum_{j \in [t-1]: i_j=i} r_{ji} \ge 0$  and  $x_{ti} = 0$  otherwise. For each  $t \in [T]$ , select  $i_t$  by Exp3, where reward in the *t*-th round is given by  $\log(1 + r_{ti_t} x_{ti_t})$ , and output  $i_t, \mathbf{x}_t^{\{i_t\}}$ .
- **Exp3\_cont** Set a parameter  $B \in \mathbb{N}$ , and consider an MAB problem instance with d(B + 1) arms in which the rewards for the d(B + 1) arms in the *t*-th round are given by  $(\log(1 + r_{ti}b/B))_{1 \le i \le d, 0 \le b \le B}$ . Apply Exp3 to this MAB problem instance.

We assess the performance of the algorithms on the basis of regrets for synthetic instances and of cumulative price relatives for real-world instances, where regrets and cumulative price relatives are averaged over 10 executions. We set parameters  $\eta$  according to Theorem 2 for Algorithm 1 and MWU\_disc, and  $\eta$  and  $\gamma$  according to Theorem 3 for Algorithm 2, Exp3\_disc, and Exp3\_cont.

Figure 1 shows average regrets for a synthetic instance with  $(d, T, C_1, C_2) = (20, 10000, -0.5, 0.5)$ . We observe that both Algorithms 1 and 2 converge faster than MWU\_disc, Exp3\_cont, and Exp3\_disc. In addition, the results empirically show that our theoretical bounds are correct.

Figures 2 and 3 show average cumulative price relatives for a real-world instance of S&P 500 stock data with  $(d, T, C_1, C_2) = (470, 1258, -0.34, 1.04)$  and for a real-world instance of crypto coin data with  $(d, T, C_1, C_2) = (7, 928, -0.7, 3.76)$ , respectively. From these figures, we observe that the cumulative price relatives of our algorithms are higher than those of baseline algorithms.

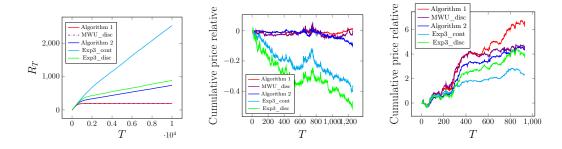


Figure 1: The average regrets over the synthetic dataset with  $(d, T, C_1, C_2) = (20, 10000, -0.5, 0.5)$ 

Figure 2: The average cumulative price relatives over S&P 500 stock dataset

Figure 3: The average cumulative price relatives over the cryptocoin historical dataset

# Acknowledgement

This work was supported by JST ERATO Grant Number JPMJER1201, Japan, and JSPS KAKENHI Grant Number JP18H05291.

<sup>&</sup>lt;sup>3</sup>https://www.kaggle.com/camnugent/sandp500

## References

- [1] A. Agarwal, E. Hazan, S. Kale, and R. E. Schapire. Algorithms for portfolio management based on the Newton method. *Proceedings of the 23rd International Conference on Machine Learning* - *ICML '06*, pages 9–16, 2006.
- [2] N. Alon and J. H. Spencer. The Probabilistic Method. John Wiley & Sons, 2004.
- [3] 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.
- [4] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002.
- [5] 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.
- [6] S. Bubeck, N. Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends* (**R**) *in Machine Learning*, 5(1):1–122, 2012.
- [7] N. Cesa-Bianchi and G. Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422, 2012.
- [8] W. Chen, Y. Wang, and Y. Yuan. Combinatorial multi-armed bandit: General framework and applications. In *International Conference on Machine Learning*, pages 151–159, 2013.
- [9] R. Combes, M. S. T. M. Shahi, A. Proutiere, et al. Combinatorial bandits revisited. In *Advances in Neural Information Processing Systems*, pages 2116–2124, 2015.
- [10] T. M. Cover. Universal portfolios. *Mathematical Finance*, 1(1):1–29, 1991.
- [11] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, 2012.
- [12] P. Das, N. Johnson, and A. Banerjee. Online portfolio selection with group sparsity. In Twenty-Eighth AAAI Conference on Artificial Intelligence, 2014.
- [13] P. M. Fenwick. A new data structure for cumulative frequency tables. Software: Practice and Experience, 24(3):327–336, 1994.
- [14] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- [15] J. E. Gentle. Computational Statistics. Springer Science & Business Media, 2009.
- [16] N. H. Hakansson and W. T. Ziemba. Capital growth theory. Handbooks in Operations Research and Management Science, 9:65–86, 1995.
- [17] E. Hazan. Introduction to online convex optimization. *Foundations and Trends*(R) *in Optimization*, 2(3-4):157–325, 2016.
- [18] E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.
- [19] A. Kalai and S. Vempala. Efficient algorithms for universal portfolios. *Journal of Machine Learning Research*, 3(Nov):423–440, 2002.
- [20] R. M. Karp. Reducibility among combinatorial problems. In Complexity of Computer Computations, pages 85–103. Springer, 1972.
- [21] J. Kelly. A new interpretation of information rate. Bell Sys. Tech. Journal, 35:917–926, 1956.
- [22] B. Li and S. C. Hoi. Online portfolio selection: A survey. ACM Computing Surveys (CSUR), 46(3):35, 2014.
- [23] B. Li and S. C. H. Hoi. On-Line Portfolio Selection with Moving Average Reversion. Proceedings of the 29th International Conference on Machine Learning (ICML-12), pages 273–280, 2012.

- [24] M. S. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret. Applications of second-order cone programming. *Linear Algebra and Its Applications*, 284(1-3):193–228, 1998.
- [25] J. Matoušek and J. Vondrák. The probabilistic method. Lecture Notes, Department of Applied Mathematics, Charles University, Prague, 2001.
- [26] E. Ordentlich and T. M. Cover. The cost of achieving the best portfolio in hindsight. *Mathematics of Operations Research*, 23(4):960–982, 1998.
- [27] S. Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends*® *in Machine Learning*, 4(2):107–194, 2012.
- [28] Y. Ye, L. Lei, and C. Ju. Hones: A fast and tuning-free homotopy method for online newton step. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics (AISTATS-18)*, pages 2008–2017, 2018.
- [29] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In Proceedings of the 20th International Conference on Machine Learning (ICML-03), pages 928–936, 2003.

# Appendix

#### A A note on the definition of the regret

**Remark 1.** When the reward  $A_T$  changes multiplicatively, the expectation of the logarithm  $\mathbf{E}[\log A_T]$  can be regarded to be a more reasonable evaluation metrics than would be the expected reward  $\mathbf{E}[A_T]$ . This is supported by the following example: suppose that  $(\mathbf{X}_t)_{t=1}^T = ((X_t^{(1)}, X_t^{(2)}))_{t=1}^T$  are Bernoulli random variables such that  $X_t^{(1)} = \begin{cases} 1.3 & \text{w. p. } 0.5 \\ 0.9 & \text{w. p. } 0.5 \end{cases}$ ,  $X_t^{(2)} = \begin{cases} 2.0 & \text{w. p. } 0.5 \\ 0.4 & \text{w. p. } 0.5 \end{cases}$ , and that  $\mathbf{X}_t$  and  $\mathbf{X}_{t'}$  are independent random variables for  $t \neq t'$ . Note that we do not assume  $X_t^{(1)}$  and  $X_t^{(2)}$  to be independent. Define  $A_T^{(1)} = \prod_{t=1}^T X_t^{(1)}$  and  $A_T^{(2)} = \prod_{t=1}^T X_t^{(2)}$ . Then, since  $\mathbf{E}[X_t^{(1)}] = 1.1$  and  $\mathbf{E}[X_t^{(1)}] = 1.2$ , we have  $\mathbf{E}[A_T^{(1)}] = 1.1^T < \mathbf{E}[A_T^{(2)}] = 1.2^T$ , which implies that we prefer  $A_T^{(1)}$  to  $A_T^{(2)}$  when determining on the basis of the expectation. However, we can show that  $\lim_{T\to\infty} A_T^{(1)} = \infty$  and  $\lim_{T\to\infty} A_T^{(2)} = 0$  with probability one, respectively. In fact, if  $A_T = \prod_{t=1}^T X_t$  is the product of i.i.d. random variables, we have

$$\lim_{T \to \infty} (A_T)^{\frac{1}{T}} = \exp\left(\lim_{T \to \infty} \frac{1}{T} \sum_{t=1}^T \log X_t\right) = \exp(\mathbf{E}[\log X_1]) \tag{9}$$

with probability one, where the last equality comes from the law of large numbers. Applying (9) to  $A_T^{(1)}$  and  $A_T^{(2)}$ , we obtain  $\lim_{T\to\infty} (A_T^{(2)})^{\frac{1}{T}} < 1 < \lim_{T\to\infty} (A_T^{(1)})^{\frac{1}{T}}$  with probability one. In general, if  $A_T^{(1)} = \prod_{t=1}^T X_t^{(1)}$  and  $A_T^{(2)} = \prod_{t=1}^T X_t^{(2)}$  are products of i.i.d. random variables, then  $\mathbf{E}[\log X_1^{(1)}] > \mathbf{E}[\log X_1^{(2)}]$  if and only if  $\lim_{t=1}^T A_T^{(1)}/A_T^{(2)} = \infty$  with probability one. These arguments imply that, in the case of a multiplicative reward model, it is reasonable to compare reward logarithms if we focus on events expected to happen with high probability.

# **B** Proof of Lemma 1

 $\begin{array}{l} \textit{Proof. Since it holds for all } x, x_0 \in [C_1, C_2] \text{ that } \frac{\mathrm{d}}{\mathrm{d}x} \log(1+x) = \frac{1}{1+x} \text{ and } \frac{\mathrm{d}^2}{\mathrm{d}x^2} \log(1+x) = \\ -\frac{1}{(1+x)^2}, \text{ we have } \log(1+x) - \log(1+x_0) \leq \frac{x-x_0}{1+x_0} - \frac{(x-x_0)^2}{2(1+C_2)^2} \leq \frac{x-x_0}{1+x_0} - \frac{C_3}{2}(\frac{x-x_0}{1+x_0})^2, \text{ where } \\ \text{we set } C_3 = \frac{(1+C_1)^2}{(1+C_2)^2}. \text{ Hence, by substituting } x = \mathbf{r}_t^\top \mathbf{x}, x_0 = \mathbf{r}_t^\top \mathbf{x}_t^S \text{ for arbitrary } t \in [T], S \in \mathcal{S} \\ \text{and } \mathbf{x} \in \Delta^S, \text{ we obtain } \log(1+\mathbf{r}_t^\top \mathbf{x}) - \log(1+\mathbf{r}_t^\top \mathbf{x}_t^S) \leq \frac{\mathbf{r}_t^\top (\mathbf{x}-\mathbf{x}_t^S)}{1+\mathbf{r}_t^\top \mathbf{x}_t^S} - \frac{C_3}{2} \left(\frac{\mathbf{r}_t^\top (\mathbf{x}-\mathbf{x}_t^S)}{1+\mathbf{r}_t^\top \mathbf{x}_t^S}\right)^2 \leq \\ \mathbf{g}_t^{S^\top} (\mathbf{x} - \mathbf{x}_t^S) - \frac{C_3}{2} (\mathbf{g}_t^{S^\top} (\mathbf{x} - \mathbf{x}_t^S))^2 = \mathbf{g}_t^{S^\top} (\mathbf{x} - \mathbf{x}_t^S) - \frac{1}{2} (\mathbf{x} - \mathbf{x}_t^S)^\top H_t^S (\mathbf{x} - \mathbf{x}_t^S), \qquad \Box \end{array}$ 

# C Proof of Theorem 2

*Proof.* In the following, we denote  $f_t(\mathbf{x}) = \log(1 + \mathbf{r}_t^{\top} \mathbf{x}) - \log(1 + C_1)$ . The regret  $R_T$  can be expressed as

$$R_{T} = \left(\sum_{t=1}^{T} f_{t}(\mathbf{x}^{*}) - \sum_{t=1}^{T} f_{t}(\mathbf{x}_{t}^{S^{*}})\right) + \left(\sum_{t=1}^{T} f_{t}(\mathbf{x}_{t}^{S^{*}}) - \sum_{t=1}^{T} f_{t}(\mathbf{x}_{t}^{S_{t}})\right).$$
(10)

Since  $S_t$  is chosen by MWU taking the input  $(F_t^S)_{S \in S} = (f_t(\mathbf{x}_t^S))_{S \in S}$ , the second term on the right-hand side of (10) can be bounded as follows (see e.g., [3]):

$$\mathbf{E}\left[\sum_{t=1}^{T} f_t(\mathbf{x}_t^{S^*}) - \sum_{t=1}^{T} f_t(\mathbf{x}_t^{S_t})\right] \le \frac{\log|\mathcal{S}|}{\eta} + C_4^2 \eta T.$$
(11)

Since  $\mathbf{x}_t^{S^*}$  is computed by FTAL, the first term on the right-hand side of (10) can be bounded as follows (see e.g., [18]):

$$\sum_{t=1}^{T} f_t(\mathbf{x}^*) - \sum_{t=1}^{T} f_t(\mathbf{x}_t^{S^*}) \le \frac{\beta}{2} + \frac{|S^*|}{C_3} \log\left(1 + \frac{C_3 C_5 T}{\beta}\right).$$
(12)

Combining (10), (11) and (12), we obtain (4).

#### **D Proof of Theorem 3**

*Proof.* The regret  $R_T$  can be expressed as (10). Since  $S_t$  is chosen by Exp3 taking the input  $(F_t^S)_{S \in S} = (f_t(\mathbf{x}_t^S))_{S \in S}$ , the second term on the right-hand side of (10) can be bounded as follows (see e.g., [5]):

$$\mathbf{E}\left[\sum_{t=1}^{T} f_t(\mathbf{x}_t^{S^*}) - \sum_{t=1}^{T} f_t(\mathbf{x}_t^{S_t})\right] \le \frac{\log|\mathcal{S}|}{\eta} + (C_4^2\eta|\mathcal{S}| + C_4\gamma)T.$$
(13)

The first term on the right-hand side of (10) can be bounded as follows:

$$\mathbf{E}\left[\sum_{t=1}^{T} f_{t}(\mathbf{x}^{*}) - \sum_{t=1}^{T} f_{t}(\mathbf{x}_{t}^{S^{*}})\right] \leq \mathbf{E}\left[\sum_{t=1}^{T} (\mathbf{g}_{t}^{S^{*}\top}(\mathbf{x}^{*} - \mathbf{x}_{t}^{S^{*}}) - \frac{1}{2}(\mathbf{x}^{*} - \mathbf{x}_{t}^{S^{*}})^{\top} H_{t}^{S^{*}}(\mathbf{x}^{*} - \mathbf{x}_{t}^{S^{*}}))\right] \\
= \mathbf{E}\left[\sum_{t=1}^{T} (\hat{\mathbf{g}}_{t}^{S^{*}\top}(\mathbf{x}^{*} - \mathbf{x}_{t}^{S^{*}}) - \frac{1}{2}(\mathbf{x}^{*} - \mathbf{x}_{t}^{S^{*}})^{\top} \hat{H}_{t}^{S^{*}}(\mathbf{x}^{*} - \mathbf{x}_{t}^{S^{*}}))\right], \tag{14}$$

where the inequality comes from (2) and the equality comes from the fact that  $\hat{\mathbf{g}}_t^S$  and  $\hat{H}_t^S$  are unbiased estimators of  $\mathbf{g}_t^S$  and  $H_t^S$ , respectively. Since  $\mathbf{x}_t^{S^*}$  is computed by FTAL as in (7), the right-hand side of can be bounded as follows (see e.g., [18]):

$$\sum_{t=1}^{T} (\hat{\mathbf{g}}_{t}^{S^{*\top}} (\mathbf{x}^{*} - \mathbf{x}_{t}^{S^{*}}) - \frac{1}{2} (\mathbf{x}^{*} - \mathbf{x}_{t}^{S^{*}})^{\top} \hat{H}_{t}^{S^{*}} (\mathbf{x}^{*} - \mathbf{x}_{t}^{S^{*}}))$$

$$\leq \frac{\beta \|\mathbf{x}^{*}\|_{2}^{2}}{2} + \sum_{t=1}^{T} \hat{\mathbf{g}}_{t}^{S^{*\top}} (\beta I + \sum_{j=1}^{t} \hat{H}_{j}^{S^{*}})^{-1} \hat{\mathbf{g}}_{t}^{S^{*}}$$

$$\leq \frac{\beta}{2} + \frac{|S|}{C_{3}\gamma} \sum_{t=1}^{T} C_{3} p_{t}^{S^{*}} \hat{\mathbf{g}}_{t}^{S^{*\top}} (\beta I + \sum_{j=1}^{t} C_{3} p_{j}^{S^{*}} \hat{\mathbf{g}}_{j}^{S^{*}} \hat{\mathbf{g}}_{j}^{S^{*}})^{-1} \hat{\mathbf{g}}_{t}^{S^{*}}$$

$$\leq \frac{\beta}{2} + \frac{|S|}{C_{3}\gamma} \log \frac{\det(\beta I + C_{3} \sum_{j=1}^{T} p_{j}^{S^{*}} \hat{\mathbf{g}}_{j}^{S^{*}} \hat{\mathbf{g}}_{j}^{S^{*}})}{\det \beta I}, \qquad (15)$$

where the first and third inequalities come from the standard analysis of FTAL, and the second inequality holds since  $p_t^S |S|/\gamma \leq 1$  from the definition of  $p_t^S$ . Denote  $M_T = C_3 \sum_{j=1}^T p_j^{S^*} \hat{\mathbf{g}}_j^{S^* \top}$ . Since  $\|\mathbf{g}_t^{S^*}\|_0 \leq |S^*| \leq k$ , the eigenvalues  $\{\lambda_1, \ldots, \lambda_d\}$  of  $M_T$  include at least d - k zero eigenvalues. From this and the fact that  $\lambda_i \geq 0$  and  $\sum_{j=1}^d \lambda_j = \operatorname{tr}(M_t)$ , we have  $\det(\beta I + M_T) = \prod_{j=1}^d (\beta + \lambda_i) \leq \beta^{d-k} (\beta + \frac{1}{k} \operatorname{tr}(M_T))^k$ . This inequality and Jensen's inequality yield  $\mathbf{E}[\log(\det(\beta I + M_T))] \leq (d-k) \log \beta + \mathbf{E}[k \log(\beta + \frac{1}{k} \operatorname{tr}(M_T))] \leq (d-k) \log \beta + k \log(\beta + \frac{1}{k} \operatorname{tr}(M_T))] \leq TkC_3C_5$ , we have  $\mathbf{E}[\log(\det(\beta I + M_T))] \leq (d-k) \log \beta + k \log(\beta + TC_3C_5)$ . Combining this with (14) and (15), we obtain (8).

## E Proof of Theorem 4

*Proof.* From a **3DM** instance  $U = \{(x_j, y_j, z_j)\}_{j=1}^d$ , we construct an input sequence  $(\mathbf{r}_t)_{t=1,...,T}$  for algorithm  $\mathcal{A}$  as follows. Let  $A = (a_{ij}) \in \{0,1\}^{3k \times d}$  be a matrix such that  $a_{ij} = 1$  if  $i = x_j$  or  $i = k + y_j$  or  $i = 2k + z_j$ , and  $a_{ij} = 0$  otherwise. From A, we construct  $B \in \mathbb{R}^{3k \times (d+1)}$  by  $B = \frac{1}{3k} [A, -\mathbf{1}_{3k}]$ , where  $\mathbf{1}_{3k}$  is an all-one vector of dimension 3k. Let  $T \ge \max\{(4 \cdot 5184k^4)^2, (5184k^4 \cdot p_2(d))^{\frac{1}{\delta}}\}$ . For each  $t \in [T]$ , take  $\mathbf{z}_t$  from the uniform random distribution on  $\{-1, 1\}^{3k}$ , independently. Then,  $\mathbf{r}_t$  can be defined by  $\mathbf{r}_t = \mathbf{1}_{d+1} + B^{\top} \mathbf{z}_t$  for each  $t \in [T]$ .

We give the sequence  $(\mathbf{r}_t)_{t=1,...,T}$  to  $\mathcal{A}$ . Let  $(\mathbf{x}_t)_{t=1,...,T}$  denote the sequence output by  $\mathcal{A}$ . We determine that  $U \in \mathbf{3DM}$  if  $\sum_{t=1}^T \log(1 + \mathbf{r}_t^\top \mathbf{x}_t) \ge T(\log 2 - \frac{1}{5184k^4})$  holds, while otherwise we

determine that  $U \notin \mathbf{3DM}$  holds. Below, we prove that this determination is correct with a probability of at least 2/3.

Assume that  $U \in 3DM$ . Then, there exists  $\mathbf{y}^* \in \{0,1\}^d$  such that  $\|\mathbf{y}^*\|_0 = k$  and  $A\mathbf{y}^* = \mathbf{1}_{3k}$ , and there exists  $\mathbf{y}^* \in \{0,1\}^d$  such that  $\|\mathbf{y}^*\|_0 = \|\mathbf{y}^*\|_1 = k$  and  $A\mathbf{y}^* = \mathbf{1}_{3k}$ . Define  $\mathbf{x}^* := \frac{1}{k+1} \begin{bmatrix} \mathbf{y}^* \\ 1 \end{bmatrix}$ . The vector  $\mathbf{x}^*$  satisfies  $\mathbf{x}^* \in \Delta^S$  for some  $S \in \mathcal{S}_{k+1}$ . Moreover, it holds that  $\mathbf{r}_t^\top \mathbf{x}^* = \mathbf{1}_{d+1}^\top \mathbf{x}^* + \mathbf{z}_t^\top B\mathbf{x}^* = 1 + \frac{1}{3k(k+1)} \mathbf{z}_t^\top (A\mathbf{y}^* - \mathbf{1}_{3k}) = 1$ . Hence, we obtain

$$\max_{S \in \mathcal{S}_{k+1}, \mathbf{x} \in \Delta^S} \sum_{t=1}^{T} \log(1 + \mathbf{r}_t^{\top} \mathbf{x}) \ge \sum_{t=1}^{T} \log(1 + \mathbf{r}_t^{\top} \mathbf{x}^*) = T \log 2.$$

From this inequality and  $R_T \leq p_2(d)T^{1-\delta}$  (with a probability  $\geq 2/3$ ), we obtain

$$\sum_{t=1}^{T} \log(1 + \mathbf{r}_t^{\top} \mathbf{x}_t) \ge \max_{S \in \mathcal{S}_{k+1}, \mathbf{x} \in \Delta^S} \sum_{t=1}^{T} \log(1 + \mathbf{r}_t^{\top} \mathbf{x}) - R_T$$
$$\ge T \log 2 - p_2(d) T^{1-\delta} \ge T \left(\log 2 - \frac{1}{5184k^4}\right)$$

where the last inequality comes from  $T \ge (5184k^4 \cdot p_2(d))^{\frac{1}{\delta}}$ . This inequality means that the decision is correct with a probability  $\geq 2/3$  if  $U \in 3DM$ .

For the remainder of the proof, we assume that  $U \notin \mathbf{3DM}$ . This assumption implies that, for all  $\mathbf{y} \in \mathbb{R}^d_{\geq 0}$  satisfying  $\|\mathbf{y}\|_0 \leq k$ , we have  $\min_{1 \leq i \leq 3k} (A\mathbf{y})_i = 0$ . Moreover, since each column of Ahas at least one entry of value 1, we have  $\max_{1 \le i \le 3k} (A\mathbf{y})_i \ge \|\mathbf{y}\|_{\infty}$  for all  $\mathbf{y} \in \mathbb{R}^d_{\ge 0}$ .

We first prove that  $||B\mathbf{x}||_2 \ge \frac{1}{12k^2} ||\mathbf{x}||_1$  holds for all  $\mathbf{x} \in \mathbb{R}^{d+1}_{\ge 0}$  satisfying  $||\mathbf{x}||_0 \le k+1$ . We consider the following two cases: the last entry of  $\mathbf{x}$  is either positive or zero. The former case is when  $\mathbf{x}$  is expressed as  $\mathbf{x} = \begin{vmatrix} \mathbf{y} \\ y_0 \end{vmatrix}$  with  $\mathbf{y} \in \mathbb{R}^d_{\geq 0}$ ,  $\|\mathbf{y}\|_0 \leq k$  and  $y_0 > 0$ . In this case, we have

$$\begin{split} \|B\mathbf{x}\|_{\infty} &\geq \frac{1}{3k} \max\left\{ |\min_{1 \leq i \leq 3k} (A\mathbf{y})_{i} - y_{0}|, |\max_{1 \leq i \leq 3k} (A\mathbf{y})_{i} - y_{0}| \right\} \\ &= \frac{1}{3k} \max\left\{ |y_{0}|, |\max_{1 \leq i \leq 3k} (A\mathbf{y})_{i} - y_{0}| \right\} \\ &\geq \frac{1}{3k} \max\left\{ |y_{0}|, \frac{1}{2}|\max_{1 \leq i \leq 3k} (A\mathbf{y})_{i}| \right\} \geq \frac{1}{3k} \max\left\{ |y_{0}|, \frac{1}{2} \|\mathbf{y}\|_{\infty} \right\} \geq \frac{1}{6k} \|\mathbf{x}\|_{\infty}, \end{split}$$

where the second inequality comes from the fact that arbitrary  $y_0$  satisfies  $\max\{|y_0|, |a-y_0|\} \ge |a|/2$ . In the latter case, namely, when  $\mathbf{x} = \begin{bmatrix} \mathbf{y} \\ 0 \end{bmatrix}$  with some  $\mathbf{x} \in \mathbb{R}^d_{\geq 0}$  such that  $\|\mathbf{y}\|_0 \leq k+1$ , we have  $\|B\mathbf{x}\|_{\infty} \geq \frac{1}{3k} \|\mathbf{y}\|_{\infty} = \frac{1}{3k} \|\mathbf{x}\|_{\infty}.$  Accordingly, in both of these cases, we have  $\|B\mathbf{x}\|_{\infty} \geq \frac{1}{6k} \|\mathbf{x}\|_{\infty}$ , and hence, we have  $\|B\mathbf{x}\|_{2} \geq \|B\mathbf{x}\|_{\infty} \geq \frac{1}{6k} \|\mathbf{x}\|_{\infty} \geq \frac{1}{6k(k+1)} \|\mathbf{x}\|_{1} \geq \frac{1}{12k^{2}} \|\mathbf{x}\|_{1}.$ 

Then, since  $\log(1+y) \leq \log 2 + \frac{1}{2}(y-1) - \frac{1}{18}(y-1)^2$  for  $y \in [0,2]$ , and since  $\mathbf{z}_t$  are statistically independent of  $\mathbf{x}_t$  and  $\mathbf{E}[\mathbf{z}_t] = \mathbf{0}$ ,  $\mathbf{E}[\mathbf{z}_t\mathbf{z}_t^{\top}] = I$ , we have

$$\begin{split} & \underset{\mathbf{z}_{t},\mathbf{x}_{t}}{\mathbf{E}} \left[ \log(1 + \mathbf{r}_{t}^{\top} \mathbf{x}_{t}) \right] \\ & \leq \log 2 + \underset{\mathbf{z}_{t},\mathbf{x}_{t}}{\mathbf{E}} \left[ \frac{1}{2} (\|\mathbf{x}_{t}\|_{1} + \mathbf{z}_{t}^{\top} B \mathbf{x}_{t} - 1) - \frac{1}{18} (\|\mathbf{x}_{t}\|_{1} + \mathbf{z}_{t}^{\top} B \mathbf{x}_{t} - 1)^{2} \right] \\ & = \log 2 - \frac{1}{2} + \underset{\mathbf{x}_{t},\mathbf{z}_{t}}{\mathbf{E}} \left[ \frac{1}{2} \|\mathbf{x}_{t}\|_{1} - \frac{1}{18} \left( \mathbf{x}_{t}^{\top} B^{\top} \mathbf{z}_{t} \mathbf{z}_{t}^{\top} B \mathbf{x}_{t} - 2 \mathbf{z}_{t}^{\top} B \mathbf{x}_{t} (\|\mathbf{x}_{t}\|_{1} - 1) + (\|\mathbf{x}_{t}\|_{1} - 1)^{2} \right) \right] \\ & \leq \log 2 - \frac{1}{2} + \underset{\mathbf{x}_{t}}{\mathbf{E}} \left[ \frac{1}{2} \|\mathbf{x}_{t}\|_{1} - \frac{1}{18} \|B \mathbf{x}_{t}\|_{2}^{2} \right] \\ & \leq \log 2 - \frac{1}{2} + \underset{\mathbf{x}_{t}}{\mathbf{E}} \left[ \frac{1}{2} \|\mathbf{x}_{t}\|_{1} - \frac{1}{18(12k^{2})^{2}} \|\mathbf{x}_{t}\|_{1}^{2} \right] \leq \log 2 - \frac{1}{2592k^{4}}. \end{split}$$

This inequality means that the stochastic process  $\{X_t\}_{t=1}^T$  defined by  $X_t = \sum_{j=1}^t \log(1 + \mathbf{r}_t \mathbf{x}_t) - t(\log 2 - \frac{1}{2592k^4})$  is a sub-martingale. From the definition,  $\{X_t\}_{t=1}^T$  satisfies  $|X_t - X_{t+1}| < \log 3$  for all t. Hence, from the Azuma-Hoeffding inequality [2],  $X_T$  is bounded as  $X_T < 4\sqrt{T}$  with a probability of at least 2/3. Consequently, we have

$$\sum_{t=1}^{T} \log(1 + \mathbf{r}_t^{\top} \mathbf{x}_t) < T(\log 2 - \frac{1}{2592k^4}) + 4\sqrt{T} \le T\left(\log 2 - \frac{1}{5184k^4}\right),$$

where the last inequality comes from  $T \ge (4 \cdot 5184k^4)^2$ . This means that the decision is correct with a probability of at least 2/3.

# F Proof of Theorem 6

Let us first consider the following lemma.

**Lemma 8.** If  $0 \le p_1 \le p_2 \le 1$  and random variables  $X_1, X_2$  follow the binomial random distributions  $B(k, p_1), B(k, p_2)$ , respectively, then we have

$$\mathbf{E}_{X_2 \sim B(k,p_2)} \left[ \log \left( 1 + \frac{1}{k} X_2 \right) \right] - \mathbf{E}_{X_1 \sim B(k,p_1)} \left[ \log \left( 1 + \frac{1}{k} X_1 \right) \right] \ge \frac{p_2 - p_1}{2}$$
(16)

*Proof.* Define  $Y_1 = k - X_1$  and  $Y_2 = k - X_2$ . Then we have  $Y_1 \sim B(k, 1 - p_1)$  and  $Y_2 \sim B(k, 1 - p_2)$ . From the Maclaurin series of  $\log(2 - x) = \log 2 - \frac{1}{2}x - \frac{1}{2 \cdot 2^2}x^2 - \cdots = \log 2 - \sum_{n=1}^{\infty} \frac{x^n}{n2^n}$ , we have

$$\begin{split} & \underset{X_{2} \sim B(k,p_{2})}{\mathbf{E}} \left[ \log \left( 1 + \frac{1}{k} X_{2} \right) \right] - \underset{X_{1} \sim B(k,p_{1})}{\mathbf{E}} \left[ \log \left( 1 + \frac{1}{k} X_{1} \right) \right] \\ &= \underset{Y_{2} \sim B(k,1-p_{2})}{\mathbf{E}} \left[ \log \left( 2 - \frac{1}{k} Y_{2} \right) \right] - \underset{Y_{1} \sim B(k,1-p_{1})}{\mathbf{E}} \left[ \log \left( 2 - \frac{1}{k} Y_{1} \right) \right] \\ &= \underset{n=1}{\overset{\infty}{\sum}} \frac{1}{n(2k)^{n}} \left( \underset{Y_{1} \sim B(k,1-p_{1})}{\mathbf{E}} \left[ Y_{1}^{n} \right] - \underset{Y_{2} \sim B(k,1-p_{2})}{\mathbf{E}} \left[ Y_{2}^{n} \right] \right) \\ &\geq \frac{1}{2k} \left( \underset{Y_{1} \sim B(k,1-p_{1})}{\mathbf{E}} \left[ Y_{1} \right] - \underset{Y_{2} \sim B(k,1-p_{2})}{\mathbf{E}} \left[ Y_{2} \right] \right) = \frac{p_{2} - p_{1}}{2}. \end{split}$$

We are now ready to prove Theorem 6.

*Proof of Theorem 6.* We construct an input sequence  $\{\mathbf{r}_t\}_{t=1,2,...}$  so that entries  $r_{ti}$  follow a uniform random distribution over  $\{0,1\}$  independently. We can show that

$$\mathbf{E}_{\mathbf{r},\mathbf{x}}[R_T(\mathbf{r})] = \mathbf{E}_{\mathbf{r},\mathbf{x}}\left[\max_{S\in\mathcal{S}_k,\mathbf{x}\in\Delta^S}\sum_{t=1}^T \log(1+\mathbf{r}_t^{\top}\mathbf{x}) - \sum_{t=1}^T \log(1+\mathbf{r}_t^{\top}\mathbf{x}_t)\right] = \Omega\left(\sqrt{T\log\frac{d}{k}}\right) \quad (17)$$

for all algorithms, by means of considering the following two inequalities:

$$\mathbf{E}_{\mathbf{r}_{t},\mathbf{x}_{t}}[\log(1+\mathbf{r}_{t}^{\top}\mathbf{x}_{t})] \leq \mathbf{E}_{X_{1}}\left[\log\left(1+\frac{1}{k}X_{1}\right)\right],\tag{18}$$

$$\mathbf{E}_{\mathbf{r}_{t},\mathbf{x}_{t}}\left[\max_{S\in\mathcal{S}_{k},\mathbf{x}\in\Delta^{S}}\sum_{t=1}^{T}\log(1+\mathbf{r}_{t}^{\top}\mathbf{x})\right] \geq T\cdot\mathbf{E}_{X_{1}}\left[\log\left(1+\frac{1}{k}X_{1}\right)\right] + \Omega\left(\sqrt{T\log\frac{d}{k}}\right), \quad (19)$$

where  $X_1$  is a binomial random variable following B(k, 1/2).

First, let us prove the inequality (18). Consider a function  $\mathbf{x} \mapsto \mathbf{E}[\log(1 + \mathbf{r}_t^\top \mathbf{x})]$ , and suppose  $S \in S_k$ . We can then confirm that this is a concave function and that, for the optimization problem

 $\underset{\mathbf{x}\in\Delta^{S}}{\arg\max} \underset{\mathbf{r}_{t}}{\mathbf{E}}[\log(1+\mathbf{r}_{t}^{\top}\mathbf{x})], \text{ the vector } \frac{1}{k}\mathbf{1}_{S} \text{ is the unique point satisfying KKT conditions, where } \mathbf{1}_{S}$ 

stands for the indicator vector of S, i.e.,  $\mathbf{1}_{S} = [\chi_{1}, \dots, \chi_{d}]^{\top}$  where  $\chi_{i} = 1$  if  $i \in S$  and  $\chi_{i} = 0$  if  $i \in [d] \setminus S$ . Consequently, we have  $\max_{\mathbf{x} \in \Delta^{S}} \log(1 + \mathbf{r}_{t}^{\top}\mathbf{x}) = \underset{\mathbf{r}_{t}}{\mathbf{E}}[\log(1 + \frac{1}{k}\mathbf{1}_{S}^{\top}\mathbf{r}_{t})] = \underset{X_{1}}{\mathbf{E}}[\log(1 + \frac{1}{k}X_{1})]$ since  $\mathbf{1}_{S}^{\top}\mathbf{r}_{t}$  follows the binomial distribution B(k, 1/2). Since  $\mathbf{x}_{t} \in \Delta^{S_{t}}$  for some  $S_{t} \in S_{k}$  and  $S_{t}, \mathbf{x}_{t}$  are stochastically independent of  $\mathbf{r}_{t}$ , we obtain  $\underset{\mathbf{r}_{t},\mathbf{x}_{t}}{\mathbf{E}}[\log(1 + \mathbf{r}_{t}^{\top}\mathbf{x}_{t})] \leq \underset{\mathbf{r}_{t}}{\mathbf{E}}[\log(1 + \frac{1}{k}\mathbf{r}_{t}^{\top}\mathbf{1}S_{t})] =$ 

 $\mathop{\mathbf{E}}_{X_1}\left[\log\left(1+\frac{1}{k}X_1\right)\right].$ 

Next, let us prove the inequality (19). For each  $i \in [d]$ , define  $r_i := \sum_{t=1}^{T} r_{ti}$ . Since  $r_{ti}$  follows a Bernoulli distribution with parameter 1/2 independently,  $r_i$  follows the binomial distribution B(T, 1/2). Let  $\sigma : [d-k] \to [d-k]$  be a permutation such that  $r_{\sigma(1)} \ge r_{\sigma(2)} \ge \cdots \ge r_{\sigma(d-k)}$ . Since the posterior random distribution of  $r_{ti}$  given  $r_i$  is the Bernoulli distribution of parameter  $r_i/T$ , for  $\mathbf{x}_2 = \frac{1}{k} \mathbf{1}_{\{\sigma(1), \sigma(2), \dots, \sigma(k)\}}$  and for arbitrary constant  $s \ge T/2$ , we have

$$\begin{aligned} \mathbf{E}_{\mathbf{r}} \left[ \chi_{\{r_{\sigma(k)} \geq s\}} \cdot \sum_{t=1}^{T} \log(1 + \mathbf{r}_{t}^{\top} \mathbf{x}_{2}) \right] &\geq \mathbf{E}_{X_{2} \sim B(k, \frac{s}{T})} \left[ \chi_{\{r_{\sigma(k)} \geq s\}} \cdot \sum_{t=1}^{T} \log(1 + \frac{1}{k} X_{2}) \right] \\ &= T \cdot \operatorname{Prob}[r_{\sigma(k)} \geq s] \cdot \mathbf{E}_{X_{2} \sim B(k, \frac{s}{T})} \left[ \log(1 + \frac{1}{k} X_{2}) \right], \end{aligned}$$

where  $\chi_A$  stands for the indicator function for arbitrary events A. Moreover, since  $r_{d-k+1}, \ldots, r_d$  are independent of  $r_{\sigma(k)}$ , for  $\mathbf{x}_1 = \frac{1}{k} \mathbf{1}_{\{d-k+1,\ldots,d\}}$ , we have

$$\mathbf{E}_{\mathbf{r}}\left[\chi_{\{r_{\sigma(k)} < s\}} \sum_{t=1}^{T} \log(1 + \mathbf{r}_{t}^{\top} \mathbf{x}_{1})\right] = T \cdot \operatorname{Prob}[r_{\sigma(k)} < s] \cdot \mathbf{E}_{X_{1} \sim B(k, \frac{1}{2})} \left[\log(1 + \frac{1}{k} X_{1})\right].$$

Hence, we obtain

$$\begin{split} & \mathbf{F}_{\mathbf{r}} \left[ \max_{S \in \mathcal{S}_{k}, \mathbf{x} \in \Delta^{S}} \sum_{t=1}^{T} \log(1 + \mathbf{r}_{t}^{\top} \mathbf{x}) \right] \\ &= \mathbf{F}_{\mathbf{r}} \left[ \chi_{\{r_{\sigma(k)} \geq s\}} \max_{S \in \mathcal{S}_{k}, \mathbf{x} \in \Delta^{S}} \sum_{t=1}^{T} \log(1 + \mathbf{r}_{t}^{\top} \mathbf{x}) \right] + \mathbf{F}_{\mathbf{r}} \left[ \chi_{\{r_{\sigma(k)} < s\}} \max_{S \in \mathcal{S}_{k}, \mathbf{x} \in \Delta^{S}} \sum_{t=1}^{T} \log(1 + \mathbf{r}_{t}^{\top} \mathbf{x}) \right] \\ &\geq \mathbf{F}_{\mathbf{r}} \left[ \chi_{\{r_{\sigma(k)} \geq s\}} \sum_{t=1}^{T} \log(1 + \mathbf{r}_{t}^{\top} \mathbf{x}_{2}) \right] + \mathbf{F}_{\mathbf{r}} \left[ \chi_{\{r_{\sigma(k)} < s\}} \sum_{t=1}^{T} \log(1 + \mathbf{r}_{t}^{\top} \mathbf{x}_{1}) \right] \\ &\geq T \cdot \operatorname{Prob}[r_{\sigma(k)} \geq s] \cdot \mathbf{E}_{X_{2}} \left[ \log(1 + \frac{1}{k} X_{2}) \right] + T \cdot \operatorname{Prob}[r_{\sigma(k)} \leq s] \cdot \mathbf{E}_{X_{1}} \left[ \sum_{t=1}^{T} \log(1 + \frac{1}{k} X_{1}) \right] \\ &\geq T \cdot \mathbf{E}_{X_{1}} \left[ \log(1 + \frac{1}{k} X_{1}) \right] + T \cdot \operatorname{Prob}[r_{\sigma(k)} \geq s] \cdot \frac{1}{2} \cdot \left( \frac{s}{T} - \frac{1}{2} \right), \end{split}$$

where  $X_1 \sim B(k, \frac{1}{2})$ ,  $X_2 \sim B(k, \frac{s}{T})$  and the last inequality comes from Lemma 8. We now can show that we have  $\operatorname{Prob}[r_{\sigma(k)} \geq s] = \Omega(1)$  for  $s = \frac{T}{2} + \Omega(\sqrt{T \log \frac{d}{k}})$ , which proves (19). Let  $F : \mathbb{R} \to [0, 1]$  denote the cumulative distribution function of B(T, 1/2), i.e.,  $F(x) = \operatorname{Prob}[r_i \leq x]$ . From a standard concentration lemma of a binomial distribution (see, e.g., Proposition 7.3.2 in [25]), we have  $F(\frac{T}{2} + t) \leq 1 - \frac{1}{15} \exp\left(-16\frac{t^2}{T}\right)$ . Hence, setting  $t = \frac{1}{4}\sqrt{T \log \frac{d-k}{15k}}$ , we obtain

$$\operatorname{Prob}\left[r_{\sigma(k)} \geq \frac{T}{2} + t\right] = \operatorname{Prob}\left[F(r_{\sigma(k)}) \geq F\left(\frac{T}{2} + t\right)\right]$$
$$\geq \operatorname{Prob}\left[F(r_{\sigma(k)}) \geq 1 - \frac{1}{15}\exp\left(-16\frac{t^2}{T}\right)\right] = \operatorname{Prob}\left[F(r_{\sigma(k)}) \geq 1 - \frac{k}{d-k}\right]$$

Since  $F(r_i)$  follows the uniform distribution on [0, 1] independently,  $F(r_{\sigma(k)})$  follows the probability distribution of the order statistic sampled from the standard uniform distribution, which is the beta

distribution Beta(d-2k+1,k) (see, e.g., [15]). This means that  $\text{Prob}[1-F(r_{\sigma(k)}) \leq \frac{k}{d-k}] \geq 1/2$ . Combining the above two inequalities, for  $s = \frac{T}{2} + \frac{1}{4}\sqrt{T\log\frac{d-k}{15k}}$ , we have

$$\begin{split} & \mathbf{E}_{\mathbf{r}} \left[ \max_{S \in \mathcal{S}_{k}, \mathbf{x} \in \Delta^{S}} \sum_{t=1}^{T} \log(1 + \mathbf{r}_{t}^{\top} \mathbf{x}) \right] \\ & \geq T \cdot \mathbf{E}_{X_{1}} \left[ \log(1 + \frac{1}{k} X_{1}) \right] + \frac{1}{2} \operatorname{Prob}[r_{\sigma(k)} \geq s] \cdot \left(s - \frac{T}{2}\right) \\ & \geq T \cdot \mathbf{E}_{X_{1}} \left[ \log(1 + \frac{1}{k} X_{1}) \right] + \frac{1}{16} \left( \sqrt{T \log \frac{d - k}{15k}} \right) \\ & \geq T \cdot \mathbf{E}_{X_{1}} \left[ \log(1 + \frac{1}{k} X_{1}) \right] + \Omega \left( \sqrt{T \log \frac{d}{k}} \right), \end{split}$$

where the last inequality comes from  $d \ge 17k$ . Consequently, we obtain (19). From (18) and (19) we have (17).

# G Proof of Theorem 7

Lemma 9. For arbitrary  $\epsilon \in [0, 1/2]$ , let  $z_0, z_1, \dots, z_k \in \{-1, 1\}$  be independent random variables such that  $z_0 = \begin{cases} 1 & \text{w.p. } 1/2 - \epsilon \\ -1 & \text{w.p. } 1/2 + \epsilon \end{cases}$  and  $z_i = \begin{cases} 1 & \text{w.p. } 1/2 \\ -1 & \text{w.p. } 1/2 \end{cases}$  for  $i = 1, \dots, d$ . Set  $X_1 = \sum_{i=1}^{k} z_i$  and  $X_2 = \sum_{i=1}^{k-1} z_i + z_0 \prod_{j=1}^{k-1} z_j$ . We then have  $\mathbf{E} \left[ \log \left( 2 - \frac{1}{k} X_2 \right) \right] - \mathbf{E} \left[ \log \left( 2 - \frac{1}{k} X_1 \right) \right] \ge \frac{2\epsilon}{k(2k)^k}.$ (20)

*Proof.* Denote  $w = z_0 \prod_{j=1}^{k-1} z_j$ . Let  $n_1, n_2, \ldots, n_k$  be arbitrary non-negative integers. Set  $m_i$  to be  $n_i$  modulo 2, i.e.,  $m_i = 0$  if  $n_i$  is even and  $m_i = 1$  if  $n_i$  is odd, for  $i = 1, \ldots, k$ . We then have

$$\mathbf{E}[z_1^{n_1} z_2^{n_2} \cdots z_{k-1}^{n_{k-1}} z_k^{n_k}] = \mathbf{E}[z_1^{m_1} z_2^{m_2} \cdots z_{k-1}^{m_{k-1}} z_k^{m_k}] = \begin{cases} 1 & \text{if } m_1 = \cdots = m_k = 0\\ 0 & \text{otherwise} \end{cases}$$
$$\mathbf{E}[z_1^{n_1} z_2^{n_2} \cdots z_{k-1}^{n_{k-1}} w^{n_k}] = \mathbf{E}[z_1^{m_1} z_2^{m_2} \cdots z_{k-1}^{m_{k-1}} w^{m_k}] = \begin{cases} 1 & \text{if } m_1 = \cdots = m_k = 0\\ -2\epsilon & \text{if } m_1 = \cdots = m_k = 1\\ 0 & \text{otherwise} \end{cases} ,$$

which means  $\mathbf{E}[z_1^{n_1}\cdots z_{k-1}^{n_{k-1}}z_k^{n_k}] \geq \mathbf{E}[z_1^{n_1}\cdots z_{k-1}^{n_{k-1}}w^{n_k}]$ . Hence,  $X_1 = \sum_{i=1}^k z_i$  and  $X_2 = \sum_{i=1}^{k-1} z_i + w$  satisfies  $\mathbf{E}[X_1^n] \geq \mathbf{E}[X_2^n]$  for all non-negative integers n and  $\mathbf{E}[X_1^k] - \mathbf{E}[X_2^k] = 2\epsilon$ . From the Maclaurin series of  $\log(2-x) = \log 2 - \sum_{n=1}^{\infty} \frac{x^n}{n2^n}$ , we have

$$\mathbf{E}\left[\log\left(2-\frac{1}{k}X_{2}\right)\right] - \mathbf{E}\left[\log\left(2-\frac{1}{k}X_{1}\right)\right] = \sum_{n=1}^{\infty} \frac{1}{n(2k)^{n}} (\mathbf{E}[X_{1}^{n}] - \mathbf{E}[X_{2}^{n}])$$
$$\geq \frac{1}{k(2k)^{k}} (\mathbf{E}[X_{1}^{k}] - \mathbf{E}[X_{2}^{k}]) = \frac{2\epsilon}{k(2k)^{k}}.$$

#### **Proof of Theorem 7**

*Proof.* For each  $S^* \in S_k$ , we define a random distribution  $D_{S^*}$  on  $\{-1,1\}^d$  so that  $\mathbf{z} = [z_1, \ldots, z_d]^\top \sim D_{S^*}$  satisfies

$$\prod_{i \in S^*} z_i = \begin{cases} 1 & \text{w. p. } 1/2 - \epsilon \\ -1 & \text{w. p. } 1/2 + \epsilon \end{cases}, \quad \prod_{i \in S} z_i = \begin{cases} 1 & \text{w. p. } 1/2 \\ -1 & \text{w. p. } 1/2 \end{cases} \quad (S \in 2^{[d]} \setminus \{\emptyset, S^*\}).$$
(21)

Such a distribution can be constructed as follows: fix an index  $i^* \in S^*$  and, for  $i \in [d] \setminus \{i^*\}$ , let  $z_i = \begin{cases} 1 & \text{w. p. } 1/2 \\ -1 & \text{w. p. } 1/2 \end{cases}$  and  $z_0 = \begin{cases} 1 & \text{w. p. } 1/2 - \epsilon \\ -1 & \text{w. p. } 1/2 + \epsilon \end{cases}$  independently. Define  $z_{i^*} = z_0 \prod_{i \in S^* \setminus \{i^*\}} z_i$ . Suppose that the input sequence  $\mathbf{r}_t$  is given by  $\mathbf{r}_t = \mathbf{1} - \mathbf{z}_t$ , where  $\mathbf{z}_t \sim D_S^*$  independently for  $t = 1, 2, \ldots, T$ . If  $\mathbf{z}$  follows  $D_S^*$ , for any  $S \in \mathcal{S}_k \setminus \{\emptyset S^*\}$ ,  $\mathbf{z}|_S$  follows the uniform distribution on  $\{-1, 1\}^S$ , and hence, we have  $\max_{\mathbf{x} \in \Delta^{S'}} \mathbf{E}_{\mathbf{z}_t \sim D_S^*} \left[ \log \left(1 + \mathbf{r}_t^\top \mathbf{x} \right) \right] = \sum_{i=1}^{N} \sum_{j=1}^{N} \sum_{i=1}^{N} \sum_{j=1}^{N} \sum_{j=1}$ 

$$\underbrace{\mathbf{E}}_{\mathbf{z}_{t}\sim D_{S^{*}}} \left[ \log\left(1 + \frac{1}{k}\mathbf{r}_{t}^{\top}\mathbf{1}_{S}\right) \right]. \text{ From Lemma 9, for } S \in \mathcal{S}_{k} \setminus \{\emptyset S^{*}\}, \text{ we have} \\
 \underbrace{\mathbf{E}}_{\mathbf{z}_{t}\sim D_{S^{*}}} \left[ \log\left(1 + \frac{1}{k}\mathbf{r}_{t}^{\top}\mathbf{1}_{S^{*}}\right) \right] - \max_{\mathbf{x}\in\Delta^{S}} \underbrace{\mathbf{E}}_{\mathbf{z}_{t}\sim D_{S^{*}}} \left[ \log\left(1 + \mathbf{r}_{t}^{\top}\mathbf{x}\right) \right] \\
 \geq \underbrace{\mathbf{E}}_{\mathbf{z}_{t}\sim D_{S^{*}}} \left[ \log\left(1 + \frac{1}{k}\mathbf{r}_{t}^{\top}\mathbf{1}_{S^{*}}\right) \right] - \underbrace{\mathbf{E}}_{\mathbf{z}_{t}\sim D_{S^{*}}} \left[ \log\left(1 + \frac{1}{k}\mathbf{r}_{t}^{\top}\mathbf{1}_{S}\right) \right] \geq \frac{2\epsilon}{k(2k)^{k}}. \quad (22)$$

Since any randomized algorithm is equivalent to an a priori random choice from the set of all deterministic strategies, and since the input defined above is oblivious to the output of the algorithm, it suffices to prove a lower bound on the expected regret of any *deterministic* algorithm (this is not crucial for the proof but simplifies the notation). We consider an arbitrary deterministic algorithm and let  $\{(S_t, \mathbf{x}_t)\}_{t=1}^T$  denote the output for the random input sequence  $\{\mathbf{r}_t\}_{t=1}^T$  given by  $\mathbf{r}_t = \mathbf{1} + \mathbf{z}_t$  and  $\mathbf{z}_t \sim D_{S^*}$ . Let  $N_S$  be a random variable denoting the number of  $t \in [T]$  such that  $S_t = S$ , i.e.,  $N_S = |\{t \in [T] \mid S_t = S\}|$ . From the equation (22), we have

$$\mathbf{E}_{\mathbf{z}_{t}\sim D_{S^{*}}}[R_{T}] \geq \sum_{t=1}^{T} \mathbf{E}_{\mathbf{z}_{t}\sim D_{S^{*}}} \left[ \log \left( 1 + \frac{1}{k} \mathbf{r}_{t}^{\top} \mathbf{1}_{S} \right) \right] - \sum_{t=1}^{T} \mathbf{E}_{S_{t}} \left[ \max_{\mathbf{x}\in\Delta^{S_{t}}} \mathbf{E}_{\mathbf{z}_{t}\sim D_{S^{*}}} \left[ \log \left( 1 + \mathbf{r}_{t}^{\top} \mathbf{x} \right) \right] \right] \\
\geq \left( T - \mathbf{E}_{\mathbf{z}_{t}\sim D_{S^{*}}}[N_{S^{*}}] \right) \frac{2\epsilon}{k(2k)^{k}}.$$
(23)

Let us evalute  $\underset{\mathbf{z}_t \sim D_{S^*}}{\mathbf{E}} [N_{S^*}]$ . Define  $D_0$  to be the uniform probabilistic distribution on  $\{-1, 1\}^d$ . Then, for all  $S \in \mathcal{S}_k \setminus \{\emptyset, S^*\}$ , we have  $D_{S^*}|_S = D_0|_S$ , i.e., if  $\mathbf{z} \sim D_{S^*}$  and  $\mathbf{z}' \sim D_0$ , then  $\mathbf{z}|_S$  and  $\mathbf{z}'|_S$  follows the same distribution (a uniform distribution on  $\{-1, 1\}^S$ ). Hence, in the same way as in Lemma A.1. of [5], we can show that

$$\mathbf{\underline{E}}_{\mathbf{z}_t \sim D_{S^*}}[N_{S^*}] - \mathbf{\underline{E}}_{\mathbf{z}_t \sim D_0}[N_{S^*}] \le \frac{T}{\sqrt{2}} \sqrt{\mathbf{\underline{E}}_{t \sim D_0}[N_{S^*}] \cdot \mathrm{KL}\left(D_0|_{S^*} \mid \mid D_{S^*}|_{S^*}\right)}$$
(24)

where  $\operatorname{KL}(P||Q) = \underset{P}{\operatorname{E}}(\log \frac{\mathrm{d}P}{\mathrm{d}Q})$  is the Kullback-Leibler divergence. The chain rule for relative entropy (see, e.g., Theorem 2.5.3 of [11]) gives, for  $S^* = \{i_1, \ldots i_k\}$ ,

$$\operatorname{KL}\left(D_{0}|_{S^{*}} \left| \left| D_{S^{*}}|_{S^{*}}\right) = \operatorname{KL}(\operatorname{Prob}_{\mathbf{z}\sim D_{0}}[(z_{i})_{i\in S^{*}}] \left| \operatorname{Prob}_{\mathbf{z}\sim D_{S^{*}}}[(z_{i})_{i\in S^{*}}]\right) \right|$$

$$= \sum_{j=1}^{k} \underbrace{\mathbf{E}}_{(z_{i_{s}})_{s

$$= \underbrace{\mathbf{E}}_{(z_{i_{s}})_{s

$$= -\frac{1}{2} \log(1 - 4\epsilon^{2}). \tag{25}$$$$$$

In the above equations, the third equality holds because  $\operatorname{Prob}_{\mathbf{z}\sim D_0}[z_{i_j} | (z_{i_s})_{s<j}]$  and  $\operatorname{Prob}_{\mathbf{z}\sim D_{S^*}}[z_{i_j} | (z_{i_s})_{s<j}]$  are equal to the Bernoulli distribution of parameter 1/2 for j < k. The last equality holds because  $\operatorname{Prob}_{\mathbf{z}\sim D_0}[z_{i_k} | (z_{i_s})_{s<k}]$  follows Bernoulli distribution of parameter 1/2 and  $\operatorname{Prob}_{\mathbf{z}\sim D_{S^*}}[z_{i_k} | (z_{i_s})_{s<k}]$  follows Bernoulli distribution of parameter  $1/2 + \epsilon$  or  $1/2 - \epsilon$ . Combining (23), (24) and (25), we have

$$\mathbf{E}_{\mathbf{z}_t \sim D_{S^*}}[R_T] \ge \left(T - \mathbf{E}_{\mathbf{z}_t \sim D_0}[N_{S^*}] - \frac{T}{2}\sqrt{-\mathbf{E}_{\mathbf{z}_t \sim D_0}[N_{S^*}]\log(1-\epsilon^2)}\right) \frac{2\epsilon}{k(2k)^k}.$$

Suppose that  $S^*$  is chosen at random uniformly from  $S_k$ , before play begins. Then, from the above inequality, the expected regret is bounded as

$$\mathbf{E}_{S^*, \mathbf{z}_t \sim D_{S^*}}[R_T] \ge \frac{1}{|\mathcal{S}_k|} \sum_{S^* \in \mathcal{S}_k} \left( T - \mathbf{E}_{\mathbf{z}_t \sim D_0}[N_{S^*}] - \frac{T}{2} \sqrt{-\frac{\mathbf{E}}{\mathbf{z}_t \sim D_0}[N_{S^*}] \log(1 - 4\epsilon^2)} \right) \frac{2\epsilon}{k(2k)^k}$$
$$\ge \left( T - \frac{T}{|\mathcal{S}_k|} - \frac{T}{2} \sqrt{-\frac{T}{|\mathcal{S}_k|} \log(1 - 4\epsilon^2)} \right) \frac{2\epsilon}{k(2k)^k},$$

where the second inequality comes from  $\sum_{S \in S_k} \mathbf{E}_{\mathbf{z}_t \sim D_0}[N_S] = T$  and  $\sum_{S \in S_k} \sqrt{\mathbf{E}_{\mathbf{z}_t \sim D_0}[N_S]} \leq \sqrt{T|S_k|}$ . Using the inequality  $-\log(1-x) \leq x/2$  for  $x \in [0, 1/4]$ , we have

$$\mathbf{E}_{S^*, \mathbf{z}_t \sim D_{S^*}}[R_T] \ge T\left(1 - \frac{1}{|\mathcal{S}_k|} - \epsilon \sqrt{\frac{T}{2|\mathcal{S}_k|}}\right) \frac{2\epsilon}{k(2k)^k}$$

for  $\epsilon \in [0, 1/4]$ . By setting  $\epsilon = \min\{1/4, \frac{1}{2}\sqrt{\frac{|S_k|}{T}}\}$ , we obtain

$$\mathbf{E}_{S^*, \mathbf{z}_t \sim D_{S^*}}[R_T] = \Omega\left(\min\left\{\frac{T}{k(2k)^k}, \sqrt{\frac{T|S_k|}{k^2(2k)^{2k}}}\right\}\right) \ge \Omega\left(\min\left\{\frac{T}{k(2k)^k}, \sqrt{T\left(\frac{d}{5k^3}\right)^k}\right\}\right),\tag{26}$$

where the second inequality follows from  $|\mathcal{S}_k| = {d \choose k} \ge {(\frac{d}{k})^k}$  and  $k^2 = O((\frac{5}{4})^k)$ .

Consider an arbitrary randomized algorithm and let  $\lambda$  denote the algorithm's internal randomization. Then, since  $\lambda$  is probabilistically independent from  $S^*$ , **r** and (26) for all deterministic algorithms, we have

$$\mathop{\mathbf{E}}_{S^*,\{\mathbf{r}\}}\left[\mathop{\mathbf{E}}_{\lambda}[R_T]\right] = \mathop{\mathbf{E}}_{\lambda}\left[\mathop{\mathbf{E}}_{S^*,\{\mathbf{r}\}}[R_T]\right] = \Omega\left(\min\left\{\frac{T}{k(2k)^k}, \sqrt{T\left(\frac{d}{5k^3}\right)^k}\right\}\right).$$