# Almost Optimal Algorithms for Linear Stochastic Bandits with Heavy-Tailed Payoffs

Han Shao\* Xiaotian Yu\* Irwin King Michael R. Lyu Department of Computer Science and Engineering The Chinese University of Hong Kong {hshao,xtyu,king,lyu}@cse.cuhk.edu.hk

## Abstract

In linear stochastic bandits, it is commonly assumed that payoffs are with sub-Gaussian noises. In this paper, under a weaker assumption on noises, we study the problem of linear stochastic bandits with heavy-tailed payoffs (LinBET), where the distributions have finite moments of order  $1 + \epsilon$ , for some  $\epsilon \in (0, 1]$ . We rigorously analyze the regret lower bound of LinBET as  $\Omega(T^{\frac{1}{1+\epsilon}})$ , implying that finite moments of order 2 (i.e., finite variances) yield the bound of  $\Omega(\sqrt{T})$ , with T being the total number of rounds to play bandits. The provided lower bound also indicates that the state-of-the-art algorithms for LinBET are far from optimal. By adopting median of means with a well-designed allocation of decisions and truncation based on historical information, we develop two novel bandit algorithms, where the regret upper bounds match the lower bound up to polylogarithmic factors. To the best of our knowledge, we are the first to solve LinBET optimally in the sense of the polynomial order on T. Our proposed algorithms are evaluated based on synthetic datasets, and outperform the state-of-the-art results.

# 1 Introduction

The decision-making model named Multi-Armed Bandits (MAB), where at each time step an algorithm chooses an arm among a given set of arms and then receives a stochastic payoff with respect to the chosen arm, elegantly characterizes the tradeoff between exploration and exploitation in sequential learning. The algorithm usually aims at maximizing cumulative payoffs over a sequence of rounds. A natural and important variant of MAB is linear stochastic bandits with the expected payoff of each arm satisfying a linear mapping from the arm information to a real number. The model of linear stochastic bandits enjoys some good theoretical properties, e.g., there exists a closed-form solution of the linear mapping at each time step in light of ridge regression. Many practical applications take advantage of MAB and its variants to control decision performance, e.g., online personalized recommendations (Li et al., 2010) and resource allocations (Lattimore et al., 2014).

In most previous studies of MAB and linear stochastic bandits, a common assumption is that noises in observed payoffs are sub-Gaussian conditional on historical information (Abbasi-Yadkori et al., 2011; Bubeck et al., 2012), which encompasses cases of all bounded payoffs and many unbounded payoffs, e.g., payoffs of an arm following a Gaussian distribution. However, there do exist practical scenarios of non-sub-Gaussian noises in observed payoffs for sequential decisions, such as high-probability extreme returns in investments for financial markets (Cont and Bouchaud, 2000) and fluctuations of neural oscillations (Roberts et al., 2015), which are called heavy-tailed noises. Thus, it is significant to completely study theoretical behaviours of sequential decisions in the case of heavy-tailed noises.

<sup>\*</sup>The first two authors contributed equally.

Many practical distributions, e.g., Pareto distributions and Weibull distributions, are heavy-tailed, which perform high tail probabilities compared with exponential family distributions. We consider a general characterization of heavy-tailed payoffs in bandits, where the distributions have finite moments of order  $1 + \epsilon$ , where  $\epsilon \in (0, 1]$ . When  $\epsilon = 1$ , stochastic payoffs are generated from distributions with finite variances. When  $0 < \epsilon < 1$ , stochastic payoffs are generated from distributions with infinite variances (Shao and Nikias, 1993). Note that, different from sub-Gaussian noises in the traditional bandit setting, noises from heavy-tailed distributions do not enjoy exponentially decaying tails, and thus make it more difficult to learn a parameter of an arm.

The regret of MAB with heavy-tailed payoffs has been well addressed by Bubeck et al. (2013), where stochastic payoffs have bounds on raw or central moments of order  $1 + \epsilon$ . For MAB with finite variances (i.e.,  $\epsilon = 1$ ), the regret of truncation algorithms or median of means recovers the optimal regret for MAB under the sub-Gaussian assumption. Recently, Medina and Yang (2016) investigated theoretical guarantees for the problem of <u>lin</u>ear stochastic <u>b</u>andits with <u>heavy-tailed</u> payoffs (LinBET). It is surprising to find that, for  $\epsilon = 1$ , the regret of bandit algorithms by Medina and Yang (2016) to solve LinBET is  $\tilde{O}(T^{\frac{3}{4}})^2$ , which is far away from the regret of the state-of-the-art algorithms (i.e.,  $\tilde{O}(\sqrt{T})$ ) in linear stochastic bandits under the sub-Gaussian assumption (Dani et al., 2008a; Abbasi-Yadkori et al., 2011). Thus, the most interesting and non-trivial question is

## Is it possible to recover the regret of $\widetilde{O}(\sqrt{T})$ when $\epsilon = 1$ for LinBET?

In this paper, we answer this question affirmatively. Specifically, we investigate the problem of LinBET characterized by finite  $(1 + \epsilon)$ -th moments, where  $\epsilon \in (0, 1]$ . The problem of LinBET raises several interesting challenges. The first challenge is the lower bound of the problem, which remains unknown. The technical issues come from the construction of an elegant setting for LinBET, and the derivation of a lower bound with respect to  $\epsilon$ . The second challenge is how to develop a robust estimator for the parameter in LinBET, because heavy-tailed noises greatly affect errors of the conventional least-squares estimator. It is worth mentioning that Medina and Yang (2016) has tried to tackle this challenge, but their estimators do not make full use of the contextual information of chosen arms to eliminate the effect from heavy-tailed noises, which eventually leads to large regrets. The third challenge is how to successfully adopt median of means and truncation to solve LinBET with regret upper bounds matching the lower bound as closely as possible.

**Our Results.** First of all, we rigorously analyze the lower bound on the problem of LinBET, which enjoys a polynomial order on T as  $\Omega(T^{\frac{1}{1+\epsilon}})$ . The lower bound provides two essential hints: one is that finite variances in LinBET yield a bound of  $\Omega(\sqrt{T})$ , and the other is that algorithms by Medina and Yang (2016) are sub-optimal. Then, we develop two novel bandit algorithms to solve LinBET based on the basic techniques of median of means and truncation. Both the algorithms adopt the optimism in the face of uncertainty principle, which is common in bandit problems (Abbasi-Yadkori et al., 2011; Munos et al., 2014). The regret upper bounds of the proposed two algorithms, which are  $\tilde{O}(T^{\frac{1}{1+\epsilon}})$ , match the lower bound up to polylogarithmic factors. To the best of our knowledge, we are the first to solve LinBET almost optimally. We conduct experiments based on synthetic datasets, which are generated by Student's *t*-distribution and Pareto distribution, to demonstrate the effectiveness of our algorithms. Experimental results show that our algorithms outperform the state-of-the-art results. The contributions of this paper are summarized as follows:

- We provide the lower bound for the problem of LinBET characterized by finite  $(1 + \epsilon)$ -th moments, where  $\epsilon \in (0, 1]$ . In the analysis, we construct an elegant setting of LinBET, which results in a regret bound of  $\Omega(T^{\frac{1}{1+\epsilon}})$  in expectation for any bandit algorithm.
- We develop two novel bandit algorithms, which are named as MENU and TOFU (with details shown in Section 4). The MENU algorithm adopts median of means with a well-designed allocation of decisions and the TOFU algorithm adopts truncation via historical information. Both algorithms achieve the regret  $\tilde{O}(T^{\frac{1}{1+\epsilon}})$  with high probability.
- We conduct experiments based on synthetic datasets to demonstrate the effectiveness of our proposed algorithms. By comparing our algorithms with the state-of-the-art results, we show improvements on cumulative payoffs for MENU and TOFU, which are strictly consistent with theoretical guarantees in this paper.

<sup>&</sup>lt;sup>2</sup>We omit polylogarithmic factors of T for  $\widetilde{O}(\cdot)$ .

## **2** Preliminaries and Related Work

In this section, we first present preliminaries, i.e., notations and learning setting of LinBET. Then, we give a detailed discussion on the line of research for bandits with heavy-tailed payoffs.

#### 2.1 Notations

For a positive integer K,  $[K] \triangleq \{1, 2, \dots, K\}$ . Let the  $\ell$ -norm of a vector  $x \in \mathbb{R}^d$  be  $||x||_{\ell} \triangleq (x_1^{\ell} + \dots + x_d^{\ell})^{\frac{1}{\ell}}$ , where  $\ell \ge 1$  and  $x_i$  is the *i*-th element of x with  $i \in [d]$ . For  $r \in \mathbb{R}$ , its absolute value is |r|, its ceiling integer is  $\lceil r \rceil$ , and its floor integer is  $\lfloor r \rfloor$ . The inner product of two vectors x, y is denoted by  $x^{\top}y = \langle x, y \rangle$ . Given a positive definite matrix  $A \in \mathbb{R}^{d \times d}$ , the weighted Euclidean norm of a vector  $x \in \mathbb{R}^d$  is  $||x||_A = \sqrt{x^{\top}Ax}$ .  $\mathbb{B}(x, r)$  denotes a Euclidean ball centered at x with radius  $r \in \mathbb{R}_+$ , where  $\mathbb{R}_+$  is the set of positive numbers. Let e be Euler's number, and  $I_d \in \mathbb{R}^{d \times d}$  an identity matrix. Let  $\mathbb{1}_{\{\cdot\}}$  be an indicator function, and  $\mathbb{E}[X]$  the expectation of X.

#### 2.2 Learning Setting

For a bandit algorithm  $\mathcal{A}$ , we consider sequential decisions with the goal to maximize cumulative payoffs, where the total number of rounds for playing bandits is T. For each round  $t = 1, \dots, T$ , the bandit algorithm  $\mathcal{A}$  is given a decision set  $D_t \subseteq \mathbb{R}^d$  such that  $||x||_2 \leq D$  for any  $x \in D_t$ .  $\mathcal{A}$  has to choose an arm  $x_t \in D_t$  and then observes a stochastic payoff  $y_t(x_t)$ . For notation simplicity, we also write  $y_t = y_t(x_t)$ . The expectation of the observed payoff for the chosen arm satisfies a linear mapping from the arm to a real number as  $y_t(x_t) \triangleq \langle x_t, \theta_* \rangle + \eta_t$ , where  $\theta_*$  is an underlying parameter with  $||\theta_*||_2 \leq S$  and  $\eta_t$  is a random noise. Without loss of generality, we assume  $\mathbb{E} [\eta_t | \mathcal{F}_{t-1}] = 0$ , where  $\mathcal{F}_{t-1} \triangleq \{x_1, \dots, x_t\} \cup \{\eta_1, \dots, \eta_{t-1}\}$  is a  $\sigma$ -filtration and  $\mathcal{F}_0 = \emptyset$ . Clearly, we have  $\mathbb{E}[y_t(x_t)|\mathcal{F}_{t-1}] = \langle x_t, \theta_* \rangle$ . For an algorithm  $\mathcal{A}$ , to maximize cumulative payoffs is equivalent to minimizing the regret as

$$R(\mathcal{A},T) \triangleq \left(\sum_{t=1}^{T} \langle x_t^*, \theta_* \rangle\right) - \left(\sum_{t=1}^{T} \langle x_t, \theta_* \rangle\right) = \sum_{t=1}^{T} \langle x_t^* - x_t, \theta_* \rangle, \tag{1}$$

where  $x_t^*$  denotes the optimal decision at time t for  $\theta_*$ , i.e.,  $x_t^* \in \arg \max_{x \in D_t} \langle x, \theta_* \rangle$ . In this paper, we will provide high-probability upper bound of  $R(\mathcal{A}, T)$  with respect to  $\mathcal{A}$ , and provide the lower bound for LinBET in expectation for any algorithm. The problem of LinBET is defined as below.

**Definition 1** (LinBET). Given a decision set  $D_t$  for time step  $t = 1, \dots, T$ , an algorithm  $\mathcal{A}$ , of which the goal is to maximize cumulative payoffs over T rounds, chooses an arm  $x_t \in D_t$ . With  $\mathcal{F}_{t-1}$ , the observed stochastic payoff  $y_t(x_t)$  is conditionally heavy-tailed, i.e.,  $\mathbb{E}\left[|y_t|^{1+\epsilon}|\mathcal{F}_{t-1}\right] \leq b$  or  $\mathbb{E}\left[|y_t - \langle x_t, \theta_* \rangle|^{1+\epsilon}|\mathcal{F}_{t-1}\right] \leq c$ , where  $\epsilon \in (0, 1]$ , and  $b, c \in (0, +\infty)$ .

#### 2.3 Related Work

The model of MAB dates back to 1952 with the original work by Robbins et al. (1952), and its inherent characteristic is the trade-off between exploration and exploitation. The asymptotic lower bound of MAB was developed by Lai and Robbins (1985), which is logarithmic with respect to the total number of rounds. An important technique called upper confidence bound was developed to achieve the lower bound (Lai and Robbins, 1985; Agrawal, 1995). Other related techniques to solve the problem of sequential decisions include Thompson sampling (Thompson, 1933; Chapelle and Li, 2011; Agrawal and Goyal, 2012) and Gittins index (Gittins et al., 2011).

The problem of MAB with heavy-tailed payoffs characterized by finite  $(1 + \epsilon)$ -th moments has been well investigated (Bubeck et al., 2013; Vakili et al., 2013; Yu et al., 2018). Bubeck et al. (2013) pointed out that finite variances in MAB are sufficient to achieve regret bounds of the same order as the optimal regret for MAB under the sub-Gaussian assumption, and the order of T in regret bounds increases when  $\epsilon$  decreases. The lower bound of MAB with heavy-tailed payoffs has been analyzed (Bubeck et al., 2013), and robust algorithms by Bubeck et al. (2013) are optimal. Theoretical guarantees by Bubeck et al. (2013); Vakili et al. (2013) are for the setting of finite arms. In Vakili et al. (2013), primary theoretical results were presented for the case of  $\epsilon > 1$ . We notice that the case of  $\epsilon > 1$  is not interesting, because it reduces to the case of finite variances in MAB. For the problem of linear stochastic bandits, which is also named linear reinforcement learning by Auer (2002), the lower bound is  $\Omega(d\sqrt{T})$  when contextual information of arms is from a *d*dimensional space (Dani et al., 2008b). Bandit algorithms matching the lower bound up to polylogarithmic factors have been well developed (Auer, 2002; Dani et al., 2008a; Abbasi-Yadkori et al., 2011; Chu et al., 2011). Notice that all these studies assume that stochastic payoffs contain sub-Gaussian noises. More variants of MAB can be discussed by Bubeck et al. (2012).

It is surprising to find that the lower bound of LinBET remains unknown. In Medina and Yang (2016), bandit algorithms based on truncation and median of means were presented. When  $\epsilon$  is finite for LinBET, the algorithms by Medina and Yang (2016) cannot recover the bound of  $\tilde{O}(\sqrt{T})$  which is the regret of the state-of-the-art algorithms in linear stochastic bandits under the sub-Gaussian assumption. Medina and Yang (2016) conjectured that it is possible to recover  $\tilde{O}(\sqrt{T})$  with  $\epsilon$  being a finite number for LinBET. Thus, it is urgent to conduct a thorough analysis of the conjecture in consideration of the importance of heavy-tailed noises in real scenarios. Solving the conjecture generalizes the practical applications of bandit models. Practical motivating examples for bandits with heavy-tailed payoffs include delays in end-to-end network routing (Liebeherr et al., 2012) and sequential investments in financial markets (Cont and Bouchaud, 2000).

Recently, the assumption in stochastic payoffs of MAB was relaxed from sub-Gaussian noises to bounded kurtosis (Lattimore, 2017), which can be viewed as an extension of Bubeck et al. (2013). The interesting point of Lattimore (2017) is the scale free algorithm, which might be practical in applications. Besides, Carpentier and Valko (2014) investigated extreme bandits, where stochastic payoffs of MAB follow Fréchet distributions. The setting of extreme bandits fits for the real scenario of anomaly detection without contextual information. The order of regret in extreme bandits is characterized by distributional parameters, which is similar to the results by Bubeck et al. (2013).

It is worth mentioning that, for linear regression with heavy-tailed noises, several interesting studies have been conducted. Hsu and Sabato (2016) proposed a generalized method in light of median of means for loss minimization with heavy-tailed noises. Heavy-tailed noises in Hsu and Sabato (2016) might come from contextual information, which is more complicated than the setting of stochastic payoffs in this paper. Therefore, linear regression with heavy-tailed noises usually requires a finite fourth moment. In Audibert et al. (2011), the basic technique of truncation was adopted to solve robust linear regression in the absence of exponential moment condition. The related studies in this line of research are not directly applicable for the problem of LinBET.

# 3 Lower Bound

In this section, we provide the lower bound for LinBET. We consider heavy-tailed payoffs with finite  $(1 + \epsilon)$ -th raw moments in the analysis. In particular, we construct the following setting. Assume  $d \ge 2$  is even (when d is odd, similar results can be easily derived by considering the first d - 1 dimensions). For  $D_t \subseteq \mathbb{R}^d$  with  $t \in [T]$ , we fix the decision set as  $D_1 = \cdots = D_T = D_{(d)}$ . Then, the fixed decision set is constructed as  $D_{(d)} \triangleq \{(x_1, \cdots, x_d) \in \mathbb{R}^d_+ : x_1 + x_2 = \cdots = x_{d-1} + x_d = 1\}$ , which is a subset of intersection of the cube  $[0, 1]^d$  and the hyperplane  $x_1 + \cdots + x_d = d/2$ . We define a set  $S_d \triangleq \{(\theta_1, \cdots, \theta_d) : \forall i \in [d/2], (\theta_{2i-1}, \theta_{2i}) \in \{(2\Delta, \Delta), (\Delta, 2\Delta)\}\}$  with  $\Delta \in (0, 1/d]$ . The payoff functions take values in  $\{0, (1/\Delta)^{\frac{1}{\epsilon}}\}$  such that, for every  $x \in D_{(d)}$ , the expected payoff is  $\theta_*^{\top} x$ . To be more specific, we have the payoff function of x as

$$y(x) = \begin{cases} \left(\frac{1}{\Delta}\right)^{\frac{1}{\epsilon}} & \text{with a probability of } \Delta^{\frac{1}{\epsilon}} \theta_*^\top x, \\ 0 & \text{with a probability of } 1 - \Delta^{\frac{1}{\epsilon}} \theta_*^\top x. \end{cases}$$
(2)

We have the theorem for the lower bound of LinBET as below.

**Theorem 1** (Lower Bound of LinBET). If  $\theta_*$  is chosen uniformly at random from  $S_d$ , and the payoff for each  $x \in D_{(d)}$  is in  $\{0, (1/\Delta)^{\frac{1}{e}}\}$  with mean  $\theta_*^{\top} x$ , then for any algorithm  $\mathcal{A}$  and every  $T \ge (d/12)^{\frac{e}{1+\epsilon}}$ , we have

$$\mathbb{E}\left[R(\mathcal{A},T)\right] \ge \frac{d}{192} T^{\frac{1}{1+\epsilon}}.$$
(3)

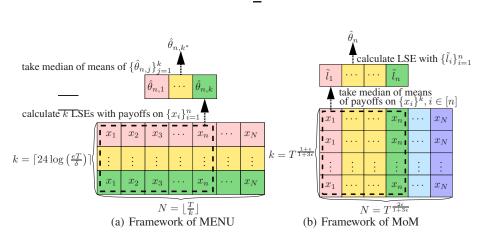


Figure 1: Framework comparison between our MENU and MoM by Medina and Yang (2016).

In the proof of Theorem 1, we first prove the lower bound when d = 2, and then generalize the argument to any d > 2. We notice that the parameter in the original *d*-dimensional space is rearranged to d/2 tuples, each of which is a 2-dimensional vector as  $(\theta_{2i-1}, \theta_{2i}) \in \{(2\Delta, \Delta), (\Delta, 2\Delta)\}$  with  $i \in [d/2]$ . If the *i*-th tuple of the parameter is selected as  $(2\Delta, \Delta)$ , then the *i*-th tuple of the optimal arm is  $(x_{*,2i-1}, x_{*,2i}) = (1,0)$ . In this case, if we define the *i*-th tuple of the chosen arm as  $(x_{t,2i-1}, x_{t,2i})$ , the instantaneous regret is  $\Delta(1 - x_{t,2i-1})$ . Then, the regret can be represented as an integration of  $\Delta(1 - x_{t,2i-1})$  over  $D_{(d)}$ . Finally, with common inequalities in information theory, we obtain the regret lower bound by setting  $\Delta = T^{-\frac{\epsilon}{1+\epsilon}}/12$ .

We notice that martingale differences to prove the lower bound for linear stochastic bandits in (Dani et al., 2008a) are not directly feasible for the proof of lower bound in LinBET, because under our construction of heavy-tailed payoffs (i.e., Eq. (4)), the information of  $\epsilon$  is excluded. Besides, our proof is partially inspired by Bubeck (2010). We show the detailed proof of Theorem 1 in Appendix A.

**Remark 1.** The above lower bound provides two essential hints: one is that finite variances in LinBET yield a bound of  $\Omega(\sqrt{T})$ , and the other is that algorithms proposed by Medina and Yang (2016) are far from optimal. The result in Theorem 1 strongly indicates that it is possible to design bandit algorithms recovering  $\tilde{O}(\sqrt{T})$  with finite variances.

# 4 Algorithms and Upper Bounds

In this section, we develop two novel bandit algorithms to solve LinBET, which turns out to be almost optimal. We rigorously prove regret upper bounds for the proposed algorithms. In particular, our core idea is based on the <u>optimism</u> in the <u>face</u> of <u>uncertainty</u> principle (OFU). The first algorithm is <u>me</u>dian of means under OF<u>U</u> (MENU) shown in Algorithm 1, and the second algorithm is <u>truncation</u> under <u>OFU</u> (TOFU) shown in Algorithm 2. For comparisons, we directly name the bandit algorithm based on median of means in Medina and Yang (2016) as MoM, and name the bandit algorithm based on confidence region with truncation in Medina and Yang (2016) as CRT.

Both algorithms in this paper adopt the tool of ridge regression. At time step t, let  $\hat{\theta}_t$  be the  $\ell^2$ -regularized least-squares estimate (LSE) of  $\theta_*$  as  $\hat{\theta}_t = V_t^{-1}X_t^{\top}Y_t$ , where  $X_t \in \mathbb{R}^{t \times d}$  is a matrix of which rows are  $x_1^{\top}, \dots, x_t^{\top}, V_t \triangleq X_t^{\top}X_t + \lambda I_d, Y_t \triangleq (y_1, \dots, y_t)$  is a vector of the historical observed payoffs until time t and  $\lambda > 0$  is a regularization parameter.

## 4.1 MENU and Regret

**Description of MENU.** To conduct median of means in LinBET, it is common to allocate T pulls of bandits among  $N \leq T$  epochs, and for each epoch the same arm is played multiple times to obtain an estimate of  $\theta_*$ . We find that there exist different ways to construct the epochs. We design the framework of MENU in Figure 1(a), and show the framework of MoM designed

#### Algorithm 1 Median of means under OFU (MENU)

1: input  $d, c, \epsilon, \delta, \lambda, S, T, \{D_n\}_{n=1}^N$ 2: initialization:  $k = \lfloor 24 \log \left(\frac{eT}{\delta}\right) \rfloor$ ,  $N = \lfloor \frac{T}{k} \rfloor$ ,  $V_0 = \lambda I_d$ ,  $C_0 = \mathbb{B}(\mathbf{0}, S)$ 3: for  $n = 1, 2, \dots, N$  do  $(x_n, \tilde{\theta}_n) = \arg \max_{(x,\theta) \in D_n \times C_{n-1}} \langle x, \theta \rangle$ Play  $x_n$  with k times and observe payoffs  $y_{n,1}, y_{n,2}, \cdots, y_{n,k}$ 4:  $\triangleright$  to select an arm 5: 6:  $V_n = V_{n-1} + x_n x_n^{\dagger}$ For  $j \in [k]$ ,  $\hat{\theta}_{n,j} = V_n^{-1} \sum_{i=1}^n y_{i,j} x_i$ For  $j \in [k]$ , let  $r_j$  be the median of  $\{\|\hat{\theta}_{n,j} - \hat{\theta}_{n,s}\|_{V_n} : s \in [k] \setminus j\}$ 7:  $\triangleright$  to calculate LSE for the *j*-th group 8:  $k^* = \arg\min_{j \in [k]} r_j$  $\beta_n = 3\left((9dc)^{\frac{1}{1+\epsilon}} n^{\frac{1-\epsilon}{2(1+\epsilon)}} + \lambda^{\frac{1}{2}}S\right)$ 9: ▷ to take median of means of estimates 10:  $C_n = \{ \theta : \|\theta - \hat{\theta}_{n,k^*}\|_{V_n} \le \beta_n \}$ 11: ▷ to update the confidence region 12: end for

by Medina and Yang (2016) in Figure 1(b). For MENU and MoM, we have the following three differences. First, for each epoch  $n = 1, \dots, N$ , MENU plays the same arm  $x_n$  by  $O(\log(T))$  times, while MoM plays the same arm by  $O(T^{\frac{1+\epsilon}{1+3\epsilon}})$  times. Second, at epoch n with historical payoffs, MENU conducts LSEs by  $O(\log(T))$  times, each of which is based on  $\{x_i\}_{i=1}^n$ , while MoM conducts LSE by one time based on intermediate payoffs calculated via median of means of observed payoffs. Third, MENU adopts median of means of LSEs, while MoM adopts median of means of the observed payoffs. Intuitively, the execution of multiple LSEs will lead to the improved regret of MENU. With a better trade-off between k and N in Figure 1(a), we derive an improved upper bound of regret in Theorem 2.

In light of Figure 1(a), we develop algorithmic procedures in Algorithm 1 for MENU. We notice that, in order to guarantee the median of means of LSEs not far away from the true underlying parameter with high probability, we construct the confidence interval in Line 10 of Algorithm 1. Now we have the following theorem for the regret upper bound of MENU.

**Theorem 2** (Regret Analysis for the MENU Algorithm). Assume that for all t and  $x_t \in D_t$  with  $||x_t||_2 \leq D$ ,  $||\theta_*||_2 \leq S$ ,  $|x_t^\top \theta_*| \leq L$  and  $\mathbb{E}[|\eta_t|^{1+\epsilon} |\mathcal{F}_{t-1}] \leq c$ . Then, with probability at least  $1 - \delta$ , for every  $T \geq 256 + 24 \log (e/\delta)$ , the regret of the MENU algorithm satisfies

$$R(\textit{MENU},T) \le 6\left((9dc)^{\frac{1}{1+\epsilon}} + \lambda^{\frac{1}{2}}S + L\right)T^{\frac{1}{1+\epsilon}} \left(24\log\left(\frac{eT}{\delta}\right) + 1\right)^{\frac{\epsilon}{1+\epsilon}} \sqrt{2d\log\left(1 + \frac{TD^2}{\lambda d}\right)}$$

The technical challenges in MENU (i.e., Algorithm 1) and its proofs are discussed as follows. Based on the common techniques in linear stochastic bandits (Abbasi-Yadkori et al., 2011), to guarantee the instantaneous regret in LinBET, we need to guarantee  $\|\theta_* - \hat{\theta}_{n,k^*}\|_{V_n} \leq \beta_n$  with high probability. We attack this issue by guaranteeing  $\|\theta_* - \hat{\theta}_{n,j}\|_{V_n} \leq \beta_n/3$  with a probability of 3/4, which could reduce to a problem of bounding a weighted sum of historical noises. Interestingly, by conducting singular value decomposition on  $X_n$  (of which rows are  $x_1^\top, \dots, x_n^\top$ ), we find that 2-norm of the weights is no greater than 1. Then the weighted sum can be bounded by a term as  $O\left(n^{\frac{1-\epsilon}{2(1+\epsilon)}}\right)$ . With a standard analysis in linear stochastic bandits from the instantaneous regret to the regret, we achieve the above results for MENU. We show the detailed proof of Theorem 2 in Appendix B.

**Remark 2.** For MENU, we adopt the assumption of heavy-tailed payoffs on central moments, which is required in the basic technique of median of means (Bubeck et al., 2013). Besides, there exists an implicit mild assumption in Algorithm 1 that, at each epoch n, the decision set must contain the selected arm  $x_n$  at least k times, which is practical in applications, e.g., online personalized recommendations (Li et al., 2010). The condition of  $T \ge 256 + 24 \log (e/\delta)$  is required for  $T \ge k$ . The regret upper bound of MENU is  $\tilde{O}(T^{\frac{1}{1+\epsilon}})$ , which implies that finite variances in LinBET are sufficient to achieve  $\tilde{O}(\sqrt{T})$ .

#### 4.2 TOFU and Regret

**Description of TOFU.** We demonstrate the algorithmic procedures of TOFU in Algorithm 2. We point out two subtle differences between our TOFU and the algorithm of CRT as follows. In TOFU,

Algorithm 2 Truncation under OFU (TOFU)

1: input  $d, b, \epsilon, \delta, \lambda, T, \{D_t\}_{t=1}^T$ 2: initialization:  $V_0 = \lambda I_d, C_0 = \mathbb{B}(\mathbf{0}, S)$ 3: for  $t = 1, 2, \cdots, T$  do  $b_t = \left(\frac{b}{\log\left(\frac{2T}{\delta}\right)}\right)^{\frac{1}{\epsilon}} t^{\frac{1-\epsilon}{2(1+\epsilon)}}$ 4:  $(x_t, \tilde{\theta}_t) = \arg \max_{(x,\theta) \in D_t \times C_{t-1}} \langle x, \theta \rangle$ 5: ▷ to select an arm Play  $x_t$  and observe a payoff  $y_t$ 6: 7:  $V_t = V_{t-1} + x_t x_t^{\top} \text{ and } X_t^{\top} = [x_1, \cdots, x_t]$  $\begin{bmatrix} u_1, \cdots, u_d \end{bmatrix}^\top = V_t^{-1/2} X_t^\top$  for  $i = 1, \cdots, d$  do 8: 9: 10:  $Y_i^{\dagger} = (y_1 \mathbb{1}_{u_{i,1}y_1 \le b_t}, \cdots, y_t \mathbb{1}_{u_{i,t}y_t \le b_t})$ ▷ to truncate the payoffs end for  $\theta_t^{\dagger} = V_t^{-1/2} (u_1^{\top} Y_1^{\dagger}, \cdots, u_d^{\top} Y_d^{\dagger})$   $\beta_t = 4\sqrt{d} b^{\frac{1}{1+\epsilon}} \left( \log\left(\frac{2dT}{\delta}\right) \right)^{\frac{\epsilon}{1+\epsilon}} t^{\frac{1-\epsilon}{2(1+\epsilon)}} + \lambda^{\frac{1}{2}} S$ 11: 12: 13: Update  $C_t = \{\theta : \|\theta - \theta_t^{\dagger}\|_{V_t} \leq \beta_t\}$ 14: ▷ to update the confidence region 15: end for

Table 1: Statistics of synthetic datasets in experiments. For Student's *t*-distribution,  $\nu$  denotes the degree of freedom,  $l_p$  denotes the location,  $s_p$  denotes the scale. For Pareto distribution,  $\alpha$  denotes the shape and  $s_m$  denotes the scale. NA denotes not available.

dataset	$D_t$ {#arms,#dimensions}	distribution {parameters}	$\{\epsilon, b, c\}$	mean of the optimal arm
S1	{20,10}	Student's <i>t</i> -distribution $\{\nu = 3, l_p = 0, s_p = 1\}$	{1.00, NA, 3.00}	4.00
S2	{100,20}	Student's <i>t</i> -distribution $\{\nu = 3, l_p = 0, s_p = 1\}$	{1.00, NA, 3.00}	7.40
S3	{20,10}	Pareto distribution $\{\alpha = 2, s_m = \frac{x_t^\top \theta_*}{2}\}$	{0.50, 7.72, NA}	3.10
S4	{100,20}	Pareto distribution $\{\alpha = 2, s_m = \frac{x_t^\top \theta_*}{2}\}$	{0.50, 54.37, NA}	11.39

to obtain the accurate estimate of  $\theta_*$ , we need to trim all historical payoffs for each dimension individually. Besides, the truncating operations depend on the historical information of arms. By contrast, in CRT, the historical payoffs are trimmed once, which is controlled only by the number of rounds for playing bandits. Compared to CRT, our TOFU achieves a tighter confidence interval, which can be found from the setting of  $\beta_t$ . Now we have the following theorem for the regret upper bound of TOFU.

**Theorem 3** (Regret Analysis for the TOFU Algorithm). Assume that for all t and  $x_t \in D_t$  with  $||x_t||_2 \leq D$ ,  $||\theta_*||_2 \leq S$ ,  $|x_t^\top \theta_*| \leq L$  and  $\mathbb{E}[|y_t|^{1+\epsilon} | \mathcal{F}_{t-1}] \leq b$ . Then, with probability at least  $1 - \delta$ , for every  $T \geq 1$ , the regret of the TOFU algorithm satisfies

$$R(TOFU,T) \le 2T^{\frac{1}{1+\epsilon}} \left( 4\sqrt{d}b^{\frac{1}{1+\epsilon}} \left( \log\left(\frac{2dT}{\delta}\right) \right)^{\frac{\epsilon}{1+\epsilon}} + \lambda^{\frac{1}{2}}S + L \right) \sqrt{2d\log\left(1 + \frac{TD^2}{\lambda d}\right)}$$

Similarly to the proof in Theorem 2, we can achieve the above results for TOFU. Due to space limitation, we show the detailed proof of Theorem 3 in Appendix C.

**Remark 3.** For TOFU, we adopt the assumption of heavy-tailed payoffs on raw moments. It is worth pointing out that, when  $\epsilon = 1$ , we have regret upper bound for TOFU as  $\tilde{O}(d\sqrt{T})$ , which implies that we recover the same order of d as that under sub-Gaussian assumption (Abbasi-Yadkori et al., 2011). A weakness in TOFU is high time complexity, because for each round TOFU needs to truncate all historical payoffs. The time complexity might be reasonably reduced by dividing T into multiple epochs, each of which contains only one truncation.

# **5** Experiments

In this section, we conduct experiments based on synthetic datasets to evaluate the performance of our proposed bandit algorithms: MENU and TOFU. For comparisons, we adopt two baselines: MoM and CRT proposed by Medina and Yang (2016). We run multiple independent repetitions for each dataset in a personal computer under Windows 7 with Intel CPU@3.70GHz and 16GB memory.

## 5.1 Datasets and Setting

To show effectiveness of bandit algorithms, we will demonstrate cumulative payoffs with respect to number of rounds for playing bandits over a fixed finite-arm decision set. For verifications, we adopt four synthetic datasets (named as S1–S4) in the experiments, of which statistics are shown in Table 1. The experiments on heavy tails require  $\epsilon$ , *b* or  $\epsilon$ , *c* to be known, which corresponds to the assumptions of Theorem 2 or Theorem 3. According to the required information, we can apply MENU or TOFU into practical applications. We adopt Student's *t* and Pareto distributions because they are common in practice. For Student's *t*-distributions, we easily estimate *c*, while for Pareto distributions, we easily estimate *b*. Besides, we can choose different parameters (e.g., larger values) in the distributions, and recalculate the parameters of *b* and *c*.

For S1 and S2, which contain different numbers of arms and different dimensions for the contextual information, we adopt standard Student's *t*-distribution to generate heavy-tailed noises. For the chosen arm  $x_t \in D_t$ , the expected payoff is  $x_t^{\top} \theta_*$ , and the observed payoff is added a noise generated from a standard Student's *t*-distribution. We generate each dimension of contextual information for an arm, as well as the underlying parameter, from a uniform distribution over [0, 1]. The standard Student's *t*-distribution implies that the bound for the second central moment of S1 and S2 is 3.

For S3 and S4, we adopt Pareto distribution, where the shape parameter is set as  $\alpha = 2$ . We know  $x_t^{\top} \theta_* = \alpha s_m / (\alpha - 1)$  implying  $s_m = x_t^{\top} \theta_* / 2$ . Then, we set  $\epsilon = 0.5$  leading to the bound of raw moment as  $\mathbb{E}\left[|y_t|^{1.5}\right] = \alpha s_m^{1.5} / (\alpha - 1.5) = 4s_m^{1.5}$ . We take the maximum of  $4s_m^{1.5}$  among all arms as the bound of the 1.5-th raw moment. We generate arms and the parameter similar to S1 and S2.

In figures, we show the average of cumulative payoffs with time evolution over ten independent repetitions for each dataset, and show error bars of a standard variance for comparing the robustness of algorithms. For S1 and S2, we run MENU and MoM and set  $T = 2 \times 10^4$ . For S3 and S4, we run TOFU and CRT and set  $T = 1 \times 10^4$ . For all algorithms, we set  $\lambda = 1.0$ , and  $\delta = 0.1$ .

## 5.2 Results and Discussions

We show experimental results in Figure 2. From the figure, we clearly find that our proposed two algorithms outperform MoM and CRT, which is consistent with the theoretical results in Theorems 2 and 3. We also evaluate our algorithms with other synthetic datasets, as well as different  $\lambda$  and  $\delta$ , and observe similar superiority of MENU and TOFU. Finally, for further comparison on regret, complexity and storage of four algorithms, we list the results shown in Table 2.

algorithm	MoM	MENU	CRT	TOFU
regret	$\widetilde{O}(T^{\frac{1+2\epsilon}{1+3\epsilon}})$	$\widetilde{O}(T^{\frac{1}{1+\epsilon}})$	$\widetilde{O}(T^{\frac{1}{2} + \frac{1}{2(1+\epsilon)}})$	$\widetilde{O}(T^{\frac{1}{1+\epsilon}})$
complexity	O(T)	$O(T\log T)$	O(T)	$O(T^2)$
storage	O(1)	$O(\log T)$	O(1)	O(T)

Table 2: Comparison on regret, complexity and storage of four algorithms.

# 6 Conclusion

We have studied the problem of LinBET, where stochastic payoffs are characterized by finite  $(1 + \epsilon)$ -th moments with  $\epsilon \in (0, 1]$ . We broke the traditional assumption of sub-Gaussian noises in payoffs of bandits, and derived theoretical guarantees based on the prior information of bounds on finite moments. We rigorously analyzed the lower bound of LinBET, and developed two novel

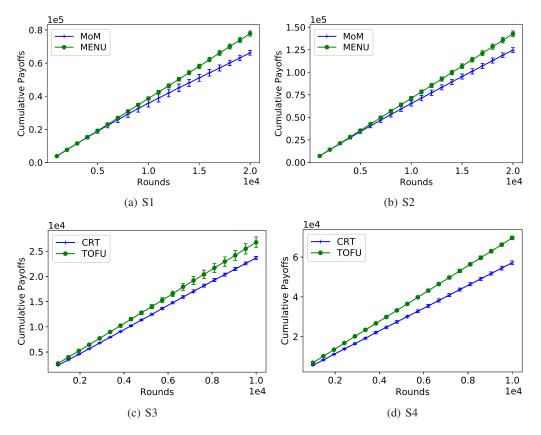


Figure 2: Comparison of cumulative payoffs for synthetic datasets S1-S4 with four algorithms.

bandit algorithms with regret upper bounds matching the lower bound up to polylogarithmic factors. Two novel algorithms were developed based on median of means and truncation. In the sense of polynomial dependence on T, we provided optimal algorithms for the problem of LinBET, and thus solved an open problem, which has been pointed out by Medina and Yang (2016). Finally, our proposed algorithms have been evaluated based on synthetic datasets, and outperformed the state-of-the-art results. Since both algorithms in this paper require a priori knowledge of  $\epsilon$ , future directions in this line of research include automatic learning of LinBET without information of distributional moments, and evaluation of our proposed algorithms in real-world scenarios.

## Acknowledgments

The work described in this paper was partially supported by the Research Grants Council of the Hong Kong Special Administrative Region, China (No. CUHK 14208815 and No. CUHK 14210717 of the General Research Fund), and Microsoft Research Asia (2018 Microsoft Research Asia Collaborative Research Award).

## References

- 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.
- R. Agrawal. Sample mean based index policies by  $O(\log n)$  regret for the multi-armed bandit problem. Advances in Applied Probability, 27(4):1054–1078, 1995.
- S. Agrawal and N. Goyal. Analysis of Thompson sampling for the multi-armed bandit problem. In *Conference on Learning Theory*, pages 39–1, 2012.
- J.-Y. Audibert, O. Catoni, et al. Robust linear least squares regression. *The Annals of Statistics*, 39 (5):2766–2794, 2011.
- P. Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learn*ing Research, 3(Nov):397–422, 2002.

- S. Bubeck. *Bandits games and clustering foundations*. PhD thesis, Université des Sciences et Technologie de Lille-Lille I, 2010.
- 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.
- S. Bubeck, N. Cesa-Bianchi, and G. Lugosi. Bandits with heavy tail. *IEEE Transactions on Infor*mation Theory, 59(11):7711–7717, 2013.
- A. Carpentier and M. Valko. Extreme bandits. In *Advances in Neural Information Processing Systems*, pages 1089–1097, 2014.
- O. Chapelle and L. Li. An empirical evaluation of Thompson sampling. In Advances in Neural Information Processing Systems, pages 2249–2257, 2011.
- W. Chu, L. Li, L. Reyzin, and R. Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 208–214, 2011.
- R. Cont and J.-P. Bouchaud. Herd behavior and aggregate fluctuations in financial markets. *Macroeconomic Dynamics*, 4(2):170–196, 2000.
- V. Dani, T. P. Hayes, and S. M. Kakade. Stochastic linear optimization under bandit feedback. In *Conference on Learning Theory*, pages 355–366, 2008a.
- 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, 2008b.
- J. Gittins, K. Glazebrook, and R. Weber. *Multi-armed bandit allocation indices*. John Wiley & Sons, 2011.
- D. Hsu and S. Sabato. Heavy-tailed regression with a generalized median-of-means. In *International Conference on Machine Learning*, pages 37–45, 2014.
- D. Hsu and S. Sabato. Loss minimization and parameter estimation with heavy tails. *The Journal of Machine Learning Research*, 17(1):543–582, 2016.
- T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- T. Lattimore. A scale free algorithm for stochastic bandits with bounded kurtosis. In Advances in Neural Information Processing Systems, pages 1583–1592, 2017.
- T. Lattimore, K. Crammer, and C. Szepesvári. Optimal resource allocation with semi-bandit feedback. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, pages 477–486. AUAI Press, 2014.
- L. Li, W. Chu, J. Langford, and R. E. Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the Nineteenth International Conference on World Wide Web*, pages 661–670. ACM, 2010.
- J. Liebeherr, A. Burchard, and F. Ciucu. Delay bounds in communication networks with heavy-tailed and self-similar traffic. *IEEE Transactions on Information Theory*, 58(2):1010–1024, 2012.
- A. M. Medina and S. Yang. No-regret algorithms for heavy-tailed linear bandits. In *International Conference on Machine Learning*, pages 1642–1650, 2016.
- R. Munos et al. From bandits to monte-carlo tree search: The optimistic principle applied to optimization and planning. *Foundations and Trends*(R) *in Machine Learning*, 7(1):1–129, 2014.
- H. Robbins et al. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952.
- J. A. Roberts, T. W. Boonstra, and M. Breakspear. The heavy tail of the human brain. *Current Opinion in Neurobiology*, 31:164–172, 2015.
- Y. Seldin, F. Laviolette, N. Cesa-Bianchi, J. Shawe-Taylor, and P. Auer. PAC-Bayesian inequalities for martingales. *IEEE Transactions on Information Theory*, 58(12):7086–7093, 2012.
- M. Shao and C. L. Nikias. Signal processing with fractional lower order moments: stable processes and their applications. *Proceedings of the IEEE*, 81(7):986–1010, 1993.
- W. R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- S. Vakili, K. Liu, and Q. Zhao. Deterministic sequencing of exploration and exploitation for multiarmed bandit problems. *IEEE Journal of Selected Topics in Signal Processing*, 7(5):759–767, 2013.
- X. Yu, H. Shao, M. R. Lyu, and I. King. Pure exploration of multi-armed bandits with heavy-tailed payoffs. In *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence*, pages 937–946. AUAI Press, 2018.

## Appendix A Proof of Theorem 1 (Lower Bound of LinBET)

We prove the lower bound for  $d \geq 2$ . Assume d is even (when d is odd, similar results can be easily derived by considering the first d-1 dimensions). For  $D_t \subseteq \mathbb{R}^d$  with  $t \in [T]$ , we fix the decision set as  $D_1 = \cdots = D_T = D_{(d)}$ . Then, the fixed decision set is constructed as  $D_{(d)} \triangleq \{(x_1, \cdots, x_d) \in \mathbb{R}^d : x_1 + x_2 = \cdots = x_{d-1} + x_d = 1\}$ , which is a subset of intersection of the cube  $[0, 1]^d$  and the hyperplane  $x_1 + \cdots + x_d = d/2$ . We define a set  $S_d \triangleq \{(\theta_1, \cdots, \theta_d) : \forall i \in [d/2], (\theta_{2i-1}, \theta_{2i}) \in \{(2\Delta, \Delta), (\Delta, 2\Delta)\}\}$  with  $\Delta \in (0, 1/d]$ . The payoff functions take values in  $\{0, (1/\Delta)^{\frac{1}{\epsilon}}\}$  with  $\epsilon \in (0, 1]$ , for every  $x \in D_{(d)}$ , the expected payoff is  $\theta_*^\top x$ , where  $\theta_*$  is the underlying parameter drawn from  $S_d$ . To be more specific, we have the payoff function of x as

$$y(x) = \begin{cases} \left(\frac{1}{\Delta}\right)^{\frac{1}{\epsilon}} & \text{with a probability of } \Delta^{\frac{1}{\epsilon}} \theta_*^\top x, \\ 0 & \text{with a probability of } 1 - \Delta^{\frac{1}{\epsilon}} \theta_*^\top x. \end{cases}$$
(4)

In this setting, the  $(1 + \epsilon)$ -th raw moments of payoffs are bounded by d and  $|\theta_*^\top x| \le 1$ . We start the proof with the 2-dimensional case in Subsection A.1. Its extension to the general case (i.e., d > 2) is provided in Subsection A.2. Though we set a fixed decision set in the proofs, we can easily extend the lower bound here to the setting of time-varying decision sets, as discussed by Dani et al. (2008a).

#### **A.1** d = 2 Case

Let  $\mu_0 = (\Delta, \Delta)$ ,  $\mu_1 = (2\Delta, \Delta)$  and  $\mu_2 = (\Delta, 2\Delta)$ . The 2-dimensional decision set is  $D_{(2)} = \{(x_1, x_2) \in \mathbb{R}^2_+ : x_1 + x_2 = 1\}$ . Our payoff functions take values in  $\{0, (1/\Delta)^{\frac{1}{\epsilon}}\}$ , and for every  $x \in D_{(2)}$ , the expected payoff is  $\theta_*^\top x$ , where  $\theta_*$  is chosen uniformly at random from  $\{\mu_1, \mu_2\}$ . It is easy to find  $\mu_j^\top x = \Delta(1 + x_j)$  which is maximized at  $x_j = 1$  for  $j \in \{1, 2\}$ , and  $\mu_0^\top x = \Delta$  for any  $x \in D_{(2)}$ .

**Lemma 1.** If  $\theta_*$  is chosen uniformly at random from  $\{\mu_1, \mu_2\}$ , and the payoff for each  $x \in D_{(2)}$  is in  $\{0, (1/\Delta)^{\frac{1}{\epsilon}}\}$  with mean  $\theta_*^{\top} x$ , then for every algorithm  $\mathcal{A}$  and every  $T \ge 1$ , the regret satisfies

$$\mathbb{E}[R(\mathcal{A},T)] \ge \frac{1}{96} T^{\frac{1}{1+\epsilon}}.$$
(5)

*Proof.* We consider a deterministic algorithm  $\mathcal{A}$  first. Let  $q_{x,T} = T(x)/T$ , where T(x) denotes the number of pulls of arm x.  $\mathbb{Q}_T$  is the empirical distribution of arms with respect to  $q_{x,T}$  and X is drawn from  $\mathbb{Q}_T$ . We let  $\mathbb{P}_j$  and  $\mathbb{E}_j$  denote, respectively, the probability distribution of X conditional on  $\theta_* = \mu_j$  and the expectation conditional on  $\theta_* = \mu_j$ , where  $j \in \{0, 1, 2\}$ . Thus, we have  $\mathbb{P}_j(X \in \mathcal{E}) = \mathbb{E}_j[\sum_{x \in \mathcal{E}} T(x)]/T$  for any  $\mathcal{E} \subseteq D_{(2)}$ . At each time step  $t, x_t = (x_{t,1}, x_{t,2})$  is selected. We let  $y_t^* = \langle x_t^*, \theta_* \rangle$ . Hence, for  $j \in \{1, 2\}$ , we have

$$\mathbb{E}_{j}\left[\sum_{t=1}^{T} (y_{t}^{*} - y_{t}(x_{t}))\right] = \sum_{t=1}^{T} \mathbb{E}_{j}\left[\Delta(1 - x_{t,j})\right] = T \int_{D_{(2)}} \Delta(1 - x_{j}) d\mathbb{P}_{j}(x)$$
$$= T\Delta\left(1 - \int_{D_{(2)}} x_{j} d\mathbb{P}_{j}(x)\right) = T\Delta\left(1 - \left(\int_{0 \le x_{j} \le \frac{1}{2}} x_{j} d\mathbb{P}_{j}(x) + \int_{\frac{1}{2} < x_{j} \le 1} x_{j} d\mathbb{P}_{j}(x)\right)\right)$$
$$\geq T\Delta\left(1 - \left(\frac{1}{2}\mathbb{P}_{j}\left(0 \le X_{j} \le \frac{1}{2}\right) + \mathbb{P}_{j}\left(\frac{1}{2} < X_{j} \le 1\right)\right)\right), \tag{6}$$

which implies

$$\mathbb{E}[R(\mathcal{A},T)] = \mathbb{E}_{\theta_*} \left[ \mathbb{E}_j \left[ \sum_{t=1}^T (y_t^* - y_t(x_t)) \right] \right]$$
  

$$\geq T\Delta \left( 1 - \frac{1}{2} \sum_{j=1}^2 \left( \frac{1}{2} \mathbb{P}_j \left( 0 \le X_j \le \frac{1}{2} \right) + \mathbb{P}_j \left( \frac{1}{2} < X_j \le 1 \right) \right) \right).$$
(7)

According to Pinsker's inequality, for any  $\mathcal{E} \subseteq D_{(2)}$ , we have

$$\mathbb{P}_{j}(X \in \mathcal{E}) \leq \mathbb{P}_{0}(X \in \mathcal{E}) + \sqrt{\frac{1}{2}\mathrm{KL}(\mathbb{P}_{0}, \mathbb{P}_{j})},\tag{8}$$

where  $KL(\mathbb{P}_0, \mathbb{P}_j)$  denotes the Kullback-Leibler divergence (simply KL divergence). Hence,

$$\mathbb{E}[R(\mathcal{A},T)] \ge T\Delta \left(1 - \frac{1}{2}\sum_{j=1}^{2} \left(\frac{1}{2}\mathbb{P}_{0}\left(0 \le X_{j} \le \frac{1}{2}\right) + \mathbb{P}_{0}\left(\frac{1}{2} < X_{j} \le 1\right) + \frac{3}{2}\sqrt{\frac{1}{2}\mathrm{KL}(\mathbb{P}_{0},\mathbb{P}_{j})}\right)$$
$$= T\Delta \left(\frac{1}{4} - \frac{3}{4}\sum_{j=1}^{2}\sqrt{\frac{1}{2}\mathrm{KL}(\mathbb{P}_{0},\mathbb{P}_{j})}\right). \tag{9}$$

Since  $\mathcal{A}$  is deterministic, the sequence of received rewards  $W_T \triangleq (y_1, y_2, \cdots, y_T) \in \{0, (1/\Delta)^{\frac{1}{\epsilon}}\}^T$ uniquely determines the empirical distribution  $\mathbb{Q}_T$  and thus,  $\mathbb{Q}_T$  conditional on  $W_T$  is the same for any  $\theta_*$ . We let  $\mathbb{P}_j^t$  be the probability distribution of  $W_t = (y_1, y_2, \cdots, y_t)$  conditional on  $\theta_* = \mu_j$ . Based on the chain rule for KL divergence, we have

$$\mathrm{KL}(\mathbb{P}_0, \mathbb{P}_j) \le \mathrm{KL}(\mathbb{P}_0^T, \mathbb{P}_j^T).$$
(10)

Further, iteratively using the chain rule for KL divergence, we have

$$KL(\mathbb{P}_{0}^{T},\mathbb{P}_{j}^{T}) = KL(\mathbb{P}_{0}^{1},\mathbb{P}_{j}^{1}) + \sum_{t=2}^{T} \int_{W_{t-1}} KL\left(\mathbb{P}_{0}^{t}(\cdot|w_{t-1}),\mathbb{P}_{j}^{t}(\cdot|w_{t-1})\right) d\mathbb{P}_{0}^{t-1}(W_{t-1})$$
  
=  $KL(\mathbb{P}_{0}^{1},\mathbb{P}_{j}^{1}) +$ (11)

$$\sum_{t=2}^{T} \int_{x_t \in D_{(2)}} \int_{W_{t-1}|x_{t,j}=x_j} \operatorname{KL}\left(\Delta^{\frac{1+\epsilon}{\epsilon}}, \Delta^{\frac{1+\epsilon}{\epsilon}}(1+x_j)\right) d\mathbb{P}_0^{t-1}(W_{t-1}|x_{t,j}=x_j) d\mathbb{P}_0^{t-1}(x_{t,j}=x_j)$$
(12)

$$\leq 2\Delta^{\frac{1+\epsilon}{\epsilon}} + \sum_{t=2}^{T} \int_{x_t \in D_{(2)}} \int_{W_{t-1}|x_{t,j}=x_j} 2\Delta^{\frac{1+\epsilon}{\epsilon}} d\mathbb{P}_0^{t-1} (W_{t-1}|x_{t,j}=x_j) d\mathbb{P}_0^{t-1} (x_{t,j}=x_j)$$
(13)

$$=2T\Delta^{\frac{1+\epsilon}{\epsilon}},\tag{14}$$

where Eq. (13) could be derived by setting  $\Delta \leq (1/2)^{\frac{\epsilon}{1+\epsilon}}$ . Note that for any  $p, q \in (0, 1)$ , let  $\mathbb{P}$  and  $\mathbb{Q}$  denote the Bernoulli distribution with parameters p and q respectively. We denote  $\mathrm{KL}(\mathbb{P}, \mathbb{Q})$  as  $\mathrm{KL}(p, q)$  in Eq. (12). Therefore, we have

$$\mathbb{E}[R(\mathcal{A},T)] \ge T\Delta\left(\frac{1}{4} - \frac{3}{2}\sqrt{T\Delta^{\frac{1+\epsilon}{\epsilon}}}\right) \ge \frac{1}{96}T^{\frac{1}{1+\epsilon}},\tag{15}$$

where we set  $\Delta = T^{-\frac{\epsilon}{1+\epsilon}}/12$ .

So far we have discussed the case where  $\mathcal{A}$  is a deterministic algorithm. When  $\mathcal{A}$  is a randomized algorithm, the result is the same. In particular, let  $\mathbb{E}_{\mathcal{A}}$  denote the expectation with respect to the randomness of  $\mathcal{A}$ . Then, we have

$$\mathbb{E}[R(\mathcal{A},T)] = \mathbb{E}_{\mathcal{A}}\left[\mathbb{E}_{\theta_*}\left[\mathbb{E}_j\left[\sum_{t=1}^T \left(y_t^* - y_t(x_t)\right)\right]\right]\right].$$
(16)

If we fix the realization of the algorithm's randomization, the results of the previous steps for a deterministic algorithm apply and  $\mathbb{E}_{\theta_*}\left[\mathbb{E}_i\left[\sum_{t=1}^T (y_t^* - y_t(x_t))\right]\right]$  could be lower bounded as before. Hence,  $\mathbb{E}[R(\mathcal{A}, T)]$  is lower bounded as Eq. (15).

## A.2 General Case (d > 2)

Now we suppose d > 2 is even. If d is odd, we just take the first d-1 dimensions into consideration. Then we consider the contribution to the total expected regret from the choice of  $(x_{2i-1}, x_{2i})$ , for all  $i \in [d/2]$ . We call  $(x_{2i-1}, x_{2i})$  the *i*-th component of x. Analogously to the d = 2 case, we set  $(\theta_{*,2i-1}, \theta_{*,2i}) \in {\mu_1, \mu_2}$ . The decision region is  $D_{(d)} = {(x_1, \dots, x_d) \in \mathbb{R}^d_+ : x_1 + x_2 = \dots = x_{d-1} + x_d = 1}$ . Then, by following the proof for d = 2 case, we could derive the regret due to the *i*-th component of x as

$$\mathbb{E}\left[R^{(i)}(\mathcal{A},T)\right] \ge \frac{1}{96}T^{\frac{1}{1+\epsilon}},\tag{17}$$

where  $i \in [d/2]$ . Summing over the d/2 components of Eq. (17) completes the proof for Theorem 1.

# Appendix B Proof of Theorem 2 (Regret Analysis for the MENU Algorithm)

To prove Theorem 2, we start with proving the following two lemmas. Recall that the algorithm in the paper is based on least-squares estimate (LSE).

**Lemma 2** (Confidence Ellipsoid of LSE). Let  $\hat{\theta}_n$  denote the LSE of  $\theta_*$  with the sequence of decisions  $x_1, \dots, x_n$  and observed payoffs  $y_1, \dots, y_n$ . Assume that for all  $\tau \in [n]$  and all  $x_\tau \in D_\tau \subseteq \mathbb{R}^d$ ,  $\mathbb{E}[|\eta_\tau|^{1+\epsilon}|\mathcal{F}_{\tau-1}] \leq c$  and  $\|\theta_*\|_2 \leq S$ . Then  $\hat{\theta}_n$  satisfies

$$Pr\left(\|\hat{\theta}_{n} - \theta_{*}\|_{V_{n}} \le (9dc)^{\frac{1}{1+\epsilon}} n^{\frac{1-\epsilon}{2(1+\epsilon)}} + \lambda^{\frac{1}{2}}S\right) \ge \frac{3}{4},\tag{18}$$

where  $\lambda > 0$  is a regularization parameter and  $V_n = \lambda I_d + \sum_{\tau=1}^n x_{\tau} x_{\tau}^{\top}$ .

*Proof.* The singular value decomposition of  $X_n$  is  $U\Sigma_n V^{\top}$ , where U is an  $n \times d$  matrix with orthonormal columns, V is a  $d \times d$  unitary matrix and  $\Sigma_n$  is an  $n \times n$  diagonal matrix with non-negative entries. We calculate  $V_n = V(\Sigma_n^2 + \lambda I_d)V^{\top}$  and

$$V_n^{-\frac{1}{2}}X_n^{\top} = V\left(\Sigma_n^2 + \lambda I_d\right)^{-\frac{1}{2}}\Sigma_n U^{\top}.$$
(19)

Let  $u_i^{\top}$  denote the *i*-th row of  $V\left(\Sigma_n^2 + \lambda I_d\right)^{-\frac{1}{2}} \Sigma_n U^{\top}$ , which leads to  $||u_i||_2 \leq 1$ . More importantly, by optimization, we have  $||u_i||_{1+\epsilon} \leq n^{\frac{1-\epsilon}{2(1+\epsilon)}}$ . By letting  $Y_n = (y_1, \cdots, y_n)$ , we have

$$\begin{aligned} \|\hat{\theta}_{n} - \theta_{*}\|_{V_{n}} &= \|V_{n}^{-1}X_{n}^{\top}(Y_{n} - X_{n}\theta_{*}) - \lambda V_{n}^{-1}\theta_{*}\|_{V_{n}} \\ &\leq \|V_{n}^{-\frac{1}{2}}X_{n}^{\top}(Y_{n} - X_{n}\theta_{*})\|_{2} + \lambda \|\theta_{*}\|_{V_{n}^{-1}} \leq \sqrt{\sum_{i=1}^{d} \left(u_{i}^{\top}(Y_{n} - X_{n}\theta_{*})\right)^{2}} + \lambda^{\frac{1}{2}}S. \end{aligned}$$
(20)

Inspired by Bubeck et al. (2013); Medina and Yang (2016), we bound the desired probability by using a union bound as

$$\Pr\left(\sum_{i=1}^{d} \left(\sum_{\tau=1}^{n} u_{i,\tau} \eta_{\tau}\right)^{2} > \gamma^{2}\right) \leq \Pr\left(\exists i, \tau, |u_{i,\tau} \eta_{\tau}| > \gamma\right) + \Pr\left(\sum_{i=1}^{d} \left(\sum_{\tau=1}^{n} u_{i,\tau} \eta_{\tau} \mathbb{1}_{|u_{i,\tau} \eta_{\tau}| \le \gamma}\right)^{2} > \gamma^{2}\right)$$

$$(21)$$

where  $\mathbb{1}_{\{\cdot\}}$  is the indicator function. By using a union bound and Markov's inequality, the first term could be bounded as

$$\Pr\left(\exists i, \tau, |u_{i,\tau}\eta_{\tau}| > \gamma\right) \le \sum_{i=1}^{d} \sum_{\tau=1}^{n} \Pr(|u_{i,\tau}\eta_{\tau}| > \gamma) \le \frac{\sum_{i=1}^{d} \sum_{\tau=1}^{n} \mathbb{E}[|u_{i,\tau}\eta_{\tau}|^{1+\epsilon}]}{\gamma^{1+\epsilon}} \qquad (22)$$
$$\le \frac{\sum_{i=1}^{d} \sum_{\tau=1}^{n} |u_{i,\tau}|^{1+\epsilon} c}{\gamma^{1+\epsilon}} \le \frac{dcn^{\frac{1-\epsilon}{2}}}{\gamma^{1+\epsilon}}.$$

Based on Markov's inequality, we bound the second term as

$$\Pr\left(\sum_{i=1}^{d} \left(\sum_{\tau=1}^{n} u_{i,\tau} \eta_{\tau} \mathbb{1}_{|u_{i,\tau}\eta_{\tau}| \leq \gamma}\right)^{2} > \gamma^{2}\right) \leq \frac{\mathbb{E}\left[\sum_{i=1}^{d} \left(\sum_{\tau=1}^{n} u_{i,\tau} \eta_{\tau} \mathbb{1}_{|u_{i,\tau}\eta_{\tau}| \leq \gamma}\right)^{2}\right]}{\gamma^{2}}$$

$$= \sum_{i=1}^{d} \left(\frac{\mathbb{E}\left[\sum_{\tau=1}^{n} (u_{i,\tau}\eta_{\tau})^{2} \mathbb{1}_{|u_{i,\tau}\eta_{\tau}| \leq \gamma}\right]}{\gamma^{2}} + 2\frac{\mathbb{E}\left[\sum_{\tau'>\tau} (u_{i,\tau}\eta_{\tau}) \mathbb{1}_{|u_{i,\tau}\eta_{\tau}| \leq \gamma}\right]}{\gamma^{2}}\right)$$

$$\leq \sum_{i=1}^{d} \left(\frac{\mathbb{E}\left[\sum_{\tau=1}^{n} (u_{i,\tau}\eta_{\tau})^{2} \mathbb{1}_{|u_{i,\tau}\eta_{\tau}| \leq \gamma}\right]}{\gamma^{2}} + 2\frac{\sum_{\tau'>\tau} \mathbb{E}[(u_{i,\tau}\eta_{\tau}) \mathbb{1}_{|u_{i,\tau}\eta_{\tau}| \leq \gamma}]\mathbb{E}[(u_{i,\tau'}\eta_{\tau'})\mathbb{1}_{|u_{i,\tau'}\eta_{\tau'}| \leq \gamma}|\mu_{i,\tau}\eta_{\tau}|]}{\gamma^{2}}\right)$$

$$\leq \sum_{i=1}^{d} \left(\frac{\sum_{\tau=1}^{n} |u_{i,\tau}|^{1+\epsilon}c}{\gamma^{1+\epsilon}} + \left(\frac{\sum_{\tau=1}^{n} |u_{i,\tau}|^{1+\epsilon}c}{\gamma^{1+\epsilon}}\right)^{2}\right)\right)$$

$$(23)$$

$$\leq \frac{dcn^{\frac{1-\epsilon}{2}}}{\gamma^{1+\epsilon}} + d\left(\frac{n^{\frac{1-\epsilon}{2}}c}{\gamma^{1+\epsilon}}\right)^{2}.$$

Note that Eq. (23) uses the fact that 
$$\mathbb{E}[(u_{i,\tau}\eta_{\tau})\mathbb{1}_{|u_{i,\tau}\eta_{\tau}|\leq\gamma}|\mathcal{F}_{\tau-1}] = -\mathbb{E}[(u_{i,\tau}\eta_{\tau})\mathbb{1}_{|u_{i,\tau}\eta_{\tau}|>\gamma}|\mathcal{F}_{\tau-1}].$$
  
Finally, setting  $\gamma = (9dc)^{\frac{1}{1+\epsilon}}n^{\frac{1-\epsilon}{2(1+\epsilon)}}$  completes the proof.

**Lemma 3.** Recall  $\hat{\theta}_{n,j}$ ,  $\hat{\theta}_{n,k^*}$  and  $V_n$  in MENU (i.e., Algorithm 1). If there exists a  $\gamma > 0$  such that  $Pr\left(\|\hat{\theta}_{n,j} - \theta_*\|_{V_n} \leq \gamma\right) \geq \frac{3}{4}$  holds for all  $j \in [k]$  with  $k \geq 1$ , then with probability at least  $1 - e^{-\frac{k}{24}}, \|\hat{\theta}_{n,k^*} - \theta_*\|_{V_n} \leq 3\gamma$ .

*Proof.* The proof is inspired by Hsu and Sabato (2014). We define  $b_j \triangleq \mathbb{1}_{\|\hat{\theta}_{n,j}-\theta_*\|_{V_n} > \gamma}$ ,  $p_j \triangleq \Pr(b_j = 1)$  and  $\mathbb{B}_{V_n}(\theta_*, \gamma) \triangleq \{\theta : \|\theta - \theta_*\|_{V_n} \leq \gamma\}$ . We know that  $p_j < 1/4$ . By Azuma-Hoeffding's inequality, we have

$$\Pr\left(\sum_{j=1}^{k} b_j \ge \frac{k}{3}\right) < \Pr\left(\sum_{j=1}^{k} b_j - p_j \ge \frac{k}{12}\right) \le e^{-\frac{k}{24}},\tag{25}$$

which means that more than 2/3 of  $\{\hat{\theta}_{n,1}, \dots, \hat{\theta}_{n,k}\}$  are contained in  $\mathbb{B}_{V_n}(\theta_*, \gamma)$  (denoting by this event  $\mathcal{E}$ ) with probability at least  $1 - e^{-\frac{k}{24}}$ . Note that the value k/3 in Eq. (25) could also be set as other values in (k/4, k/2). Conditional on the event  $\mathcal{E}$ , by letting  $r_j$  be the median of  $\{\|\hat{\theta}_{n,j} - \hat{\theta}_{n,s}\|_{V_n} : s \in [k] \setminus j\}$ , we have

- If  $\hat{\theta}_{n,j} \in \mathbb{B}_{V_n}(\theta_*, \gamma)$ ,  $\|\hat{\theta}_{n,j} \hat{\theta}_{n,s}\|_{V_n} \le 2\gamma$  for all  $\hat{\theta}_{n,s} \in \mathbb{B}_{V_n}(\theta_*, \gamma)$  by triangle inequality. Therefore,  $r_j \le 2\gamma$ .
- If  $\hat{\theta}_{n,j} \notin \mathbb{B}_{V_n}(\theta_*, 3\gamma)$ ,  $\|\hat{\theta}_{n,j} \hat{\theta}_{n,s}\|_{V_n} > 2\gamma$  for all  $\hat{\theta}_{n,s} \in \mathbb{B}_{V_n}(\theta_*, \gamma)$  by triangle inequality. Therefore,  $r_j > 2\gamma$ .

Combining the above two cases completes proof.

Based on Lemmas 2 and 3, by setting  $k = \lceil 24 \log (eT/\delta) \rceil$ , we have  $\|\hat{\theta}_{n,k^*} - \theta_*\|_{V_n} \leq 3\left((9dc)^{\frac{1}{1+\epsilon}}n^{\frac{1-\epsilon}{2(1+\epsilon)}} + \lambda^{\frac{1}{2}}S\right)$  with probability at least  $1 - \delta/T$ . The following part of proof is standard (Dani et al., 2008a; Abbasi-Yadkori et al., 2011). We include it for the sake of completeness. By letting  $\beta_n = 3\left((9dc)^{\frac{1}{1+\epsilon}}n^{\frac{1-\epsilon}{2(1+\epsilon)}} + \lambda^{\frac{1}{2}}S\right)$ , we can decompose the instantaneous regret as follows:

$$r_{n} = \theta_{*}^{\top} x_{*} - \theta_{*}^{\top} x_{n} \leq \tilde{\theta}_{n}^{\top} x_{n} - \theta_{*}^{\top} x_{n} \leq \left( \|\tilde{\theta}_{n} - \hat{\theta}_{n-1,k^{*}}\|_{V_{n-1}} + \|\hat{\theta}_{n-1,k^{*}} - \theta_{*}\|_{V_{n-1}} \right) \|x_{n}\|_{V_{n-1}^{-1}}$$

$$\leq 2\beta_{n-1} \|x_{n}\|_{V_{n-1}^{-1}}, \tag{26}$$

where we recall that  $(x_n, \tilde{\theta}_n)$  is optimistic in MENU. Note that, for n = 1, the above inequality also holds with  $V_0 = \lambda I_d$ . On the other hand, by considering  $|x_t^{\top} \theta_*| \leq L$ , we always have

$$r_n \le 2L. \tag{27}$$

We can get that

γ

$$x_n \le 2\min\{\beta_{n-1} \| x_n \|_{V_{n-1}^{-1}}, L\} \le 2(\beta_{n-1} + L)\min\{\| x_n \|_{V_{n-1}^{-1}}, 1\}.$$
(28)

Following Lemma 11 of Abbasi-Yadkori et al. (2011), we know that

$$\sum_{n=1}^{N} \min\{\|x_n\|_{V_{n-1}^{-1}}^2, 1\} \le 2\sum_{n=1}^{N} \log(1 + \|x_n\|_{V_{n-1}^{-1}}^2) = 2\log\left(\frac{\det(V_N)}{\det(V_0)}\right) \le 2d\log\left(1 + \frac{ND^2}{\lambda d}\right), \quad (29)$$

where N is the number of epochs in MENU. Therefore, the total regret can be upper bounded by

$$R(\text{MENU}, T) \le k \sum_{n=1}^{N} r_n \le k \sqrt{N \sum_{n=1}^{N} r_n^2} \le 2k N^{\frac{1}{2}} (\beta_N + L) \sqrt{\sum_{n=1}^{N} \min\{\|x_n\|_{V_{n-1}^{-1}}^2, 1\}}$$
  
$$\le 6 \left( (12dc)^{\frac{1}{1+\epsilon}} + \lambda^{\frac{1}{2}} S + L \right) T^{\frac{1}{1+\epsilon}} \left( 24 \log\left(\frac{eT}{\delta}\right) + 1 \right)^{\frac{\epsilon}{1+\epsilon}} \sqrt{2d \log\left(1 + \frac{TD^2}{\lambda d}\right)}.$$
(30)

The condition of  $T \ge 256 + 24 \log(e/\delta)$  is required for  $T \ge k$ , which completes the proof.

## Appendix C Proof of Theorem 3 (Regret Analysis for the TOFU Algorithm)

**Lemma 4** (Confidence Ellipsoid of Truncated Estimate). With the sequence of decisions  $x_1, \dots, x_t$ , the truncated payoffs  $\{Y_i^{\dagger}\}_{i=1}^d$  and the parameter estimate  $\theta_t^{\dagger}$  are defined in TOFU (i.e., Algorithm 2). Assume that for all  $\tau \in [t]$  and all  $x_{\tau} \in D_{\tau} \subseteq \mathbb{R}^d$ ,  $\mathbb{E}[|y_{\tau}|^{1+\epsilon}|\mathcal{F}_{\tau-1}] \leq b$  and  $\|\theta_*\|_2 \leq S$ . With probability at least  $1 - \delta$ , we have

$$\|\theta_t^{\dagger} - \theta_*\|_{V_t} \le 4\sqrt{d}b^{\frac{1}{1+\epsilon}} \left(\log\left(\frac{2d}{\delta}\right)\right)^{\frac{\epsilon}{1+\epsilon}} t^{\frac{1-\epsilon}{2(1+\epsilon)}} + \lambda^{\frac{1}{2}}S,\tag{31}$$

where  $\lambda > 0$  is a regularization parameter and  $V_t = \lambda I_d + \sum_{\tau=1}^t x_\tau x_\tau^\top$ .

Proof. Similarly to Eq. (20), we have

$$\|\theta_t^{\dagger} - \theta_*\|_{V_t} \le \sqrt{\sum_{i=1}^d \left(u_i^{\top}(Y_i^{\dagger} - X_t\theta_*)\right)^2} + \lambda^{\frac{1}{2}}S.$$
(32)

We let  $y_{\tau}^{\dagger}$  denote  $Y_{i,\tau}^{\dagger}$  for notation simplicity as the following proof holds for all  $i \in [d]$ . Then with probability at least  $1 - \delta/d$ , we have

$$\begin{aligned} u_{i}^{\top} \left(Y_{i}^{\dagger} - X_{t} \theta_{*}\right) &= \sum_{\tau=1}^{t} u_{i,\tau} \left(y_{\tau}^{\dagger} - \mathbb{E}[y_{\tau} | \mathcal{F}_{\tau-1}]\right) \end{aligned} \tag{33} \\ &= \sum_{\tau=1}^{t} u_{i,\tau} \left(y_{\tau}^{\dagger} - \mathbb{E}\left[y_{\tau}^{\dagger} | \mathcal{F}_{\tau-1}\right] - \mathbb{E}\left[y_{\tau} \mathbb{1}_{|u_{i,\tau} y_{\tau}| > b_{t}} | \mathcal{F}_{\tau-1}\right]\right) \\ &\leq \left|\sum_{\tau=1}^{t} u_{i,\tau} \left(y_{\tau}^{\dagger} - \mathbb{E}[y_{\tau}^{\dagger} | \mathcal{F}_{\tau-1}]\right)\right| + \left|\sum_{\tau=1}^{t} u_{i,\tau} \mathbb{E}[y_{\tau} \mathbb{1}_{|u_{i,\tau} y_{\tau}| > b_{t}} | \mathcal{F}_{\tau-1}]\right| \\ &\leq \left|2b_{t} \log\left(\frac{2d}{\delta}\right) + \frac{1}{2b_{t}} \sum_{\tau=1}^{t} \mathbb{E}\left[u_{i,\tau}^{2} \left(y_{\tau}^{\dagger} - \mathbb{E}\left[y_{\tau}^{\dagger} | \mathcal{F}_{\tau-1}\right]\right)^{2} | \mathcal{F}_{\tau-1}\right]\right| + \left|\sum_{\tau=1}^{t} \mathbb{E}[u_{i,\tau} y_{\tau} \mathbb{1}_{|u_{i,\tau} y_{\tau}| > b_{t}} | \mathcal{F}_{\tau-1}]\right| \end{aligned} \tag{34} \\ &\leq 2b_{t} \log\left(\frac{2d}{\delta}\right) + \frac{\sum_{\tau=1}^{t} |u_{i,\tau}|^{1+\epsilon}b}{2b_{t}^{\epsilon}} + \frac{\sum_{\tau=1}^{t} |u_{i,\tau}|^{1+\epsilon}b}{b_{t}^{\epsilon}} \\ &\leq 4b^{\frac{1}{1+\epsilon}} \left(\log\left(\frac{2d}{\delta}\right)\right)^{\frac{\epsilon}{1+\epsilon}} t^{\frac{1-\epsilon}{2(1+\epsilon)}}, \end{aligned} \tag{35}$$

 $\leq 4b^{\frac{1}{1+\epsilon}} \left( \log\left(\frac{2a}{\delta}\right) \right)^{1+\epsilon} t^{\frac{1-\epsilon}{2(1+\epsilon)}},$ (35) where Eq. (34) is obtained by applying Bernstein's inequality for martingales (Seldin et al., 2012) and Eq. (35) is obtained by the fact that  $||u_i||_{1+\epsilon} \leq t^{\frac{1-\epsilon}{2(1+\epsilon)}}$  and by setting  $b_t = (b/\log(2d/\delta))^{\frac{1}{1+\epsilon}} t^{\frac{1-\epsilon}{2(1+\epsilon)}}$ . Combining Eq. (32) and Eq. (35) completes the proof. With similar procedures to the proof of Theorem 2, we have the regret of TOFU as:

$$R(\text{TOFU},T) \le 2T^{\frac{1}{1+\epsilon}} \left( 4\sqrt{d}b^{\frac{1}{1+\epsilon}} \left( \log\left(\frac{2dT}{\delta}\right) \right)^{\frac{\epsilon}{1+\epsilon}} + \lambda^{\frac{1}{2}}S + L \right) \sqrt{2d\log\left(1 + \frac{TD^2}{\lambda d}\right)}, \quad (36)$$

which completes the proof.