

---

# Continuous Mean-Covariance Bandits

---

**Yihan Du**

IIS, Tsinghua University  
Beijing, China  
duyh18@mails.tsinghua.edu.cn

**Siwei Wang**

CST, Tsinghua University  
Beijing, China  
wangsw2020@mail.tsinghua.edu.cn

**Zhixuan Fang**

IIS, Tsinghua University, Beijing, China  
Shanghai Qi Zhi Institute, Shanghai, China  
zfang@mail.tsinghua.edu.cn

**Longbo Huang\***

IIS, Tsinghua University  
Beijing, China  
longbohuang@mail.tsinghua.edu.cn

## Abstract

Existing risk-aware multi-armed bandit models typically focus on risk measures of individual options such as variance. As a result, they cannot be directly applied to important real-world online decision making problems with correlated options. In this paper, we propose a novel Continuous Mean-Covariance Bandit (CMCB) model to explicitly take into account option correlation. Specifically, in CMCB, there is a learner who sequentially chooses weight vectors on given options and observes random feedback according to the decisions. The agent’s objective is to achieve the best trade-off between reward and risk, measured with option covariance. To capture different reward observation scenarios in practice, we consider three feedback settings, i.e., full-information, semi-bandit and full-bandit feedback. We propose novel algorithms with optimal regrets (within logarithmic factors), and provide matching lower bounds to validate their optimality. The experimental results also demonstrate the superiority of our algorithms. To the best of our knowledge, this is the first work that considers option correlation in risk-aware bandits and explicitly quantifies how arbitrary covariance structures impact the learning performance. The novel analytical techniques we developed for exploiting the estimated covariance to build concentration and bounding the risk of selected actions based on sampling strategy properties can likely find applications in other bandit analysis and be of independent interests.

## 1 Introduction

The stochastic Multi-Armed Bandit (MAB) [3, 30, 2] problem is a classic online learning model, which characterizes the exploration-exploitation trade-off in decision making. Recently, due to the increasing requirements of risk guarantees in practical applications, the Mean-Variance Bandits (MVB) [28, 32, 36] which aim at balancing the rewards and performance variances have received extensive attention. While MVB provides a successful risk-aware model, it only considers discrete decision space and focuses on the variances of individual arms (assuming independence among arms).

However, in many real-world scenarios, a decision often involves multiple options with certain correlation structure, which can heavily influence risk management and cannot be ignored. For instance, in finance, investors can select portfolios on multiple correlated assets, and the investment risk is closely related to the correlation among the chosen assets. The well-known “risk diversification” strategy [5] embodies the importance of correlation to investment decisions. In clinical trials, a

---

\*Corresponding author.

treatment often consists of different drugs with certain ratios, and the correlation among drugs plays an important role in the treatment risk. Failing to handle the correlation among multiple options, existing MVB results cannot be directly applied to these important real-world tasks.

Witnessing the above limitation of existing risk-aware results, in this paper, we propose a novel Continuous Mean-Covariance Bandit (CMCB) model, which considers a set of options (base arms) with continuous decision space and measures the risk of decisions with the option correlation. Specifically, in this model, a learner is given  $d$  base arms, which are associated with an unknown joint reward distribution with a mean vector and covariance. At each timestep, the environment generates an underlying random reward for each base arm according to the joint distribution. Then, the learner selects a weight vector of base arms and observes the rewards. The goal of the learner is to minimize the expected cumulative regret, i.e., the total difference of the reward-risk (mean-covariance) utilities between the chosen actions and the optimal action, where the optimal action is defined as the weight vector that achieves the best trade-off between the expected reward and covariance-based risk. To capture important observation scenarios in practice, we consider three feedback settings in this model, i.e., full-information (CMCB-FI), semi-bandit (CMCB-SB) and full-bandit (CMCB-FB) feedback, which vary from seeing rewards of all options to receiving rewards of the selected options to only observing a weighted sum of rewards.

The CMCB framework finds a wide range of real-world applications, including finance [25], company operation [26] and online advertising [29]. For example, in stock markets, investors choose portfolios based on the observed prices of all stocks (full-information feedback), with the goal of earning high returns and meanwhile minimizing risk. In company operation, managers allocate investment budgets to several correlated business and only observe the returns of the invested business (semi-bandit feedback), with the objective of achieving high returns and low risk. In clinical trials, clinicians select a treatment comprised of different drugs and only observe an overall therapeutic effect (full-bandit feedback), where good therapeutic effects and high stability are both desirable.

For both CMCB-FI and CMCB-SB, we propose optimal algorithms (within logarithmic factors) and establish matching lower bounds for the problems, and contribute novel techniques in analyzing the risk of chosen actions and exploiting the covariance information. For CMCB-FB, we develop a novel algorithm which adopts a carefully designed action set to estimate the expected rewards and covariance, with non-trivial regret guarantees. Our theoretical results offer an explicit quantification of the influences of arbitrary covariance structures on learning performance, and our empirical evaluations also demonstrate the superior performance of our algorithms.

Our work differs from previous works on bandits with covariance [34, 35, 13, 27] in the following aspects. (i) We consider the reward-risk objective under continuous decision space and stochastic environment, while existing works study either combinatorial bandits, where the decision space is discrete and risk is not considered in the objective, or adversarial online optimization. (ii) We do not assume a prior knowledge or direct feedback on the covariance matrix as in [34, 35, 13]. (iii) Our results for full-information and full-bandit feedback explicitly characterize the impacts of arbitrary covariance structures, whereas prior results, e.g., [13, 27], only focus on independent or positively-correlated cases. These differences pose new challenges in algorithm design and analysis, and demand new analytical techniques.

We summarize the main contributions as follows.

- We propose a novel risk-aware bandit model called continuous mean-covariance bandit (CMCB), which considers correlated options with continuous decision space, and characterizes the trade-off between reward and covariance-based risk. Motivated by practical reward observation scenarios, three feedback settings are considered under CMCB, i.e., full-information (CMCB-FI), semi-bandit (CMCB-SB) and full-bandit (CMCB-FB).
- We design an algorithm `MC-Empirical` for CMCB-FI with an optimal  $O(\sqrt{T})$  regret (within logarithmic factors), and develop a novel analytical technique to build a relationship on risk between chosen actions and the optimal one using properties of the sampling strategy. We also derive a matching lower bound, by analyzing the gap between hindsight knowledge and available empirical information under a Bayesian environment.
- For CMCB-SB, we develop `MC-UCB`, an algorithm that exploits the estimated covariance information to construct confidence intervals and achieves the optimal  $O(\sqrt{T})$  regret (up to

logarithmic factors). A matching regret lower bound is also established, by investigating the necessary regret paid to differentiate two well-chosen distinct instances.

- We propose a novel algorithm MC-ETE for CMCB-FB, which employs a well-designed action set to carefully estimate the reward means and covariance, and achieves an  $O(T^{\frac{2}{3}})$  regret guarantee under the severely limited feedback.

To our best knowledge, our work is the *first* to explicitly characterize the influences of *arbitrary* covariance structures on learning performance in risk-aware bandits. Our results shed light into optimal risk management in online decision making with correlated options. Due to space limitation, we defer all detailed proofs to the supplementary material.

## 2 Related Work

**(Risk-aware Bandits)** Sani et al. [28] initiate the classic mean-variance paradigm [25, 16] in bandits, and formulate the mean-variance bandit problem, where the learner plays a single arm each time and the risk is measured by the variances of individual arms. Vakili & Zhao [31, 32] further study this problem under a different metric and complete the regret analysis. Zhu & Tan [36] provide a Thompson Sampling-based algorithm for mean-variance bandits. In addition to variance, several works consider other risk criteria. The VaR measure is studied in [12], and CVaR is also investigated to quantify the risk in [15, 18]. Cassel et al. [6] propose a general risk measure named empirical distributions performance measure (EDPM) and present an algorithmic framework for EDPM. All existing studies on risk-aware bandits only consider discrete decision space and assume independence among arms, and thus they cannot be applied to our CMCB problem.

**(Bandits with Covariance)** In the stochastic MAB setting, while there have been several works [13, 27] on covariance, they focus on the combinatorial bandit problem without considering risk. Degenne & Perchet [13] study the combinatorial semi-bandits with correlation, which assume a known upper bound on the covariance, and design an algorithm with this prior knowledge of covariance. Perrault et al. [27] further investigate this problem without the assumption on covariance under the sub-exponential distribution framework, and propose an algorithm with a tight asymptotic regret analysis. In the adversarial setting, Warmuth & Kuzmin [34, 35] consider an online variance minimization problem, where at each timestep the learner chooses a weight vector and receives a covariance matrix, and propose the exponentiated gradient based algorithms. Our work differs from the above works in the following aspects: compared to [13, 27], we consider a continuous decision space instead of combinatorial space, study the reward-risk objective instead of only maximizing the expected reward, and investigate two more feedback settings other than the semi-bandit feedback. Compared to [34, 35], we consider the stochastic environment and in our case, the covariance cannot be directly observed and needs to be estimated.

## 3 Continuous Mean-Covariance Bandits (CMCB)

Here we present the formulation for the Continuous Mean-Covariance Bandits (CMCB) problem. Specifically, a learner is given  $d$  base arms labeled  $1, \dots, d$  and a decision (action) space  $\mathcal{D} \subseteq \Delta_d$ , where  $\Delta_d = \{\mathbf{w} \in \mathbb{R}^d : 0 \leq w_i \leq 1, \forall i \in [d], \sum_i w_i = 1\}$  denotes the probability simplex in  $\mathbb{R}^d$ . The base arms are associated with an unknown  $d$ -dimensional joint reward distribution with mean vector  $\boldsymbol{\theta}^*$  and positive semi-definite covariance matrix  $\Sigma^*$ , where  $\Sigma_{ii}^* \leq 1$  for any  $i \in [d]$  without loss of generality. For any action  $\mathbf{w} \in \mathcal{D}$ , which can be regarded as a weight vector placed on the base arms, the instantaneous reward-risk utility is given by the following *mean-covariance* function

$$f(\mathbf{w}) = \mathbf{w}^\top \boldsymbol{\theta}^* - \rho \mathbf{w}^\top \Sigma^* \mathbf{w}, \quad (1)$$

where  $\mathbf{w}^\top \boldsymbol{\theta}^*$  denotes the expected reward,  $\mathbf{w}^\top \Sigma^* \mathbf{w}$  represents the risk, i.e., reward variance, and  $\rho > 0$  is a risk-aversion parameter that controls the weight placed on the risk. We define the optimal action as  $\mathbf{w}^* = \operatorname{argmax}_{\mathbf{w} \in \mathcal{D}} f(\mathbf{w})$ . Compared to linear bandits [1, 19], the additional quadratic term in  $f(\mathbf{w})$  raises significant challenges in estimating the covariance, bounding the risk of chosen actions and deriving covariance-dependent regret bounds.

At each timestep  $t$ , the environment generates an underlying (unknown to the learner) random reward vector  $\boldsymbol{\theta}_t = \boldsymbol{\theta}^* + \boldsymbol{\eta}_t$  according to the joint distribution, where  $\boldsymbol{\eta}_t$  is a zero-mean noise vector and it is

---

**Algorithm 1** MC-Empirical

---

- 1: **Input:** Risk-aversion parameter  $\rho > 0$ .
  - 2: Initialization: Pull action  $\mathbf{w}_1 = (\frac{1}{d}, \dots, \frac{1}{d})$ , and observe  $\boldsymbol{\theta}_1 = (\theta_{1,1}, \dots, \theta_{d,1})^\top$ .  $\hat{\boldsymbol{\theta}}_{1,i} \leftarrow \theta_{1,i}$ ,  $\forall i \in [d]$ .  $\hat{\Sigma}_{1,ij} = (\theta_{1,i} - \hat{\boldsymbol{\theta}}_{1,i}^*)(\theta_{1,j} - \hat{\boldsymbol{\theta}}_{1,j}^*)$ ,  $\forall i, j \in [d]$ .
  - 3: **for**  $t = 2, 3, \dots$  **do**
  - 4:      $\mathbf{w}_t = \operatorname{argmax}_{\mathbf{w} \in \Delta_d} (\mathbf{w}^\top \hat{\boldsymbol{\theta}}_{t-1}^* - \rho \mathbf{w}^\top \hat{\Sigma}_{t-1} \mathbf{w})$
  - 5:     Pull  $\mathbf{w}_t$ , observe  $\boldsymbol{\theta}_t = (\theta_{t,1}, \dots, \theta_{t,d})^\top$
  - 6:      $\hat{\boldsymbol{\theta}}_{t,i}^* \leftarrow \frac{1}{t} \sum_{s=1}^t \theta_{s,i}$ ,  $\forall i \in [d]$
  - 7:      $\hat{\Sigma}_{t,ij} = \frac{1}{t} \sum_{s=1}^t (\theta_{s,i} - \hat{\boldsymbol{\theta}}_{t,i}^*)(\theta_{s,j} - \hat{\boldsymbol{\theta}}_{t,j}^*)$ ,  $\forall i, j \in [d]$
  - 8: **end for**
- 

independent among different timestep  $t$ . Note that here we consider an additive vector noise to the parameter  $\boldsymbol{\theta}^*$ , instead of the simpler scalar noise added in the observation (i.e.,  $y_t = \mathbf{w}_t^\top \boldsymbol{\theta}^* + \eta_t$ ) as in linear bandits [1, 19]. Our noise setting better models the real-world scenarios where distinct actions incur different risk, and enables us to explicitly quantify the correlation effects. Following the standard assumption in the bandit literature [24, 13, 36], we assume the noise is sub-Gaussian, i.e.,  $\forall \mathbf{u} \in \mathbb{R}^d$ ,  $\mathbb{E}[\exp(\mathbf{u}^\top \boldsymbol{\eta}_t)] \leq \exp(\frac{1}{2} \mathbf{u}^\top \Sigma^* \mathbf{u})$ , where  $\Sigma^*$  is unknown. The learner selects an action  $\mathbf{w}_t \in \mathcal{D}$  and observes the feedback according to a certain structure (specified later). For any time horizon  $T > 0$ , define the expected cumulative regret as

$$\mathbb{E}[\mathcal{R}(T)] = \sum_{t=1}^T \mathbb{E}[f(\mathbf{w}^*) - f(\mathbf{w}_t)].$$

The objective of the learner is to minimize  $\mathbb{E}[\mathcal{R}(T)]$ . Note that our mean-covariance function Eq. (1) extends the popular mean-variance measure [28, 32, 36] to the continuous decision space.

In the following, we consider three feedback settings motivated by reward observation scenarios in practice, including (i) full-information (CMCB-FI), observing random rewards of all base arms after a pull, (ii) semi-bandit (CMCB-SB), only observing random rewards of the selected base arms, and (iii) full-bandit (CMCB-FB), only seeing a weighted sum of the random rewards from base arms. We will present the formal definitions of these three feedback settings in the following sections.

**Notations.** For action  $\mathbf{w} \in \mathcal{D}$ , let  $I_{\mathbf{w}}$  be a diagonal matrix such that  $I_{\mathbf{w},ii} = \mathbb{I}\{w_i > 0\}$ . For a matrix  $A$ , let  $A_{\mathbf{w}} = I_{\mathbf{w}} A I_{\mathbf{w}}$  and  $\Lambda_A$  be a diagonal matrix with the same diagonal as  $A$ .

## 4 CMCB with Full-Information Feedback (CMCB-FI)

We start with CMCB with full-information feedback (CMCB-FI). In this setting, at each timestep  $t$ , the learner selects  $\mathbf{w}_t \in \Delta_d$  and observes the random reward  $\theta_{t,i}$  for all  $i \in [d]$ . CMCB-FI provides an online learning model for the celebrated Markowitz [25, 16] problem in finance, where investors select portfolios and can observe the prices of all stocks at the end of the trading days.

Below, we propose an optimal Mean-Covariance Empirical algorithm (MC-Empirical) for CMCB-FI, and provide a novel regret analysis that fully characterizes how an arbitrary covariance structure affects the regret performance. We also present a matching lower bound for CMCB-FI to demonstrate the optimality of MC-Empirical.

### 4.1 Algorithm for CMCB-FI

Algorithm 1 shows the detailed steps of MC-Empirical. Specifically, at each timestep  $t$ , we use the empirical mean  $\hat{\boldsymbol{\theta}}_t$  and covariance  $\hat{\Sigma}_t$  to estimate  $\boldsymbol{\theta}^*$  and  $\Sigma^*$ , respectively. Then, we form  $\hat{f}_t(\mathbf{w}) = \mathbf{w}^\top \hat{\boldsymbol{\theta}}_t - \rho \mathbf{w}^\top \hat{\Sigma}_t \mathbf{w}$ , an empirical mean-covariance function of  $\mathbf{w} \in \Delta_d$ , and always choose the action with the maximum empirical objective value.

Although MC-Empirical appears to be intuitive, its analysis is highly non-trivial due to covariance-based risk in the objective. In this case, a naive universal bound cannot characterize the impact of covariance, and prior gap-dependent analysis (e.g., [13, 27]) cannot be applied to solve our continuous space analysis with gap approximating to zero. Instead, we develop two novel techniques to handle the covariance, including using the actual covariance to analyze the confidence region of the expected rewards, and exploiting the empirical information of the sampling strategy to bound the

risk gap between selected actions and the optimal one. Different from prior works [13, 27], which assume a prior knowledge on covariance or only focus on the independent and positively-related cases, our analysis does not require extra knowledge of covariance and explicitly quantifies the effects of arbitrary covariance structures. The regret performance of MC-Empirical is summarized in Theorem 1.

**Theorem 1** (Upper Bound for CMCB-FI). *Consider the continuous mean-covariance bandits with full-information feedback (CMCB-FI). For any  $T > 0$ , algorithm MC-Empirical (Algorithm 1) achieves an expected cumulative regret bounded by*

$$O\left(\left(\min\left\{\sqrt{\mathbf{w}^{*\top}\Sigma^*\mathbf{w}^*} + \rho^{-\frac{1}{2}}\sqrt{\theta_{\max}^* - \theta_{\min}^*}, \sqrt{\Sigma_{\max}^*}\right\} + \rho\right)\ln T\sqrt{T}\right), \quad (2)$$

where  $\theta_{\max}^* = \max_{i \in [d]} \theta_i^*$ ,  $\theta_{\min}^* = \min_{i \in [d]} \theta_i^*$  and  $\Sigma_{\max}^* = \max_{i \in [d]} \Sigma_{ii}^*$ .

*Proof sketch.* Let  $D_t$  be the diagonal matrix which takes value  $t$  at each diagonal entry. We first build confidence intervals for the expected rewards of actions and the covariance as  $|\mathbf{w}^\top \boldsymbol{\theta}^* - \mathbf{w}^\top \hat{\boldsymbol{\theta}}_{t-1}| \leq p_t(\mathbf{w}) \triangleq c_1 \sqrt{\ln t} \sqrt{\mathbf{w}^\top D_{t-1}^{-1} (\lambda \Lambda_{\Sigma^*} D_{t-1} + \sum_{s=1}^{t-1} \Sigma^*) D_{t-1}^{-1} \mathbf{w}}$  and  $|\Sigma_{ij}^* - \hat{\Sigma}_{ij,t-1}| \leq q_t \triangleq c_2 \frac{\ln t}{\sqrt{t-1}}$ . Here  $\lambda = \frac{\mathbf{w}^{*\top} \Sigma^* \mathbf{w}^*}{\Sigma_{\max}^*}$  and  $c_1, c_2$  are positive constants. Then, we obtain the confidence interval of  $f(\mathbf{w})$  as  $|\hat{f}_{t-1}(\mathbf{w}) - f(\mathbf{w})| \leq r_t(\mathbf{w}) \triangleq p_t(\mathbf{w}) + \rho \mathbf{w}^\top Q_t \mathbf{w}$ , where  $Q_t$  is a matrix with all entries equal to  $q_t$ . Since algorithm MC-Empirical always plays the empirical best action, we have  $f(\mathbf{w}^*) - f(\mathbf{w}_t) \leq \hat{f}_{t-1}(\mathbf{w}^*) + r_t(\mathbf{w}^*) - f(\mathbf{w}_t) \leq \hat{f}_{t-1}(\mathbf{w}_t) + r_t(\mathbf{w}^*) - f(\mathbf{w}_t) \leq r_t(\mathbf{w}^*) + r_t(\mathbf{w}_t)$ . Plugging the definitions of  $f(\mathbf{w})$  and  $r_t(\mathbf{w})$ , we have

$$-\Delta_{\theta^*} + \rho \left( \mathbf{w}_t^\top \Sigma^* \mathbf{w}_t - \mathbf{w}^{*\top} \Sigma^* \mathbf{w}^* \right) \leq f(\mathbf{w}^*) - f(\mathbf{w}_t) \stackrel{(a)}{\leq} \frac{c_3 \ln t \left( \sqrt{\mathbf{w}^{*\top} \Sigma^* \mathbf{w}^*} + \sqrt{\mathbf{w}_t^\top \Sigma^* \mathbf{w}_t} + \rho \right)}{\sqrt{t-1}}, \quad (3)$$

where  $\Delta_{\theta^*} = \theta_{\max}^* - \theta_{\min}^*$  and  $c_3$  is a positive constant. Since our goal is to bound the regret  $f(\mathbf{w}^*) - f(\mathbf{w}_t)$  and in inequality (a) only the  $\sqrt{\mathbf{w}_t^\top \Sigma^* \mathbf{w}_t}$  term is a variable, the challenge falls on bounding  $\mathbf{w}_t^\top \Sigma^* \mathbf{w}_t$ . Note that the left-hand-side of Eq. (3) is linear with respect to  $\mathbf{w}_t^\top \Sigma^* \mathbf{w}_t$  and the right-hand-side only contains  $\sqrt{\mathbf{w}_t^\top \Sigma^* \mathbf{w}_t}$ . Then, using the property of sampling strategy on  $\mathbf{w}_t$ , i.e., Eq. (3), again, after some algebraic analysis, we obtain  $\mathbf{w}_t^\top \Sigma^* \mathbf{w}_t \leq c_4 (\mathbf{w}^{*\top} \Sigma^* \mathbf{w}^* + \frac{1}{\rho} \Delta_{\theta^*} + \frac{1}{\rho} \sqrt{\frac{\ln t}{t-1}} \sqrt{\mathbf{w}^{*\top} \Sigma^* \mathbf{w}^*} + \frac{\ln t}{\sqrt{t-1}} + \frac{\ln t}{\rho^2(t-1)})$  for some constant  $c_4$ . Plugging it into inequality (a) and doing a summation over  $t$ , we obtain the theorem.  $\square$

**Remark 1.** As we will show in Section 4.2, this  $O(\sqrt{T})$  regret matches the lower bound up to a logarithmic factor. Moreover, Theorem 1 fully characterizes how an *arbitrary* covariance structure impacts the regret bound. To see this, note that in Eq. (2), under the  $\min$  operation, the first  $\sqrt{\mathbf{w}^{*\top} \Sigma^* \mathbf{w}^*}$ -related term dominates under reasonable  $\rho$ , and shrinks from positive to negative correlation, which implies that the more the base arms are negatively (positively) correlate, the lower (higher) regret the learner suffers. The intuition behind is that the negative (positive) correlation diversifies (intensifies) the risk of estimation error and narrows (enlarges) the confidence region for the expected reward of an action, which leads to a reduction (an increase) of regret.

Also note that when  $\rho = 0$ , the CMCB-FI problem reduces to a  $d$ -armed bandit problem with full-information feedback, and Eq. (2) becomes  $\tilde{O}(\sqrt{\Sigma_{\max}^* T})$ . For this degenerated case, the optimal gap-dependent regret is  $O(\frac{\Sigma_{\max}^*}{\Delta})$  for constant gap  $\Delta > 0$ . By setting  $\Delta = \sqrt{\Sigma_{\max}^* / T}$  at this gap-dependent result, one obtains the optimal gap-independent regret  $O(\sqrt{\Sigma_{\max}^* T})$ . Hence, when  $\rho = 0$ , Eq. (2) still offers a tight gap-independent regret bound.

## 4.2 Lower Bound for CMCB-FI

Now we provide a regret lower bound for CMCB-FI, which demonstrates that the  $O(\sqrt{T})$  regret of MC-Empirical is in fact optimal (up to a logarithmic factor).

Since CMCB-FI considers full-information feedback and continuous decision space where the reward gap  $\Delta$  (between the optimal action and the nearest optimal action) approximates to zero, existing lower bound analysis for linear [10, 11] or discrete [21, 13, 27] bandit problems cannot be applied to this problem.

---

**Algorithm 2** MC-UCB

---

- |  |   |
|--|---|
| 1: <b>Input:</b> $\rho > 0$ , $c \in (0, \frac{1}{2}]$ and regularization parameter $\lambda > 0$ .  | 6: $\mathbf{w}_t \leftarrow \underset{\mathbf{w} \in \Delta_d^c}{\operatorname{argmax}} (\mathbf{w}^\top \hat{\boldsymbol{\theta}}_{t-1} + E_t(\mathbf{w}) - \rho \mathbf{w}^\top \hat{\Sigma}_t \mathbf{w})$ |
| 2: Initialize: $\forall i \in [d]$ , pull $e_i$ that has 1 at the $i$ -th entry and 0 elsewhere. $\forall i, j \in [d], i \neq j$ , pull $e_{ij}$ that has $\frac{1}{2}$ at the $i$ -th and the $j$ -th entries, and 0 elsewhere. Update $N_{ij}(d^2)$ , $\forall i, j \in [d]$ , $\hat{\boldsymbol{\theta}}_{d^2}$ and $\hat{\Sigma}_{d^2}$ . | 7: Pull $\mathbf{w}_t$ and observe all $\theta_{t,i}$ s.t. $w_{t,i} > 0$  |
| 3: <b>for</b> $t = d^2 + 1, \dots$ <b>do</b>   | 8: $J_{t,ij} \leftarrow \mathbb{I}\{w_{t,i}, w_{t,j} > 0\}$ , $\forall i, j \in [d]$  |
| 4: $\underline{\Sigma}_{t,ij} \leftarrow \hat{\Sigma}_{t-1,ij} - g_{ij}(t)$  | 9: $N_{ij}(t) \leftarrow N_{ij}(t-1) + J_{t,ij}$ , $\forall i, j \in [d]$   |
| 5: $\bar{\Sigma}_{t,ij} \leftarrow \hat{\Sigma}_{t-1,ij} + g_{ij}(t)$  | 10: $\hat{\theta}_{t,i}^* \leftarrow \frac{\sum_{s=1}^t J_{t,ii} \theta_{s,i}}{N_{ii}(t)}$ , $\forall i \in [d]$  |
|  | 11: $\hat{\Sigma}_{t,ij} \leftarrow \frac{\sum_{s=1}^t J_{t,ij} (\theta_{s,i} - \hat{\theta}_{t,i}^*) (\theta_{s,j} - \hat{\theta}_{t,j}^*)}{N_{ij}(t)}$ , $\forall i, j \in [d]$                             |
|  | 12: <b>end for</b>  |
- 

To tackle this challenge, we contribute a new analytical procedure to establish the lower bound for continuous and full-information bandit problems from the Bayesian perspective. The main idea is to construct an instance distribution, where  $\boldsymbol{\theta}^*$  is drawn from a well-chosen prior Gaussian distribution. After  $t$  pulls the posterior of  $\boldsymbol{\theta}^*$  is still Gaussian with a mean vector  $\mathbf{u}_t$  related to sample outcomes. Since the hindsight strategy simply selects the action which maximizes the mean-covariance function with respect to  $\boldsymbol{\theta}^*$  while a feasible strategy can only utilize the sample information ( $\mathbf{u}_t$ ), we show that any algorithm must suffer  $\Omega(\sqrt{T})$  regret due to the gap between random  $\boldsymbol{\theta}^*$  and its mean  $\mathbf{u}_t$ . Theorem 2 below formally states this lower bound.

**Theorem 2** (Lower Bound for CMCB-FI). *There exists an instance distribution of the continuous mean-covariance bandits with full-information feedback problem (CMCB-FI), for which any algorithm has an expected cumulative regret bounded by  $\Omega(\sqrt{T})$ .*

**Remark 2.** This parameter-free lower bound demonstrates that the regret upper bound (Theorem 1) of MC-Empirical is optimal (within a logarithmic factor), since under the constructed instance distribution, Theorem 1 also implies a matching  $\tilde{O}(\sqrt{T})$  parameter-free result, i.e., when  $\rho = 1/\sqrt{T}$ , Eq. (2) becomes  $\tilde{O}((\sqrt{\Sigma_{\max}^*} + 1/\sqrt{T})\sqrt{T}) = \tilde{O}(\sqrt{T})$ . Unlike discrete bandit problems [21, 13, 27] where the optimal regret is usually  $\frac{\log T}{\Delta}$  for constant gap  $\Delta > 0$ , CMCB-FI has continuous decision space with gap  $\Delta \rightarrow 0$  and a polylogarithmic regret is not achievable in general. In such continuous bandit literature [20, 10, 11], the parameter ( $\boldsymbol{\theta}^*$ ,  $\Sigma^*$  and  $\rho$ ) dependent lower bound is an open problem.

## 5 CMCB with Semi-Bandit Feedback (CMCB-SB)

In many practical tasks, the learner may not be able to simultaneously select (place positive weights on) all options and observe full information. Instead, the weight of each option is usually lower bounded and cannot be arbitrarily small. As a result, the learner only selects a subset of options and obtains their feedback, e.g., company investments [14] on multiple business.

Motivated by such tasks, in this section we consider the CMCB problem with semi-bandit feedback (CMCB-SB), where the decision space is a restricted probability simplex  $\Delta_d^c = \{\mathbf{w} \in \mathbb{R}^d : w_i = 0 \text{ or } c \leq w_i \leq 1, \forall i \in [d] \text{ and } \sum_i w_i = 1\}$  for some constant  $0 < c \leq \frac{1}{2}$ .<sup>2</sup> In this scenario, at timestep  $t$ , the learner selects  $\mathbf{w}_t \in \Delta_d^c$  and only observes the rewards  $\{\theta_{t,i} : w_{t,i} \geq c\}$  from the base arms that are placed positive weights on. Below, we propose the Mean-Covariance Upper Confidence Bound algorithm (MC-UCB) for CMCB-SB, and provide a regret lower bound, which shows that MC-UCB achieves the optimal performance with respect to  $T$ .

### 5.1 Algorithm for CMCB-SB

Algorithm MC-UCB for CMCB-SB is described in Algorithm 2. The main idea is to use the optimistic covariance to construct a confidence region for the expected reward of an action and calculate an upper confidence bound of the mean-covariance function, and then select the action with the maximum optimistic mean-covariance value.

---

<sup>2</sup>When  $c > \frac{1}{2}$ , the learner can only place all weight on one option, and the problem trivially reduces to the mean-variance bandit setting [28, 36]. In this case, our Theorem 3 still provides a tight gap-independent bound.



In Algorithm 2,  $N_{ij}(t)$  denotes the number of times  $w_{s,i}, w_{s,j} > 0$  occurs among timestep  $s \in [t]$ .  $J_{t,ij}$  is an indicator variable that takes value 1 if  $w_{t,i}, w_{t,j} > 0$  and 0 otherwise.  $D_t$  is a diagonal matrix such that  $D_{t,ii} = N_{ii}(t)$ . In Line 2, we update the number of observations by  $N_{ii}(d^2) \leftarrow 2d-1$  for all  $i \in [d]$  and  $N_{ij}(d^2) \leftarrow 2$  for all  $i, j \in [d], i \neq j$  (due to the initialized  $d^2$  pulls), and calculate the empirical mean  $\hat{\theta}_{d^2}^*$  and empirical covariance  $\hat{\Sigma}_{d^2}^*$  using the equations in Lines 10,11.

For any  $t > 1$  and  $i, j \in [d]$ , we define the confidence radius of covariance  $\Sigma_{ij}^*$  as  $g_{ij}(t) \triangleq 16 \left( \frac{3 \ln t}{N_{ij}(t-1)} \vee \sqrt{\frac{3 \ln t}{N_{ij}(t-1)}} \right) + \sqrt{\frac{48 \ln^2 t}{N_{ij}(t-1)N_{ii}(t-1)}} + \sqrt{\frac{36 \ln^2 t}{N_{ij}(t-1)N_j(t-1)}}$ , and the confidence region for the expected reward  $\mathbf{w}^\top \boldsymbol{\theta}^*$  of action  $\mathbf{w}$  as

$$E_t(\mathbf{w}) \triangleq \sqrt{2\beta(\delta_t) \left( \mathbf{w}^\top D_{t-1}^{-1} \left( \lambda \Lambda_{\bar{\Sigma}_t} D_{t-1} + \sum_{s=1}^{t-1} \bar{\Sigma}_{s, \mathbf{w}_s} \right) D_{t-1}^{-1} \mathbf{w} \right)},$$

where  $\lambda > 0$  is the regularization parameter,  $\beta(\delta_t) = \ln\left(\frac{1}{\delta_t}\right) + d \ln \ln t + \frac{d}{2} \ln\left(1 + \frac{e}{\lambda}\right)$  is the confidence term and  $\delta_t = \frac{1}{t \ln^2 t}$  is the confidence parameter. At each timestep  $t$ , algorithm MC-UCB calculates the upper confidence bound of  $f(\mathbf{w})$  using  $g_{ij}(t)$  and  $E_t(\mathbf{w})$ , and selects the action  $\mathbf{w}_t$  that maximizes this upper confidence bound. Then, the learner observes rewards  $\theta_{t,i}$  with  $w_{t,i} > 0$  and update the statistical information according to the feedback.

In regret analysis, unlike [13] which uses a universal upper bound to analyze confidence intervals, we incorporate the estimated covariance into the confidence region for the expected reward of an action, which enables us to derive tighter regret bound and explicitly quantify the impact of the covariance structure on algorithm performance. We also contribute a new technique for handling the challenge raised by having different numbers of observations among base arms, in order to obtain an optimal  $O(\sqrt{T})$  regret (here prior gap-dependent analysis [13, 27] still cannot be applied to solve this continuous problem). Theorem 3 gives the regret upper bound of algorithm MC-UCB.

**Theorem 3** (Upper Bound for CMCB-SB). *Consider the continuous mean-covariance bandits with semi-bandit feedback problem (CMCB-SB). Then, for any  $T > 0$ , algorithm MC-UCB (Algorithm 2) with regularization parameter  $\lambda > 0$  has an expected cumulative regret bounded by*

$$O\left(\sqrt{L(\lambda)(\|\Sigma^*\|_+ + d^2)d \ln^2 T \cdot T} + \rho d \ln T \sqrt{T}\right),$$

where  $L(\lambda) = (\lambda + 1)(\ln(1 + \lambda^{-1}) + 1)$  and  $\|\Sigma^*\|_+ = \sum_{i,j \in [d]} (\Sigma_{ij}^* \vee 0)$  for any  $i, j \in [d]$ .

**Remark 3.** Theorem 3 captures the effects of covariance structures in CMCB-SB, i.e., positive correlation renders a larger  $\|\Sigma^*\|_+$  factor than the negative correlation or independent case, since the covariance influences the rate of estimate concentration for the expected rewards of actions. The regret bound for CMCB-SB has a heavier dependence on  $d$  than that for CMCB-FI. This matches the fact that semi-bandit feedback only reveals rewards of the queried dimensions, and provides less information than full-information feedback in terms of observable dimensions.

## 5.2 Lower Bound for CMCB-SB

In this subsection, we establish a lower bound for CMCB-SB, and show that algorithm MC-UCB achieves the optimal regret with respect to  $T$  up to logarithmic factors.

The insight of the lower bound analysis is to construct two instances with a gap in the expected reward vector  $\boldsymbol{\theta}^*$ , where the optimal actions under these two instances place positive weights on different base arms. Then, when the gap is set to  $\sqrt{\ln T/T}$ , any algorithm must suffer  $\Omega\left(\sqrt{T \ln T}\right)$  regret for differentiating these two instances. Theorem 4 summarizes the lower bound for CMCB-SB.

**Theorem 4** (Lower Bound for CMCB-SB). *There exists an instance distribution of the continuous mean-covariance bandits with semi-bandit feedback (CMCB-SB) problem, for which any algorithm has an expected cumulative regret bounded by  $\Omega\left(\sqrt{cdT}\right)$ .*

**Remark 4.** Theorem 4 demonstrates that the regret upper bound (Theorem 3) of MC-UCB is optimal with respect to  $T$  (up to logarithmic factors). Similar to CMCB-FI, CMCB-SB considers continuous

---

**Algorithm 3** MC-ETE

---

1: <b>Input:</b> $\rho > 0$ , $\tilde{d} = \frac{d(d+1)}{2}$ and design action	12: $\mathbf{y}_{N_\pi(t)} \leftarrow (y_{N_\pi(t),1}, \dots, y_{N_\pi(t),\tilde{d}})^\top$
set $\pi = \{\mathbf{v}_1, \dots, \mathbf{v}_{\tilde{d}}\}$ .	13: $\hat{\mathbf{y}}_t \leftarrow \frac{\sum_{s=1}^{N_\pi(t)} \mathbf{y}_s}{N_\pi(t)}$
2: Initialize: $N_\pi(0) \leftarrow 0$ . $t \leftarrow 1$ .	14: $\hat{z}_{t,k} = \frac{\sum_{s=1}^{N_\pi(t)} (y_{s,k} - \hat{y}_{t,k})^2}{N_\pi(t)}, \forall k \in [\tilde{d}]$
3: <b>Repeat</b> lines 4-22:	15: $\hat{\mathbf{z}}_t \leftarrow (\hat{z}_{t,1}, \dots, \hat{z}_{t,\tilde{d}})^\top$
4: <b>if</b> $N_\pi(t-1) > t^{\frac{2}{3}}/d$ <b>then</b>	16: $\hat{\boldsymbol{\theta}}_t \leftarrow B_\pi^+ \hat{\mathbf{y}}_t$
5: $\mathbf{w}_t = \operatorname{argmax}_{\mathbf{w} \in \Delta_d} (\mathbf{w}^\top \hat{\boldsymbol{\theta}}_{t-1}^* - \rho \mathbf{w}^\top \hat{\Sigma}_{t-1} \mathbf{w})$	17: $\hat{\boldsymbol{\sigma}}_t \leftarrow C_\pi^+ \hat{\mathbf{z}}_t$
6: $t \leftarrow t + 1$	18: Reshape $\hat{\boldsymbol{\sigma}}_t$ to $d \times d$ matrix $\hat{\Sigma}_t$
7: <b>else</b>	19: <b>end if</b>
8: $N_\pi(t) \leftarrow N_\pi(t-1) + 1$	20: $t \leftarrow t + 1$
9: <b>for</b> $k = 1, \dots, \tilde{d}$ <b>do</b>	21: <b>end for</b>
10:     Pull $\mathbf{v}_k$ and observe $y_{N_\pi(t),k}$	22: <b>end if</b>
11: <b>if</b> $k = \tilde{d}$ <b>then</b>	

---

decision space with  $\Delta \rightarrow 0$ , and thus the lower bound differs from those gap-dependent results  $\frac{\log T}{\Delta}$  in discrete bandit problems [21, 13, 27]. Our lower bound shows that for CMCB-SB, no improvement upon  $O(\sqrt{T})$  regret is possible in general.

## 6 CMCB with Full-Bandit Feedback (CMCB-FB)

In this section, we further study the CMCB problem with full-bandit feedback (CMCB-FB), where at timestep  $t$ , the learner selects  $\mathbf{w}_t \in \Delta_d$  and only observes the weighted sum of random rewards, i.e.,  $y_t = \mathbf{w}_t^\top \boldsymbol{\theta}_t$ . This setting models many real-world decision making tasks, where the learner can only attain an aggregate feedback from the chosen options, such as clinical trials [33].

### 6.1 Algorithm for CMCB-FB

We propose the Mean-Covariance Exploration-Then-Exploitation algorithm (MC-ETE) for CMCB-FB in Algorithm 3. Specifically, we first choose a design action set  $\pi = \{\mathbf{v}_1, \dots, \mathbf{v}_{\tilde{d}}\}$  which contains  $\tilde{d} = d(d+1)/2$  actions and satisfies that  $B_\pi = (\mathbf{v}_1^\top; \dots; \mathbf{v}_{\tilde{d}}^\top)$  and  $C_\pi = (v_{1,1}^2, \dots, v_{1,d}^2, 2v_{1,1}v_{1,2}, \dots, 2v_{1,d-1}v_{1,d}; \dots; v_{\tilde{d},1}^2, \dots, v_{\tilde{d},d}^2, 2v_{\tilde{d},1}v_{\tilde{d},2}, \dots, 2v_{\tilde{d},d-1}v_{\tilde{d},d})$  are of full column rank. We also denote their Moore-Penrose inverses by  $B_\pi^+$  and  $C_\pi^+$ , and it holds that  $B_\pi^+ B_\pi = I^{d \times d}$  and  $C_\pi^+ C_\pi = I^{\tilde{d} \times \tilde{d}}$ . There exist more than one feasible  $\pi$ , and for simplicity and good performance we choose  $\mathbf{v}_1, \dots, \mathbf{v}_d$  as standard basis vectors in  $\mathbb{R}^d$  and  $\{\mathbf{v}_{d+1}, \dots, \mathbf{v}_{\tilde{d}}\}$  as the set of all  $\binom{d}{2}$  vectors where each vector has two entries equal to  $\frac{1}{2}$  and others equal to 0.

In an exploration round (Lines 8-21), we pull the designed actions in  $\pi$  and maintain their empirical rewards and variances. Through linear transformation by  $B_\pi^+$  and  $C_\pi^+$ , we obtain the estimators of the expected rewards and covariance of base arms (Lines 16-17). When the estimation confidence is high enough, we exploit the attained information to select the empirical best action (Lines 5). Theorem 5 presents the regret guarantee of MC-ETE.

**Theorem 5** (Upper Bound for CMCB-FB). *Consider the continuous mean-covariance bandits with full-bandit feedback problem (CMCB-FB). Then, for any  $T > 0$ , algorithm MC-ETE (Algorithm 3) achieves an expected cumulative regret bounded by*

$$O\left(Z(\rho, \pi) \sqrt{d(\ln T + d^2)} \cdot T^{\frac{2}{3}} + d\Delta_{\max} \cdot T^{\frac{2}{3}}\right),$$

where  $Z(\rho, \pi) = \max_{\mathbf{w} \in \Delta_d} (\sqrt{\mathbf{w}^\top B_\pi^+ \Sigma_\pi^* (B_\pi^+)^\top \mathbf{w} + \rho \|C_\pi^+\|})$ ,  $\Sigma_\pi^* = \operatorname{diag}(\mathbf{v}_1^\top \Sigma^* \mathbf{v}_1, \dots, \mathbf{v}_{\tilde{d}}^\top \Sigma^* \mathbf{v}_{\tilde{d}})$ ,  $\|C_\pi^+\| = \max_{i \in [\tilde{d}]} \sum_{j \in [\tilde{d}]} |C_{\pi,ij}^+|$  and  $\Delta_{\max} = f(\mathbf{w}^*) - \min_{\mathbf{w} \in \Delta_d} f(\mathbf{w})$ .

**Remark 5.** The choice of  $\pi$  will affect the regret factor  $\Sigma_\pi^*$  contained in  $Z(\rho, \pi)$ . Under our construction,  $\Sigma_\pi^*$  can be regarded as a uniform representation of covariance  $\Sigma^*$ , and thus our regret bound demonstrates how the learning performance is influenced by the covariance structure, i.e., negative (positive) correlation shrinks (enlarges) the factor and leads to a lower (higher) regret.






Figure 1: Experiments for CMCB-FI, CMCB-SB and CMCB-FB on the synthetic and real-world datasets.

**Discussion on the ETE strategy.** In contrast to common ETE-type algorithms, MC-ETE requires novel analytical techniques in handling the transformed estimate concentration while preserving the covariance information in regret bounds. In analysis, we build a novel concentration using key matrices  $B_\pi^+$  and  $C_\pi^+$  to adapt to the actual covariance structure, and construct a super-martingale which takes the aggregate noise in an exploration round as analytical basis to prove the concentration. These techniques allow us to capture the correlations in the results, and are new compared to both the former FI/SB settings and covariance-related bandit literature [35, 13, 27].

In fact, under the full-bandit feedback, it is highly challenging to estimate the covariance without using a fixed exploration (i.e., ETE) strategy. Note that even for its simplified offline version, where one uses given (non-fixed) full-bandit data to estimate the covariance, there is *no available solution* in the statistics literature to our best knowledge. Hence, for such online tasks with severely limited feedback, ETE is the most viable strategy currently available, as used in many partial observation works [23, 7, 8]. We remark that our contribution in this setting focuses on designing a practical solution and deriving regret guarantees which explicitly characterize the correlation impacts. The lower bound for CMCB-FB remains open, which we leave for future work.

## 7 Experiments

In this section, we present experimental results for our algorithms on both synthetic and real-world [22] datasets. For the synthetic dataset, we set  $\theta^* = [0.2, 0.3, 0.2, 0.2, 0.2]^\top$ , and  $\Sigma^*$  has all diagonal entries equal to 1 and all off-diagonal entries equal to  $-0.05$ . For the real-world dataset, we use an open dataset *US Funds from Yahoo Finance* on Kaggle [22], which provides financial data of 1680 ETF funds in 2010-2017. We select five funds and generate a stochastic distribution ( $\theta^*$  and  $\Sigma^*$ ) from the data of returns (since we study a stochastic bandit problem). For both datasets, we set  $d = 5$  and  $\rho \in \{0.1, 10\}$ . The random reward  $\theta_t$  is drawn i.i.d. from Gaussian distribution  $\mathcal{N}(\theta^*, \Sigma^*)$ . We perform 50 independent runs for each algorithm and show the average regret and 95% confidence interval across runs,<sup>3</sup> with logarithmic y-axis for clarity of magnitude comparison.

<sup>3</sup>In some cases, since algorithms are doing similar procedures (e.g., in Figures 1(c),1(f), the algorithms are exploring the designed actions) and have low performance variance, the confidence intervals are narrow and indistinguishable.

**(CMCB-FI)** We compare our algorithm MC-Empirical with two algorithms OGD [17] and LinearFI. OGD (Online Gradient Descent) [17] is designed for general online convex optimization with also a  $O(\sqrt{T})$  regret guarantee, but its result cannot capture the covariance impacts as ours. LinearFI is a linear adaption of MC-Empirical that only aims to maximize the expected rewards. Figures 1(a),1(d) show that our MC-Empirical enjoys multiple orders of magnitude reduction in regret compared to the benchmarks, since it efficiently exploits the empirical observations to select actions and well handles the covariance-based risk. In particular, the performance superiority of MC-Empirical over OGD demonstrates that our sample strategy sufficiently utilize the observed information than conventional gradient descent based policy.

**(CMCB-SB)** For CMCB-SB, we compare MC-UCB with two adaptations of OLS-UCB [13] (state-of-the-art for combinatorial bandits with covariance), named MC-UCB- $\Gamma$  and OLS-UCB-C. MC-UCB- $\Gamma$  uses the confidence region with a universal covariance upper bound  $\Gamma$ , instead of the adapting one used in our MC-UCB. OLS-UCB-C directly adapts OLS-UCB [13] to the continuous decision space and only considers maximizing the expected rewards in its objective. As shown in Figures 1(b),1(e), MC-UCB achieves the lowest regret since it utilizes the covariance information to accelerate the estimate concentration. Due to lack of a covariance-adapting confidence interval, MC-UCB- $\Gamma$  shows an inferior regret performance than MC-UCB, and OLS-UCB-C suffers the highest regret due to its ignorance of risk.

**(CMCB-FB)** We compare MC-ETE with two baselines, OGD-ETE, which adopts OGD [17] during the exploitation phase, and LinearFB, which only investigates the expected reward maximization. From Figures 1(c),1(f), one can see that, MC-ETE achieves the best regret performance due to its effective estimation of the covariance-based risk and efficiency in exploitation. Due to the inefficiency of gradient descent based policy in utilizing information, OGD-ETE has a higher regret than MC-ETE, whereas LinearFB shows the worst performance owing to the unawareness of the risk.

## 8 Conclusion and Future Work

In this paper, we propose a novel continuous mean-covariance bandit (CMCB) model, which investigates the reward-risk trade-off measured by option correlation. Under this model, we consider three feedback settings, i.e., full-information, semi-bandit and full-bandit feedback, to formulate different real-world reward observation scenarios. We propose novel algorithms for CMCB to achieve the optimal regrets (within logarithmic factors), and provide lower bounds for the problems to demonstrate our optimality. We also present empirical evaluations to show the superior performance of our algorithms. To our best knowledge, this is the first work to fully characterize the impacts of arbitrary covariance structures on learning performance for risk-aware bandits. There are several interesting directions for future work. For example, how to design an adaptive algorithm for CMCB-FB is a challenging open problem, and the lower bound for CMCB-FB is also worth further investigation.

## Acknowledgments and Disclosure of Funding

The work of Yihan Du and Longbo Huang is supported in part by the Technology and Innovation Major Project of the Ministry of Science and Technology of China under Grant 2020AAA0108400 and 2020AAA0108403.

## References

- [1] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in Neural Information Processing Systems*, 24:2312–2320, 2011.
- [2] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on Learning Theory*, pages 39–1, 2012.
- [3] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002.
- [4] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.
- [5] Kirt C Butler and Dale L Domian. Risk, diversification, and the investment horizon. *Journal of Portfolio Management*, 17(3):41, 1991.

- [6] Asaf Cassel, Shie Mannor, and Assaf Zeevi. A general approach to multi-armed bandits under risk criteria. In *Conference on Learning Theory*, pages 1295–1306, 2018.
- [7] Sougata Chaudhuri and Ambuj Tewari. Phased exploration with greedy exploitation in stochastic combinatorial partial monitoring games. In *Advances in Neural Information Processing Systems*, pages 2433–2441. 2016.
- [8] Lixing Chen, Jie Xu, and Zhuo Lu. Contextual combinatorial multi-armed bandits with volatile arms and submodular reward. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [9] Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999.
- [10] Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. In *Conference on Learning Theory*, 2008.
- [11] Varsha Dani, Sham M Kakade, and Thomas Hayes. The price of bandit information for online optimization. In *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2008.
- [12] Yahel David and Nahum Shimkin. Pure exploration for max-quantile bandits. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 556–571. Springer, 2016.
- [13] Rémy Degenne and Vianney Perchet. Combinatorial semi-bandit with known covariance. In *Advances in Neural Information Processing Systems*, pages 2972–2980, 2016.
- [14] A Diadin. Analysis of a trade company investment project. *Business and Economics*, page 104, 2019.
- [15] Nicolas Galichet, Michele Sebag, and Olivier Teytaud. Exploration vs exploitation vs safety: Risk-aware multi-armed bandits. In *Asian Conference on Machine Learning*, pages 245–260, 2013.
- [16] Stephan M Gasser, Margarethe Rammerstorfer, and Karl Weinmayer. Markowitz revisited: Social portfolio engineering. *European Journal of Operational Research*, 258(3):1181–1190, 2017.
- [17] Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.
- [18] Anmol Kagrecha, Jayakrishnan Nair, and Krishna Jagannathan. Distribution oblivious, risk-aware algorithms for multi-armed bandits with unbounded rewards. *Advances in Neural Information Processing Systems*, 32:11272–11281, 2019.
- [19] Abbas Kazerouni, Mohammad Ghavamzadeh, Yasin Abbasi-Yadkori, and Benjamin Van Roy. Conservative contextual linear bandits. In *Advances in Neural Information Processing Systems*, pages 3913–3922, 2017.
- [20] Robert Kleinberg. Nearly tight bounds for the continuum-armed bandit problem. In *Advances in Neural Information Processing Systems*, volume 17. MIT Press, 2005.
- [21] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- [22] Stefano Leone. Dataset: US funds dataset from yahoo finance. Kaggle, 2020. <https://www.kaggle.com/stefanoleone992/mutual-funds-and-etfs?select=ETFs.csv>.
- [23] Tian Lin, Bruno Abrahao, Robert Kleinberg, John Lui, and Wei Chen. Combinatorial partial monitoring game with linear feedback and its applications. In *International Conference on Machine Learning*, pages 901–909, 2014.
- [24] Andrea Locatelli, Maurilio Gutzeit, and Alexandra Carpentier. An optimal algorithm for the thresholding bandit problem. In *International Conference on Machine Learning*, pages 1690–1698. PMLR, 2016.

- [25] Harry M Markowitz et al. Portfolio selection. *Journal of Finance*, 7(1):77–91, 1952.
- [26] Joanne M McInerney and Tim S Roberts. Online learning: Social interaction and the creation of a sense of community. *Journal of Educational Technology & Society*, 7(3):73–81, 2004.
- [27] Pierre Perrault, Michal Valko, and Vianney Perchet. Covariance-adapting algorithm for semi-bandits with application to sparse outcomes. In *Conference on Learning Theory*, pages 3152–3184, 2020.
- [28] Amir Sani, Alessandro Lazaric, and Rémi Munos. Risk-aversion in multi-armed bandits. In *Advances in Neural Information Processing Systems*, pages 3275–3283, 2012.
- [29] Eric M Schwartz, Eric T Bradlow, and Peter S Fader. Customer acquisition via display advertising using multi-armed bandit experiments. *Marketing Science*, 36(4):500–522, 2017.
- [30] William 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.
- [31] Sattar Vakili and Qing Zhao. Mean-variance and value at risk in multi-armed bandit problems. In *The 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1330–1335. IEEE, 2015.
- [32] Sattar Vakili and Qing Zhao. Risk-averse multi-armed bandit problems under mean-variance measure. *IEEE Journal of Selected Topics in Signal Processing*, 10(6):1093–1111, 2016.
- [33] Sofía S Villar, Jack Bowden, and James Wason. Multi-armed bandit models for the optimal design of clinical trials: benefits and challenges. *Statistical Science: A Review Journal of the Institute of Mathematical Statistics*, 30(2):199, 2015.
- [34] Manfred K Warmuth and Dima Kuzmin. Online variance minimization. In *International Conference on Computational Learning Theory*, pages 514–528. Springer, 2006.
- [35] Manfred K Warmuth and Dima Kuzmin. Online variance minimization. *Machine Learning*, 87(1):1–32, 2012.
- [36] Qiuyu Zhu and Vincent YF Tan. Thompson sampling algorithms for mean-variance bandits. *International Conference on Machine Learning*, 2020.

## Appendix

### A Technical Lemmas

In this section, we introduce two technical lemmas which will be used in our analysis.

Lemmas 1 and 2 give the concentration guarantees of algorithm MC-UCB for CMCB-SB, which sets up a foundation for the concentration guarantees in CMCB-FI. For ease of notation, we use  $\Sigma$  for a shorthand of the covariance matrix  $\Sigma^*$  in Appendix.

**Lemma 1** (Concentration of Covariance for CMCB-SB). *Consider the CMCB-SB problem and algorithm MC-UCB (Algorithm 2). Define the event*

$$\mathcal{G}_t \triangleq \left\{ |\Sigma_{ij} - \hat{\Sigma}_{ij,t-1}| \leq 16 \left( \frac{3 \ln t}{N_{ij}(t-1)} \vee \sqrt{\frac{3 \ln t}{N_{ij}(t-1)}} \right) + \sqrt{\frac{61 \ln^2 t}{N_{ij}(t-1)N_i(t-1)}} + \sqrt{\frac{36 \ln^2 t}{N_{ij}(t-1)N_j(t-1)}} \right\}, \forall i, j \in [d]$$

For any  $t \geq 2$ , we have

$$\Pr[\mathcal{G}_t] \geq 1 - \frac{10d^2}{t^2}.$$

*Proof.* According to Proposition 2 in [27], we have that for any  $t \geq 2$  and  $i, j \in [d]$ ,

$$\Pr \left[ |\Sigma_{ij} - \hat{\Sigma}_{ij,t-1}| \leq 16 \left( \frac{3 \ln t}{N_{ij}(t-1)} \vee \sqrt{\frac{3 \ln t}{N_{ij}(t-1)}} \right) + \sqrt{\frac{61 \ln^2 t}{N_{ij}(t-1)N_i(t-1)}} + \sqrt{\frac{36 \ln^2 t}{N_{ij}(t-1)N_j(t-1)}} \right] \leq 1 - \frac{10}{t^2}.$$

Using a union bound on  $i, j \in [d]$ , we obtain Lemma 1.  $\square$

**Lemma 2** (Concentration of Means for CMCB-SB). *Consider the CMCB-SB problem and algorithm MC-UCB (Algorithm 2). Let  $0 < \lambda < 1$ , and define  $\delta_t = \frac{1}{t \ln^2 t}$  and  $\beta(\delta_t) = \ln(1/\delta_t) + d \ln \ln t + \frac{d}{2} \ln(1 + e/\lambda)$  for  $t \geq 2$ . Then, for any  $t \geq 2$  and  $\mathbf{w} \in \Delta_d^c$ , with probability at least  $1 - \delta_t$ , we have*

$$\left| \mathbf{w}^\top \boldsymbol{\theta}^* - \mathbf{w}^\top \hat{\boldsymbol{\theta}}_{t-1} \right| \leq \sqrt{2\beta(\delta_t)} \sqrt{\mathbf{w}^\top D_{t-1}^{-1} \left( \lambda \Lambda_\Sigma D_{t-1} + \sum_{s=1}^{t-1} \Sigma_{\mathbf{w}_s} \right) D_{t-1}^{-1} \mathbf{w}}.$$

Further define  $E_t(\mathbf{w}) = \sqrt{2\beta(\delta_t)} \sqrt{\mathbf{w}^\top D_{t-1}^{-1} (\lambda \Lambda_{\bar{\Sigma}_t} D_{t-1} + \sum_{s=1}^{t-1} \bar{\Sigma}_{s, \mathbf{w}_s}) D_{t-1}^{-1} \mathbf{w}}$ . Then, for any  $t \geq 2$  and  $\mathbf{w} \in \Delta_d^c$ , the event  $\mathcal{H}_t \triangleq \{|\mathbf{w}^\top \boldsymbol{\theta}^* - \mathbf{w}^\top \hat{\boldsymbol{\theta}}_{t-1}| \leq E_t(\mathbf{w})\}$  satisfies  $\Pr[\mathcal{H}_t \mid \mathcal{G}_t] \geq 1 - \delta_t$ .

*Proof.* The proof of Lemma 2 follows the analysis procedure in [13]. Specifically, assuming that event  $\mathcal{G}_t$  occurs, we have  $\bar{\Sigma}_{t,ij} \geq \Sigma_{ij}$  for any  $i, j \in [d]$  and

$$E_t(\mathbf{w}) \geq \sqrt{2\beta(\delta_t)} \sqrt{\mathbf{w}^\top D_{t-1}^{-1} \left( \lambda \Lambda_\Sigma D_{t-1} + \sum_{s=1}^{t-1} \Sigma_{\mathbf{w}_s} \right) D_{t-1}^{-1} \mathbf{w}}.$$

Hence, to prove Lemma 2, it suffices to prove that

$$\Pr \left[ \left| \mathbf{w}^\top \boldsymbol{\theta}^* - \mathbf{w}^\top \hat{\boldsymbol{\theta}}_{t-1} \right| > \sqrt{2\beta(\delta_t)} \sqrt{\mathbf{w}^\top D_{t-1}^{-1} \left( \lambda \Lambda_\Sigma D_{t-1} + \sum_{s=1}^{t-1} \Sigma_{\mathbf{w}_s} \right) D_{t-1}^{-1} \mathbf{w}} \right] \leq \delta_t. \quad (4)$$

Recall that  $N_i(t) = \sum_{s=1}^t \mathbb{I}\{w_{s,i} \geq c\}$  and  $D_t$  is a diagonal matrix such that  $D_{t,ii} = N_i(t)$  for any  $t > 0$ . For any  $\mathbf{w} \in \Delta_d^c$ , let  $I_{\mathbf{w}}$  denote the diagonal matrix such that  $I_{ii} = 1$  for any  $w_i \geq c$  and  $I_{jj} = 0$  for any  $w_j = 0$ , and let  $\Sigma_{\mathbf{w}} = I_{\mathbf{w}} \Sigma I_{\mathbf{w}}$ . Let  $\boldsymbol{\varepsilon}_t$  be the vector such that  $\boldsymbol{\eta}_t = \Sigma^{\frac{1}{2}} \boldsymbol{\varepsilon}_t$  for any  $t > 0$ . Then, we have for any  $\mathbf{w} \in \Delta_d^c$  that

$$\begin{aligned} \left| \mathbf{w}^\top (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{t-1}) \right| &= \left| -\mathbf{w}^\top D_{t-1}^{-1} \sum_{s=1}^{t-1} I_{\mathbf{w}_s} \Sigma^{\frac{1}{2}} \boldsymbol{\varepsilon}_s \right| \\ &= \left| -\mathbf{w}^\top D_{t-1}^{-1} \left( D + \sum_{s=1}^{t-1} \Sigma_{\mathbf{w}_s} \right)^{\frac{1}{2}} \left( D + \sum_{s=1}^{t-1} \Sigma_{\mathbf{w}_s} \right)^{-\frac{1}{2}} \sum_{s=1}^{t-1} I_{\mathbf{w}_s} \Sigma^{\frac{1}{2}} \boldsymbol{\varepsilon}_s \right| \\ &\leq \sqrt{\mathbf{w}^\top D_{t-1}^{-1} \left( D + \sum_{s=1}^{t-1} \Sigma_{\mathbf{w}_s} \right) D_{t-1}^{-1} \mathbf{w}} \cdot \left\| \sum_{s=1}^{t-1} I_{\mathbf{w}_s} \Sigma^{\frac{1}{2}} \boldsymbol{\varepsilon}_s \right\|_{(D + \sum_{s=1}^{t-1} \Sigma_{\mathbf{w}_s})^{-1}} \end{aligned}$$

Let  $S_t = \sum_{s=1}^{t-1} I_{\mathbf{w}_s} \Sigma^{\frac{1}{2}} \boldsymbol{\varepsilon}_s$ ,  $V_t = \sum_{s=1}^{t-1} \Sigma_{\mathbf{w}_s}$  and  $I_{D+V_t} = \frac{1}{2} \|S_t\|_{(D+V_t)^{-1}}^2$ . We get

$$\left\| \sum_{s=1}^{t-1} I_{\mathbf{w}_s} \Sigma^{\frac{1}{2}} \boldsymbol{\varepsilon}_s \right\|_{(D + \sum_{s=1}^{t-1} \Sigma_{\mathbf{w}_s})^{-1}} = \|S_t\|_{(D+V_t)^{-1}} = \sqrt{2I_{D+V_t}}.$$

Since  $D \preceq \lambda \Lambda_\Sigma D_{t-1}$ , we have

$$\begin{aligned} \left| \mathbf{w}^\top (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{t-1}) \right| &\leq \sqrt{\mathbf{w}^\top D_{t-1}^{-1} D D_{t-1}^{-1} \mathbf{w} + \mathbf{w}^\top D_{t-1}^{-1} \left( \sum_{s=1}^{t-1} \Sigma_{\mathbf{w}_s} \right) D_{t-1}^{-1} \mathbf{w}} \cdot \sqrt{2I_{D+V_t}} \\ &\leq \sqrt{\lambda \mathbf{w}^\top D_{t-1}^{-1} \Lambda_\Sigma \mathbf{w} + \mathbf{w}^\top D_{t-1}^{-1} \left( \sum_{s=1}^{t-1} \Sigma_{\mathbf{w}_s} \right) D_{t-1}^{-1} \mathbf{w}} \cdot \sqrt{2I_{D+V_t}} \\ &= \sqrt{\mathbf{w}^\top D_{t-1}^{-1} \left( \lambda \Lambda_\Sigma D_{t-1} + \sum_{s=1}^{t-1} \Sigma_{\mathbf{w}_s} \right) D_{t-1}^{-1} \mathbf{w}} \cdot \sqrt{2I_{D+V_t}} \end{aligned}$$

Thus,

$$\begin{aligned} &\Pr \left[ \left| \mathbf{w}^\top (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{t-1}) \right| > \sqrt{2\beta(\delta_t)} \sqrt{\mathbf{w}^\top D_{t-1}^{-1} \left( \lambda \Lambda_{\bar{\Sigma}_t} D_{t-1} + \sum_{s=1}^{t-1} \bar{\Sigma}_{s, \mathbf{w}_s} \right) D_{t-1}^{-1} \mathbf{w}} \right] \\ &\leq \Pr \left[ \sqrt{\mathbf{w}^\top D_{t-1}^{-1} \left( \lambda \Lambda_\Sigma D_{t-1} + \sum_{s=1}^{t-1} \Sigma_{\mathbf{w}_s} \right) D_{t-1}^{-1} \mathbf{w}} \cdot \sqrt{2I_{D+V_t}} \right. \\ &\quad \left. > \sqrt{2\beta(\delta_t)} \sqrt{\mathbf{w}^\top D_{t-1}^{-1} \left( \lambda \Lambda_\Sigma D_{t-1} + \sum_{s=1}^{t-1} \Sigma_{\mathbf{w}_s} \right) D_{t-1}^{-1} \mathbf{w}} \right] \end{aligned}$$



$$= \Pr [I_{D+V_t} > \beta(\delta_t)]$$

Hence, to prove Eq. (4), it suffices to prove

$$\Pr [I_{D+V_t} > \beta(\delta_t)] \leq \delta_t. \quad (5)$$

To do so, we introduce some notions. Let  $\mathcal{J}_t$  be the  $\sigma$ -algebra  $\sigma(\mathbf{w}_1, \varepsilon_1, \dots, \mathbf{w}_{t-1}, \varepsilon_{t-1}, \mathbf{w}_t)$ . Let  $\mathbf{u} \in \mathbb{R}^d$  be a multivariate Gaussian random variable with mean  $\mathbf{0}$  and covariance  $D^{-1}$ , which is independent of all the other random variables, and use  $\varphi(\mathbf{u})$  denote its probability density function. Define

$$P_s^{\mathbf{u}} = \exp \left( \mathbf{u}^\top I_{\mathbf{w}_s} \Sigma^{\frac{1}{2}} \varepsilon_s - \frac{1}{2} \mathbf{u}^\top \Sigma_{\mathbf{w}_s} \mathbf{u} \right),$$

$$M_t^{\mathbf{u}} \triangleq \exp \left( \mathbf{u}^\top S_t - \frac{1}{2} \|\mathbf{u}\|_{V_t}^2 \right),$$

and

$$M_t \triangleq \mathbb{E}_{\mathbf{u}} [M_t^{\mathbf{u}}] = \int_{\mathbb{R}^d} \exp \left( \mathbf{u}^\top S_t - \frac{1}{2} \|\mathbf{u}\|_{V_t}^2 \right) \varphi(\mathbf{u}) d\mathbf{u}.$$

We have  $M_t^{\mathbf{u}} = \prod_{s=1}^{t-1} P_s^{\mathbf{u}}$ . In the following, we prove  $\mathbb{E}[M_t] \leq 1$ .

For any  $s > 0$ , according to the sub-Gaussian property,  $\eta_s = \Sigma^{\frac{1}{2}} \varepsilon_s$  satisfies

$$\forall \mathbf{v} \in \mathbb{R}^d, \mathbb{E} \left[ e^{\mathbf{v}^\top \Sigma^{\frac{1}{2}} \varepsilon_s} \right] \leq e^{\frac{1}{2} \mathbf{v}^\top \Sigma \mathbf{v}},$$

which is equivalent to

$$\forall \mathbf{v} \in \mathbb{R}^d, \mathbb{E} \left[ e^{\mathbf{v}^\top \Sigma^{\frac{1}{2}} \varepsilon_s - \frac{1}{2} \mathbf{v}^\top \Sigma \mathbf{v}} \right] \leq 1.$$

Thus, we have

$$\mathbb{E} [P_s^{\mathbf{u}} | \mathcal{J}_s] = \mathbb{E} \left[ \exp \left( \mathbf{u}^\top I_{\mathbf{w}_s} \Sigma^{\frac{1}{2}} \varepsilon_s - \frac{1}{2} \mathbf{u}^\top \Sigma_{\mathbf{w}_s} \mathbf{u} \right) | \mathcal{J}_s \right] \leq 1.$$

Then, we can obtain

$$\begin{aligned} \mathbb{E}[M_t^{\mathbf{u}} | \mathcal{J}_{t-1}] &= \mathbb{E} [\prod_{s=1}^{t-1} P_s^{\mathbf{u}} | \mathcal{J}_{t-1}] \\ &= (\prod_{s=1}^{t-2} P_s^{\mathbf{u}}) \mathbb{E} [P_{t-1}^{\mathbf{u}} | \mathcal{J}_{t-1}] \\ &\leq M_{t-1}^{\mathbf{u}}, \end{aligned}$$

which implies that  $M_t^{\mathbf{u}}$  is a super-martingale and  $\mathbb{E}[M_t^{\mathbf{u}} | \mathbf{u}] \leq 1$ . Thus,

$$\mathbb{E}[M_t] = \mathbb{E}_{\mathbf{u}} [\mathbb{E}[M_t^{\mathbf{u}} | \mathbf{u}]] \leq 1.$$

According to Lemma 9 in [1], we have

$$M_t \triangleq \int_{\mathbb{R}^d} \exp \left( \mathbf{u}^\top S_t - \frac{1}{2} \|\mathbf{u}\|_{V_t}^2 \right) \varphi(\mathbf{u}) d\mathbf{u} = \sqrt{\frac{\det D}{\det(D + V_t)}} \exp(I_{D+V_t}).$$

Thus,

$$\mathbb{E} \left[ \sqrt{\frac{\det D}{\det(D + V_t)}} \exp(I_{D+V_t}) \right] \leq 1.$$

Now we prove Eq. (5). First, we have

$$\begin{aligned}
\Pr [I_{D+V_t} > \beta(\delta_t)] &= \Pr \left[ \sqrt{\frac{\det D}{\det(D+V_t)}} \exp(I_{D+V_t}) > \sqrt{\frac{\det D}{\det(D+V_t)}} \exp(\beta(\delta_t)) \right] \\
&= \Pr \left[ M_t > \frac{1}{\sqrt{\det(I + D^{-\frac{1}{2}} V_t D^{-\frac{1}{2}})}} \exp(\beta(\delta_t)) \right] \\
&\leq \frac{\mathbb{E}[M_t] \sqrt{\det(I + D^{-\frac{1}{2}} V_t D^{-\frac{1}{2}})}}{\exp(\beta(\delta_t))} \\
&\leq \frac{\sqrt{\det(I + D^{-\frac{1}{2}} V_t D^{-\frac{1}{2}})}}{\exp(\beta(\delta_t))}
\end{aligned} \tag{6}$$

Then, for some constant  $\gamma > 0$  and for any  $\mathbf{a} = (a_1, \dots, a_d) \in \mathbb{N}^d$ , we define the set of timesteps  $\mathcal{K}_{\mathbf{a}} \subseteq [T]$  such that

$$t \in \mathcal{K}_{\mathbf{a}} \Leftrightarrow \forall i \in d, (1 + \gamma)^{a_i} \leq N_i(t) < (1 + \gamma)^{a_i + 1}.$$

Define  $D_{\mathbf{a}}$  a diagonal matrix with  $D_{\mathbf{a},ii} = (1 + \gamma)^{a_i}$ .

Suppose  $t \in \mathcal{K}_{\mathbf{a}}$  for some fixed  $\mathbf{a}$ . We have

$$\frac{1}{1 + \gamma} D_t \preceq D_{\mathbf{a}} \preceq D_t.$$

Let  $D = \lambda \Lambda_{\Sigma} D_{\mathbf{a}} \succeq \frac{\lambda}{1 + \gamma} \Lambda_{\Sigma} D_t$ . Then, we have

$$D^{-\frac{1}{2}} V_t D^{-\frac{1}{2}} \preceq \frac{1 + \gamma}{\lambda} D_t^{-\frac{1}{2}} \Lambda_{\Sigma}^{-\frac{1}{2}} V_t \Lambda_{\Sigma}^{-\frac{1}{2}} D_t^{-\frac{1}{2}},$$

where matrix  $D_t^{-\frac{1}{2}} \Lambda_{\Sigma}^{-\frac{1}{2}} V_t \Lambda_{\Sigma}^{-\frac{1}{2}} D_t^{-\frac{1}{2}}$  has  $d$  ones on the diagonal. Since the determinant of a positive definite matrix is smaller than the product of its diagonal terms, we have

$$\begin{aligned}
\det(I + D^{-\frac{1}{2}} V_t D^{-\frac{1}{2}}) &\leq \det\left(I + \frac{1 + \gamma}{\lambda} D_t^{-\frac{1}{2}} \Lambda_{\Sigma}^{-\frac{1}{2}} V_t \Lambda_{\Sigma}^{-\frac{1}{2}} D_t^{-\frac{1}{2}}\right) \\
&\leq \left(1 + \frac{1 + \gamma}{\lambda}\right)^d
\end{aligned} \tag{7}$$

Let  $0 < \lambda < 1$  and  $\gamma = e - 1$ . Using Eqs. (6) and (7),  $\beta(\delta_t) = \ln(1/\delta_t) + d \ln \ln t + \frac{d}{2} \ln(1 + e/\lambda) = \ln(t \ln^2 t) + d \ln \ln t + \frac{d}{2} \ln(1 + e/\lambda)$ , and a union bound over  $\mathbf{a}$ , we have

$$\begin{aligned}
\Pr [I_{D+V_t} > \beta(\delta_t)] &\leq \sum_{\mathbf{a}} \Pr [I_{D+V_t} > \beta(\delta_t) | t \in \mathcal{K}_{\mathbf{a}}, D = \lambda \Lambda_{\Sigma} D_{\mathbf{a}}] \\
&\leq \sum_{\mathbf{a}} \frac{\sqrt{\det(I + D^{-\frac{1}{2}} V_t D^{-\frac{1}{2}})}}{\exp(\beta(\delta_t))} \\
&\leq \left(\frac{\ln t}{\ln(1 + \gamma)}\right)^d \cdot \frac{\left(1 + \frac{1 + \gamma}{\lambda}\right)^{\frac{d}{2}}}{\exp(\ln(t \ln^2 t) + d \ln \ln t + \frac{d}{2} \ln(1 + \frac{e}{\lambda}))} \\
&= (\ln t)^d \cdot \frac{\left(1 + \frac{e}{\lambda}\right)^{\frac{d}{2}}}{t \ln^2 t \cdot (\ln t)^d \cdot \left(1 + \frac{e}{\lambda}\right)^{\frac{d}{2}}}
\end{aligned}$$

$$\begin{aligned}
&= \frac{1}{t \ln^2 t} \\
&= \delta_t
\end{aligned}$$

Thus, Eq. (5) holds and we complete the proof of Lemma 2.  $\square$

## B Proof for CMCB-FI

### B.1 Proof of Theorem 1

In order to prove Theorem 1, we first have the following Lemmas 3 and 4, which are adaptations of Lemmas 1 and 2 to CMCB-FI.

**Lemma 3** (Concentration of Covariance for CMCB-FI). *Consider the CMCB-FI problem and algorithm MC-Empirical (Algorithm 1). For any  $t \geq 2$ , the event*

$$\mathcal{E}_t \triangleq \left\{ |\Sigma_{ij} - \hat{\Sigma}_{ij,t-1}| \leq 16 \left( \frac{3 \ln t}{t-1} \vee \sqrt{\frac{3 \ln t}{t-1}} \right) + (6 + 4\sqrt{3}) \frac{\ln t}{t-1}, \forall i, j \in [d] \right\}$$

satisfies

$$\Pr[\mathcal{E}_t] \geq 1 - \frac{10d^2}{t^2}$$

*Proof.* In CMCB-FI, we have  $N_{ij}(t-1) = t-1$  for any  $t \geq 2$  and  $i, j \in [d]$ . Then, Lemma 3 can be obtained by applying Lemma 1 with  $N_{ij}(t-1) = t-1$  for any  $i, j \in [d]$ .  $\square$

**Lemma 4** (Concentration of Means for CMCB-FI). *Consider the CMCB-FI problem and algorithm MC-Empirical (Algorithm 1). Let  $0 < \lambda < 1$ . Define  $\delta_t = \frac{1}{t \ln^2 t}$  and  $\beta(\delta_t) = \ln(1/\delta_t) + \ln \ln t + \frac{d}{2} \ln(1 + e/\lambda)$  for  $t \geq 2$ . Define  $E_t(\mathbf{w}) = \sqrt{2\beta(\delta_t)} \sqrt{\mathbf{w}^\top D_{t-1}^{-1} (\lambda \Lambda_\Sigma D_{t-1} + \sum_{s=1}^{t-1} \Sigma) D_{t-1}^{-1} \mathbf{w}}$ . Then, for any  $t \geq 2$  and  $\mathbf{w} \in \Delta_d$ , the event  $\mathcal{F}_t \triangleq \{|\mathbf{w}^\top \boldsymbol{\theta}^* - \mathbf{w}^\top \hat{\boldsymbol{\theta}}_{t-1}| \leq E_t(\mathbf{w})\}$  satisfies  $\Pr[\mathcal{F}_t] \geq 1 - \delta_t$ .*

*Proof.* In CMCB-FI,  $D_t$  is a diagonal matrix such that  $D_{t,ii} = N_i(t) = t$ . Then, Lemma 4 can be obtained by applying Lemma 2 with  $D_t = tI$  and that the union bound on the number of samples only needs to consider one dimension. Specifically, in the proof of Lemma 2, we replace the set of timesteps  $\mathcal{K}_a$  with  $\mathcal{K}_a \subseteq [T]$  for  $a \in \mathbb{N}$ , which stands for

$$t \in \mathcal{K}_a \Leftrightarrow (1 + \gamma)^a \leq t < (1 + \gamma)^{a+1}.$$

This completes the proof.  $\square$

Now we are ready to prove Theorem 1.

*Proof.* (Theorem 1) Let  $\Delta_t = f(\mathbf{w}^*) - f(\mathbf{w}_t)$ ,  $g(t) = 16 \left( \frac{3 \ln t}{t-1} \vee \sqrt{\frac{3 \ln t}{t-1}} \right) + (6 + 4\sqrt{3}) \frac{\ln t}{t-1}$  denote the confidence radius of covariance  $\Sigma_{ij}^*$  for any  $i, j \in [d]$ , and  $G(t)$  be the matrix with all entries equal to  $g(t)$ . For any  $\mathbf{w} \in \Delta_d^c$ , define  $\hat{f}_t(\mathbf{w}) = \mathbf{w}^\top \hat{\boldsymbol{\theta}}_t - \rho \mathbf{w}^\top \hat{\Sigma}_t \mathbf{w}$  and  $h_t(\mathbf{w}) = E_t(\mathbf{w}) + \rho \mathbf{w}^\top G_t \mathbf{w}$ .

For any  $t \geq 2$ , suppose that event  $\mathcal{E}_t \cap \mathcal{F}_t$  occurs. Then,

$$|\hat{f}_{t-1}(\mathbf{w}) - f(\mathbf{w})| \leq h_t(\mathbf{w}).$$

Therefore, we have

$$\Delta_t \leq |\hat{f}_{t-1}(\mathbf{w}^*) - f(\mathbf{w}^*)| + |\hat{f}_{t-1}(\mathbf{w}_t) - f(\mathbf{w}_t)|.$$

This is because if instead  $\Delta_t > |\hat{f}_{t-1}(\mathbf{w}^*) - f(\mathbf{w}^*)| + |\hat{f}_{t-1}(\mathbf{w}_t) - f(\mathbf{w}_t)|$ , we have

$$\begin{aligned}
& \hat{f}_{t-1}(\mathbf{w}^*) - \hat{f}_{t-1}(\mathbf{w}_t) \\
&= \hat{f}_{t-1}(\mathbf{w}^*) - \hat{f}_{t-1}(\mathbf{w}_t) + (f(\mathbf{w}^*) - f(\mathbf{w}_t)) - (f(\mathbf{w}^*) - f(\mathbf{w}_t)) \\
&\geq \Delta_t - (f(\mathbf{w}^*) - \hat{f}_{t-1}(\mathbf{w}^*)) - (\hat{f}_{t-1}(\mathbf{w}_t) - f(\mathbf{w}_t)) \\
&\geq \Delta_t - |(f(\mathbf{w}^*) - \hat{f}_{t-1}(\mathbf{w}^*))| - |(\hat{f}_{t-1}(\mathbf{w}_t) - f(\mathbf{w}_t))| \\
&> 0,
\end{aligned}$$

which contradicts the selection strategy of  $\mathbf{w}_t$  in algorithm MC-Empirical. Thus, we obtain

$$\begin{aligned}
\Delta_t &\leq |\hat{f}_{t-1}(\mathbf{w}^*) - f(\mathbf{w}^*)| + |\hat{f}_{t-1}(\mathbf{w}_t) - f(\mathbf{w}_t)| \\
&\leq h_t(\mathbf{w}^*) + h_t(\mathbf{w}_t) \\
&= E_t(\mathbf{w}^*) + \rho \mathbf{w}^{*\top} G_t \mathbf{w}^* + E_t(\mathbf{w}_t) + \rho \mathbf{w}_t^\top G_t \mathbf{w}_t
\end{aligned} \tag{8}$$

Now, for any  $\mathbf{w} \in \Delta_d^c$ , we have

$$\begin{aligned}
\mathbf{w}^\top G_t \mathbf{w} &= \sum_{i,j \in [d]} g(t) w_i w_j \\
&= g(t) \sum_{i,j \in [d]} w_i w_j \\
&= g(t) \left( \sum_i w_i \right)^2 \\
&= g(t)
\end{aligned}$$

and

$$\begin{aligned}
E_t(\mathbf{w}) &= \sqrt{2\beta(\delta_t)} \sqrt{\mathbf{w}^\top D_{t-1}^{-1} \left( \lambda \Lambda_\Sigma D_{t-1} + \sum_{s=1}^{t-1} \Sigma \right) D_{t-1}^{-1} \mathbf{w}} \\
&= \sqrt{2\beta(\delta_t)} \sqrt{\lambda \mathbf{w}^\top D_{t-1}^{-1} \Lambda_\Sigma \mathbf{w} + \mathbf{w}^\top D_{t-1}^{-1} \left( \sum_{s=1}^{t-1} \Sigma \right) D_{t-1}^{-1} \mathbf{w}} \\
&= \sqrt{2\beta(\delta_t)} \sqrt{\lambda \mathbf{w}^\top D_{t-1}^{-1} \Lambda_\Sigma \mathbf{w} + \mathbf{w}^\top D_{t-1}^{-1} \Sigma \mathbf{w}} \\
&= \sqrt{2\beta(\delta_t)} \sqrt{\frac{1}{t-1} \lambda \mathbf{w}^\top \Lambda_\Sigma \mathbf{w} + \frac{1}{t-1} \mathbf{w}^\top \Sigma \mathbf{w}} \\
&\leq \sqrt{\frac{2\beta(\delta_t)}{t-1}} \sqrt{\lambda \Sigma_{\max} + \mathbf{w}^\top \Sigma \mathbf{w}},
\end{aligned}$$

where  $\Sigma_{\max}$  denotes the maximum diagonal entry of  $\Sigma$ .

For any  $t \geq 7$ ,  $\frac{3 \ln t}{t-1} < \sqrt{\frac{3 \ln t}{t-1}}$  and  $g(t) \leq (6 + 20\sqrt{3}) \frac{\ln t}{\sqrt{t-1}}$ , Eq. (8) can be written as

$$\begin{aligned}
\Delta_t &\leq E_t(\mathbf{w}^*) + \rho \mathbf{w}^{*\top} G_t \mathbf{w}^* + E_t(\mathbf{w}_t) + \rho \mathbf{w}_t^\top G_t \mathbf{w}_t \\
&\leq \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( \sqrt{\lambda \Sigma_{\max} + \mathbf{w}_t^\top \Sigma \mathbf{w}_t} + \sqrt{\lambda \Sigma_{\max} + \mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) + 82\rho \frac{\ln t}{\sqrt{t-1}}
\end{aligned}$$

$$\begin{aligned}
&\leq \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda\Sigma_{\max}} + \sqrt{\mathbf{w}_t^\top \Sigma \mathbf{w}_t} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) + 82\rho \frac{\ln t}{\sqrt{t-1}} \\
&= \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda\Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) + 82\rho \frac{\ln t}{\sqrt{t-1}} + \sqrt{\frac{2\beta(\delta_t)}{t-1}} \cdot \sqrt{\mathbf{w}_t^\top \Sigma \mathbf{w}_t}
\end{aligned} \tag{9}$$

Next, we investigate the upper bound of  $\mathbf{w}_t^\top \Sigma \mathbf{w}_t$ . According to Eq. (8), we have that  $\mathbf{w}_t^\top \Sigma \mathbf{w}_t$  satisfies

$$\Delta_t \leq h_t(\mathbf{w}^*) + h_t(\mathbf{w}_t) \tag{10}$$

In Eq. (10), we have

$$\begin{aligned}
\Delta_t &= f(\mathbf{w}^*) - f(\mathbf{w}_t) \\
&\geq \theta_{\min}^* - \rho \mathbf{w}^{*\top} \Sigma \mathbf{w}^* - \theta_{\max}^* + \rho \mathbf{w}_t^\top \Sigma \mathbf{w}_t
\end{aligned}$$

and

$$\begin{aligned}
h_t(\mathbf{w}^*) + h_t(\mathbf{w}_t) &= E_t(\mathbf{w}^*) + \rho \mathbf{w}^{*\top} G_t \mathbf{w}^* + E_t(\mathbf{w}_t) + \rho \mathbf{w}_t^\top G_t \mathbf{w}_t \\
&\leq \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda\Sigma_{\max}} + \sqrt{\mathbf{w}_t^\top \Sigma \mathbf{w}_t} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) + 82\rho \frac{\ln t}{\sqrt{t-1}}.
\end{aligned}$$

Thus,  $\mathbf{w}_t^\top \Sigma \mathbf{w}_t$  satisfies

$$\begin{aligned}
\theta_{\min}^* - \rho \mathbf{w}^{*\top} \Sigma \mathbf{w}^* - \theta_{\max}^* + \rho \mathbf{w}_t^\top \Sigma \mathbf{w}_t &\leq \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda\Sigma_{\max}} + \sqrt{\mathbf{w}_t^\top \Sigma \mathbf{w}_t} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) \\
&\quad + 82\rho \frac{\ln t}{\sqrt{t-1}}
\end{aligned}$$

Rearranging the terms, we have

$$\begin{aligned}
\rho \mathbf{w}_t^\top \Sigma \mathbf{w}_t - \sqrt{\frac{2\beta(\delta_t)}{t-1}} \sqrt{\mathbf{w}_t^\top \Sigma \mathbf{w}_t} - \left( \theta_{\max}^* - \theta_{\min}^* + \rho \mathbf{w}^{*\top} \Sigma \mathbf{w}^* \right. \\
\left. + \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda\Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) + 82\rho \frac{\ln t}{\sqrt{t-1}} \right) \leq 0
\end{aligned} \tag{11}$$

Let  $x = \mathbf{w}_t^\top \Sigma \mathbf{w}_t$  and  $0 < \lambda < 1$ . Define function

$$y(x) = \rho x - c_1 \sqrt{x} - c_2 \leq 0,$$

where  $c_1 = \sqrt{\frac{2\beta(\delta_t)}{t-1}} > 0$  and  $c_2 = \theta_{\max}^* - \theta_{\min}^* + \rho \mathbf{w}^{*\top} \Sigma \mathbf{w}^* + \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda\Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) + 82\rho \frac{\ln t}{\sqrt{t-1}} > 0$ . When  $t \geq t_0 \triangleq \max \left\{ (1 + e/\lambda)^{\frac{d}{2}}, 7 \right\}$ , we have

$$\beta(\delta_t) = \ln(t \ln^2 t) + \ln \ln t + \frac{d}{2} \ln(1 + e/\lambda) \leq 5 \ln t$$

Now since

$$y(x) = \rho x - c_1 \sqrt{x} - c_2$$

$$= \rho \left( \sqrt{x} - \frac{c_1}{2\rho} \right)^2 - \frac{c_1^2}{4\rho} - c_2,$$

by letting  $y(x) \leq 0$ , we have

$$\begin{aligned} x &\leq \left( \frac{c_1}{2\rho} + \sqrt{\frac{c_1^2}{4\rho^2} + \frac{c_2}{\rho}} \right)^2 \\ &\leq 2\frac{c_1^2}{4\rho^2} + 2\frac{c_1^2}{4\rho^2} + 2\frac{c_2}{\rho} \\ &= \frac{c_1^2}{\rho^2} + \frac{2c_2}{\rho} \end{aligned}$$

Therefore

$$\begin{aligned} \mathbf{w}_t^\top \Sigma \mathbf{w}_t &\leq \frac{1}{\rho^2} \frac{2\beta(\delta_t)}{t-1} + \frac{2}{\rho} \left( \theta_{\max}^* - \theta_{\min}^* + \rho \mathbf{w}^{*\top} \Sigma \mathbf{w}^* + \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda \Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) \right. \\ &\quad \left. + 82\rho \frac{\ln t}{\sqrt{t-1}} \right) \\ &\leq 2\mathbf{w}^{*\top} \Sigma \mathbf{w}^* + \frac{2}{\rho} (\theta_{\max}^* - \theta_{\min}^*) + \frac{2}{\rho} \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda \Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) + 164 \frac{\ln t}{\sqrt{t-1}} \\ &\quad + \frac{1}{\rho^2} \frac{2\beta(\delta_t)}{t-1} \end{aligned}$$

Thus, we have that,  $\mathbf{w}_t^\top \Sigma \mathbf{w}_t$  satisfies

$$\begin{aligned} \mathbf{w}_t^\top \Sigma \mathbf{w}_t &\leq \min \left\{ 2\mathbf{w}^{*\top} \Sigma \mathbf{w}^* + \frac{2}{\rho} (\theta_{\max}^* - \theta_{\min}^*) + \frac{2}{\rho} \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda \Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) \right. \\ &\quad \left. + 164 \frac{\ln t}{\sqrt{t-1}} + \frac{1}{\rho^2} \frac{2\beta(\delta_t)}{t-1}, \mathbf{w}_{\max}^\top \Sigma \mathbf{w}_{\max} \right\}, \end{aligned} \quad (12)$$

where  $\mathbf{w}_{\max}^\top \triangleq \operatorname{argmax}_{\mathbf{w} \in \Delta_d^c} \mathbf{w}^\top \Sigma \mathbf{w}$ .

Below we discuss the two terms in Eq. (12) separately.

Case (i): Plugging the first term of the upper bound of  $\mathbf{w}_t^\top \Sigma \mathbf{w}_t$  in Eq. (12) into Eq. (9), we have that for  $t \geq t_0$ ,

$$\begin{aligned} \Delta_t &\leq \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda \Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) + 82\rho \frac{\ln t}{\sqrt{t-1}} + \sqrt{\frac{2\beta(\delta_t)}{t-1}} \cdot \sqrt{\mathbf{w}_t^\top \Sigma \mathbf{w}_t} \\ &\leq \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda \Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) + 82\rho \frac{\ln t}{\sqrt{t-1}} + \sqrt{\frac{2\beta(\delta_t)}{t-1}} \\ &\quad \sqrt{2\mathbf{w}^{*\top} \Sigma \mathbf{w}^* + \frac{2}{\rho} (\theta_{\max}^* - \theta_{\min}^*) + \frac{2}{\rho} \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda \Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) + 164 \frac{\ln t}{\sqrt{t-1}} + \frac{1}{\rho^2} \frac{2\beta(\delta_t)}{t-1}} \\ &\leq \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda \Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) + 82\rho \frac{\ln t}{\sqrt{t-1}} + \sqrt{\frac{2\beta(\delta_t)}{t-1}}. \end{aligned}$$



$$\begin{aligned}
& \left( \sqrt{2\mathbf{w}^{*\top}\Sigma\mathbf{w}^*} + \frac{\sqrt{2}}{\sqrt{\rho}}\sqrt{\theta_{\max}^* - \theta_{\min}^*} + \frac{\sqrt{2}}{\sqrt{\rho}}\left(\frac{2\beta(\delta_t)}{t-1}\right)^{\frac{1}{4}} \left( \sqrt{2}(\lambda\Sigma_{\max})^{\frac{1}{4}} + (\mathbf{w}^{*\top}\Sigma\mathbf{w}^*)^{\frac{1}{4}} \right) \right. \\
& \left. + 13\frac{\sqrt{\ln t}}{(t-1)^{\frac{1}{4}}} + \frac{1}{\rho}\sqrt{\frac{2\beta(\delta_t)}{t-1}} \right) \\
& \leq \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda\Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top}\Sigma\mathbf{w}^*} + \sqrt{2\mathbf{w}^{*\top}\Sigma\mathbf{w}^*} + \frac{\sqrt{2}}{\sqrt{\rho}}\sqrt{\theta_{\max}^* - \theta_{\min}^*} \right) + 82\rho\frac{\ln t}{\sqrt{t-1}} \\
& \quad + \frac{\sqrt{2}}{\sqrt{\rho}}\left(\frac{2\beta(\delta_t)}{t-1}\right)^{\frac{3}{4}} \left( \sqrt{2}(\lambda\Sigma_{\max})^{\frac{1}{4}} + (\mathbf{w}^{*\top}\Sigma\mathbf{w}^*)^{\frac{1}{4}} \right) + 42\frac{\ln t}{(t-1)^{\frac{3}{4}}} + \frac{1}{\rho}\frac{2\beta(\delta_t)}{t-1} \\
& \leq \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda\Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top}\Sigma\mathbf{w}^*} + \sqrt{2\mathbf{w}^{*\top}\Sigma\mathbf{w}^*} + \frac{\sqrt{2}}{\sqrt{\rho}}\sqrt{\theta_{\max}^* - \theta_{\min}^*} \right) + 82\rho\frac{\ln t}{\sqrt{t-1}} \\
& \quad + \frac{\sqrt{2}}{\sqrt{\rho}}\left(\frac{2\beta(\delta_t)}{t-1}\right)^{\frac{3}{4}} \left( \sqrt{2}(\lambda\Sigma_{\max})^{\frac{1}{4}} + (\mathbf{w}^{*\top}\Sigma\mathbf{w}^*)^{\frac{1}{4}} \right) + 42\frac{\ln t}{(t-1)^{\frac{3}{4}}} + \frac{1}{\rho}\frac{2\beta(\delta_t)}{t-1} \\
& \leq \sqrt{\frac{\ln t}{t-1}} \left( 2\sqrt{\lambda\Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top}\Sigma\mathbf{w}^*} + \sqrt{2\mathbf{w}^{*\top}\Sigma\mathbf{w}^*} + \frac{\sqrt{2}}{\sqrt{\rho}}\sqrt{\theta_{\max}^* - \theta_{\min}^*} \right) + 82\rho\frac{\ln t}{\sqrt{t-1}} \\
& \quad + 42\frac{\ln t}{(t-1)^{\frac{3}{4}}} \left( \frac{1}{\sqrt{\rho}} + 1 \right) + \frac{1}{\rho}\frac{10\ln t}{t-1}
\end{aligned}$$

According to Lemmas 3 and 4, for any  $t \geq 2$ , we bound the probability of event  $\neg(\mathcal{E}_t \cap \mathcal{F}_t)$  as follows.

$$\begin{aligned}
\Pr[\neg(\mathcal{E}_t \cap \mathcal{F}_t)] & \leq \frac{10d^2}{t^2} + \frac{1}{t \ln^2 t} \\
& \leq \frac{10d^2}{t \ln^2 t} + \frac{1}{t \ln^2 t} \\
& = \frac{11d^2}{t \ln^2 t}
\end{aligned}$$

Recall that  $t_0 = \max\left\{(1 + e/\lambda)^{\frac{d}{2}}, 7\right\}$ . For any horizon  $T$ , summing over  $t = 1, \dots, t_0$  and  $t = t_0, \dots, T$ , we obtain the regret upper bound

$$\begin{aligned}
\mathbb{E}[\mathcal{R}(T)] & = O(t_0\Delta_{\max}) + \sum_{t=t_0}^T O(\Delta_{\max} \cdot \Pr[\neg(\mathcal{E}_t \cap \mathcal{F}_t)] + \Delta_t \cdot \mathbb{I}\{\mathcal{E}_t \cap \mathcal{F}_t\}) \\
& = O(\lambda^{-\frac{d}{2}}\Delta_{\max}) + \sum_{t=t_0}^T O\left(\Delta_{\max} \cdot \frac{d^2}{t \ln^2 t}\right) + \sum_{t=t_0}^T O(\Delta_t \cdot \mathbb{I}\{\mathcal{E}_t \cap \mathcal{F}_t\}) \\
& = O(\lambda^{-\frac{d}{2}}\Delta_{\max}) + \sum_{t=t_0}^T O\left(\sqrt{\frac{\ln t}{t-1}} \left( \sqrt{\lambda\Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top}\Sigma\mathbf{w}^*} + \frac{1}{\sqrt{\rho}}\sqrt{\theta_{\max}^* - \theta_{\min}^*} \right) \right. \\
& \quad \left. + \rho\frac{\ln t}{\sqrt{t-1}} + \frac{\ln t}{(t-1)^{\frac{3}{4}}} \left( \frac{1}{\sqrt{\rho}} + 1 \right) + \frac{1}{\rho}\frac{\ln t}{t-1} \right)
\end{aligned}$$

$$\begin{aligned}
&= O(\lambda^{-\frac{d}{2}} \Delta_{\max}) + O\left(\ln T \sqrt{T} \left( \sqrt{\lambda \Sigma_{\max}} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} + \frac{1}{\sqrt{\rho}} \sqrt{\theta_{\max}^* - \theta_{\min}^*} + \rho \right)\right) \\
&\quad + \ln T \cdot T^{\frac{1}{4}} \left( \frac{1}{\sqrt{\rho}} + 1 \right) + \ln^2 T \cdot \frac{1}{\rho}
\end{aligned}$$

Case (ii): Plugging the second term of the upper bound of  $\mathbf{w}_t^\top \Sigma \mathbf{w}_t$  in Eq. (12) into Eq. (9), we have that for  $t \geq t_0$ ,

$$\begin{aligned}
\Delta_t &\leq \sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( 2\sqrt{\lambda \Sigma_{\max}} + \sqrt{\mathbf{w}_{\max}^\top \Sigma \mathbf{w}_{\max}} + \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} \right) + 82\rho \frac{\ln t}{\sqrt{t-1}} \\
&\leq 2\sqrt{\frac{2\beta(\delta_t)}{t-1}} \left( \sqrt{\lambda \Sigma_{\max}} + \sqrt{\mathbf{w}_{\max}^\top \Sigma \mathbf{w}_{\max}} \right) + 82\rho \frac{\ln t}{\sqrt{t-1}}
\end{aligned}$$

For any horizon  $T$ , summing over  $t = 1, \dots, t_0$  and  $t = t_0, \dots, T$ , we obtain the regret upper bound

$$\begin{aligned}
\mathbb{E}[\mathcal{R}(T)] &= O(t_0 \Delta_{\max}) + \sum_{t=t_0}^T O(\Delta_{\max} \cdot \Pr[-(\mathcal{E}_t \cap \mathcal{F}_t)] + \Delta_t \cdot \mathbb{I}\{\mathcal{E}_t \cap \mathcal{F}_t\}) \\
&= O(\lambda^{-\frac{d}{2}} \Delta_{\max}) + \sum_{t=t_0}^T O\left(\Delta_{\max} \cdot \frac{d^2}{t \ln^2 t}\right) + \sum_{t=t_0}^T O(\Delta_t \cdot \mathbb{I}\{\mathcal{E}_t \cap \mathcal{F}_t\}) \\
&= O(\lambda^{-\frac{d}{2}} \Delta_{\max}) + \sum_{t=t_0}^T O\left(\sqrt{\frac{\beta(\delta_t)}{t-1}} \left( \sqrt{\lambda \Sigma_{\max}} + \sqrt{\mathbf{w}_{\max}^\top \Sigma \mathbf{w}_{\max}} \right) + \rho \frac{\ln t}{\sqrt{t-1}}\right) \\
&= O(\lambda^{-\frac{d}{2}} \Delta_{\max}) + \sum_{t=t_0}^T O\left(\sqrt{\frac{\ln t}{t-1}} \left( \sqrt{\lambda \Sigma_{\max}} + \sqrt{\mathbf{w}_{\max}^\top \Sigma \mathbf{w}_{\max}} \right) + \rho \frac{\ln t}{\sqrt{t-1}}\right) \\
&= O(\lambda^{-\frac{d}{2}} \Delta_{\max}) + O\left(\ln T \sqrt{T} \left( \sqrt{\lambda \Sigma_{\max}} + \sqrt{\mathbf{w}_{\max}^\top \Sigma \mathbf{w}_{\max}} + \rho \right)\right)
\end{aligned}$$

Combining both cases (i) and (ii), we can obtain

$$\begin{aligned}
\mathbb{E}[\mathcal{R}(T)] &= O\left(\min \left\{ \ln T \sqrt{T} \left( \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} + \rho^{-\frac{1}{2}} \sqrt{\theta_{\max}^* - \theta_{\min}^*} \right) + \ln T \cdot T^{\frac{1}{4}} \left( \rho^{-\frac{1}{2}} + 1 \right) \right. \right. \\
&\quad \left. \left. + \ln^2 T \cdot \rho^{-1}, \ln T \sqrt{T} \sqrt{\mathbf{w}_{\max}^\top \Sigma \mathbf{w}_{\max}} \right\} + \ln T \sqrt{T} \left( \sqrt{\lambda \Sigma_{\max}} + \rho \right) + \lambda^{-\frac{d}{2}} \Delta_{\max} \right)
\end{aligned}$$

Let  $\Sigma_{\max}^* = \max_{i \in [d]} \Sigma_{ii}^*$ . Setting  $\lambda = \frac{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*}{\Sigma_{\max}^*}$  and ignoring the terms of  $o(\ln T \sqrt{T})$  order, we obtain

$$\begin{aligned}
\mathbb{E}[\mathcal{R}(T)] &= O\left(\left(\min \left\{ \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} + \rho^{-\frac{1}{2}} \sqrt{\theta_{\max}^* - \theta_{\min}^*}, \sqrt{\mathbf{w}_{\max}^\top \Sigma \mathbf{w}_{\max}} \right\} + \rho\right) \ln T \sqrt{T}\right) \\
&= O\left(\left(\min \left\{ \sqrt{\mathbf{w}^{*\top} \Sigma \mathbf{w}^*} + \rho^{-\frac{1}{2}} \sqrt{\theta_{\max}^* - \theta_{\min}^*}, \sqrt{\Sigma_{\max}^*} \right\} + \rho\right) \ln T \sqrt{T}\right)
\end{aligned}$$

This completes the proof.  $\square$

## B.2 Proof of Theorem 2

In order to prove Theorem 2, we first analyze the offline problem of CMCB-FI. Suppose that the covariance matrix  $\Sigma$  is positive definite.

**(Offline Problem of CMCB-FI)** We define the quadratic optimization  $\text{QuadOpt}(\boldsymbol{\theta}^*, \Sigma)$  as

$$\begin{aligned} \min_{\mathbf{w}} \quad & f(\mathbf{w}) = \rho \mathbf{w}^\top \Sigma \mathbf{w} - \mathbf{w}^\top \boldsymbol{\theta}^* \\ \text{s.t.} \quad & w_i \geq 0, \quad \forall i \in [d] \\ & \sum_{i=1}^d w_i = 1 \end{aligned}$$

and  $\mathbf{w}^*$  as the optimal solution to  $\text{QuadOpt}(\boldsymbol{\theta}^*, \Sigma)$ . We consider the KKT condition for this quadratic optimization as follows:

$$\begin{aligned} 2\rho \Sigma \mathbf{w} - \boldsymbol{\theta}^* - \mathbf{u} - v \mathbf{1} &= \mathbf{0} \\ w_i u_i &= 0, \quad \forall i \in [d] \\ u_i &\geq 0, \quad \forall i \in [d] \\ w_i &\geq 0, \quad \forall i \in [d] \\ \sum_{i=1}^d w_i &= 1 \end{aligned}$$

Let  $S \subseteq [d]$  be a subset of indexes for  $\mathbf{w}$  such that  $S = \{i \in [d] : w_i > 0\}$ . Let  $\bar{S} = [d] \setminus S$  and we have  $\bar{S} = \{i \in [d] : w_i = 0\}$ . Then, from the KKT condition, we have

$$\begin{aligned} \mathbf{w}_S &= \frac{1}{2\rho} \Sigma_S^{-1} \boldsymbol{\theta}_S^* + \frac{1 - \|\frac{1}{2\rho} \Sigma_S^{-1} \boldsymbol{\theta}_S^*\|}{\|\Sigma_S^{-1} \mathbf{1}\|} \Sigma_S^{-1} \mathbf{1} \succ \mathbf{0} \\ \mathbf{w}_{\bar{S}} &= \mathbf{0} \\ v &= \frac{2\rho(1 - \|\frac{1}{2\rho} \Sigma_S^{-1} \boldsymbol{\theta}_S^*\|)}{\|\Sigma_S^{-1} \mathbf{1}\|} \\ \mathbf{u} &= 2\rho \Sigma \mathbf{w} - \boldsymbol{\theta}^* - v \mathbf{1} \succeq \mathbf{0} \end{aligned} \tag{13}$$

Since this problem is a quadratic optimization and the covariance matrix  $\Sigma$  is positive-definite, there is a unique feasible  $S$  satisfying the above inequalities and the solution  $\mathbf{w}_S, \mathbf{w}_{\bar{S}}$  is the optimal solution  $\mathbf{w}^*$ .

**Main Proof.** Now, we give the proof of Theorem 2.

*Proof.* (Theorem 2) First, we choose prior distributions for  $\boldsymbol{\theta}^*$  and  $\Sigma$ . We assume that  $\boldsymbol{\theta}^* \sim \mathcal{N}(0, \frac{1}{\omega} I)$  and  $\Sigma \sim \pi_I$ , where  $\omega > 0$  and  $\pi_I$  takes probability 1 at the support  $I$  and probability 0 anywhere else. Define  $\hat{\boldsymbol{\theta}}_t \triangleq \frac{1}{t} \sum_{i=1}^t \theta_i$  and  $\boldsymbol{\mu}_t = \frac{t}{t+\omega} \hat{\boldsymbol{\theta}}_t$ . Then, we see that  $\hat{\boldsymbol{\theta}}_t \sim \mathcal{N}(0, \frac{\omega+t}{t\omega} I)$  and  $\boldsymbol{\mu}_t \sim \mathcal{N}(0, \frac{t}{(t+\omega)\omega} I)$ .

Thus, the posterior of  $\boldsymbol{\theta}^*$  is given by

$$\boldsymbol{\theta}^* | \theta_1, \dots, \theta_t, \Sigma \sim \mathcal{N}\left(\frac{t}{t+\omega} \hat{\boldsymbol{\theta}}_t, \frac{1}{t+\omega} I\right) = \mathcal{N}\left(\boldsymbol{\mu}_t, \frac{1}{t+\omega} I\right).$$

The posterior of  $\Sigma$  is still  $\Sigma \sim \pi_I$ , i.e.,  $\Sigma$  is always a fixed identity matrix. Under the Bayesian setting, the expected regret is given by

$$\sum_{t=1}^T \mathbb{E}_{\boldsymbol{\mu}_t \sim \mathcal{N}(0, \frac{t}{(t+\omega)\omega} I)} \left[ \mathbb{E}_{\boldsymbol{\theta}^* | \boldsymbol{\mu}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \frac{1}{t+\omega} I)} [f(\mathbf{w}^*) - f(\mathbf{w}_t)] \right].$$

Recall that  $\mathbf{w}^*$  is the optimal solution to  $\text{QuadOpt}(\boldsymbol{\theta}^*, \Sigma)$ . It can be seen that the best strategy of  $\mathbf{w}_t$  at timestep  $t$  is to select the optimal solution to  $\text{QuadOpt}(\boldsymbol{\mu}_t, \Sigma)$  and we use algorithm  $\mathcal{A}$  to denote this strategy. Thus, to obtain a regret lower bound for the problem, it suffices to prove a regret lower bound of algorithm  $\mathcal{A}$  for the problem.

Below we prove a regret lower bound of algorithm  $\mathcal{A}$  for the problem.

**Step (i).** We consider the case when  $\mathbf{w}^*$  and  $\mathbf{w}_t$  both lie in the interior of the  $d$ -dimensional probability simplex  $\Delta_d$ , i.e.,  $w_i^* > 0, \forall i \in [d]$  and  $w_{t,i} > 0, \forall i \in [d]$ . From Eq. (13),  $\mathbf{w}^*$  satisfies

$$\frac{1}{2\rho} I^{-1} \boldsymbol{\theta}^* + \frac{1 - \|\frac{1}{2\rho} I^{-1} \boldsymbol{\theta}^*\|}{\|I^{-1} \mathbf{1}\|} I^{-1} \mathbf{1} \succ \mathbf{0}.$$

Rearranging the terms, we have

$$\|\boldsymbol{\theta}^*\| \mathbf{1} - d\boldsymbol{\theta}^* \prec 2\rho \mathbf{1},$$

which is equivalent to

$$\begin{cases} \theta_2^* + \cdots + \theta_d^* - (d-1)\theta_1^* < 2\rho \\ \vdots \\ \theta_1^* + \cdots + \theta_{d-1}^* - (d-1)\theta_d^* < 2\rho \end{cases} \quad (14)$$

Similarly,  $\mathbf{w}_t$  satisfies

$$\begin{cases} \mu_2 + \cdots + \mu_d - (d-1)\mu_1 < 2\rho \\ \vdots \\ \mu_1 + \cdots + \mu_{d-1} - (d-1)\mu_d < 2\rho \end{cases} \quad (15)$$

We first derive a condition that makes  $\boldsymbol{\mu}_t$  lie in the interior of  $\Delta_d$ . Recall that  $\boldsymbol{\mu}_t \sim \mathcal{N}(0, \frac{t}{(t+\omega)\omega} I)$ . Define event

$$\mathcal{E}_t \triangleq \left\{ -3\sqrt{\frac{t}{(t+\omega)\omega}} \leq \mu_t \leq 3\sqrt{\frac{t}{(t+\omega)\omega}} \right\}.$$

According to the  $3 - \sigma$  principle for Gaussian distributions, we have

$$\Pr[\mathcal{E}_t] \geq (99.7\%)^d.$$

Conditioning on  $\mathcal{E}_t$ , under which Eq. (15) hold, it suffices to let

$$3(d-1)\sqrt{\frac{t}{(t+\omega)\omega}} - \left( -3(d-1)\sqrt{\frac{t}{(t+\omega)\omega}} \right) < \rho,$$

which is equivalent to

$$\left(1 + \frac{\omega}{t}\right) \omega > \frac{36(d-1)^2}{\rho^2}, \quad (16)$$

when  $t > 0, m > 0, t + \omega > 0$ . Let  $t_1 > 0$  be the smallest timestep that satisfies Eq. (16). Thus, when  $\mathcal{E}_t$  occurs and  $t \geq t_1$ ,  $\boldsymbol{\mu}_t$  lie in the interior of  $\Delta_d$ .

Next, we derive some condition that make  $\boldsymbol{\theta}^*$  lie in the interior of  $\Delta_d$ . Recall that  $\boldsymbol{\theta}^* | \boldsymbol{\mu}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \frac{1}{t+\omega} I)$ .

Fix  $\boldsymbol{\mu}_t$ , and then we define event

$$\mathcal{F}_t \triangleq \left\{ -3\sqrt{\frac{t}{(t+\omega)\omega}} \leq \boldsymbol{\theta}^* - \boldsymbol{\mu}_t \leq 3\sqrt{\frac{t}{(t+\omega)\omega}} \right\}.$$

According to the 3 –  $\sigma$  principle for Gaussian distributions, we have

$$\Pr[\mathcal{F}_t] \geq (99.7\%)^d.$$

Conditioning on  $\mathcal{F}_t$ , in order to let Eq. (14) hold, it suffices to let

$$3(d-1)\frac{1}{\sqrt{t+\omega}} - \left( -3(d-1)\frac{1}{\sqrt{t+\omega}} \right) < \rho,$$

which is equivalent to

$$t > \frac{36(d-1)^2}{\rho^2} - \omega, \quad (17)$$

when  $t + \omega > 0$ . Let  $t_2 > 0$  be the smallest timestep that satisfies Eq. (17). Thus, when  $\mathcal{F}_t$  occurs and  $t \geq t_2$ ,  $\boldsymbol{\theta}^*$  lie in the interior of  $\Delta_d$ .

**Step (ii).** Suppose that  $\mathcal{E}_t \cap \mathcal{F}_t \cap \mathcal{G}_t$  occurs and consider  $t \geq \tilde{t} \triangleq \max\{t_1, t_2\}$ . Then,  $\boldsymbol{w}^*$  and  $\boldsymbol{w}_t$  both lie in the interior of  $\Delta_d$ , i.e.,  $w_i^* > 0, \forall i \in [d]$  and  $w_{t,i} > 0, \forall i \in [d]$ . We have

$$\begin{aligned} \boldsymbol{w}^* &= \frac{1}{2\rho}\Sigma^{-1}\boldsymbol{\theta}^* + \frac{1 - \|\frac{1}{2\rho}\Sigma^{-1}\boldsymbol{\theta}^*\|}{\|\Sigma^{-1}\mathbf{1}\|}\Sigma^{-1}\mathbf{1} \\ \boldsymbol{w}_t &= \frac{1}{2\rho}\Sigma^{-1}\boldsymbol{\mu}_t + \frac{1 - \|\frac{1}{2\rho}\Sigma^{-1}\boldsymbol{\mu}_t\|}{\|\Sigma^{-1}\mathbf{1}\|}\Sigma^{-1}\mathbf{1} \end{aligned}$$

Let  $\Delta\boldsymbol{\theta}_t \triangleq \boldsymbol{\mu}_t - \boldsymbol{\theta}^*$  and thus  $\Delta\boldsymbol{\theta}_t|\boldsymbol{\mu}_t \sim \mathcal{N}(0, \frac{1}{t+\omega}I)$ . Let  $\Delta\boldsymbol{w}_t \triangleq \boldsymbol{w}_t - \boldsymbol{w}^* = \frac{1}{2\rho}\Sigma^{-1}\Delta\boldsymbol{\theta}_t - \frac{\|\frac{1}{2\rho}\Sigma^{-1}\Delta\boldsymbol{\theta}_t\|}{\|\Sigma^{-1}\mathbf{1}\|}\Sigma^{-1}\mathbf{1} = \frac{1}{2\rho}\Delta\boldsymbol{\theta}_t - \frac{1}{2\rho d}\|\Delta\boldsymbol{\theta}_t\|\mathbf{1}$ . Then, we have

$$\begin{aligned} f(\boldsymbol{w}^*) - f(\boldsymbol{w}_t) &= f(\boldsymbol{w}^*) - f(\boldsymbol{w}^* + \Delta\boldsymbol{w}_t) \\ &= ((\boldsymbol{w}^*)^\top \boldsymbol{\theta}^* - \rho(\boldsymbol{w}^*)^\top \Sigma \boldsymbol{w}^*) \\ &\quad - ((\boldsymbol{w}^*)^\top \boldsymbol{\theta}^* + (\Delta\boldsymbol{w}_t)^\top \boldsymbol{\theta}^* - \rho(\boldsymbol{w}^*)^\top \Sigma \boldsymbol{w}^* - 2\rho(\Delta\boldsymbol{w}_t)^\top \Sigma \boldsymbol{w}^* - \rho(\Delta\boldsymbol{w}_t)^\top \Sigma \Delta\boldsymbol{w}_t) \\ &= -(\Delta\boldsymbol{w}_t)^\top \boldsymbol{\theta}^* + 2\rho(\Delta\boldsymbol{w}_t)^\top \Sigma \boldsymbol{w}^* + \rho(\Delta\boldsymbol{w}_t)^\top \Sigma \Delta\boldsymbol{w}_t \\ &= (\Delta\boldsymbol{w}_t)^\top \nabla f(\boldsymbol{\theta}^*) + \rho(\Delta\boldsymbol{w}_t)^\top \Sigma \Delta\boldsymbol{w}_t \\ &\geq \rho(\Delta\boldsymbol{w}_t)^\top \Sigma \Delta\boldsymbol{w}_t \\ &= \rho(\Delta\boldsymbol{w}_t)^\top \Delta\boldsymbol{w}_t \\ &= \rho \left( \frac{1}{4\rho^2} (\Delta\boldsymbol{\theta}_t)^\top \Delta\boldsymbol{\theta}_t - \frac{1}{4\rho^2 d^2} \|\Delta\boldsymbol{\theta}_t\|^2 d \right) \\ &= \frac{1}{4\rho} \left( \sum_{i=1}^d \Delta\theta_{t,i}^2 - \frac{(\sum_{i=1}^d \Delta\theta_{t,i})^2}{d} \right) \\ &= \frac{1}{4\rho d} \sum_{1 \leq i < j \leq d} (\Delta\theta_{t,i} - \Delta\theta_{t,j})^2 \end{aligned}$$

Let  $i_1 = 1, j_1 = 2, i_2 = 3, j_2 = 4, \dots, i_{\lceil \frac{d}{2} \rceil} = 2 \lceil \frac{d}{2} \rceil - 1, j_{\lceil \frac{d}{2} \rceil} = 2 \lceil \frac{d}{2} \rceil$ . For any  $i_k, j_k$  ( $k \in [\lceil \frac{d}{2} \rceil]$ ),  $\Delta\theta_{t,i_k} - \Delta\theta_{t,j_k} | \boldsymbol{\mu}_t \sim \mathcal{N}(0, \frac{2}{t+\omega})$  and they are mutually independent among  $k$ .

Fix  $\boldsymbol{\mu}_t$ , and then we define event

$$\mathcal{G}_t \triangleq \left\{ |\Delta\theta_{t,i_k} - \Delta\theta_{t,j_k}| \geq 0.3 \sqrt{\frac{2}{t+\omega}}, \forall k \in \left[ \left\lceil \frac{d}{2} \right\rceil \right] \right\}.$$

From the c.d.f. of Gaussian distributions, we have

$$\Pr[\mathcal{G}_t] \geq (75\%)^{\lceil \frac{d}{2} \rceil}.$$

From this, we get

$$\begin{aligned} \Pr[\mathcal{F}_t \cap \mathcal{G}_t] &\geq 1 - \Pr[\bar{\mathcal{F}}_t] - \Pr[\bar{\mathcal{G}}_t] \\ &\geq 1 - (1 - (99.7\%)^d) - (1 - (75\%)^{\lceil \frac{d}{2} \rceil}) \\ &\geq (99.7\%)^d + (75\%)^{\lceil \frac{d}{2} \rceil} - 1. \end{aligned}$$

When  $d \leq 18$ ,  $\Pr[\mathcal{F}_t \cap \mathcal{G}_t] \geq (99.7\%)^d + (75\%)^{\lceil \frac{d}{2} \rceil} - 1 > 0$ .

**Step (iii).** We bound the expected regret by considering the event  $\mathcal{E}_t \cap \mathcal{F}_t \cap \mathcal{G}_t$  and  $t \geq \tilde{t}$ . Specifically,

$$\begin{aligned} &\sum_{t=1}^T \mathbb{E}_{\boldsymbol{\mu}_t \sim \mathcal{N}(0, \frac{t}{(t+\omega)\omega} I)} \left[ \mathbb{E}_{\boldsymbol{\theta}^* | \boldsymbol{\mu}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \frac{1}{t+\omega} I)} [f(\boldsymbol{w}^*) - f(\boldsymbol{w}_t)] \right] \\ &\geq \sum_{t=\tilde{t}}^T \mathbb{E}_{\boldsymbol{\mu}_t \sim \mathcal{N}(0, \frac{t}{(t+\omega)\omega} I)} \left[ \mathbb{E}_{\boldsymbol{\theta}^* | \boldsymbol{\mu}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \frac{1}{t+\omega} I)} [f(\boldsymbol{w}^*) - f(\boldsymbol{w}_t)] | \mathcal{E}_t \right] \Pr[\mathcal{E}_t] \\ &\geq \sum_{t=\tilde{t}}^T \mathbb{E}_{\boldsymbol{\mu}_t \sim \mathcal{N}(0, \frac{t}{(t+\omega)\omega} I)} \left[ \mathbb{E}_{\boldsymbol{\theta}^* | \boldsymbol{\mu}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \frac{1}{t+\omega} I)} [f(\boldsymbol{w}^*) - f(\boldsymbol{w}_t) | \mathcal{F}_t \cap \mathcal{G}_t] \Pr[\mathcal{F}_t \cap \mathcal{G}_t | \mathcal{E}_t] \right] \Pr[\mathcal{E}_t] \\ &\geq \sum_{t=\tilde{t}}^T \mathbb{E}_{\boldsymbol{\mu}_t \sim \mathcal{N}(0, \frac{t}{(t+\omega)\omega} I)} \left[ \mathbb{E}_{\boldsymbol{\theta}^* | \boldsymbol{\mu}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \frac{1}{t+\omega} I)} \left[ \frac{1}{4\rho d} \sum_{1 \leq i < j \leq d} (\Delta\theta_{t,i} - \Delta\theta_{t,j})^2 | \mathcal{F}_t \cap \mathcal{G}_t \cap \mathcal{E}_t \right] \cdot \Pr[\mathcal{F}_t \cap \mathcal{G}_t] \right] \Pr[\mathcal{E}_t] \\ &\geq \sum_{t=\tilde{t}}^T \mathbb{E}_{\boldsymbol{\mu}_t \sim \mathcal{N}(0, \frac{t}{(t+\omega)\omega} I)} \left[ \mathbb{E}_{\boldsymbol{\theta}^* | \boldsymbol{\mu}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \frac{1}{t+\omega} I)} \left[ \frac{\lceil \frac{d}{2} \rceil}{4\rho d} \cdot \frac{0.3^2 \cdot 2}{t+\omega} | \mathcal{F}_t \cap \mathcal{G}_t \cap \mathcal{E}_t \right] \cdot \left( (99.7\%)^d + (75\%)^{\lceil \frac{d}{2} \rceil} - 1 \right) \right] (99.7\%)^d \\ &= \sum_{t=\tilde{t}}^T \frac{0.01125}{\rho(t+\omega)} \left( (99.7\%)^d + (75\%)^{\lceil \frac{d}{2} \rceil} - 1 \right) (99.7\%)^d \\ &= \frac{0.01125 \left( (99.7\%)^d + (75\%)^{\lceil \frac{d}{2} \rceil} - 1 \right) (99.7\%)^d}{\rho} \ln \left( \frac{t+\omega}{\tilde{t}+\omega} \right). \end{aligned}$$

In the following, we consider an intrinsic bound of the expected regret and set the problem parameters to proper quantities. Since the expected regret is upper bounded by  $\Delta_{\max} T$  and

$$\Delta_{\max} = f_{\max} - f_{\min}$$



$$\begin{aligned}
&\geq f\left(\frac{1}{d}\mathbf{1}\right) - (\theta_{\min}^* - \rho) \\
&= \frac{1}{d} \sum_{i=1}^d \theta_i^* - \rho \frac{1}{d} - (\theta_{\min}^* - \rho) \\
&= \frac{1}{d} \sum_{i=1}^d \theta_i^* - \theta_{\min}^* + \frac{d-1}{d} \rho \\
&> 0,
\end{aligned}$$

we conclude that the expected regret is lower bounded by

$$\min \left\{ \frac{0.01125 \left( (99.7\%)^d + (75\%)^{\lceil \frac{d}{2} \rceil} - 1 \right) (99.7\%)^d}{\rho} \ln \left( \frac{t + \omega}{\tilde{t} + m} \right), \left( \frac{1}{d} \sum_{i=1}^d \theta_i^* - \theta_{\min}^* + \frac{d-1}{d} \rho \right) T \right\}.$$

Choose  $\omega = 36(d-1)^2 T$  and  $\rho = \frac{1}{\sqrt{T}}$ . According to Eqs. (16) and (17), we have  $t_1 = t_2 = \tilde{t} = 1$ . Therefore, for  $d \leq 18$ , the expected regret is lower bounded by  $\Omega(\sqrt{T})$ . □

## C Proof for CMCB-SB

### C.1 Proof of Theorem 3

*Proof.* (Theorem 3) Denote  $\Delta_t = f(\mathbf{w}^*) - f(\mathbf{w}_t)$ ,  $\bar{f}_t(\mathbf{w}) = \mathbf{w}^\top \hat{\boldsymbol{\theta}}_{t-1} + E_t(\mathbf{w}) - \rho \mathbf{w}^\top \bar{\Sigma}_{t-1} \mathbf{w}$  and  $G_t$  the matrix whose  $ij$ -th entry is  $g_{ij}(t)$ .

For  $t \geq d^2 + 1$ , suppose that event  $\mathcal{G}_t \cap \mathcal{H}_t$  occurs. Define  $h_t(\mathbf{w}) \triangleq E_t(\mathbf{w}) + \rho \mathbf{w}^\top G_t \mathbf{w}$  for any  $\mathbf{w} \in \Delta_d^c$ . Then, we have

$$0 \leq \bar{f}_t(\mathbf{w}) - f(\mathbf{w}) \leq 2h_t(\mathbf{w}).$$

According to the selection strategy of  $\mathbf{w}_t$ , we obtain

$$\begin{aligned}
f(\mathbf{w}^*) - f(\mathbf{w}_t) &\leq \bar{f}_t(\mathbf{w}^*) - f(\mathbf{w}_t) \\
&\leq \bar{f}_t(\mathbf{w}_t) - f(\mathbf{w}_t) \\
&\leq 2h_t(\mathbf{w}_t)
\end{aligned}$$

Thus, for any  $T \geq d^2 + 1$ , we have that

$$\begin{aligned}
&\sum_{t=d^2+1}^T (f(\mathbf{w}^*) - f(\mathbf{w}_t)) \\
&\leq 2 \sum_{t=d^2+1}^T h_t(\mathbf{w}_t) \\
&= 2 \sum_{t=d^2+1}^T \left( \sqrt{2\beta(\delta_t)} \sqrt{\mathbf{w}^\top D_{t-1}^{-1} (\lambda \Lambda_{\bar{\Sigma}_t} D_{t-1} + \sum_{s=1}^{t-1} \bar{\Sigma}_{s, \mathbf{w}_s}) D_{t-1}^{-1} \mathbf{w} + \rho \mathbf{w}^\top G_t \mathbf{w}} \right)
\end{aligned}$$

$$\begin{aligned}
&\leq 2\sqrt{2\beta(\delta_T)} \sum_{t=d^2+1}^T \sqrt{\mathbf{w}^\top D_{t-1}^{-1} (\lambda \Lambda_{\bar{\Sigma}_t} D_{t-1} + \sum_{s=1}^{t-1} \bar{\Sigma}_{s, \mathbf{w}_s}) D_{t-1}^{-1} \mathbf{w} + 2\rho \sum_{t=d^2+1}^T \mathbf{w}^\top G_t \mathbf{w}} \\
&\leq 2\sqrt{2\beta(\delta_T)} \sqrt{T \cdot \sum_{t=d^2+1}^T \left( \mathbf{w}^\top D_{t-1}^{-1} (\lambda \Lambda_{\bar{\Sigma}_t} D_{t-1} + \sum_{s=1}^{t-1} \bar{\Sigma}_{s, \mathbf{w}_s}) D_{t-1}^{-1} \mathbf{w} \right)} \\
&\quad + 2\rho \sum_{t=d^2+1}^T \mathbf{w}^\top G_t \mathbf{w} \\
&\leq 2\sqrt{2\beta(\delta_T)} \sqrt{T} \sqrt{\underbrace{\lambda \sum_{t=d^2+1}^T (\mathbf{w}^\top D_{t-1}^{-1} \Lambda_{\bar{\Sigma}_t} \mathbf{w})}_{\Gamma_1} + \underbrace{\sum_{t=d^2+1}^T \left( \mathbf{w}^\top D_{t-1}^{-1} \sum_{s=1}^{t-1} \bar{\Sigma}_{s, \mathbf{w}_s} D_{t-1}^{-1} \mathbf{w} \right)}_{\Gamma_2}} \\
&\quad + 2\rho \underbrace{\sum_{t=d^2+1}^T \mathbf{w}^\top G_t \mathbf{w}}_{\Gamma_3}
\end{aligned}$$

Let  $\Sigma_{ij}^+ = \Sigma_{ij} \vee 0$  for any  $i, j, \in [d]$ . We first address  $\Gamma_3$ . Since for any  $t \geq d^2 + 1$  and  $i, j \in [d]$ ,

$$\begin{aligned}
g_{ij}(t) &= 16 \left( \frac{3 \ln t}{N_{ij}(t-1)} \vee \sqrt{\frac{3 \ln t}{N_{ij}(t-1)}} \right) + \sqrt{\frac{48 \ln^2 t}{N_{ij}(t-1) N_i(t-1)}} + \sqrt{\frac{36 \ln^2 t}{N_{ij}(t-1) N_j(t-1)}} \\
&\leq 48 \frac{\ln t}{\sqrt{N_{ij}(t-1)}} + \sqrt{\frac{48 \ln^2 t}{N_{ij}(t-1)}} + \sqrt{\frac{36 \ln^2 t}{N_{ij}(t-1)}} \\
&\leq 61 \frac{\ln t}{\sqrt{N_{ij}(t-1)}},
\end{aligned}$$

we can bound  $\Gamma_3$  as follows

$$\begin{aligned}
\Gamma_3 &= \sum_{t=d^2+1}^T \mathbf{w}^\top G_t \mathbf{w} \\
&= \sum_{t=d^2+1}^T \sum_{i, j \in [d]} g_{ij}(t) w_{t,i} w_{t,j} \\
&\leq 61 \sum_{i, j \in [d]} \sum_{t=d^2+1}^T \frac{\ln t}{\sqrt{N_{ij}(t-1)}} w_{t,i} w_{t,j} \\
&\leq 61 \ln T \sum_{i, j \in [d]} \sum_{t=d^2+1}^T \frac{w_{t,i} w_{t,j}}{\sqrt{\sum_{s=1}^{t-1} w_{s,i} w_{s,j}}} \\
&\leq 61 \ln T \sum_{i, j \in [d]} \sum_{t=d^2+1}^T \frac{w_{t,i} w_{t,j}}{\sqrt{\sum_{s=1}^{t-1} w_{s,i} w_{s,j}}}
\end{aligned}$$

$$\begin{aligned}
&\leq 122 \ln T \sum_{i,j \in [d]} \sqrt{\sum_{t=1}^T w_{t,i} w_{t,j}} \\
&\leq 122 \ln T \sqrt{d^2 \sum_{i,j \in [d]} \sum_{t=1}^T w_{t,i} w_{t,j}} \\
&= 122 \ln T \sqrt{d^2 \sum_{t=1}^T \sum_{i,j \in [d]} w_{t,i} w_{t,j}} \\
&= 122 \ln T \sqrt{d^2 \sum_{t=1}^T \left( \sum_{i \in [d]} w_{t,i} \right)^2} \\
&\leq 122d \ln T \sqrt{T}
\end{aligned}$$

Next, we obtain a bound for  $\Gamma_1$ .

$$\begin{aligned}
\Gamma_1 &= \sum_{t=d^2+1}^T (\mathbf{w}^\top D_{t-1}^{-1} \Lambda_{\Sigma_t} \mathbf{w}) \\
&= \sum_{t=d^2+1}^T \sum_{i \in [d]} \frac{\bar{\Sigma}_{t,ii}}{N_i(t-1)} w_{t,i}^2 \\
&\leq \sum_{t=d^2+1}^T \sum_{i \in [d]} \frac{\bar{\Sigma}_{t,ii}}{\sum_{s=1}^{t-1} w_{s,i}} w_{t,i}^2 \\
&\leq \sum_{i \in [d]} \sum_{t=d^2+1}^T \frac{\Sigma_{ii} + 2g_{ii}(t)}{\sum_{s=1}^{t-1} w_{s,i}} w_{t,i}^2 \\
&\leq \sum_{i \in [d]} \left( \sum_{t=d^2+1}^T \frac{\Sigma_{ii}}{\sum_{s=1}^{t-1} w_{s,i}} w_{t,i} + 122 \sum_{t=d^2+1}^T \frac{\frac{\ln t}{\sqrt{N_i(t-1)}}}{\sum_{s=1}^{t-1} w_{s,i}} w_{t,i} \right) \\
&\leq \sum_{i \in [d]} \left( \Sigma_{ii} \sum_{t=d^2+1}^T \frac{1}{\sum_{s=1}^{t-1} w_{s,i}} w_{t,i} + 122 \ln T \sum_{t=d^2+1}^T \frac{1}{(\sum_{s=1}^{t-1} w_{s,i})^{\frac{3}{2}}} w_{t,i} \right) \\
&\leq \sum_{i \in [d]} \left( \Sigma_{ii}^+ \ln \left( \sum_{t=1}^T w_{t,i} \right) + 244 \ln T \frac{1}{\sqrt{\sum_{t=1}^{d^2} w_{t,i}}} \right) \\
&\leq \sum_{i \in [d]} \left( \Sigma_{ii}^+ \ln T + 244 \ln T \frac{1}{\sqrt{\sum_{t=1}^{d^2} w_{t,i}}} \right)
\end{aligned}$$

Finally, we bound  $\Gamma_2$ .

$\Gamma_2$

$$\begin{aligned}
&= \sum_{t=d^2+1}^T \left( \mathbf{w}_t^\top D_{t-1}^{-1} \left( \sum_{s=1}^{t-1} \bar{\Sigma}_{s, \mathbf{w}_s} \right) D_{t-1}^{-1} \mathbf{w}_t \right) \\
&= \sum_{i,j \in [d]} \sum_{t=d^2+1}^T \frac{\sum_{s=1}^{t-1} \bar{\Sigma}_{s,ij} \mathbb{I}\{w_{s,i}, w_{s,j} > 0\}}{N_i(t-1)N_j(t-1)} w_{t,i} w_{t,j} \\
&\leq \sum_{i,j \in [d]} \sum_{t=d^2+1}^T \frac{\sum_{s=1}^{t-1} \bar{\Sigma}_{s,ij} \mathbb{I}\{w_{s,i}, w_{s,j} > 0\}}{N_{ij}^2(t-1)} w_{t,i} w_{t,j} \\
&\leq \sum_{t=d^2+1}^T \sum_{i,j \in [d]} \frac{\sum_{s=1}^{t-1} (\Sigma_{ij} + 2g_{ij}(s)) \mathbb{I}\{w_{s,i}, w_{s,j} > 0\}}{N_{ij}^2(t-1)} w_{t,i} w_{t,j} \\
&= \sum_{i,j \in [d]} \left( \sum_{t=d^2+1}^T \frac{\Sigma_{ij} \sum_{s=1}^{t-1} \mathbb{I}\{w_{s,i}, w_{s,j} > 0\}}{N_{ij}^2(t-1)} w_{t,i} w_{t,j} + \sum_{t=d^2+1}^T \frac{2 \sum_{s=1}^{t-1} g_{ij}(s) \mathbb{I}\{w_{s,i}, w_{s,j} > 0\}}{N_{ij}^2(t-1)} w_{t,i} w_{t,j} \right) \\
&\leq \sum_{i,j \in [d]} \left( \Sigma_{ij} \sum_{t=d^2+1}^T \frac{1}{N_{ij}(t-1)} w_{t,i} w_{t,j} + 122 \sum_{t=d^2+1}^T \frac{\sum_{s=1}^{t-1} \frac{\ln s}{\sqrt{N_{ij}(s-1)}} \mathbb{I}\{w_{s,i}, w_{s,j} > 0\}}{N_{ij}^2(t-1)} w_{t,i} w_{t,j} \right) \\
&\leq \sum_{i,j \in [d]} \left( \Sigma_{ij} \sum_{t=d^2+1}^T \frac{1}{\sum_{s=1}^{t-1} w_{s,i} w_{s,j}} w_{t,i} w_{t,j} + 122 \sum_{t=d^2+1}^T \frac{\ln t \sum_{s=1}^{t-1} \frac{1}{\sqrt{\sum_{\ell=1}^{s-1} \mathbb{I}\{w_{\ell,i}, w_{\ell,j} > 0\}}} \mathbb{I}\{w_{s,i}, w_{s,j} > 0\}}{N_{ij}^2(t-1)} w_{t,i} w_{t,j} \right) \\
&\leq \sum_{i,j \in [d]} \left( \Sigma_{ij} \sum_{t=d^2+1}^T \frac{1}{\sum_{s=1}^{t-1} w_{s,i} w_{s,j}} w_{t,i} w_{t,j} + 244 \sum_{t=d^2+1}^T \frac{\ln t \sqrt{\sum_{s=1}^{t-1} \mathbb{I}\{w_{s,i}, w_{s,j} > 0\}}}{N_{ij}^2(t-1)} w_{t,i} w_{t,j} \right) \\
&= \sum_{i,j \in [d]} \left( \Sigma_{ij} \sum_{t=d^2+1}^T \frac{1}{\sum_{s=1}^{t-1} w_{s,i} w_{s,j}} w_{t,i} w_{t,j} + 244 \ln T \sum_{t=d^2+1}^T \frac{1}{N_{ij}^{\frac{3}{2}}(t-1)} w_{t,i} w_{t,j} \right) \\
&\leq \sum_{i,j \in [d]} \left( \Sigma_{ij} \sum_{t=d^2+1}^T \frac{1}{\sum_{s=1}^{t-1} w_{s,i} w_{s,j}} w_{t,i} w_{t,j} + 244 \ln T \sum_{t=d^2+1}^T \frac{1}{\left( \sum_{s=1}^{t-1} w_{s,i} w_{s,j} \right)^{\frac{3}{2}}} w_{t,i} w_{t,j} \right) \\
&\leq \sum_{i,j \in [d]} \left( \Sigma_{ij}^+ \ln \left( \sum_{t=1}^T w_{t,i} w_{t,j} \right) + 488 \ln T \frac{1}{\sqrt{\sum_{t=1}^{d^2} w_{t,i} w_{t,j}}} \right) \\
&\leq \sum_{i,j \in [d]} \left( \Sigma_{ij}^+ \ln T + 488 \ln T \frac{1}{\sqrt{\sum_{t=1}^{d^2} w_{t,i} w_{t,j}}} \right)
\end{aligned}$$

Recall that  $\beta(\delta_T) = \ln(T \ln^2 T) + d \ln \ln T + \frac{d}{2} \ln(1 + e/\lambda) = O(\ln T + d \ln \ln T + d \ln(1 + \lambda^{-1}))$ . Combining the bounds of  $\Gamma_1, \Gamma_2$  and  $\Gamma_3$ , we obtain

$$\sum_{t=d^2+1}^T (f(\mathbf{w}^*) - f(\mathbf{w}_t))$$

$$\begin{aligned}
&\leq 2\sqrt{2\beta(\delta_T)}\sqrt{T} \sqrt{\underbrace{\lambda \sum_{t=d^2+1}^T (\mathbf{w}^\top D_{t-1}^{-1} \Lambda_{\bar{\Sigma}_t} \mathbf{w})}_{\Gamma_1} + \underbrace{\sum_{t=d^2+1}^T \left( \mathbf{w}^\top D_{t-1}^{-1} \sum_{s=1}^{t-1} \bar{\Sigma}_{s, \mathbf{w}_s} D_{t-1}^{-1} \mathbf{w} \right)}_{\Gamma_2} + 2\rho \underbrace{\sum_{t=d^2+1}^T \mathbf{w}^\top G_t \mathbf{w}}_{\Gamma_3}} \\
&\leq 2\sqrt{2\beta(\delta_T)}\sqrt{T} \sqrt{\lambda \sum_{i \in [d]} \left( \Sigma_{ii}^+ \ln T + 244 \ln T \frac{1}{\sqrt{\sum_{t=1}^{d^2} w_{t,i}}} \right) + \sum_{i,j \in [d]} \left( \Sigma_{ij}^+ \ln T + 488 \ln T \frac{1}{\sqrt{\sum_{t=1}^{d^2} w_{t,i} w_{t,j}}} \right)} \\
&\quad + 244\rho d \ln T \sqrt{T} \\
&\leq 2\sqrt{2\beta(\delta_T)}\sqrt{T} \sqrt{\lambda \ln T \sum_{i \in [d]} \Sigma_{ii}^+ + \ln T \sum_{i,j \in [d]} \Sigma_{ij}^+ + (244\lambda + 488)d^2 \ln T + 244\rho d \ln T \sqrt{T}} \\
&= O \left( \sqrt{(\ln T + d \ln \ln T + d \ln(1 + \lambda^{-1})) \cdot T} \sqrt{\lambda \ln T \sum_{i \in [d]} \Sigma_{ii}^+ + \ln T \sum_{i,j \in [d]} \Sigma_{ij}^+ + (\lambda + 1)d^2 \ln T + \rho d \ln T \sqrt{T}} \right) \\
&= O \left( \sqrt{d(\ln T + \ln(1 + \lambda^{-1})) \cdot T} \sqrt{\lambda \ln T \sum_{i \in [d]} \Sigma_{ii}^+ + \ln T \sum_{i,j \in [d]} \Sigma_{ij}^+ + (\lambda + 1)d^2 \ln T + \rho d \ln T \sqrt{T}} \right) \\
&= O \left( \sqrt{d \ln T (\ln T + \ln(1 + \lambda^{-1})) \cdot T} \sqrt{\lambda \sum_{i \in [d]} \Sigma_{ii}^+ + \sum_{i,j \in [d]} \Sigma_{ij}^+ + (\lambda + 1)d^2 + \rho d \ln T \sqrt{T}} \right)
\end{aligned}$$

According to Lemmas 1 and 2, for any  $t \geq 2$ , the probability of event  $\neg(\mathcal{G}_t \cap \mathcal{H}_t)$  satisfies

$$\begin{aligned}
\Pr[\neg(\mathcal{G}_t \cap \mathcal{H}_t)] &\leq \frac{10d^2}{t^2} + \frac{1}{t \ln^2 t} \\
&\leq \frac{10d^2}{t \ln^2 t} + \frac{1}{t \ln^2 t} \\
&= \frac{11d^2}{t \ln^2 t}
\end{aligned}$$

Therefore, for any horizon  $T$ , we obtain the regret upper bound

$$\begin{aligned}
\mathbb{E}[\mathcal{R}(T)] &= O(\Delta_{\max}) + \sum_{t=2}^T O(\Delta_{\max} \cdot \Pr[\neg(\mathcal{G}_t \cap \mathcal{H}_t)] + \Delta_t \cdot \mathbb{I}\{\mathcal{G}_t \cap \mathcal{H}_t\}) \\
&= O(\Delta_{\max}) + \sum_{t=2}^T O\left(\Delta_{\max} \cdot \frac{d^2}{t \ln^2 t}\right) + O\left(\sqrt{d \ln T (\ln T + \ln(1 + \lambda^{-1})) T}\right) \\
&\quad \sqrt{\lambda \sum_{i \in [d]} \Sigma_{ii}^+ + \sum_{i,j \in [d]} \Sigma_{ij}^+ + (\lambda + 1)d^2 + \rho d \ln T \sqrt{T}} \\
&= O\left(\sqrt{d \ln T (\ln T + \ln(1 + \lambda^{-1})) T} \sqrt{\lambda \sum_{i \in [d]} \Sigma_{ii}^+ + \sum_{i,j \in [d]} \Sigma_{ij}^+ + (\lambda + 1)d^2 + \rho d \ln T \sqrt{T} + d^2 \Delta_{\max}}\right)
\end{aligned}$$

$$\begin{aligned}
&= O\left(\sqrt{d(\ln(1+\lambda^{-1})+1)\ln^2 T \cdot T} \sqrt{(\lambda+1)\left(\sum_{i,j \in [d]} \Sigma_{ij}^+ + d^2\right) + \rho d \ln T \sqrt{T}}\right) \\
&= O\left(\sqrt{(\lambda+1)(\ln(1+\lambda^{-1})+1)(\|\Sigma\|_+ + d^2)d \ln^2 T \cdot T + \rho d \ln T \sqrt{T}}\right) \\
&= O\left(\sqrt{L(\lambda)(\|\Sigma\|_+ + d^2)d \ln^2 T \cdot T + \rho d \ln T \sqrt{T}}\right)
\end{aligned}$$

where  $L(\lambda) = (\lambda+1)(\ln(1+\lambda^{-1})+1)$  and  $\|\Sigma\|_+ = \sum_{i,j \in [d]} (\Sigma_{ij} \vee 0)$  for any  $i, j \in [d]$ .  $\square$

## C.2 Proof of Theorem 4

*Proof.* First, we construct some instances with  $d \geq 4$ ,  $\frac{2}{d} \leq c \leq \frac{1}{2}$ ,  $\Sigma_* = I$  and  $\theta_t \sim N(\theta^*, I)$ .

Let  $I_J$  be a random instance constructed as follows: we uniformly choose a dimension  $J$  from  $[d]$ , and the expected reward vector  $\theta_J^*$  has  $\frac{1}{2} + \varepsilon$  on its  $J$ -th entry and  $\frac{1}{2}$  elsewhere, where  $\varepsilon \in (0, \frac{1}{2})$  will be specified later. Let  $I_u$  be a uniform instance, where  $\theta_u^*$  has all its entries to be  $\frac{1}{2}$ . Let  $\Pr_J[\cdot]$  and  $\Pr_u[\cdot]$  denote the probabilities under instances  $I_J$  and  $I_u$ , respectively, and let  $\Pr_j[\cdot] = \Pr_J[\cdot | J = j]$ . Analogously,  $E_J[\cdot]$ ,  $E_u[\cdot]$  and  $E_j[\cdot] = E_J[\cdot | J = j]$  denote the expectation operations.

Fix an algorithm  $\mathcal{A}$ . Let  $S_t \in \{\mathbb{R} \cup \{\perp\}\}^d$  be a random variable vector denoting the observations at timestep  $t$ , obtained by running  $\mathcal{A}$ . Here  $\perp$  denotes no observation on this dimension. Let  $Q_\perp$  denote the distribution on support  $\{\perp\}$  which takes value  $\perp$  with probability 1.

In CMCB-SB, if  $w_{t,i} > 0$ , we can observe the reward on the  $i$ -th dimension, i.e.,  $S_{t,i} = \theta_{t,i}$ ; otherwise, if  $w_{t,i} = 0$ , we cannot get observation on the  $i$ -th dimension, i.e.,  $S_{t,i} = \perp$ . Let  $D_J$  be the distribution of observation sequence  $S_1, \dots, S_t$  under instance  $I_J$ , and  $D_j = D_J|_{J=j}$  is the distribution conditioned on  $J = j$ . Let  $D_u$  be the distribution of observation sequence  $S_1, \dots, S_t$  under instance  $I_u$ . For any  $i \in [d]$ , let  $N_i = \sum_{t=1}^T \mathbb{I}\{w_{t,i} > 0\}$  be the number of pulls that has a positive weight on the  $i$ -th dimension, i.e., the number of observations on the  $i$ -th dimension.

Following the analysis procedure of Lemma A.1 in [4], we have

$$\begin{aligned}
KL(D_j \| D_u) &= \sum_{t=1}^T KL(D_u[S_t | S_1, \dots, S_{t-1}] \| D_j[S_t | S_1, \dots, S_{t-1}]) \\
&= \sum_{t=1}^T \sum_{i=1}^d (\Pr[w_{t,i} > 0] \cdot KL(N(\theta_{u,i}^*, 1) \| N(\theta_{j,i}^*, 1)) + \Pr[w_{t,i} = 0] \cdot KL(Q_\perp \| Q_\perp)) \\
&= \sum_{t=1}^T \left( \Pr[w_{t,j} > 0] \cdot KL\left(N\left(\frac{1}{2}, 1\right) \| N\left(\frac{1}{2} + \varepsilon, 1\right)\right) + \sum_{i \neq j} \Pr[w_{t,i} > 0] \cdot KL\left(N\left(\frac{1}{2}, 1\right) \| N\left(\frac{1}{2}, 1\right)\right) \right) \\
&= \frac{1}{2} \varepsilon^2 \cdot \sum_{t=1}^T \Pr[w_{t,j} > 0] \\
&= \frac{1}{2} \varepsilon^2 E_u[N_j]
\end{aligned}$$



Here the first equality comes from the chain rule of entropy [9]. The second equality is due to that given  $S_1, \dots, S_{t-1}$ , if  $w_{t,i} > 0$ , the conditional distribution of  $S_t$  is  $N(\theta_{\cdot,i}^*, 1)$ , where " $\cdot$ " refers to the subscript of instances; otherwise, if  $w_{t,i} = 0$ ,  $S_t$  is  $\perp$  deterministically. The third equality is due to that  $\theta_u^*$  and  $\theta_j^*$  only have one different entry on the  $j$ -th dimension.

Let  $\|\cdot\|$  with subscript  $TV$  denote the total variance distance, and  $KL(\cdot\|\cdot)$  denote the Kullback–Leibler divergence. Using Eq. (28) in the analysis of Lemma A.1 in [4] and Pinsker’s inequality, we have

$$\begin{aligned} E_j[N_j] &\leq E_u[N_j] + T\|D_j - D_u\|_{TV} \\ &\leq E_u[N_j] + T\sqrt{\frac{1}{2}KL(D_j\|D_u)} \\ &= E_u[N_j] + \frac{T\varepsilon}{2}\sqrt{E_u[N_j]} \end{aligned}$$

Let  $m = \lfloor \frac{1}{c} \rfloor \leq \frac{d}{2}$  denote the maximum number of positive entries for a feasible action, i.e., the maximum number of observations for a pull. Performing the above argument for all  $j \in [d]$  and using  $\sum_{j \in [d]} E_u[N_j] \leq mT$ , we have

$$\begin{aligned} \sum_{j \in [d]} E_j[N_j] &\leq \sum_{j \in [d]} E_u[N_j] + \frac{T\varepsilon}{2} \sum_{j \in [d]} \sqrt{E_u[N_j]} \\ &\leq mT + \frac{T\varepsilon}{2} \sqrt{d \sum_{j \in [d]} E_u[N_j]} \\ &\leq mT + \frac{T\varepsilon}{2} \sqrt{dmT} \end{aligned}$$

and thus

$$E_J[N_J] = \frac{1}{d} \sum_{j \in [d]} E_j[N_j] \leq \frac{mT}{d} + \frac{T\varepsilon}{2} \sqrt{\frac{mT}{d}}$$

Letting  $\rho \leq \frac{\varepsilon}{2(1-c)}$ , the expected reward (linear) term dominates  $f(\mathbf{w})$ , and the best action  $\mathbf{w}^*$  under  $I_J$  has the weight 1 on the  $J$ -th entry and 0 elsewhere.

Recall that  $m \leq \frac{1}{c}$ . For each pull that has no weight on the  $J$ -th entry, algorithm  $\mathcal{A}$  must suffer a regret at least

$$\begin{aligned} &(\frac{1}{2} + \varepsilon - \rho) - (\frac{1}{2} - \rho \cdot \frac{1}{m}) \\ &\geq \varepsilon - \frac{m-1}{m} \rho \\ &\geq \varepsilon - \frac{m-1}{m} \cdot \frac{\varepsilon}{2(1-c)} \\ &\geq \varepsilon - \frac{m-1}{m} \cdot \frac{\varepsilon}{2(1-\frac{1}{m})} \\ &= \frac{\varepsilon}{2} \end{aligned}$$

Thus, the regret is lower bounded by

$$\begin{aligned}
E[R(T)] &\geq (T - E_J[N_J]) \cdot \frac{\varepsilon}{2} \\
&\geq \left( T - \frac{mT}{d} - \frac{T\varepsilon}{2} \sqrt{\frac{mT}{d}} \right) \cdot \frac{\varepsilon}{2} \\
&= \Omega \left( T\varepsilon - T\varepsilon^2 \sqrt{\frac{mT}{d}} \right),
\end{aligned}$$

where the last equality is due to  $m \leq \frac{d}{2}$ .

Letting  $\varepsilon = a_0 \sqrt{\frac{d}{Tm}}$  for small enough constant  $a_0$ , we obtain the regret lower bound  $\Omega(\sqrt{\frac{dT}{m}}) = \Omega(\sqrt{cdT})$ . □

## D Proof for CMCB-FB

### D.1 Proof of Theorem 5

In order to prove Theorem 5, we first prove Lemmas 5 and 6, which give the concentrations of covariance and means for CMCB-FB, using different techniques than those for CMCB-SB (Lemmas 1 and 2) and CMCB-FI (Lemmas 3 and 4).

**Lemma 5** (Concentration of Covariance for CMCB-FB). *Consider the CMCB-FB problem and algorithm MC-ETE (Algorithm 3). For any  $t > 0$ , the event*

$$\mathcal{M}_t \triangleq \left\{ |\Sigma_{ij} - \hat{\Sigma}_{ij,t-1}| \leq 5 \|C_\pi^+\| \sqrt{\frac{3 \ln t}{2N_\pi(t)}} \right\}$$

satisfies

$$Pr[\mathcal{M}_t] \geq 1 - \frac{6d^2}{t^2},$$

where  $\|C_\pi^+\| \triangleq \max_{i \in [\tilde{d}]} \left\{ \sum_{j \in [\tilde{d}]} |C_{\pi,ij}^+| \right\}$ .

*Proof.* Let  $\boldsymbol{\sigma} = (\Sigma_{11}, \dots, \Sigma_{dd}, \Sigma_{12}, \dots, \Sigma_{1d}, \Sigma_{23}, \dots, \Sigma_{d,d-1})^\top \in \mathbb{R}^{\tilde{d}}$  denote the column vector that stacks the  $\tilde{d}$  distinct entries in the covariance matrix  $\Sigma$ .

Recall that the  $\tilde{d} \times \tilde{d}$  matrix

$$C_\pi = \begin{bmatrix} w_{1,1}^2 & \dots & w_{1,d}^2 & 2w_{1,1}w_{1,2} & 2w_{1,1}w_{1,3} & \dots & 2w_{1,d-1}w_{1,d} \\ w_{2,1}^2 & \dots & w_{2,d}^2 & 2w_{2,1}w_{2,2} & 2w_{2,1}w_{2,3} & \dots & 2w_{2,d-1}w_{2,d} \\ & \dots & & & & \dots & \\ & & & & & & \\ & & & & & & \\ & & & & & & \\ w_{\tilde{d},1}^2 & \dots & w_{\tilde{d},d}^2 & 2w_{\tilde{d},1}w_{\tilde{d},2} & 2w_{\tilde{d},1}w_{\tilde{d},3} & \dots & 2w_{\tilde{d},d-1}w_{\tilde{d},d} \end{bmatrix},$$

where  $w_{i,j}$  denotes the  $j$ -th entry of portfolio vector  $\mathbf{w}_i$  in design set  $W$ . We use  $C_{\pi,k}$  to denote the  $k$ -th row in matrix  $C_\pi$ .

We recall the feedback structure in algorithm MC-ETE as follows. At each timestep  $t$ , the learner plays an action  $\mathbf{w}_t \in \Delta_d$ , and observes the full-bandit feedback  $y_t = \mathbf{w}_t^\top \boldsymbol{\theta}_t$  with  $\mathbb{E}[y_t] = \mathbf{w}_t^\top \boldsymbol{\theta}^*$  and  $\text{Var}[y_t] = \mathbf{w}_t^\top \Sigma \mathbf{w}_t = \sum_{i \in [d]} w_{t,i}^2 \Sigma_{ii} +$

$\sum_{i,j \in [d], i < j} 2w_{t,i}w_{t,j}\Sigma_{ij}$ . Then, during exploration round  $s$ , where each action in design set  $\pi = \{\mathbf{v}_1, \dots, \mathbf{v}_{\tilde{d}}\}$  is pulled once, the full-bandit feedback  $\mathbf{y}_s$  has mean  $\mathbf{y}(\pi) \triangleq (\mathbf{v}_1^\top \boldsymbol{\theta}^*, \dots, \mathbf{v}_{\tilde{d}}^\top \boldsymbol{\theta}^*)^\top$  and variance  $\mathbf{z}(\pi) \triangleq (\mathbf{v}_1^\top \Sigma \mathbf{v}_1, \dots, \mathbf{v}_{\tilde{d}}^\top \Sigma \mathbf{v}_{\tilde{d}})^\top = (C_{\pi,1}^\top \boldsymbol{\sigma}, \dots, C_{\pi,\tilde{d}}^\top \boldsymbol{\sigma})^\top = C_\pi \boldsymbol{\sigma}$ . For any  $t > 0$ , denote  $\hat{\mathbf{y}}_t$  the empirical mean of  $\mathbf{y}(\pi) \triangleq (\mathbf{v}_1^\top \boldsymbol{\theta}^*, \dots, \mathbf{v}_{\tilde{d}}^\top \boldsymbol{\theta}^*)^\top$  and  $\hat{\mathbf{z}}_t$  the empirical variance of  $\mathbf{z}(\pi) \triangleq (C_{\pi,1}^\top \boldsymbol{\sigma}, \dots, C_{\pi,\tilde{d}}^\top \boldsymbol{\sigma})^\top$ .

Using the Chernoff-Hoeffding inequality for empirical variances (Lemma 1 in [28]), we have that for any  $t > 0$  and  $k \in \tilde{d}$ , with probability at least  $1 - \frac{6\tilde{d}^2}{t^2}$ ,

$$|\hat{z}_{t-1,k} - z_k| = |\hat{z}_{t-1,k} - C_{\pi,k}^\top \boldsymbol{\sigma}| \leq 5 \sqrt{\frac{3 \ln t}{2N_\pi(t-1)}}.$$

Let  $C_{\pi,i}^+$  denote the  $i$ -th row of matrix  $C_\pi^+$  and  $C_{\pi,ik}^+$  denote the  $ik$ -th entry of matrix  $C_\pi^+$ . Since  $C_\pi^+ C_\pi = I$ , for any  $i \in [\tilde{d}]$ , we have

$$\begin{aligned} |\hat{\sigma}_{t-1,i} - \sigma_i| &= |(C_{\pi,i}^+)^{\top} \hat{\mathbf{z}}_t - (C_{\pi,i}^+)^{\top} C_\pi \boldsymbol{\sigma}| \\ &\leq \sum_{k \in [\tilde{d}]} |C_{\pi,ik}^+| |\hat{z}_{t-1,k} - C_{\pi,k}^\top \boldsymbol{\sigma}| \\ &\leq 5 \sum_{k \in [\tilde{d}]} |C_{\pi,ik}^+| \sqrt{\frac{3 \ln t}{2N_\pi(t)}} \\ &\leq 5 \|C_\pi^+\| \sqrt{\frac{3 \ln t}{2N_\pi(t)}} \end{aligned}$$

In addition, since  $\boldsymbol{\sigma} = (\Sigma_{11}, \dots, \Sigma_{dd}, \Sigma_{12}, \dots, \Sigma_{1d}, \Sigma_{23}, \dots, \Sigma_{d,d-1})^\top \in \mathbb{R}^{\tilde{d}}$  is the column vector stacking the distinct entries in  $\Sigma$ , we obtain the lemma.  $\square$

**Lemma 6** (Concentration of Means for CMCB-FB). *Consider the CMCB-FB problem and algorithm MC-ETE (Algorithm 3). Let  $\delta_t > 0$ ,  $\lambda > 0$ ,  $\beta(\delta_t) = \ln(1/\delta_t) + \ln \ln t + \frac{\tilde{d}}{2} \ln(1 + e/\lambda)$  and  $E_t(\mathbf{w}) = \sqrt{2\beta(\delta_t)} \sqrt{\mathbf{w}^\top B_\pi^+ D_{t-1}^{-1} (\lambda \Lambda_{\Sigma_\pi} D_{t-1} + \sum_{s=1}^{N_\pi(t-1)} \Sigma_\pi) D_{t-1}^{-1} (B_\pi^+)^\top \mathbf{w}}$ , where  $\Sigma_\pi = \text{diag}(\mathbf{v}_1^\top \Sigma \mathbf{v}_1, \dots, \mathbf{v}_{\tilde{d}}^\top \Sigma \mathbf{v}_{\tilde{d}})$ . Then, the event  $\mathcal{N}_t \triangleq \{|\mathbf{w}^\top \boldsymbol{\theta}^* - \mathbf{w}^\top \hat{\boldsymbol{\theta}}_{t-1}| \leq E_t(\mathbf{w}), \forall \mathbf{w} \in \mathcal{D}\}$  satisfies  $\Pr[\mathcal{N}_t] \geq 1 - \delta_t$ .*

*Proof.* Recall that the  $\tilde{d} \times d$  matrix  $B_\pi = [\mathbf{v}_1^\top; \dots; \mathbf{v}_{\tilde{d}}^\top]$ , and  $B_\pi^+$  is the Moore–Penrose pseudoinverse of  $B_\pi$ . Since  $B_\pi$  is of full column rank,  $B_\pi^+$  satisfies  $B_\pi^+ B_\pi = I$ .

We recall the feedback structure in algorithm MC-ETE as follows. At each timestep  $t$ , the learner plays an action  $\mathbf{w}_t \in \Delta_d$ , and observes the full-bandit feedback  $y_t = \mathbf{w}_t^\top \boldsymbol{\theta}_t$  such that  $\mathbb{E}[y_t] = \mathbf{w}_t^\top \boldsymbol{\theta}^*$  and  $\text{Var}[y_t] = \mathbf{w}_t^\top \Sigma \mathbf{w}_t = \sum_{i \in [d]} w_{t,i}^2 \Sigma_{ii} + \sum_{i,j \in [d], i < j} 2w_{t,i}w_{t,j}\Sigma_{ij}$ . Then, during exploration round  $s$ , where each action in design set  $\pi = \{\mathbf{v}_1, \dots, \mathbf{v}_{\tilde{d}}\}$  is pulled once, the full-bandit feedback  $\mathbf{y}_s$  has mean  $\mathbf{y}(\pi) \triangleq (\mathbf{v}_1^\top \boldsymbol{\theta}^*, \dots, \mathbf{v}_{\tilde{d}}^\top \boldsymbol{\theta}^*)^\top$  and variance  $\mathbf{z}(\pi) \triangleq (\mathbf{v}_1^\top \Sigma \mathbf{v}_1, \dots, \mathbf{v}_{\tilde{d}}^\top \Sigma \mathbf{v}_{\tilde{d}})^\top = (C_{\pi,1}^\top \boldsymbol{\sigma}, \dots, C_{\pi,\tilde{d}}^\top \boldsymbol{\sigma})^\top = C_\pi \boldsymbol{\sigma}$ . For any  $t > 0$ ,  $\hat{\mathbf{y}}_t$  is the empirical mean of  $\mathbf{y}(\pi) \triangleq (\mathbf{v}_1^\top \boldsymbol{\theta}^*, \dots, \mathbf{v}_{\tilde{d}}^\top \boldsymbol{\theta}^*)^\top$  and  $\hat{\mathbf{z}}_t$  is the empirical variance of  $\mathbf{z}(\pi) \triangleq (C_{\pi,1}^\top \boldsymbol{\sigma}, \dots, C_{\pi,\tilde{d}}^\top \boldsymbol{\sigma})^\top$ .

Let  $D_t$  be the  $\tilde{d} \times \tilde{d}$  diagonal matrix such that  $D_{t,ii} = N_\pi(t)$  for any  $i \in [\tilde{d}]$  and  $t > 0$ . Let  $\Sigma_\pi$  be a  $\tilde{d} \times \tilde{d}$  diagonal matrix such that  $\Sigma_{\pi,ii} = \mathbf{w}_i^\top \Sigma \mathbf{w}_i$  for any  $i \in [\tilde{d}]$ , and thus  $\Lambda_{\Sigma_\pi} = \Sigma_\pi$ . Let  $\boldsymbol{\varepsilon}_t$  be the vector such that  $\boldsymbol{\eta}_t = \Sigma^{\frac{1}{2}} \boldsymbol{\varepsilon}_t$  for any timestep  $t > 0$ . Let  $\zeta_s = (\mathbf{v}_1^\top \Sigma^{\frac{1}{2}} \boldsymbol{\varepsilon}_{s,1}, \dots, \mathbf{v}_{\tilde{d}}^\top \Sigma^{\frac{1}{2}} \boldsymbol{\varepsilon}_{s,\tilde{d}})^\top$ , where  $\Sigma^{\frac{1}{2}} \boldsymbol{\varepsilon}_{s,k}$  denotes the noise of the  $k$ -th sample in the  $s$ -th exploration round.

Note that, the following analysis of  $|\mathbf{w}^\top(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{t-1})|$  and the constructions (definitions) of the noise  $\zeta_s$ , matrices  $S_t, V_t$  and super-martingale  $M_t^u$  are different from those in CMCB-SB (Lemmas 1 and 2) and CMCB-FI (Lemmas 3 and 4).

For any  $\mathbf{w} \in \Delta_d$ , we have

$$\begin{aligned}
|\mathbf{w}^\top(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{t-1})| &= |\mathbf{w}^\top(B_\pi^+ B_\pi \boldsymbol{\theta}^* - B_\pi^+ \hat{\mathbf{y}}_{t-1})| \\
&= |\mathbf{w}^\top B_\pi^+ (\mathbf{y}(\pi) - \hat{\mathbf{y}}_{t-1})| \\
&= \left| -\mathbf{w}^\top B_\pi^+ D_{t-1}^{-1} \sum_{s=1}^{N_\pi(t-1)} \zeta_s \right| \\
&= \left| -\mathbf{w}^\top B_\pi^+ D_{t-1}^{-1} \left( D + \sum_{s=1}^{N_\pi(t-1)} \Sigma_\pi \right)^{\frac{1}{2}} \left( D + \sum_{s=1}^{N_\pi(t-1)} \Sigma_\pi \right)^{-\frac{1}{2}} \sum_{s=1}^{N_\pi(t-1)} \zeta_s \right| \\
&\leq \sqrt{\mathbf{w}^\top B_\pi^+ D_{t-1}^{-1} \left( D + \sum_{s=1}^{N_\pi(t-1)} \Sigma_\pi \right) D_{t-1}^{-1} (B_\pi^+)^\top \mathbf{w} \cdot \left\| \sum_{s=1}^{N_\pi(t-1)} \zeta_s \right\|_{(D + \sum_{s=1}^{N_\pi(t-1)} \Sigma_\pi)^{-1}}}
\end{aligned}$$

Let  $S_t = \sum_{s=1}^{N_\pi(t-1)} \zeta_s$ ,  $V_t = \sum_{s=1}^{N_\pi(t-1)} \Sigma_\pi$  and  $I_{D+V_t} = \frac{1}{2} \|S_t\|_{(D+V_t)^{-1}}^2$ . Then, we have

$$\left\| \sum_{s=1}^{N_\pi(t-1)} \zeta_s \right\|_{(D + \sum_{s=1}^{N_\pi(t-1)} \Sigma_\pi)^{-1}} = \|S_t\|_{(D+V_t)^{-1}} = \sqrt{2I_{D+V_t}}.$$

Since  $D \preceq \lambda \Lambda_{\Sigma_\pi} D_{t-1}$ , we get

$$\begin{aligned}
|\mathbf{w}^\top(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{t-1})| &\leq \sqrt{\mathbf{w}^\top B_\pi^+ D_{t-1}^{-1} D D_{t-1}^{-1} (B_\pi^+)^\top \mathbf{w} + \mathbf{w}^\top B_\pi^+ D_{t-1}^{-1} \left( \sum_{s=1}^{N_\pi(t-1)} \Sigma_\pi \right) D_{t-1}^{-1} (B_\pi^+)^\top \mathbf{w} \cdot \sqrt{2I_{D+V_t}}} \\
&\leq \sqrt{\lambda \mathbf{w}^\top B_\pi^+ D_{t-1}^{-1} \Lambda_\Sigma (B_\pi^+)^\top \mathbf{w} + \mathbf{w}^\top B_\pi^+ D_{t-1}^{-1} \left( \sum_{s=1}^{N_\pi(t-1)} \Sigma_\pi \right) D_{t-1}^{-1} (B_\pi^+)^\top \mathbf{w} \cdot \sqrt{2I_{D+V_t}}} \\
&= \sqrt{\mathbf{w}^\top B_\pi^+ D_{t-1}^{-1} \left( \lambda \Lambda_\Sigma D_{t-1} + \sum_{s=1}^{N_\pi(t-1)} \Sigma_\pi \right) D_{t-1}^{-1} (B_\pi^+)^\top \mathbf{w} \cdot \sqrt{2I_{D+V_t}}}
\end{aligned}$$

Thus,

$$\begin{aligned}
&\Pr \left[ |\mathbf{w}^\top(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{t-1})| > \sqrt{2\beta(\delta_t)} \sqrt{\mathbf{w}^\top B_\pi^+ D_{t-1}^{-1} (\lambda \Lambda_{\Sigma_\pi} D_{t-1} + \sum_{s=1}^{N_\pi(t-1)} \Sigma_\pi) D_{t-1}^{-1} (B_\pi^+)^\top \mathbf{w}} \right] \\
&\leq \Pr \left[ \sqrt{\mathbf{w}^\top B_\pi^+ D_{t-1}^{-1} \left( \lambda \Lambda_\Sigma D_{t-1} + \sum_{s=1}^{N_\pi(t-1)} \Sigma_\pi \right) D_{t-1}^{-1} (B_\pi^+)^\top \mathbf{w} \cdot \sqrt{2I_{D+V_t}}} \right]
\end{aligned}$$

$$\begin{aligned}
&> \sqrt{2\beta(\delta_t)} \sqrt{\mathbf{w}^\top B_\pi^+ D_{t-1}^{-1} \left( \lambda \Lambda_{\Sigma_\pi} D_{t-1} + \sum_{s=1}^{N_\pi(t-1)} \Sigma_\pi \right) D_{t-1}^{-1} (B_\pi^+)^\top \mathbf{w}} \\
&= \Pr [I_{D+V_t} > \beta(\delta_t)]
\end{aligned}$$

Hence, to prove Eq. (4), it suffices to prove

$$\Pr [I_{D+V_t} > \beta(\delta_t)] \leq \delta_t. \quad (18)$$

To prove Eq. (18), we introduce some notions. Let  $\mathbf{u} \in \mathbb{R}^d$  be a multivariate Gaussian random variable with mean  $\mathbf{0}$  and covariance  $D^{-1}$ , which is independent of all the other random variables and we use  $\varphi(\mathbf{u})$  denote its probability density function. Let

$$\begin{aligned}
P_s^{\mathbf{u}} &= \exp \left( \mathbf{u}^\top \zeta_s - \frac{1}{2} \mathbf{u}^\top \Sigma_\pi \mathbf{u} \right), \\
M_t^{\mathbf{u}} &\triangleq \exp \left( \mathbf{u}^\top S_t - \frac{1}{2} \|\mathbf{u}\|_{V_t}^2 \right),
\end{aligned}$$

and

$$M_t \triangleq \mathbb{E}_{\mathbf{u}} [M_t^{\mathbf{u}}] = \int_{\mathbb{R}^d} \exp \left( \mathbf{u}^\top S_t - \frac{1}{2} \|\mathbf{u}\|_{V_t}^2 \right) \varphi(\mathbf{u}) d\mathbf{u},$$

where  $s = 1, \dots, N_\pi(t-1)$  is the index of exploration round. We have  $M_t^{\mathbf{u}} = \prod_{s=1}^{N_\pi(t-1)} P_s^{\mathbf{u}}$ . In the following, we prove  $\mathbb{E}[M_t] \leq 1$ .

For any timestep  $t$ ,  $\eta_t = \Sigma^{\frac{1}{2}} \boldsymbol{\varepsilon}_t$  is  $\Sigma$ -sub-Gaussian and  $\eta_t$  is independent among different timestep  $t$ . Then,  $\zeta_s = (\mathbf{v}_1^\top \Sigma^{\frac{1}{2}} \boldsymbol{\varepsilon}_{s,1}, \dots, \mathbf{v}_{\bar{d}}^\top \Sigma^{\frac{1}{2}} \boldsymbol{\varepsilon}_{s,\bar{d}})^\top$  is  $\Sigma_\pi$ -sub-Gaussian. According to the sub-Gaussian property,  $\zeta_s$  satisfies

$$\forall \mathbf{v} \in \mathbb{R}^d, \mathbb{E} \left[ e^{\mathbf{v}^\top \zeta_s} \right] \leq e^{\frac{1}{2} \mathbf{v}^\top \Sigma_\pi \mathbf{v}},$$

which is equivalent to

$$\forall \mathbf{v} \in \mathbb{R}^d, \mathbb{E} \left[ e^{\mathbf{v}^\top \zeta_s - \frac{1}{2} \mathbf{v}^\top \Sigma_\pi \mathbf{v}} \right] \leq 1.$$

Let  $\mathcal{J}_s$  be the  $\sigma$ -algebra  $\sigma(W, \zeta_1, \dots, W, \zeta_{s-1}, \pi)$ . Thus, we have

$$\mathbb{E} [P_s^{\mathbf{u}} | \mathcal{J}_s] = \mathbb{E} \left[ \exp \left( \mathbf{u}^\top \zeta_s - \frac{1}{2} \mathbf{u}^\top \Sigma_\pi \mathbf{u} \right) | \mathcal{J}_s \right] \leq 1.$$

Then, we can obtain

$$\begin{aligned}
\mathbb{E} [M_t^{\mathbf{u}} | \mathcal{J}_{N_\pi(t-1)}] &= \mathbb{E} \left[ \prod_{s=1}^{N_\pi(t-1)} P_s^{\mathbf{u}} | \mathcal{J}_{N_\pi(t-1)} \right] \\
&= \left( \prod_{s=1}^{N_\pi(t-2)} P_s^{\mathbf{u}} \right) \mathbb{E} \left[ P_{N_\pi(t-1)}^{\mathbf{u}} | \mathcal{J}_{N_\pi(t-1)} \right] \\
&\leq M_{t-1}^{\mathbf{u}},
\end{aligned}$$

which implies that  $M_t^{\mathbf{u}}$  is a super-martingale and  $\mathbb{E}[M_t^{\mathbf{u}} | \mathbf{u}] \leq 1$ . Thus,

$$\mathbb{E}[M_t] = \mathbb{E}_{\mathbf{u}} [\mathbb{E}[M_t^{\mathbf{u}} | \mathbf{u}]] \leq 1.$$

According to Lemma 9 in [1], we have

$$M_t \triangleq \int_{\mathbb{R}^d} \exp \left( \mathbf{u}^\top S_t - \frac{1}{2} \|\mathbf{u}\|_{V_t}^2 \right) \varphi(\mathbf{u}) d\mathbf{u} = \sqrt{\frac{\det D}{\det(D + V_t)}} \exp(I_{D+V_t}).$$

Thus,

$$\mathbb{E} \left[ \sqrt{\frac{\det D}{\det(D + V_t)}} \exp(I_{D+V_t}) \right] \leq 1.$$

Now we prove Eq. (18). First, we have

$$\begin{aligned} \Pr [I_{D+V_t} > \beta(\delta_t)] &= \Pr \left[ \sqrt{\frac{\det D}{\det(D + V_t)}} \exp(I_{D+V_t}) > \sqrt{\frac{\det D}{\det(D + V_t)}} \exp(\beta(\delta_t)) \right] \\ &= \Pr \left[ M_t > \frac{1}{\sqrt{\det(I + D^{-\frac{1}{2}} V_t D^{-\frac{1}{2}})}} \exp(\beta(\delta_t)) \right] \\ &\leq \frac{\mathbb{E}[M_t] \sqrt{\det(I + D^{-\frac{1}{2}} V_t D^{-\frac{1}{2}})}}{\exp(\beta(\delta_t))} \\ &\leq \frac{\sqrt{\det(I + D^{-\frac{1}{2}} V_t D^{-\frac{1}{2}})}}{\exp(\beta(\delta_t))} \end{aligned} \quad (19)$$

Then, for some constant  $\gamma > 0$  and for any  $a \in \mathbb{N}$ , we define the set of timesteps  $\mathcal{K}_a \subseteq [T]$  such that

$$t \in \mathcal{K}_a \Leftrightarrow (1 + \gamma)^a \leq N_\pi(t - 1) < (1 + \gamma)^{a+1}.$$

Define  $D_a$  as a diagonal matrix such that  $D_{a,ii} = (1 + \gamma)^a$ ,  $\forall i \in \tilde{d}$ . Suppose  $t \in \mathcal{K}_a$  for some fixed  $a$ . Then, we have

$$\frac{1}{1 + \gamma} D_t \preceq D_a \preceq D_t.$$

Let  $D = \lambda \Lambda_{\Sigma_\pi} D_a \succeq \frac{\lambda}{1 + \gamma} \Lambda_{\Sigma_\pi} D_t$ . Then, we have

$$D^{-\frac{1}{2}} V_t D^{-\frac{1}{2}} \preceq \frac{1 + \gamma}{\lambda} D_t^{-\frac{1}{2}} \Lambda_{\Sigma_\pi}^{-\frac{1}{2}} V_t \Lambda_{\Sigma_\pi}^{-\frac{1}{2}} D_t^{-\frac{1}{2}},$$

where matrix  $D_t^{-\frac{1}{2}} \Lambda_{\Sigma_\pi}^{-\frac{1}{2}} V_t \Lambda_{\Sigma_\pi}^{-\frac{1}{2}} D_t^{-\frac{1}{2}}$  has  $\tilde{d}$  ones on the diagonal. Since the determinant of a positive definite matrix is smaller than the product of its diagonal terms, we have

$$\begin{aligned} \det(I + D^{-\frac{1}{2}} V_t D^{-\frac{1}{2}}) &\leq \det\left(I + \frac{1 + \gamma}{\lambda} D_t^{-\frac{1}{2}} \Lambda_{\Sigma_\pi}^{-\frac{1}{2}} V_t \Lambda_{\Sigma_\pi}^{-\frac{1}{2}} D_t^{-\frac{1}{2}}\right) \\ &\leq \left(1 + \frac{1 + \gamma}{\lambda}\right)^{\tilde{d}} \end{aligned} \quad (20)$$

Let  $\gamma = e - 1$ . Using Eqs. (19) and (20),  $\beta(\delta_t) = \ln(1/\delta_t) + \ln \ln t + \frac{\tilde{d}}{2} \ln(1 + e/\lambda) = \ln(t \ln^2 t) + \ln \ln t + \frac{\tilde{d}}{2} \ln(1 + \frac{e}{\lambda})$  and a union bound over  $a$ , we have

$$\begin{aligned} \Pr [I_{D+V_t} > \beta(\delta_t)] &\leq \sum_a \Pr [I_{D+V_t} > \beta(\delta_t) | t \in \mathcal{K}_a, D = \lambda \Lambda_{\Sigma_\pi} D_a] \\ &\leq \sum_a \frac{\sqrt{\det(I + D^{-\frac{1}{2}} V_t D^{-\frac{1}{2}})}}{\exp(\beta(\delta_t))} \end{aligned}$$

$$\begin{aligned}
&\leq \frac{\ln t}{\ln(1+\gamma)} \cdot \frac{(1 + \frac{1+\gamma}{\lambda})^{\frac{d}{2}}}{\exp(\ln(t \ln^2 t) + \ln \ln t + \frac{d}{2} \ln(1 + \frac{e}{\lambda}))} \\
&= \ln t \cdot \frac{(1 + \frac{e}{\lambda})^{\frac{d}{2}}}{t \ln^2 t \cdot \ln t \cdot (1 + \frac{e}{\lambda})^{\frac{d}{2}}} \\
&= \frac{1}{t \ln^2 t} \\
&= \delta_t
\end{aligned}$$

Thus, Eq. (18) holds and we complete the proof of Lemma 2.  $\square$

Now, we give the proof of Theorem 5.

*Proof.* (Theorem 5) First, we bound the number of exploration rounds up to time  $T$ . Let  $\psi(t) = t^{\frac{2}{3}}/d$ . According to the condition of exploitation (Line 4 in Algorithm 3), we have that if at timestep  $t$  algorithm MC-E TE starts an exploration round, then  $t$  satisfies

$$N_\pi(t-1) \leq \psi(t).$$

Let  $t_0$  denote the timestep at which algorithm MC-E TE starts the last exploration round. Then, we have

$$N_\pi(t_0-1) \leq \psi(t_0)$$

and thus

$$\begin{aligned}
N_\pi(T) &= N_\pi(t_0) \\
&= N_\pi(t_0-1) + 1 \\
&\leq \psi(t_0) + 1 \\
&\leq \psi(T) + 1.
\end{aligned}$$

Next, for each timestep  $t$ , we bound the estimation error of  $f(\mathbf{w})$ . For any  $t > 0$ , let  $\Delta_t \triangleq f(\mathbf{w}^*) - f(\mathbf{w}_t)$  and  $\hat{f}_t(\mathbf{w}) \triangleq \mathbf{w}^\top \hat{\boldsymbol{\theta}}_t - \rho \mathbf{w}^\top \hat{\Sigma}_t \mathbf{w}$ . Suppose that event  $\mathcal{M}_t \cap \mathcal{N}_t$  occurs. Then, according to Lemmas 5 and 6, we have that for any  $\mathbf{w} \in \Delta_d$ ,

$$\begin{aligned}
|f(\mathbf{w}) - \hat{f}_{t-1}(\mathbf{w})| &\leq \sqrt{2\beta(\delta_t)} \sqrt{\mathbf{w}^\top B_\pi^+ D_{t-1}^{-1} \left( \lambda \Lambda_{\Sigma_\pi} D_{t-1} + \sum_{s=1}^{N_\pi(t-1)} \Sigma_\pi \right) D_{t-1}^{-1} (B_\pi^+)^{\top} \mathbf{w}} \\
&\quad + 5\rho \|C_\pi^+\| \sqrt{\frac{3 \ln t}{2N_\pi(t-1)}} \\
&\leq \sqrt{2 \ln(t \ln^2 t) + \ln \ln t + \frac{\tilde{d}}{2} \ln(1 + \frac{e}{\lambda})} \cdot \sqrt{\mathbf{w}^\top B_\pi^+ D_{t-1}^{-1} (\lambda \Lambda_{\Sigma_\pi} + \Sigma_\pi) (B_\pi^+)^{\top} \mathbf{w}} \\
&\quad + 5\rho \|C_\pi^+\| \sqrt{\frac{3 \ln t}{2N_\pi(t-1)}} \\
&\leq 7 \sqrt{\ln t + \frac{\tilde{d}}{2} \ln(1 + \frac{e}{\lambda})} \cdot \left( \sqrt{\mathbf{w}^\top B_\pi^+ (\lambda \Lambda_{\Sigma_\pi} + \Sigma_\pi) (B_\pi^+)^{\top} \mathbf{w}} + \rho \|C_\pi^+\| \right) \cdot \sqrt{\frac{1}{N_\pi(t-1)}}
\end{aligned}$$

Let  $Z(\rho, \pi) \triangleq \max_{\mathbf{w} \in \Delta_d} \left( \sqrt{\mathbf{w}^\top B_\pi^+ (\lambda \Lambda_{\Sigma_\pi} + \Sigma_\pi) (B_\pi^+)^\top \mathbf{w}} + \rho \|C_\pi^+\| \right)$ . Then, we have

$$\left| f(\mathbf{w}) - \hat{f}_{t-1}(\mathbf{w}) \right| \leq 7 \cdot Z(\rho, \pi) \sqrt{\frac{\ln t + \tilde{d} \ln(1 + e/\lambda)}{N_\pi(t-1)}}.$$

Let  $\mathcal{L}^{\text{exploit}}(t)$  denote the event that algorithm MC-E TE does the exploitation at timestep  $t$ . For any  $t > 0$ , if  $\mathcal{L}^{\text{exploit}}(t)$  occurs, we have  $N_\pi(t-1) > \psi(t)$ . Thus

$$\begin{aligned} \left| f(\mathbf{w}) - \hat{f}_{t-1}(\mathbf{w}) \right| &< 7 \cdot Z(\rho, \pi) \sqrt{\frac{\ln t + \tilde{d} \ln(1 + e/\lambda)}{\psi(t)}} \\ &= 7 \cdot Z(\rho, \pi) \sqrt{d \left( \ln t + \tilde{d} \ln(1 + e/\lambda) \right)} \cdot t^{-\frac{1}{3}} \end{aligned}$$

According to Lemmas 5,6, for any  $t \geq 2$ , we bound the probability of event  $\neg(\mathcal{M}_t \cap \mathcal{N}_t)$  by

$$\begin{aligned} \Pr[\neg(\mathcal{M}_t \cap \mathcal{N}_t)] &\leq \frac{6d^2}{t^2} + \frac{1}{t \ln^2 t} \\ &\leq \frac{6d^2}{t \ln^2 t} + \frac{1}{t \ln^2 t} \\ &= \frac{7d^2}{t \ln^2 t} \end{aligned}$$

The expected regret of algorithm MC-E TE can be divided into two parts, one due to exploration and the other due to exploitation. Then, we can obtain

$$\begin{aligned} \mathbb{E}[\mathcal{R}(T)] &\leq N_\pi(T) \cdot \tilde{d} \Delta_{\max} + \sum_{t=1}^T \mathbb{E}[\Delta_t | \mathcal{L}^{\text{exploit}}(t)] \\ &\leq (\psi(T) + 1) \cdot \tilde{d} \Delta_{\max} + \sum_{t=1}^T (\mathbb{E}[\Delta_t | \mathcal{L}^{\text{exploit}}(t), \mathcal{M}_t \cap \mathcal{N}_t] + \mathbb{E}[\Delta_t | \mathcal{L}^{\text{exploit}}(t), \neg(\mathcal{M}_t \cap \mathcal{N}_t)] \cdot \Pr[\neg(\mathcal{M}_t \cap \mathcal{N}_t)]) \\ &= O\left( \left( \frac{T^{\frac{2}{3}}}{d} + 1 \right) \cdot \tilde{d} \Delta_{\max} + \sum_{t=1}^T \left( Z(\rho, \pi) \sqrt{d(\ln t + \tilde{d} \ln(1 + e/\lambda))} \cdot t^{-\frac{1}{3}} + \Delta_{\max} \cdot \frac{d^2}{t \ln^2 t} \right) \right) \\ &= O\left( T^{\frac{2}{3}} d \Delta_{\max} + d^2 \Delta_{\max} + Z(\rho, \pi) \sqrt{d(\ln T + d^2 \ln(1 + \lambda^{-1}))} \cdot T^{\frac{2}{3}} + d^2 \Delta_{\max} \right) \\ &= O\left( Z(\rho, \pi) \sqrt{d(\ln T + d^2 \ln(1 + \lambda^{-1}))} \cdot T^{\frac{2}{3}} + d \Delta_{\max} \cdot T^{\frac{2}{3}} + d^2 \Delta_{\max} \right) \\ &= O\left( Z(\rho, \pi) \sqrt{d(\ln T + d^2 \ln(1 + \lambda^{-1}))} \cdot T^{\frac{2}{3}} + d \Delta_{\max} \cdot T^{\frac{2}{3}} \right) \end{aligned}$$

Choosing  $\lambda = \frac{1}{2}$  and using  $\Lambda_{\Sigma_\pi} = \Sigma_\pi$ , we obtain

$$\mathbb{E}[\mathcal{R}(T)] = O\left( Z(\rho, \pi) \sqrt{d(\ln T + d^2)} \cdot T^{\frac{2}{3}} + d \Delta_{\max} \cdot T^{\frac{2}{3}} \right),$$

where  $Z(\rho, \pi) = \max_{\mathbf{w} \in \Delta_d} \left( \sqrt{\mathbf{w}^\top B_\pi^+ \Sigma_\pi (B_\pi^+)^\top \mathbf{w}} + \rho \|C_\pi^+\| \right)$ . □