

---

# Fixed-Budget Real-Valued Combinatorial Pure Exploration of Multi-Armed Bandit

---

Shintaro Nakamura  
University of Tokyo  
RIKEN AIP

Masashi Sugiyama  
RIKEN AIP  
University of Tokyo

## Abstract

We study the real-valued combinatorial pure exploration of the multi-armed bandit in the fixed-budget setting. We first introduce the Combinatorial Successive Assign (CSA) algorithm, which is the first algorithm that can identify the best action even when the size of the action class is exponentially large with respect to the number of arms. We show that the upper bound of the probability of error of the CSA algorithm matches a lower bound up to a logarithmic factor in the exponent. Then, we introduce another algorithm named the Minimax Combinatorial Successive Accepts and Rejects (Minimax-CombsAR) algorithm for the case where the size of the action class is polynomial, and show that it is optimal, which matches a lower bound. Finally, we experimentally compare the algorithms with previous methods and show that our algorithm performs better.

## 1 Introduction

The multi-armed bandit (MAB) model is an important framework in online learning since it is useful to investigate the trade-off between exploration and exploitation in decision-making problems (Auer et al., 2002; Audibert et al., 2009). Although investigating this trade-off is intrinsic in many applications, some application domains only focus on obtaining the optimal object, e.g., an arm or a set of arms, among a set of candidates, and do not care about the loss or rewards that occur during the exploration procedure. This learning problem called the pure exploration (PE)

task has received much attention (Bubeck et al., 2009; Audibert et al., 2010).

One of the important sub-fields among PE of MAB is the *combinatorial pure exploration* of the MAB (CPE-MAB)(Chen et al., 2014; Gabillon et al., 2016; Chen et al., 2017). In the CPE-MAB, we have a set of  $d$  stochastic arms, where the reward of each arm  $s \in \{1, \dots, d\}$  follows an unknown distribution with mean  $\mu_s$ , and an *action class*  $\mathcal{A}$ , which is a collection of subsets of arms with certain combinatorial structures. Then, the goal is to identify the best action from the action class  $\mathcal{A}$  by pulling a single arm each round. There are mainly two settings in the CPE-MAB. One is the *fixed confidence* setting, where the player tries to identify the optimal action with high probability with as few rounds as possible, and the other is the *fixed-budget* setting, where the player tries to identify the optimal action with a fixed number of rounds (Chen et al., 2014; Katz-Samuels et al., 2020; Wang and Zhu, 2022a; Gabillon et al., 2016). Abstractly, the goal is to identify  $\pi^*$ , which is the optimal solution for the following constraint optimization problem:

$$\begin{aligned} & \text{maximize}_{\pi} && \mu^\top \pi \\ & \text{subject to} && \pi \in \mathcal{A}, \end{aligned} \tag{1}$$

where  $\mu$  is a vector whose  $s$ -th element is the mean reward of arm  $s$  and  $\top$  denotes the transpose.

Although CPE-MAB can be applied to many models which can be formulated as (1), most of the existing works in CPE-MAB (Chen et al., 2014; Wang and Zhu, 2022b; Gabillon et al., 2016; Chen et al., 2017; Du et al., 2021; Chen et al., 2016) assume  $\mathcal{A} \subseteq \{0, 1\}^d$ . This means that although we can apply the CPE-MAB framework to the shortest path problem (Sniedovich, 2006), top- $K$  arms identification (Kalyanakrishnan and Stone, 2010), matching (Gibbons, 1985), and spanning trees (Pettie and Ramachandran, 2002), we cannot apply it to problems where  $\mathcal{A} \subset \mathbb{R}^d$ , such as the optimal transport problem (Villani, 2008), the knapsack problem (Dantzig and Mazur, 2007), and the production planning problem




Figure 1: A simple sketch of the knapsack problem. We want to know how many of each item should be included in the bag to maximize the total value. Here, the total weight of every item cannot exceed 10kg, which is the capacity of the bag.

(Pochet and Wolsey, 2010). For instance, in the knapsack problem shown in Figure 1, actions are not binary vectors since, for each item, we can put more than one in the bag, e.g., one blue one and two orange ones.

To overcome this limitation, Nakamura and Sugiyama (2023) introduced the real-valued CPE-MAB (R-CPE-MAB), where the action class is a set of real vectors, i.e.,  $\mathcal{A} \subset \mathbb{R}^d$ . Though they have investigated the R-CPE-MAB in the fixed confidence setting, there is still room for investigation in the fixed-budget setting. To the best of our knowledge, the only existing work that can be applied to the fixed-budget R-CPE-MAB is the Peace algorithm introduced by Katz-Samuels et al. (2020). However, the Peace algorithm required enumerating all of the feasible actions. This is nearly impossible since, in general, there is no efficient algorithm for enumerating all actions for a combinatorial optimization problem. Even if they could be enumerated, it would require an operation to sort an exponentially large number of actions by estimated reward, which is practically unrealistic in terms of computation time. Secondly, there is a hyper-parameter for it, and one needs to carefully choose them for the feasibility of the Peace algorithm (Yang and Tan, 2022). Finally, it needs an assumption that the rewards follow a normal distribution, which may not be satisfied in real-world applications.

In this work, we first introduce a parameter-free algorithm named the Combinatorial Successive Assign (CSA) algorithm, which is a generalized version of the Combinatorial Successive Accepts and Rejects (CSAR) algorithm proposed by Chen et al. (2014) for the ordinary CPE-MAB. Since the CSA algorithm does not require enumerating all the actions in  $\mathcal{A}$ , it can be applied to cases even when the size of the ac-

tion class  $\mathcal{A}$  is exponentially large in  $d$ . We show that the upper bound of the probability of error of the CSA algorithm matches a lower bound up to a logarithmic factor in the exponent.

Since the CSA algorithm does not match a lower bound, we introduce another algorithm named the Minimax Combinatorial Successive Accepts and Rejects (Minimax-CombSAR) algorithm inspired by Yang and Tan (2022) for the case where the size of the action class is polynomial with respect to the number of arms. In Section 4, we show that the Minimax-CombSAR algorithm is optimal, which means that the upper bound of the probability of error of the best action matches a lower bound. We also show that the Minimax-CombSAR algorithm has only one hyper-parameter, and is easily interpreted.

Finally, we report the results of numerical experiments. First, we show that the CSA algorithm can identify the best action in a knapsack problem, where the size of the action class can be exponentially large in  $d$ . Then, when the size of the action class is polynomial in  $d$ , we show that the Minimax-CombSAR algorithm performs better than the CSA algorithm and the Peace algorithm in identifying the best action in the knapsack and optimal transport problems.

## 2 Problem Formulation

In this section, we formally define our R-CPE-MAB model. Suppose we have  $d$  arms, numbered  $1, \dots, d$ . Assume that each arm  $s \in [d]$  is associated with a reward distribution  $\phi_s$ , where  $[d] = \{1, \dots, d\}$ . We assume all reward distributions have  $R$ -sub-Gaussian tails for some known constant  $R > 0$ . Formally, if  $X$  is a random variable drawn from  $\phi_s$  for some  $s \in [d]$ , then, for all  $\lambda \in \mathbb{R}$ , we have  $\mathbb{E}[\exp(\lambda X - \lambda \mathbb{E}[X])] \leq \exp(R^2 \lambda^2 / 2)$ . It is known that the family of  $R$ -sub-Gaussian tail distributions includes all distributions that are supported on  $[a, b]$ , where  $(b - a)/2 = R$ , and also many unbounded distributions such as Gaussian distributions with variance  $R^2$  (Rivasplata, 2012; Rinaldo and Bong, 2018). Let  $\boldsymbol{\mu} = (\mu_1, \dots, \mu_d)^\top$  denote the vector of expected rewards, where each element  $\mu_s = \mathbb{E}_{X \sim \phi_s}[X]$  denotes the expected reward of arm  $s$  and  $\top$  denotes the transpose. With a given  $\boldsymbol{\nu}$ , let us consider the following linear optimization problem:

$$\begin{aligned} & \text{maximize}_{\boldsymbol{\pi}} && \boldsymbol{\nu}^\top \boldsymbol{\pi} \\ & \text{subject to} && \boldsymbol{\pi} \in \mathcal{C} \subset \mathbb{R}_{\geq 0}^d. \end{aligned} \quad (2)$$

Here,  $\mathcal{C}$  is a problem-dependent feasible region of  $\boldsymbol{\pi}$ , which satisfies some combinatorial structures. Then, for any  $\boldsymbol{\nu} \in \mathbb{R}^d$ , we denote by  $\boldsymbol{\pi}^{\boldsymbol{\nu}, \mathcal{C}}$  the solution of (2). We define the action class  $\mathcal{A}$  as the set of vectors that

contains optimal solutions of (2) for any  $\nu$ , i.e.,

$$\mathcal{A} = \{\pi^{\nu, \mathcal{C}} \in \mathbb{R}_{\geq 0}^d \mid \forall \nu \in \mathbb{R}^d\}. \quad (3)$$

We assume that the size of  $\mathcal{A}$  is finite and denote it by  $K$ . This assumption is relatively mild since, for instance, in linear programming, the optimal solution can always be found at one of the vertices of the feasible region (Ahmadi, 2016).

Also, let  $\text{POSSIBLE-PI}(s) = \{x \in \mathbb{R} \mid \exists \pi \in \mathcal{A}, \pi_s = x\}$ , where  $\pi_s$  denotes the  $s$ -th element of  $\pi$ . We can see  $\text{POSSIBLE-PI}(s)$  as the set of all possible values that an action in the action set can take as the  $s$ -th element. Let us denote the size of  $\text{POSSIBLE-PI}(s)$  by  $B_s$ . Note that  $K$ , the size of  $\mathcal{A}$ , can be exponentially large in  $d$ , i.e.,  $|\mathcal{A}| = \mathcal{O}(\prod_{s=1}^d B_s)$ .

We assume we have *offline oracles* which efficiently solve the linear optimization problem (2) once  $\nu$  is given. For instance, offline oracles can be algorithms which output  $\pi^*(\nu) = \arg \max_{\pi \in \mathcal{A}} \nu^\top \pi$  in polynomial or pseudo-polynomial<sup>1</sup> time in  $d$ .

The player's objective is to identify the best action  $\pi^* = \arg \max_{\pi \in \mathcal{A}} \mu^\top \pi$  by pulling a single arm each round.

The player is given a *budget*  $T$ , and cannot pull arms more than  $T$  times. The player outputs an action  $\pi^{\text{out}}$  at the end, and she is evaluated by the *probability of error*, which is formally  $\Pr[\pi^{\text{out}} \neq \pi^*]$ .

### 3 Lower Bound of the Fixed-Budget R-CPE-MAB

In this section, we show a lower bound of the probability of error in R-CPE-MAB. As preliminaries, let us introduce some notions. First, we introduce  $\pi^{(s)}$  as follows (Nakamura and Sugiyama, 2023):

$$\pi^{(s)} = \arg \min_{\pi \in \mathcal{A} \setminus \{\pi^*\}} \frac{\mu^\top (\pi^* - \pi)}{|\pi_s^* - \pi_s|}. \quad (4)$$

Intuitively, among the actions whose  $s$ -th element is different from  $\pi^*$ ,  $\pi^{(s)}$  is the action which is the most difficult to determine whether it is the best action or not. We also introduce the notion *G-gap* (Nakamura and Sugiyama, 2023) as follows:

$$\begin{aligned} \Delta_s &= \min_{\pi \in \mathcal{A} \setminus \{\pi^*\}} \frac{\mu^\top (\pi^* - \pi)}{|\pi_s^* - \pi_s|} \\ &= \frac{\mu^\top (\pi^* - \pi^{(s)})}{|\pi_s^* - \pi_s^{(s)}|}. \end{aligned} \quad (5)$$

<sup>1</sup>A pseudo-polynomial time algorithm is a numeric algorithm whose running time is polynomial in the numeric value of the input, but not necessarily in the length of the input (Garey and Johnson, 1990)

*G-Gap* was first introduced in Nakamura and Sugiyama (2023) as a natural generalization of the notion *gap* in the CPE-MAB literature (Chen et al., 2014, 2016, 2017). This was introduced as a key notion that characterizes the difficulty of the problem instance.

In Theorem 1, we show that the sum of the inverse of squared *G-Gaps*,

$$\mathbf{H} = \sum_{s=1}^d \left( \frac{1}{\Delta_s} \right)^2, \quad (6)$$

appears in the lower bound of the probability of error of R-CPE-MAB, which implies that it characterizes the difficulty of the problem instance.

**Theorem 1.** *For any action class and any algorithm that returns an action  $\pi^{\text{out}}$  after  $T$  times of arm pulls, the probability of error is at least*

$$\mathcal{O} \left( \exp \left( -\frac{T}{\mathbf{H}} \right) \right). \quad (7)$$

We show the proof in Appendix A. If  $\mathcal{A}$  is a set of  $d$  dimensional standard basis, R-CPE-MAB becomes the standard best arm identification problem whose objective is to identify the best arm with the largest expected reward among  $d$  arms (Bubeck et al., 2009; Audibert et al., 2010; Carpentier and Locatelli, 2016). From Carpentier and Locatelli (2016), a lower bound is  $\mathcal{O} \left( \exp \left( -\frac{T}{\log(d)\mathbf{H}} \right) \right)$  for the standard best arm identification problem. It is a future work that whether the lower bound is  $\mathcal{O} \left( \exp \left( -\frac{T}{\log(d)\mathbf{H}} \right) \right)$  for general action classes.

## 4 The Combinatorial Successive Assign (CSA) Algorithm

In this section, we first introduce the CSA algorithm, which can be seen as a generalization of the CSAR algorithm (Chen et al., 2014). This algorithm can be applied to fixed-budget R-CPE-MAB even when the size of the action class  $\mathcal{A}$  is exponentially large in  $d$ . Then, we show an upper bound of the probability of error of the best action of the CSA algorithm. We also discuss the number of times offline oracles have to be called.

### 4.1 CSA Algorithm

In this subsection, we introduce the CSA algorithm, a fully parameter-free algorithm for fixed-budget R-CPE-MAB that works even when the action set  $\mathcal{A}$  can be exponentially large in  $d$ .

We first define the *constrained offline oracle* which is used in the CSA algorithm.

**Definition 1** (Constrained offline oracle). Let  $\mathbf{S} = \{(e, x) | (e, x) \in \mathbb{Z} \times \mathbb{R}\}$  be a set of tuples. A constrained offline oracle is denoted by  $\text{COOracle}: \mathbb{R}^d \times \mathbf{S} \rightarrow \mathcal{A} \cup \perp$  and satisfies

$$\text{COOracle}(\boldsymbol{\mu}, \mathbf{S}) = \begin{cases} \arg \max_{\boldsymbol{\pi} \in \mathcal{A}_{\mathbf{S}}} \boldsymbol{\mu}^\top \boldsymbol{\pi} & (\text{if } \mathcal{A}_{\mathbf{S}} \neq \emptyset), \\ \perp & (\text{if } \mathcal{A}_{\mathbf{S}} = \emptyset), \end{cases}$$

where we define  $\mathcal{A}_{\mathbf{S}} = \{\boldsymbol{\pi} \in \mathcal{A} \mid \forall (e, x) \in \mathbf{S}, \pi_e = x\}$  as the collection of feasible actions and  $\perp$  is a null symbol.

Here, we can see that a  $\text{COOracle}$  is a modification of an offline oracle specified by  $\mathbf{S}$ . In other words, for all  $(e, x)$  in  $\mathbf{S}$ , the  $\text{COOracle}$  outputs an action whose  $e$ -th element is  $x$ ; otherwise, it outputs the null symbol. In Appendix B, we discuss how to construct such  $\text{COOracles}$  for some combinatorial problems, such as the optimal transport problem and the knapsack problem.

We introduce the CSA algorithm in Algorithm 1. The CSA algorithm divides the budget into  $d$  rounds. In each round, we pull each of the remaining arms the same number of times (line 5). At each round  $t$ , the CSA outputs the empirically best action  $\hat{\boldsymbol{\pi}}(t)$  (line 6), chooses a single arm  $p(t)$  (line 14), and assigns  $\hat{\pi}_{p(t)}(t)$  for the  $e$ -th element of  $\boldsymbol{\pi}^{\text{out}}$ . Indices that are assigned are maintained in  $\mathbf{F}(t)$ , and arms  $s \in \mathbf{F}(t)$  will no longer be pulled in the next rounds. The pair of the index and the assigned value for  $\boldsymbol{\pi}^{\text{out}}$ ,  $\mathbf{S}(t)$ , is updated at every round (line 16).

## 4.2 Theoretical Analysis of CSA Algorithm

Here, we first discuss an upper bound of the probability of error. Then, we discuss the number of times we call the offline oracle.

### 4.2.1 An Upper Bound of the Probability of Error

Here, we show an upper bound of the probability of error of the CSA algorithm. Let  $\Delta_{(1)}, \dots, \Delta_{(d)}$  be a permutation of  $\Delta_1, \dots, \Delta_d$  such that  $\Delta_{(1)} \leq \dots \leq \Delta_{(d)}$ . Also, let us define

$$\mathbf{H}_2 = \max_{s \in [d]} \frac{s}{\Delta_{(s)}^2}.$$

One can verify that  $\mathbf{H}_2$  is equivalent to  $\mathbf{H}$  up to a logarithmic factor:  $\mathbf{H}_2 \leq \mathbf{H} \leq \log(2d)\mathbf{H}_2$  (Audibert et al., 2010).

**Theorem 2.** Given any  $T > d$ , action class  $\mathcal{A} \subset \mathbb{R}^d$ , and  $\boldsymbol{\mu} \in \mathbb{R}^d$ , the CSA algorithm uses at most  $T$  samples and outputs a solution  $\boldsymbol{\pi}^{\text{out}} \in \mathcal{A} \cup \{\perp\}$  such

---

**Algorithm 1** CSA: Combinatorial Successive Assign Algorithm

---

**Input:** Budget:  $T \geq 0$ ;  $\text{COOracle}: \rightarrow \mathcal{A} \cup \{\perp\}$

- 1: Define  $\tilde{\log}(n) = \sum_{i=1}^d \frac{1}{i}$
- 2:  $\tilde{T}_0 \leftarrow 0$ ,  $\mathbf{F}(t) \leftarrow \emptyset$ ,  $\mathbf{S}(t) = \emptyset$
- 3: **for**  $t = 1$  to  $d$  **do**
- 4:  $\tilde{T}(t) \leftarrow \lceil \frac{T-d}{\tilde{\log}(d)(d-t+1)} \rceil$
- 5: Pull each arm  $e \in [d] \setminus \mathbf{F}(t)$  for  $\tilde{T}(t) - \tilde{T}_{t-1}$  times
- 6:  $\hat{\boldsymbol{\pi}}(t) \leftarrow \text{COOracle}(\hat{\boldsymbol{\mu}}(t), \mathbf{S}(t))$
- 7: **if**  $\hat{\boldsymbol{\pi}}(t) = \perp$  **then**
- 8:     **Fail:** set  $\boldsymbol{\pi}^{\text{out}} \leftarrow \perp$  and return  $\boldsymbol{\pi}^{\text{out}}$
- 9: **end if**
- 10: **for all**  $e$  in  $[d] \setminus \mathbf{F}(t)$  **do**
- 11:      $\text{CHECK}(e) \leftarrow \text{POSSIBLE-PI}(e) \setminus \{\hat{\pi}_e(t)\}$
- 12:      $\tilde{\pi}^e(t) \leftarrow \arg \max_{x \in \text{CHECK}(e)} \text{COOracle}(\hat{\boldsymbol{\mu}}(t), \mathbf{S}(t) \cup (e, x))$
- 13: **end for**
- 14:  $p(t) \leftarrow \arg \max_{e \in [d] \setminus \mathbf{F}(t)} \frac{\langle \hat{\boldsymbol{\mu}}(t), \hat{\boldsymbol{\pi}}(t) - \tilde{\boldsymbol{\pi}}^e(t) \rangle}{\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)}$
- 15:  $\mathbf{F}(t+1) \leftarrow \mathbf{F}(t) \cup \{p(t)\}$
- 16:  $\mathbf{S}(t+1) \leftarrow \mathbf{S}(t) \cup \{(p(t), \hat{\pi}_{p(t)}(t))\}$
- 17: **end for**
- 18: // Convert  $\mathbf{S}(d+1)$  to  $\boldsymbol{\pi}^{\text{out}}$
- 19: **for**  $(e, x)$  in  $\mathbf{S}(d+1)$  **do**
- 20:     // The  $e$ -th element of  $\boldsymbol{\pi}^{\text{out}}$  is  $x$
- 21:      $\pi_e^{\text{out}} = x$
- 22: **end for**
- 23: **return**  $\boldsymbol{\pi}^{\text{out}}$

---

that

$$\begin{aligned} & \Pr[\boldsymbol{\pi}^{\text{out}} \neq \boldsymbol{\pi}^*] \\ & \leq d^2 \exp\left(-\frac{T-d}{2(2+L^2)^2 R^2 \tilde{\log}(d) U_{\mathcal{A}}^2 \mathbf{H}_2}\right), \end{aligned} \quad (8)$$

where  $L = \max_{e \in [d], \pi^1, \pi^2, \pi^3 \in \mathcal{A}, \pi_e^1 \neq \pi_e^3} \frac{|\pi_e^1 - \pi_e^2|}{|\pi_e^1 - \pi_e^3|}$ ,  $\tilde{\log}(d) \triangleq \sum_{s=1}^d \frac{1}{s}$ , and  $U_{\mathcal{A}} = \max_{\boldsymbol{\pi}, \boldsymbol{\pi}' \in \mathcal{A}, e \in \{s \in [d] \mid \pi_s \neq \pi'_s\}} \frac{\sum_{s=1}^d |\pi_s - \pi'_s|}{|\pi_e - \pi'_e|}$ .

We can see that the CSA algorithm is optimal up to a logarithmic factor in the exponent. Since  $L = 1$  and  $U_{\mathcal{A}} = \text{width}(\mathcal{A})$  in the ordinary CPE-MAB, we can confirm that Theorem 2 can be seen as a natural generalization of Theorem 3 in Chen et al. (2014), which shows an upper bound of the probability of error of the CSAR algorithm.

### 4.2.2 The Oracle Complexity

Next, we discuss the *oracle complexity*, the number of times we call the offline oracle (Ito et al., 2019; Xu and Li, 2021). Note that, in the CSA algorithm, we

call the COracle  $\mathcal{O}(d \sum_{s=1}^d B_s)$  times (line 12). Therefore, if each COracle call invokes the offline oracle  $N$  times, the *oracle complexity* is  $\mathcal{O}(Nd \sum_{s=1}^d B_s)$ . Finally, if the time complexity of each oracle call is  $Z$ , then the total time complexity of the CSA algorithm is  $\mathcal{O}(ZNd \sum_{s=1}^d B_s)$ .

Below, we discuss the oracle complexity of the CSA algorithm in some specific combinatorial problems such as the knapsack problem and the optimal transport problem. Note that the Peace algorithm (Katz-Samuels et al., 2020) has to enumerate all the actions in advance, where the number may be exponentially large in  $d$ . We show that the CSA algorithm mitigates the curse of dimensionality.

**The Knapsack Problem (Dantzig and Mazur, 2007).** In the knapsack problem, we have  $d$  items. Each item  $s \in [d]$  has a weight  $w_s$  and value  $v_s$ . Also, there is a knapsack whose capacity is  $W$  in which we put items. Our goal is to maximize the total value of the knapsack, not letting the total weight of the items exceed the capacity of the knapsack. Formally, the optimization problem is given as follows:

$$\begin{aligned} & \text{maximize}_{\pi \in \mathcal{A}} \quad \sum_{s=1}^d v_s \pi_s \\ & \text{subject to} \quad \sum_{s=1}^d \pi_s w_s \leq W, \end{aligned}$$

where  $\pi_s \in \mathbb{Z}_{\geq 0}$  denotes the number of item  $s$  in the knapsack. Here, if we assume the weight of each item is known, but the value is unknown, we can apply the R-CPE-MAB framework to the knapsack problem, where we estimate the values of items. The knapsack problem is NP-complete (Garey and Johnson, 1979). Hence, it is unlikely that the knapsack problem can be solved in polynomial time. However, it is well known that the knapsack problem can be solved in pseudo-polynomial time  $\mathcal{O}(dW)$  if we use dynamic programming (Kellerer et al., 2004; Fujimoto, 2016). It finds the optimal solution by constructing a table of size  $dW$  whose  $(s, w)$ -th element represents the maximum total value that can be achieved if the sum of the weights does not exceed  $w$  using up to the  $s$ th item. In some cases, it is sufficient to assume  $\mathcal{O}(dW)$  time-complexity is enough, and therefore, we use this dynamic programming method as the offline oracle. We can construct the COracle for the knapsack problem by calling this offline oracle once (see Appendix B.1 for details). Therefore, the CSA algorithm calls the offline oracle  $\mathcal{O}(1 \times d \sum_{s=1}^d B_s) = \mathcal{O}(d \sum_{s=1}^d B_s)$  times, and the total time complexity of the CSA algorithm is  $\mathcal{O}(dW \times d \sum_{s=1}^d B_s) = \mathcal{O}(d^2 W \sum_{s=1}^d B_s)$ . This is much more computationally friendly than the Peace algorithm with time complexity  $\mathcal{O}(\prod_{s=1}^d B_s)$ .

**The Optimal Transport (OT) Problem (Peyré**

**and Cuturi, 2019).** OT can be regarded as the cheapest plan to deliver resources from  $m$  suppliers to  $n$  consumers, where each supplier  $i$  and consumer  $j$  have supply  $s_i$  and demand  $d_j$ , respectively. Let  $\gamma \in \mathbb{R}_{\geq 0}^{m \times n}$  be the cost matrix, where  $\gamma_{ij}$  denotes the cost between supplier  $i$  and demander  $j$ . Our objective is to find the optimal transportation matrix,

$$\pi^* = \arg \min_{\pi \in \mathcal{G}(s, d)} \sum_{i, j} \pi_{ij} \gamma_{ij}, \quad (9)$$

where

$$\mathcal{G}(s, d) \triangleq \left\{ \Pi \in \mathbb{R}_{\geq 0}^{m \times n} \mid \Pi \mathbf{1}_n = s, \Pi^\top \mathbf{1}_m = d \right\}. \quad (10)$$

Here,  $s = (s_1, \dots, s_m)$  and  $d = (d_1, \dots, d_n)$ .  $\pi_{ij}$  represents how much resources one sends from supplier  $i$  to demander  $j$ . If we assume that the cost is unknown and changes stochastically, e.g., due to some traffic congestions, we can apply the R-CPE-MAB framework to the optimal transport problem, where we estimate the cost of each edge  $(i, j)$  between supplier  $i$  and consumer  $j$ .

Once  $\gamma$  is given, we can compute  $\pi^*$  in  $\mathcal{O}(l^3 \log l)$ , where  $l = \max(m, n)$ , by using network simplex or interior point methods (Cuturi, 2013), and we can use them as the offline oracle. It is known that the solution of linear programming can always be found at one of the vertices of the feasible region (Ahmadi, 2016), and therefore the size of the action space  $\mathcal{A} = \left\{ \pi^\nu, \mathcal{G}(s, d) \in \mathbb{R}_{\geq 0}^{m \times n} \mid \forall \nu \in \mathbb{R}^{m \times n} \right\}$  is finite. However, it is difficult to construct  $\{\text{POSSIBLE-PI}((i, j))\}_{i \in [m], j \in [n]}$  in general to run the CSA algorithm. On the other hand, if  $s$  and  $d$  are both integer vectors, we can construct  $\{\text{POSSIBLE-PI}((i, j))\}_{i \in [m], j \in [n]}$  thanks to the fact that  $\mathcal{A} = \left\{ \pi^\nu, \mathcal{G}(s, d) \in \mathbb{Z}_{\geq 0}^{m \times n} \mid \forall \nu \in \mathbb{R}^{m \times n} \right\}$  is a set of non-negative integer matrices (Goodman and O'Rourke, 1997). In this case, all actions are restricted to non-negative integers, and  $\text{POSSIBLE-PI}(i, j) = \{0, 1, \dots, \min(s_i, d_j)\}$ . In Appendix B.2, we show that the COracle can be constructed by calling the offline oracle once. Therefore, we call the offline oracle  $\mathcal{O}(1 \times mn \sum_{s=1}^{mn} B_s) = \mathcal{O}(mn \sum_{s=1}^{mn} B_s)$  times, and the total time complexity is  $\mathcal{O}(mn \sum_{s=1}^{mn} B_s \cdot l^3 \log l)$ , where  $l = \max(m, n)$ . Again, this is much more computationally friendly than the Peace algorithm with time complexity  $\mathcal{O}(\prod_{s=1}^d B_d)$ .

**A General Case When  $K(= |\mathcal{A}|) = \text{poly}(d)$ .** In some cases, we can enumerate all the possible actions in  $\mathcal{A}$ . For instance, one may use some prior knowledge of each arm, which is sometimes obtainable in the real world, to narrow down the list of actions, and make the size of the action class  $\mathcal{A}$  polynomial

in  $d$ . In Appendix B.3, we show that the time complexity of the COracle is  $\mathcal{O}(dK + K \log K)$ , and therefore the total time complexity of the CSA algorithm is  $\mathcal{O}\left((dK + K \log K) \cdot d \sum_{s=1}^d B_s\right)$ .

## 5 The Minimax-CombSAR Algorithm for R-CPE-MAB where $|\mathcal{A}| = \text{poly}(d)$

In this section, we show an algorithm for fixed-budget R-CPE-MAB named the Minimax Combinatorial Successive Accept (Minimax-CombSAR) algorithm, for the case where we can assume that the size of  $\mathcal{A}$  is polynomial in  $d$ . Let  $\mathcal{A} = \{\boldsymbol{\pi}^1 = \boldsymbol{\pi}^*, \boldsymbol{\pi}^2, \dots, \boldsymbol{\pi}^K\}$ , where  $\boldsymbol{\mu}^\top \boldsymbol{\pi}^i \geq \boldsymbol{\mu}^\top \boldsymbol{\pi}^{i+1}$ , for all  $i \in [K-1]$ . The Minimax-CombSAR algorithm is inspired by the Optimal Design-based Linear Best Arm Identification (OD-LinBAI) algorithm Yang and Tan (2022), which eliminates actions in order from those considered suboptimal and finally outputs the remaining action as the optimal action. For the Minimax-CombSAR algorithm, we show an upper bound of the probability of error.

### 5.1 Minimax-CombSAR Algorithm

We show the Minimax-CombSAR in Algorithm 2. Here, we explain it at a higher level.

We have  $\lceil \log d \rceil$  phases in the Minimax-CombSAR algorithm, and it maintains an *active* action set  $\mathbb{A}(r)$  in each phase  $r$ . In each phase  $r \in [\lceil \log d \rceil]$ , it pulls arms  $m(r) = \frac{T' - d \lceil \log_2 d \rceil}{B/2^{r-1}}$  times in total, where  $B = 2^{\lceil \log_2 d \rceil} - 1$  and  $\beta \in [0, 1]$  is a hyperparameter, and  $T' = T - \lfloor \frac{T}{d} \beta \rfloor \times d$ . In each phase, we compute an *allocation vector*  $\mathbf{p}(r) \in \left\{ \mathbf{v} \in \mathbb{R}^d \mid \sum_{s=1}^d v_s = 1 \right\} \triangleq \Pi_d$ , and pull each arm  $s \lceil p_s(r) \cdot m(r) \rceil$  times. Then, at the end of each phase  $r$ , it eliminates a subset of possibly suboptimal actions. Eventually, there is only one action  $\boldsymbol{\pi}^{\text{out}}$  in the active action set, which is the output of the algorithm.

The key to identifying the best action with high confidence is the choice of the allocation vector  $\mathbf{p}(r)$ , which determines how many times we pull each arm in phase  $r$ .

#### Choice of $\mathbf{p}(r)$

We discuss how to choose an allocation vector that is beneficial to identify the best action. Let us denote the number of times arm  $s$  is pulled before phase  $r$  starts by  $T_s(r)$ . Also, we denote by  $\hat{\mu}_s(r)$  the empirical mean of the reward from arm  $s$  before phase  $r$  starts. Then, at the end of phase  $r$ , from Hoeffding's

---

#### Algorithm 2 Minimax Combinatorial Successive Accept Algorithm

---

**Input:** Budget:  $T \geq 0$ , initialization parameter:  $\beta$ ,  
 action set:  $\mathbb{A}(1) = \mathcal{A}$

- 1: // Initialization
- 2: **for**  $s \in [d]$  **do**
- 3:   Pull arm  $s \lfloor \frac{T}{d} \beta \rfloor$  times and update  $\hat{\mu}_s(1)$
- 4:    $T_s(r) \leftarrow \lfloor \frac{T}{d} \beta \rfloor$
- 5: **end for**
- 6:  $T' \leftarrow T - \lfloor \frac{T}{d} \beta \rfloor \times d$
- 7: **for**  $r = 1$  to  $\lceil \log_2 d \rceil$  **do**
- 8:    $m(r) = \frac{T' - d \lceil \log_2 d \rceil}{B/2^{r-1}}$
- 9:   Compute  $\mathbf{p}(r)$  according to (13) or (15)
- 10:   **for**  $s \in [d]$  **do**
- 11:     Pull arm  $s \lceil p_s(r) \cdot m(r) \rceil$  times
- 12:     Update  $\hat{\mu}_s(r+1)$  with the observed samples
- 13:   **end for**
- 14:   For each action  $\boldsymbol{\pi} \in \mathbb{A}(r)$ , estimate the expected reward:  $\langle \hat{\boldsymbol{\mu}}(r+1), \boldsymbol{\pi} \rangle$
- 15:   Let  $\mathbb{A}(r+1)$  be the set of  $\lceil \frac{d}{2^r} \rceil$  actions in  $\mathbb{A}(r)$  with the largest estimates of the expected rewards
- 16: **end for**

**Output:** The only action  $\boldsymbol{\pi}^{\text{out}}$  in  $\mathbb{A}(\lceil \log_2 d \rceil + 1)$

---

inequality (Hoeffding, 1963), we have

$$\begin{aligned} & \Pr \left[ \left| \langle \hat{\boldsymbol{\mu}}(r+1) - \boldsymbol{\mu} \rangle^\top (\boldsymbol{\pi}^a - \boldsymbol{\pi}^b) \right| \geq \epsilon \right] \\ & \leq \exp \left( - \frac{\epsilon^2}{\kappa^{a,b,\mathbf{p}(r)} R^2} \right) \end{aligned} \quad (11)$$

for any  $\boldsymbol{\pi}^a, \boldsymbol{\pi}^b \in \mathcal{A}$ , and  $\epsilon \in \mathbb{R}$ . Here,  $\hat{\boldsymbol{\mu}}(r+1) = (\hat{\mu}_1(r+1), \dots, \hat{\mu}_d(r+1))^\top$  and

$$\kappa^{a,b,\mathbf{p}(r)} = \sum_{s=1}^d \frac{(\pi_s^a - \pi_s^b)^2}{T_s(r) + \lceil p_s(r) \cdot m(r) \rceil}. \quad (12)$$

(11) shows that the empirical difference  $\hat{\boldsymbol{\mu}}^\top (\boldsymbol{\pi}^a - \boldsymbol{\pi}^b)$  between  $\boldsymbol{\pi}^a$  and  $\boldsymbol{\pi}^b$  is closer to the true difference  $\boldsymbol{\mu}^\top (\boldsymbol{\pi}^a - \boldsymbol{\pi}^b)$  with high probability if we make  $\kappa^{a,b,\mathbf{p}(r)}$  small. In that case, we have a higher chance to distinguish whether  $\boldsymbol{\pi}^a$  is better than  $\boldsymbol{\pi}^b$  or not. However, when we have more than two actions in  $\mathbb{A}(r)$ ,

$$\mathbf{p}^{a,b}(r) = \arg \min_{\mathbf{p} \in \Pi_d} \kappa^{a,b,\mathbf{p}(r)}$$

is not necessarily a good allocation vector for investigating the true difference between other pairs of actions in  $\mathbb{A}(r)$ . Therefore, we propose the following allocation vector as an alternative:

$$\mathbf{p}^{\min}(r) \triangleq \arg \min_{\mathbf{p} \in \Pi_d} \max_{\boldsymbol{\pi}^a, \boldsymbol{\pi}^b \in \mathcal{A}} \kappa^{a,b,\mathbf{p}(r)}. \quad (13)$$

(13) takes a minimax approach, which computes an allocation vector that minimizes the maximum value

of the right-hand side of (11) among all of the pairs of actions in  $\mathbb{A}(r)$ . Since (13) is a  $d$ -dimensional non-linear optimization problem, it becomes computationally costly as  $d$  grows. Thus, another possible choice of the allocation vector is as follows:

$$\mathbf{q}^{\min}(r) \triangleq \arg \min_{\mathbf{p} \in \Pi_d} \max_{\boldsymbol{\pi}^a, \boldsymbol{\pi}^b \in \mathcal{A}} \lambda^{a,b,\mathbf{p}}(r), \quad (14)$$

where  $\lambda^{a,b,\mathbf{p}} = \sum_{s=1}^d \frac{(\pi_s^a - \pi_s^b)^2}{p_s(r) \cdot m(r)}$ . For specific actions  $\boldsymbol{\pi}^k, \boldsymbol{\pi}^l \in \mathcal{A}$ , from the method of Lagrange multipliers (Hoffmann and Bradley, 2009), we have

$$\begin{aligned} \mathbf{q}^{k,l}(r) &\triangleq \arg \min_{\mathbf{p} \in \Pi_d} \lambda^{k,l,\mathbf{p}}(r) \\ &= \left( \frac{|\pi_1^k - \pi_1^l|}{\sum_{s=1}^d |\pi_s^k - \pi_s^l|}, \dots, \frac{|\pi_d^k - \pi_d^l|}{\sum_{s=1}^d |\pi_s^k - \pi_s^l|} \right)^\top, \end{aligned}$$

and therefore,  $\mathbf{q}^{\min}(r)$  in (14) can be written explicitly as follows:

$$\begin{aligned} \mathbf{q}^{\min}(r) &= \mathbf{q}^{i,j}(r) \\ &= \left( \frac{|\pi_1^i - \pi_1^j|}{\sum_{s=1}^d |\pi_s^i - \pi_s^j|}, \dots, \frac{|\pi_d^i - \pi_d^j|}{\sum_{s=1}^d |\pi_s^i - \pi_s^j|} \right)^\top, \end{aligned} \quad (15)$$

where  $\boldsymbol{\pi}^i, \boldsymbol{\pi}^j = \arg \max_{\boldsymbol{\pi}^k, \boldsymbol{\pi}^l \in \mathcal{A}} \lambda^{k,l,\mathbf{p}}(r)$ . If we use (15) for the allocation vector instead of computing (13), we do not have to solve a  $d$ -dimensional non-linear optimization problem, and thus is computationally more friendly.

In some cases, actions in  $\mathcal{A}$  can be sparse, and  $\mathbf{p}(1)$  can also be sparse. If  $\mathbf{p}(1)$  is sparse, we may only have a few samples for some arms, and therefore accidentally eliminate the best action in the first phase in line 15 of Algorithm 2. To cope with this problem, we have the *initialization* phase (lines 2–5) to pull each arm  $\lfloor \frac{T}{d} \beta \rfloor$  times, where  $\beta \in [0, 1]$  is a hyperparameter. Intuitively,  $\beta$  represents how much of the total budget will be spent on the initialization phase. If  $\beta$  is too small, we may accidentally eliminate the best action in the early phase, and if it is too large, we may not have enough budget to distinguish between the best action and the next best action.

## 5.2 Theoretical Analysis of Minimax-CombSAR

Here, in Theorem 3, we show an upper bound of the probability of error of the Minimax-CombSAR algorithm.

**Theorem 3.** *For any problem instance in fixed-budget R-CPE-MAB, the Minimax-CombSAR algorithm out-*

*puts an action  $\boldsymbol{\pi}^{\text{out}}$  satisfying*

$$\begin{aligned} &\Pr [\boldsymbol{\pi}^{\text{out}} \neq \boldsymbol{\pi}^*] \\ &\leq \left( \frac{4K}{d} + 3 \log_2 d \right) \exp \left( - \frac{T' - \lceil \log_2 d \rceil}{R^2 V^2} \cdot \frac{1}{\mathbf{H}_2} \right), \end{aligned} \quad (16)$$

where  $V = \max_{\boldsymbol{\pi}^i \in \mathcal{A} \setminus \{\boldsymbol{\pi}^*\}} \left( \sum_{s=1}^d \frac{|\pi_s^1 - \pi_s^i|}{|\pi_{s(i)}^1 - \pi_{s(i)}^i|} \right)^2$  and  $s(i) = \arg \max_{s \in [d]} |\pi_s^1 - \pi_s^i|$ .

We can see that the Minimax-CombSAR algorithm is an optimal algorithm whose upper bound of the probability of error matches the lower bound shown in (7) since we have  $\mathbf{H}_2 \leq \mathbf{H} \leq \log(2d)\mathbf{H}_2$ .

## 5.3 Computational Complexity

We discuss the time complexity of the MinMax-CombSAR algorithm. We can compute the allocation vector  $\mathbf{p}(r)$  with time complexity of  $\mathcal{O}(1)$  if we use Equation (15) (line 9). The time complexity of the elimination procedure for  $\mathcal{A}$  in lines 14 and 15 is of  $\mathcal{O}(K \log K)$ . In total, since we have  $\lceil \log d \rceil$  phases, the time complexity is  $\mathcal{O}(K \log K \lceil \log d \rceil)$ .

## 5.4 Comparison with Katz-Samuels et al. (2020)

Here, we compare the Minimax-CombSAR algorithm with the Peace algorithm introduced in Katz-Samuels et al. (2020). The upper bound of the Peace algorithm can be written as follows:

$$\begin{aligned} &\Pr [\boldsymbol{\pi}^{\text{out}} \neq \boldsymbol{\pi}^*] \\ &\leq 2 \lceil \log(d) \rceil \exp \left( - \frac{T}{c' \log(d) (\gamma_* + \rho_*)} \right), \end{aligned} \quad (17)$$

where

$$\begin{aligned} \gamma_* &= \inf_{\boldsymbol{\lambda} \in \Pi_d} \mathbb{E}_{\boldsymbol{\eta} \sim \mathcal{N}(0, I)} \left[ \sup_{\boldsymbol{\pi} \in \mathcal{A} \setminus \{\boldsymbol{\pi}^*\}} \frac{(\boldsymbol{\pi}^* - \boldsymbol{\pi})^\top \mathbf{A}(\boldsymbol{\lambda})^{\frac{1}{2}} \boldsymbol{\eta}}{\boldsymbol{\mu}^\top (\boldsymbol{\pi}^* - \boldsymbol{\pi})} \right], \\ \rho_* &= \min_{\boldsymbol{\lambda} \in \Pi_d} \max_{\boldsymbol{\pi} \in \mathcal{A} \setminus \{\boldsymbol{\pi}^*\}} \frac{\sum_{s=1}^d \frac{|\pi_s^* - \pi_s|^2}{\lambda_s}}{\Delta_{\boldsymbol{\pi}^*, \boldsymbol{\pi}}^2}, \end{aligned}$$

and  $c'$  is a problem-dependent constant. Here,  $\mathbf{A}(\boldsymbol{\lambda})$  is a diagonal matrix whose  $(i, i)$  element is  $\lambda_i$ . In general, it is not clear whether the upper bound of the Minimax-CombSAR algorithm in (2) is tighter than that of the Peace algorithm shown in (17). In the experiment section, we compare the two algorithms numerically and show that our algorithm outperforms the Peace algorithm.

Table 1: The percentage of correctly identified optimal actions in the knapsack problem. The size of the action set is exponentially large in  $d$ . We conducted experiments for 100 times for each  $d$ .

$d$	10	20	30	40	50	60	70	80	90	100
Naive Baseline	93 %	82 %	63 %	65 %	49 %	47 %	33 %	30 %	24 %	15 %
CSA	98 %	84 %	77 %	70 %	55 %	49 %	51 %	50 %	35 %	32 %

 Table 2: Comparison of the percentage of the best actions correctly identified by the CSA, Minimax-CombSAR, and Peace algorithms in the knapsack problem. We assume that  $|\mathcal{A}| \leq 1000$  is guaranteed. The horizontal axis represents the number of items  $d$ . We ran the experiments a hundred times for each  $d$ .

$d$	10	15	20	25	30	35	40	45	50
Peace	97 %	86 %	88 %	78 %	74 %	71 %	58 %	58 %	55 %
CSA	98 %	96 %	93 %	87 %	83 %	81 %	80 %	73 %	73 %
MinMaxCombSAR( $\beta = 0.2$ )	98 %	93 %	95 %	89 %	93 %	87 %	83 %	85 %	83 %
MinMaxCombSAR( $\beta = 0.2$ )	99 %	94 %	95 %	92 %	93 %	85 %	86 %	79 %	83 %

## 6 Experiment

In this section, we numerically compare the CSA, Minimax-CombSAR, and Peace algorithms. First, we run an experiment where we assume  $|\mathcal{A}|$  is exponentially large in  $d$ . We see the performance of the CSA algorithm with different budgets. Also, we compare the CSA algorithm with a naive baseline method. Next, we compare the CSA, Minimax-CombSAR, and Peace algorithms, where we can assume that  $|\mathcal{A}|$  is polynomial in  $d$ .

### 6.1 When $|\mathcal{A}|$ is Exponentially Large in $d$

Here, the goal is to identify the best action for the knapsack problem. In the knapsack problem, we have  $d$  items. Each item  $s \in [d]$  has a weight  $w_s$  and value  $v_s$ . Also, there is a knapsack whose capacity is  $W$  in which we put items. Our goal is to maximize the total value of the knapsack, not letting the total weight of the items exceed the capacity of the knapsack. Formally, the optimization problem is given as follows:

$$\begin{aligned} & \text{maximize}_{\pi \in \mathcal{A}} && \sum_{s=1}^d v_s \pi_s \\ & \text{subject to} && \sum_{s=1}^d \pi_s w_s \leq W, \end{aligned}$$

where  $\pi_s$  denotes the number of item  $s$  in the knapsack. Here, the weight of each item is known, but the value is unknown, and therefore has to be estimated. In each time step, the player chooses an item  $s$  and gets an observation of value  $v_s$ , which can be regarded as a random variable from an unknown distribution with mean  $v_s$ .

Here, we assume  $|\mathcal{A}|$  is exponentially large in  $d$ . Therefore, we can not use the Minimax-CombSAR and the Peace algorithm since they have to enumerate all the possible actions. For our experiment, we gener-

ated the weight of each item  $s$ ,  $w_s$ , uniformly from  $\{1, 2, \dots, 200\}$ . For each item  $s$ , we generate  $v_s$  as  $v_s = w_s \times x$ , where  $x$  is a uniform sample from  $[1.0, 1.1]$ . The capacity of the knapsack is  $W = 200$ . Each time we choose an item  $s$ , we observe a value  $v_s + \epsilon$  where  $\epsilon$  is a noise from  $\mathcal{N}(0, 1)$ . We set  $R = 1$ .

We compare the CSA algorithm to a naive baseline, which pulls each arm equally and returns the action with the largest reward using the sample mean. The budget  $T$  is set to 10000. The results are shown in Table 1. We can see that the CSA algorithm outperforms the baseline method in all  $d$ .

### 6.2 When $|\mathcal{A}|$ is Polynomial in $d$

Here, we assume that we have prior knowledge of the rewards of arms, i.e., knowledge of the values of the items. We assume that we know each  $v_s$  is in  $[w_s, 1.1 \times w_s]$ , and use this prior knowledge to generate the action class  $\mathcal{A}$  in the following procedure. We first generate a vector  $v'$  whose  $s$ -th element  $v'_s$  is uniformly sampled from  $[w_s, 1.1 \times w_s]$ , and then solve the knapsack problem with  $v'_s$  and add the obtained solution  $\pi$  to  $\mathcal{A}$ . We repeat this 1000 times; therefore  $|\mathcal{A}| \leq 1000$ . In this experiment, the budget is  $T = 20000$ . For the Minimax-CombSAR algorithm, we set  $\beta = 0.2, 0.4$ . We run forty experiments for each  $d \in \{10, 15, \dots, 100\}$ . We show the result in Table 2. We can see that the Minimax-CombSAR algorithm outperforms the other two algorithms for almost every  $d$ . Also, the CSA algorithm outperforms the Peace algorithm for almost every  $d$ .

Next, we compare the performance of each method by conducting an experiment on the best action identification task on the optimal transport problem. We use the notation introduced in Section 4.2.2. We set the supplier vector and the demand vector as



Table 3: Comparison of the percentage of the best actions correctly identified by the CSA, Minimax-CombSAR, and Peace algorithms in the optimal transport problem. We assume that  $|\mathcal{A}| \leq 100$  is guaranteed. We show the accuracy of identifying the optimal action for  $d = \{4, 9, 16, 25, 36, 49, 64, 81, 100\}$ . For each  $d$ , we conducted the experiment for 100 times.

$d$	4	9	16	25	36	49	64	81	100
Peace	84 %	67 %	58 %	49 %	41 %	35 %	36 %	24 %	11 %
CSA	84 %	75 %	57 %	46 %	50 %	38 %	31 %	39 %	15 %
MinMax-CombSAR ( $\beta = 0.2$ )	82 %	83 %	73 %	67 %	59 %	54 %	49 %	48 %	29 %

$\mathbf{s} = (1, 2, \dots, m)$  and  $\mathbf{d} = (1, 2, \dots, n)$ , respectively. Also, we set  $m = n$ . We generate each element of the cost matrix  $\gamma$  from a uniform distribution  $U(0, 1)$ . Here, we assume that  $\gamma$  is unknown, and therefore has to be estimated. Note that the number of arms  $d$  is  $mn = m^2$  in this setting. To generate the action set, we follow the following procedure. We first generate a random matrix  $\gamma'$ , where each element is sampled from  $U(0, 1)$ . Then, we solve the optimal transport problem with the offline oracle and add the obtained solution  $\pi'$  to  $\mathcal{A}$ . We repeat this procedure for 100 times, and therefore  $|\mathcal{A}| \leq 100$ . We show the results in Table 3. We can see that the MinMax-CombSAR algorithm outperforms the other two methods.

## 7 Conclusion

In this paper, we studied the fixed-budget R-CPE-MAB. We first introduced the CSA algorithm, which is the first algorithm that can identify the best action even when the size of the action class is exponentially large with respect to the number of arms. However, it still has an extra logarithmic term in the exponent. Then, we proposed an optimal algorithm named the Minimax-CombSAR algorithm, which, although it is applicable only when the action class is polynomial, matches a lower bound. We showed that both of the algorithms outperform the existing methods.

## Acknowledgement

We thank Dr. Kevin Jamieson for his very helpful advice and comments on existing studies. SN was supported by JST SPRING, Grant Number JPMJSP2108. MS was supported by the Institute for AI and Beyond, UTokyo.

## References

- Ahmadi, A. A. (2016). Lec11p1, ORF363/COS323.
- Audibert, J.-Y., Bubeck, S., and Munos, R. (2010). Best arm identification in multi-armed bandits. In *The 23rd Conference on Learning Theory*, pages 41–53.
- Audibert, J.-Y., Munos, R., and Szepesvári, C. (2009). Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theor. Comput. Sci.*, 410:1876–1902.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256.
- Bubeck, S., Munos, R., and Stoltz, G. (2009). Pure exploration in multi-armed bandits problems. In *International Conference on Algorithmic Learning Theory*.
- Carpentier, A. and Locatelli, A. (2016). Tight (lower) bounds for the fixed budget best arm identification bandit problem. In *Annual Conference Computational Learning Theory*.
- Chen, L., Gupta, A., and Li, J. (2016). Pure exploration of multi-armed bandit under matroid constraints. In *Proceedings of the 29th Conference on Learning Theory, COLT 2016, New York, USA, June 23-26, 2016*, volume 49 of *JMLR Workshop and Conference Proceedings*, pages 647–669. JMLR.org.
- Chen, L., Gupta, A., Li, J., Qiao, M., and Wang, R. (2017). Nearly optimal sampling algorithms for combinatorial pure exploration. In Kale, S. and Shamir, O., editors, *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 482–534. PMLR.
- Chen, S., Lin, T., King, I., Lyu, M. R., and Chen, W. (2014). Combinatorial pure exploration of multi-armed bandits. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1, NIPS’14*, page 379–387, Cambridge, MA, USA. MIT Press.
- Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. In Burges, C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K., editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Dantzig, T. and Mazur, J. (2007). *Number: The Language of Science*. A Plume book. Penguin Publishing Group.

- Du, Y., Kuroki, Y., and Chen, W. (2021). Combinatorial pure exploration with bottleneck reward function. In *Advances in Neural Information Processing Systems*, volume 34, pages 23956–23967. Curran Associates, Inc.
- Fujimoto, N. (2016). A pseudo-polynomial time algorithm for solving the knapsack problem in polynomial space. In Chan, T.-H. H., Li, M., and Wang, L., editors, *Combinatorial Optimization and Applications*, pages 624–638, Cham. Springer International Publishing.
- Gabillon, V., Lazaric, A., Ghavamzadeh, M., Ortner, R., and Bartlett, P. (2016). Improved learning complexity in combinatorial pure exploration bandits. In Gretton, A. and Robert, C. C., editors, *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 1004–1012, Cadiz, Spain. PMLR.
- Garey, M. R. and Johnson, D. S. (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness (Series of Books in the Mathematical Sciences)*. W. H. Freeman, first edition edition.
- Garey, M. R. and Johnson, D. S. (1990). *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., USA.
- Gibbons, A. (1985). *Algorithmic Graph Theory*. Cambridge University Press.
- Goodman, J. E. and O’Rourke, J., editors (1997). *Handbook of Discrete and Computational Geometry*. CRC Press, Inc., USA.
- Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30.
- Hoffmann, L. and Bradley, G. (2009). *Calculus for Business, Economics, and the Social and Life Sciences, Brief*. McGraw-Hill Companies, Incorporated.
- Ito, S., Hatano, D., Sumita, H., Takemura, K., Fukunaga, T., Kakimura, N., and Kawarabayashi, K.-I. (2019). Oracle-efficient algorithms for online linear optimization with bandit feedback. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Kalyanakrishnan, S. and Stone, P. (2010). Efficient selection of multiple bandit arms: Theory and practice. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML’10, page 511–518, Madison, WI, USA. Omnipress.
- Katz-Samuels, J., Jain, L., Karnin, Z., and Jamieson, K. (2020). An empirical process approach to the union bound: Practical algorithms for combinatorial and linear bandits. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS’20, Red Hook, NY, USA. Curran Associates Inc.
- Kellerer, H., Pferschy, U., and Pisinger, D. (2004). *Knapsack Problems*. Springer, Berlin, Germany.
- Nakamura, S. and Sugiyama, M. (2023). Thompson sampling for real-valued combinatorial pure exploration of multi-armed bandit. In *The 38th Annual AAAI Conference on Artificial Intelligence*.
- Pettie, S. and Ramachandran, V. (2002). An optimal minimum spanning tree algorithm. *J. ACM*, 49(1):16–34.
- Peyré, G. and Cuturi, M. (2019). Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607.
- Pochet, Y. and Wolsey, L. A. (2010). *Production Planning by Mixed Integer Programming*. Springer Publishing Company, Incorporated, 1st edition.
- Rinaldo, A. and Bong, H. (2018). 36-710: Advanced statistical theory. lecture 3.
- Rivasplata, O. (2012). Subgaussian random variables : An expository note.
- Sniedovich, M. (2006). Dijkstra’s algorithm revisited: the dynamic programming connexion. *Control and Cybernetics*, 35:599–620.
- Villani, C. (2008). *Optimal transport – Old and new*, volume 338, pages xxii+973.
- Wang, S. and Zhu, J. (2022a). Thompson sampling for (combinatorial) pure exploration. In *Proceedings of the 39th International Conference on Machine Learning*, Baltimore, Maryland, USA.
- Wang, S. and Zhu, J. (2022b). Thompson sampling for (combinatorial) pure exploration. In *Proceedings of the 39th International Conference on Machine Learning*, Baltimore, Maryland, USA.
- Xu, H. and Li, J. (2021). Simple combinatorial algorithms for combinatorial bandits: corruptions and approximations. In de Campos, C. P., Maathuis, M. H., and Quaeghebeur, E., editors, *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI 2021, Virtual Event, 27-30 July 2021*, volume 161 of *Proceedings of Machine Learning Research*, pages 1444–1454. AUAI Press.
- Yang, J. and Tan, V. (2022). Minimax optimal fixed-budget best arm identification in linear bandits. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*.

## Checklist

The checklist follows the references. For each question, choose your answer from the three possible options: Yes, No, Not Applicable. You are encouraged to include a justification to your answer, either by referencing the appropriate section of your paper or providing a brief inline description (1-2 sentences). Please do not modify the questions. Note that the Checklist section does not count towards the page limit. Not including the checklist in the first submission won't result in desk rejection, although in such case we will ask you to upload it during the author response period and include it in camera ready (if accepted).

**In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.**

1. For all models and algorithms presented, check if you include:
  - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
  - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
  - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]
2. For any theoretical claim, check if you include:
  - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
  - (b) Complete proofs of all theoretical results. [Yes]
  - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
  - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
  - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
  - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
  - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Not Applicable]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
  - (a) Citations of the creator If your work uses existing assets. [Yes]
  - (b) The license information of the assets, if applicable. [Not Applicable]
  - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
  - (d) Information about consent from data providers/curators. [Not Applicable]
  - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
  - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
  - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
  - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

## A Proof of Theorem 1

Here, we prove Theorem 1. For the reader's convenience, we restate the Theorem.

**Theorem 1.** *For any action class and any algorithm that returns an action  $\boldsymbol{\pi}^{\text{out}}$  after  $T$  times of arm pulls, the probability of error is at least*

$$\mathcal{O}\left(\exp\left(-\frac{T}{\mathbf{H}}\right)\right). \quad (7)$$

We follow a similar discussion to Carpentier and Locatelli (2016). This is a lower bound that will hold in the much easier problem where the learner knows that the bandit setting she is facing is one of only  $d$  given bandit settings. This lower bound ensures that even in this much simpler case, the learner will make a mistake.

Let  $\mathcal{N}(\mu, 1)$  denote the Gaussian distribution of mean  $\mu$  and unit variance. For any  $k \in [d]$ , we write  $\nu_k := \mathcal{N}(\mu_k, 1)$ . Also, for any  $k \in [d]$ , we define  $\nu'_k$  as follows:

$$\nu'_k := \begin{cases} \mathcal{N}(\mu_k + 2\Delta_k, 1) & (\text{if } \pi_k^* < \pi_k^{(k)}) \\ \mathcal{N}(\mu_k - 2\Delta_k, 1) & (\text{if } \pi_k^* > \pi_k^{(k)}) \end{cases}.$$

Recall that  $\Delta_k = \frac{\boldsymbol{\mu}^\top(\boldsymbol{\pi}^* - \boldsymbol{\pi}^{(k)})}{|\pi_k^* - \pi_k^{(k)}|}$  and  $\boldsymbol{\pi}^{(k)} = \arg \min_{\boldsymbol{\pi} \neq \boldsymbol{\pi}^*} \frac{\boldsymbol{\mu}^\top(\boldsymbol{\pi}^* - \boldsymbol{\pi})}{|\pi_k^* - \pi_k|}$ .

For any  $s \in [d]$ , we define the product distributions  $\mathcal{G}^s$  as  $\nu_1^s \otimes \cdots \otimes \nu_d^s$ , where for  $1 \leq k \leq d$ ,

$$\nu_k^s := \nu_k \mathbf{1}[k \neq s] + \nu'_k \mathbf{1}[k = s]. \quad (18)$$

The bandit problem associated with distribution  $\mathcal{G}^s$ , and that we call “the bandit problem  $s$ ” is such that, for any  $1 \leq k \leq d$ , arm  $k$  has distribution  $\nu_k^s$ , i.e., all arms have distribution  $\nu_k$  except for arm  $s$  that has distribution  $\nu'_k$ . We denote by  $\boldsymbol{\mu}^{(s)}$  the vector of expected rewards of bandit problem  $s$ , i.e.,  $\mu_k^{(s)} = \mu_k$  for  $k \neq s$ ,  $\mu_k^{(s)} = \mu_k + 2\Delta_k$  for  $k = s$  and  $\pi_k^* < \pi_k^{(k)}$ ,  $\mu_k^{(s)} = \mu_k - 2\Delta_k$  for  $k = s$  and  $\pi_k^* > \pi_k^{(k)}$ . For any  $1 \leq s \leq d$ ,  $\mathbb{P}_s := \mathbb{P}_{(\mathcal{G}^s)^{\otimes T}}$  for the probability distribution of the bandit problem  $s$  according to all the samples that a strategy could possibly collect up to horizon  $T$ , i.e., according to the samples  $(X_{k,u})_{1 \leq k \leq d, 1 \leq u \leq T} \sim (\mathcal{G}^s)^{\otimes T}$ . Also, we define  $\mathbb{P}_0 := \mathbb{P}_{(\mathcal{G}^0)^{\otimes T}}$ , where  $\mathcal{G}^0 := \nu_1 \otimes \cdots \otimes \nu_d$ .

In the bandit problem  $s$ ,  $\boldsymbol{\pi}^*$  is no longer the best action since the difference of the expected reward between  $\boldsymbol{\pi}^*$  and  $\boldsymbol{\pi}^{(s)}$  is

$$\begin{aligned} & \sum_{e \in [d] \setminus s} \mu_e(\pi_e^* - \pi_e^{(s)}) - 2\Delta_s \times |\pi_s^* - \pi_s^{(s)}| \\ & \leq -\boldsymbol{\mu}^\top(\boldsymbol{\pi}^* - \boldsymbol{\pi}^{(s)}) < 0. \end{aligned} \quad (19)$$

We denote by  $\boldsymbol{\pi}^{*,(s)}$  the best action in the bandit problem  $s$ .

Now, we prove Theorem 1.

### Step 1: Definition of KL-divergence

For two distributions  $\nu$  and  $\nu'$  defined on  $\mathbb{R}$ , and that are such that  $\nu$  is absolutely continuous with respect to  $\nu'$ , we write

$$\text{KL}(\nu, \nu') = \int_{\mathbb{R}} \log\left(\frac{d\nu(x)}{d\nu'(x)}\right) d\nu(x), \quad (20)$$

for the Kullback Leibler divergence between distribution  $\nu$  and  $\nu'$ . For any  $k \in [d]$ , let us write

$$\text{KL}_k \triangleq \text{KL}(\nu_k, \nu'_k) = \frac{1}{2}(\mu_a - \mu_b)^2. \quad (21)$$

for the Kullback-Leibler divergence between two Gaussian distributions  $\mathcal{N}(\mu_a, 1)$  and  $\mathcal{N}(\mu_b, 1)$ .

Let  $1 \leq t \leq T$ . Also, let us fix  $e \in [d]$ , and think of bandit instance  $e$ . We define the quantity:

$$\begin{aligned} \widehat{\text{KL}}_{s,t} &= \frac{1}{t} \sum_{u=1}^t \log \left( \frac{d\nu_s}{d\nu'_s} (X_{s,u}) \right) \\ &= \begin{cases} \frac{1}{t} \sum_{u=1}^t -\Delta_s (X_{s,u} - \mu_s - 2\Delta_s) & \text{if } \pi_s^* > \pi_s^{(s)} \text{ and } e \neq s, \\ \frac{1}{t} \sum_{u=1}^t \Delta_s (X_{s,u} - \mu_s + 2\Delta_s) & \text{if } \pi_s^* < \pi_s^{(s)} \text{ and } e \neq s, \\ 0 & \text{if } e = s \end{cases} \end{aligned} \quad (22)$$

where  $X_{s,u} \sim_{\text{i.i.d}} \nu_s^e$  for any  $u \leq t$ .

Next, let us define the following event:

$$\xi = \left\{ \forall 1 \leq s \leq d, \forall 1 \leq t \leq T, \left| \widehat{\text{KL}}_{s,t} \right| - \text{KL}_s \leq \sqrt{2\Delta_s^2 \frac{\log(12Td)}{t}} \right\}. \quad (23)$$

The following lemma shows a concentration bound for  $\left| \widehat{\text{KL}}_{s,t} \right|$ .

**Lemma 1.** *It holds that*

$$\Pr_e [\xi] \geq \frac{5}{6} \quad (24)$$

*Proof.* When  $e = s$ ,  $\left| \widehat{\text{KL}}_{s,t} \right| - \text{KL}_s \leq \sqrt{2\Delta_s^2 \frac{\log(12Td)}{t}}$  holds for any  $t$ . Therefore, we think of  $e \neq s$ . Since  $X_{s,t}$  is a 1-sub-Gaussian random variable, we can see that  $\widehat{\text{KL}}_{s,t}$  is a  $\frac{\Delta_s^2}{t}$ -sub-Gaussian random variable from (22). We can apply the Hoeffding's inequality to this quantity, and we have that with probability larger than  $1 - \frac{1}{6dT}$ ,

$$\left| \widehat{\text{KL}}_{s,t} \right| - \text{KL}_s \leq \sqrt{\frac{2\Delta_s^2 \log(12Td)}{t}} \quad (25)$$

□

## Step2: A change of measure

Let  $\mathcal{ALG}$  denote the active strategy of the learner, that returns some action  $\boldsymbol{\pi}_{\text{out}}$  after pulling arms  $T$  times in total. Let  $(T_s)_{1 \leq s \leq T}$  denote the number of samples collected by  $\mathcal{ALG}$  on each arm. These quantities are stochastic but it holds that  $\sum_{s=1}^d T_s = T$ . For any  $1 \leq s \leq d$ , let us write

$$t_s = \mathbb{E}_0 [T_s]. \quad (26)$$

It holds also that  $\sum_{s=1}^d t_s = T$ .

We recall the change of measure identity (Audibert et al., 2010), which states that for any measurable event  $\mathcal{E}$  and for any  $1 \leq s \leq d$ ,

$$\Pr_s [\mathcal{E}] = \mathbb{E}_0 \left[ \mathbf{1} [\mathcal{E}] \exp \left( -T_s \widehat{\text{KL}}_{s,T_s} \right) \right], \quad (27)$$

where  $\mathbf{1}$  denotes the indicator function.

Then, we consider the following event:

$$\mathcal{E}_s = \{ \boldsymbol{\pi}_{\text{out}} = \boldsymbol{\pi}^* \} \cap \{ \xi \} \cap \{ T_s \leq 6t_s \}, \quad (28)$$

i.e., the event where the algorithm outputs action  $\boldsymbol{\pi}^*$  at the end, where  $\xi$  holds, and where the number of times arm  $s$  is pulled is smaller than  $6t_s$ . From (28), we have

$$\Pr_s [\mathcal{E}_s] = \mathbb{E}_0 \left[ \mathbf{1} [\mathcal{E}_s] \exp \left( -T_s \widehat{\text{KL}}_{s,T_s} \right) \right] \quad (29)$$

$$\geq \mathbb{E}_0 \left[ \mathbf{1} [\mathcal{E}_s] \exp \left( -T_s \text{KL}_s - \sqrt{2T_s \Delta_s^2 \log(12Td)} \right) \right] \quad (30)$$

$$\geq \mathbb{E}_0 \left[ \mathbf{1} [\mathcal{E}_s] \exp \left( -6t_s \text{KL}_s - \sqrt{2T \Delta_s^2 \log(12Td)} \right) \right] \quad (31)$$

$$\geq \exp \left( -6t_s \text{KL}_s - \sqrt{2T \Delta_s^2 \log(12Td)} \right) \Pr_0 [\mathcal{E}_s], \quad (32)$$

since on  $\mathcal{E}_s$ , we have that  $\xi$  holds and that  $T_s \leq 6t_s$ , and  $\mathbb{E}_0 [\widehat{\text{KL}}_{s,t}] = \text{KL}_s$  for any  $t \leq T$ .

**Step 3: Lower bound on  $\Pr_0[\mathcal{E}_1]$  for any reasonable algorithm**

Assume that the algorithm  $\mathcal{ALG}$  is a  $\delta$ -correct algorithm. Then, we have  $\Pr_0[\boldsymbol{\pi}_{\text{out}} \neq \boldsymbol{\pi}^*] \leq \delta$ . From the Markov's inequality, we have

$$\Pr_0[T_s \geq 6t_s] \leq \frac{\mathbb{E}_0 T_s}{6t_s} = \frac{1}{6}, \quad (33)$$

since  $\mathbb{E}_0 [T_s] = t_s$  for algorithm  $\mathcal{ALG}$ .

Combining Lemma 1, (33), and the fact that  $\mathcal{ALG}$  is a  $\delta$ -correct algorithm, by the union bound, we have

$$\Pr_0[\mathcal{E}_s] \geq 1 - (1/6 + \delta + 1/6) = \frac{2}{3} - \delta. \quad (34)$$

This fact combined with (32), (21), and the fact that  $\Pr_s[\boldsymbol{\pi}_{\text{out}} \neq \boldsymbol{\pi}^{*,(s)}] \geq \Pr_s[\mathcal{E}_s]$ , we have

$$\begin{aligned} \Pr_s[\boldsymbol{\pi}_{\text{out}} \neq \boldsymbol{\pi}^{*,(s)}] &\geq \left(\frac{2}{3} - \delta\right) \exp\left(-6t_s \text{KL}_s - \sqrt{2T\Delta_s^2 \log(12Td)}\right) \\ &= \left(\frac{2}{3} - \delta\right) \exp\left(-12t_s (\Delta_s)^2 - \sqrt{2T\Delta_s^2 \log(12Td)}\right) \end{aligned}$$

since  $\text{KL}_s = 2\Delta_s^2$ .

**Step 4: Conclusion.** Since  $\mathbf{H} = \sum_{s=1}^d \left(\frac{1}{\Delta_s}\right)^2$  and  $\sum_{s=1}^d t_s = T$ , there exists  $e \in [d]$  such that

$$t_e \leq \frac{T}{\mathbf{H}\Delta_e^2}, \quad (35)$$

as the contradiction yields an immediate contradiction. For this  $e$ , it holds by (35) that

$$\Pr_e[\boldsymbol{\pi}_{\text{out}} \neq \boldsymbol{\pi}^{*,(e)}] \geq \left(\frac{2}{3} - \delta\right) \exp\left(-12\frac{T}{\mathbf{H}} - 2\sqrt{T\Delta_e^2 \log(6Td)}\right). \quad (36)$$

This concludes the proof.

## B How to construct COracles

Here, we show how to construct COracles once  $\boldsymbol{\mu}$  and  $\mathbf{S}(t)$  are given for some specific combinatorial problems.

### B.1 COracle for the Knapsack Problem

Here, we show that we can construct the COracle for the knapsack problem by calling the *offline oracle* once. Let  $\mathbf{S}^1(t)$  be the set that collects only the first component of each of all elements in  $\mathbf{S}(t)$ . Also, let  $W_{\text{done}} = \sum_{(e,x) \in \mathbf{S}(t)} xw_e$ . If  $W - W_{\text{done}} < 0$ , the COracle outputs  $\perp$ . Otherwise, we call the *offline oracle* and solve the following optimization problem:

$$\begin{aligned} &\text{maximize}_{\boldsymbol{\xi}} \quad \sum_{s \in [d] \setminus \mathbf{S}^1(t)} \mu_s \xi_s \\ &\text{subject to} \quad \sum_{s \in \mathbf{S}^1(t)} \xi_s w_s \leq W - W_{\text{done}}. \end{aligned}$$

Then, we return  $\text{COracle}(\boldsymbol{\mu}, \mathbf{S}(t)) = \mathbf{S}(t) \cup \boldsymbol{\xi}$ .

## B.2 COracle for the Optimal Transport Problem

Here, in Algorithm 3, we show the *COracle* for the optimal transport problem, which calls the *offline oracle* once. First, let us define  $\mathbf{S}^1(t)$  as the set that collects only the first element of each of all elements in  $\mathbf{S}(t)$ . Then, let us define  $\pi^{\text{CO}}$  as the final output of the COracle. Also, let us define  $\mathbf{s}'$  and  $\mathbf{d}'$  as follows:

$$\forall i \in [m], j \in [n], x \in \text{POSSIBLE-PI}((i, j)), s'_i = \begin{cases} x & (\text{if } ((i, j), x) \in \mathbf{S}(t), \\ 0 & (\text{otherwise}), \end{cases}$$

and

$$\forall i \in [m], j \in [n], x \in \text{POSSIBLE-PI}((i, j)), d'_j = \begin{cases} x & (\text{if } ((i, j), x) \in \mathbf{S}(t), \\ 0 & (\text{otherwise}). \end{cases}$$

Also, let us define  $\pi' \in \mathbb{R}^{m \times n}$  whose  $(i, j)$ -th element is defined as follows:

$$\forall i \in [m], j \in [n], x \in \text{POSSIBLE-PI}((i, j)), \pi'_{ij} = \begin{cases} x & (\text{if } ((i, j), x) \in \mathbf{S}(t), \\ 0 & (\text{otherwise}). \end{cases}$$

Intuitively, for any  $((i, j), x) \in \mathbf{S}(t)$ , we have to send resources from supplier  $i$  to demander  $j$  for a value of  $s'_i (= d'_j)$ . This means that we have resources of  $s_i - s'_i$  left to send from supplier  $i$ , and  $d_j - d'_j$  left to send to demander  $j$ . Since we can not send a negative value of resources,  $s_i - s'_i \geq 0$  and  $d_j - d'_j \geq 0$  have to be satisfied.

From the above discussion, we can formally write  $\pi^{\text{CO}}$  as follows:

$$\pi^{\text{CO}} = \begin{cases} \sum_{i,j} \pi'_{ij} \gamma_{ij} + \arg \min_{\pi'' \in \mathcal{G}(\mathbf{s} - \mathbf{s}', \mathbf{d} - \mathbf{d}')} \sum_{i,j} \pi''_{ij} \gamma'_{ij} & (\text{if there is no negative value in neither } \mathbf{s} - \mathbf{s}' \text{ nor } \mathbf{d} - \mathbf{d}'), \\ \perp & (\text{otherwise}). \end{cases}$$

Here,  $\gamma'$  is defined as

$$\forall i \in [m], j \in [n], x \in \text{POSSIBLE-PI}((i, j)), \gamma'_{ij} = \begin{cases} -\infty & (\text{if } ((i, j), x) \in \mathbf{S}(t), \\ \gamma_{ij} & (\text{otherwise}), \end{cases}$$

to not let  $\pi^{\text{CO}}$  send materials more than the amount determined by  $\mathbf{S}(t)$ .

---

### Algorithm 3 COracle for the Optimal Transport Problem

---

**Input:** Cost matrix:  $\gamma$ , supply vector  $\mathbf{s} \in \mathbb{R}^m$ , demand vector  $\mathbf{d} \in \mathbb{R}^n$ ,  $\mathbf{S}(t)$

- 1:  $\pi' \leftarrow \mathbf{0}^{m \times n}$
  - 2: **for**  $((i, j), x)$  in  $\mathbf{S}(t)$  **do**
  - 3:    $\pi'_{ij} \leftarrow x$
  - 4:    $s_i \leftarrow s_i - x$
  - 5:    $d_j \leftarrow d_j - x$
  - 6:    $\gamma_{ij} \leftarrow -\infty$
  - 7:   **if**  $s_i < 0$  or  $d_j < 0$  **then**
  - 8:     Return  $\perp$
  - 9:   **end if**
  - 10: **end for**
  - 11: Compute  $\pi'' = \arg \min_{\pi \in \mathcal{G}(\mathbf{s}, \mathbf{d})} \sum_{i,j} \pi_{ij} \gamma_{ij}$  using the offline oracle
  - 12: Return  $\pi^{\text{CO}} = \pi' + \pi''$
- 

## B.3 COracle for a General Case When $K(= |\mathcal{A}|) = \text{poly}(d)$

Here, we show the time complexity of the COracle when the size of the action class is polynomial in  $d$ . If  $\mathcal{A}_S$  is an empty set, the COracle returns  $\perp$ . Otherwise, we return  $\pi^{\text{return}} = \arg \max_{\pi \in \mathcal{A}_S} \boldsymbol{\mu}^\top \pi$ . The time complexity to construct  $\mathcal{A}_S$  is  $\mathcal{O}(dK)$ . This is because, for each action  $\pi \in \mathcal{A}$ , we check every dimension to see if  $\pi \in \mathcal{A}_S$ . Then, the time complexity of finding  $\pi^{\text{return}}$  from  $\mathcal{A}_S$  is  $\mathcal{O}(K \log K)$ . Therefore the time complexity of the COracle is  $\mathcal{O}(dK + K \log K)$ .

## C Proof of Theorem 2

Here, we prove Theorem 2. We first introduce some notions that are useful to prove it. Then, we show some preparatory lemmas that are needed to prove Theorem 2.

### C.1 Preparatories

Let us introduce some notions that are useful to prove the theorem.

#### C.1.1 Arm-Value Pair

We define an *arm-value pair set*  $\mathbf{M}(\boldsymbol{\pi}) \subset [d] \times \mathbb{R}$  of  $\boldsymbol{\pi}$  as follows:

$$\mathbf{M}(\boldsymbol{\pi}) = \{(e, \pi_e) \mid \forall e \in [d]\}.$$

Also, we define the *arm-value pair family*  $\mathcal{M}$  as follows:

$$\mathcal{M} = \{\mathbf{M}(\boldsymbol{\pi}) \mid \forall \boldsymbol{\pi} \in \mathcal{A}\}.$$

For any arm-value pair set  $\mathbf{M} \in \mathcal{M}$ , we denote by  $M_e$  the second component of the ordered pair whose first component is  $e$ , i.e.,  $M_e = x$  for any  $(e, x) \in \mathbf{M}$ . Also, we call  $M_e$  the  $e$ -th element of  $\mathbf{M}$ .

For any two different arm-value pair set  $\mathbf{M}, \mathbf{M}' \in \mathcal{M}$ , we define an operator  $\boxminus$  such that the  $e$ -th element of  $\mathbf{M}' \boxminus \mathbf{M} \subset [d] \times \mathbb{R}$ ,  $(\mathbf{M}' \boxminus \mathbf{M})_e$ , is defined as follows:

$$(\mathbf{M}' \boxminus \mathbf{M})_e = \begin{cases} M'_e - M_e & (\text{if } M'_e > M_e) \\ 0 & (\text{if } M'_e \leq M_e) \end{cases}.$$

#### C.1.2 Exchange Class

We define an *exchange set*  $b$  as an ordered pair of sets  $b = (\mathbf{b}^+, \mathbf{b}^-)$ , where  $\mathbf{b}^+, \mathbf{b}^- \subset [d] \times \mathbb{R} \setminus \{0\}$ . We say  $\mathbf{b}^+$  (or  $\mathbf{b}^-$ ) has an  $e$ -changer if it has an element whose first component is  $e$ . For any  $\mathbf{b}^+$  (or  $\mathbf{b}^-$ ), we denote by  $b_e^+$  the second component of the ordered pair whose first component is  $e$ , i.e.,  $b_e^+ = x$  for any  $(e, x) \in \mathbf{b}^+$ . Also, for any  $b = (\mathbf{b}^+, \mathbf{b}^-)$ , we do not let both  $\mathbf{b}^+$  and  $\mathbf{b}^-$  have  $e$ -changers.

Next, for any arm-value pair set  $\mathbf{M} \in \mathcal{M}$ , exchange set  $b = (\mathbf{b}^+, \mathbf{b}^-)$ , and  $e \in [d]$ , we define operator  $\oplus$  such that the  $e$ -th element of  $\mathbf{M} \oplus b \subset [d] \times \mathbb{R}$ ,  $(\mathbf{M} \oplus b)_e$ , is defined as follows:

$$(\mathbf{M} \oplus b)_e = \begin{cases} M_e + b_e^+ & (\text{if } \mathbf{b}^+ \text{ has an } e\text{-changer}) \\ M_e - b_e^- & (\text{if } \mathbf{b}^- \text{ has an } e\text{-changer}) \\ M_e & (\text{if neither } \mathbf{b}^+ \text{ nor } \mathbf{b}^- \text{ has an } e\text{-changer}) \end{cases}. \quad (37)$$

Similarly, for any arm-value pair set  $\mathbf{M} \in \mathcal{M}$ , exchange set  $b = (\mathbf{b}^+, \mathbf{b}^-)$ , and  $e \in [d]$ , we define operator  $\ominus$  such that the  $e$ -th element of  $\mathbf{M} \ominus b \subset [d] \times \mathbb{R}$ ,  $(\mathbf{M} \ominus b)_e$ , is defined as follows:

$$(\mathbf{M} \ominus b)_e = \begin{cases} M_e - b_e^+ & (\text{if } \mathbf{b}^+ \text{ has an } e\text{-changer}) \\ M_e + b_e^- & (\text{if } \mathbf{b}^- \text{ has an } e\text{-changer}) \\ M_e & (\text{if neither } \mathbf{b}^+ \text{ nor } \mathbf{b}^- \text{ has an } e\text{-changer}) \end{cases}. \quad (38)$$

We call a collection of exchange sets  $\mathcal{B}$  an *exchange class* for  $\mathcal{M}$  if  $\mathcal{B}$  satisfies the following property. For any  $\mathbf{M}, \mathbf{M}' \in \mathcal{M}$  and  $e \in [d]$ , where  $(\mathbf{M} \boxminus \mathbf{M}')_e > 0$ , there exists an exchange set  $b = (\mathbf{b}^+, \mathbf{b}^-) \in \mathcal{B}$  that satisfies the following five constraints: (a)  $b_e^- = (\mathbf{M} \boxminus \mathbf{M}')_e$  (b)  $\mathbf{b}^+ \subseteq \mathbf{M}' \boxminus \mathbf{M}$ , (c)  $\mathbf{b}^- \subseteq \mathbf{M} \boxminus \mathbf{M}'$ , (d)  $(\mathbf{M} \oplus b) \in \mathcal{M}$  and (e)  $(\mathbf{M}' \ominus b) \in \mathcal{M}$ .

For any  $\mathbf{a} \subset [d] \times \mathbb{R}$ , let  $a_e$  denote the second component of the ordered pair whose first component is  $e$ , i.e.,  $a_e = x$  for any  $(e, x) \in \mathbf{a}$ . Then, let  $\boldsymbol{\chi}(\mathbf{a}) \in \mathbb{R}^d$  denote the vector defined as follows:

$$\chi_e(\mathbf{a}) = \begin{cases} a_e & (\text{if } \mathbf{a} \text{ has an element whose first component is } e) \\ 0 & (\text{otherwise}) \end{cases}.$$

Also, for any exchange set  $b$ , we define  $\boldsymbol{\chi}(b) = \boldsymbol{\chi}(\mathbf{b}^+) - \boldsymbol{\chi}(\mathbf{b}^-)$ .



## C.2 Preparatory Lemmas

Here, let us introduce some lemmas that are useful to prove the theorem. Below, we define  $\mathbf{M}^* = \mathbf{M}(\boldsymbol{\pi}^*)$ .

**Lemma 2** (Interpolation Lemma). *Let  $\mathcal{B}$  be an exchange class of  $\mathcal{M}$ , and  $\mathbf{M}, \mathbf{M}'$  be two different members of  $\mathcal{M}$ . Then, for any  $e \in \{s \in [d] \mid M_s \neq M'_s\}$ , there exists an exchange set  $b = (\mathbf{b}^+, \mathbf{b}^-) \in \mathcal{B}$ , which satisfies the following five constraints: (a)  $b_e^- = (\mathbf{M} \boxminus \mathbf{M}')_e$  or  $b_e^+ = (\mathbf{M}' \boxminus \mathbf{M})_e$ , (b)  $\mathbf{b}^- \subseteq (\mathbf{M} \boxminus \mathbf{M}')$ , (c)  $\mathbf{b}^+ \subseteq \mathbf{M}' \boxminus \mathbf{M}$ , (d)  $(\mathbf{M} \oplus b) \in \mathcal{M}$  and (e)  $(\mathbf{M}' \ominus b) \in \mathcal{M}$ . Moreover, if  $\mathbf{M}' = \mathbf{M}^*$ , then we have*

$$\frac{\langle \boldsymbol{\mu}, \boldsymbol{\chi}(b) \rangle}{\chi_e(b)} \geq \Delta_e. \quad (39)$$

*Proof.* We decompose our proof into two cases.

**Case (1):**  $(\mathbf{M} \boxminus \mathbf{M}')_e > 0$

By the definition of the exchange class, we know that there exists  $b = (\mathbf{b}^+, \mathbf{b}^-) \in \mathcal{B}$  which satisfies that there is an  $e \in [d]$  where (a)  $b_e^- = (\mathbf{M} \boxminus \mathbf{M}')_e$ , (b)  $\mathbf{b}^- \subseteq (\mathbf{M} \boxminus \mathbf{M}')$ , (c)  $\mathbf{b}^+ \subseteq \mathbf{M}' \boxminus \mathbf{M}$ , (d)  $(\mathbf{M} \oplus b) \in \mathcal{M}$  and (e)  $(\mathbf{M}' \ominus b) \in \mathcal{M}$ . Therefore the five constraints are satisfied.

**Case (2):**  $(\mathbf{M}' \boxminus \mathbf{M})_e > 0$

Using the definition of the exchange class, we see that there exists  $b = (\mathbf{c}^+, \mathbf{c}^-) \in \mathcal{B}$  such that (a)  $c_e^- = (\mathbf{M}' \boxminus \mathbf{M})_e$ , (b)  $\mathbf{c}^- \subseteq (\mathbf{M}' \boxminus \mathbf{M})$ , (c)  $\mathbf{c}^+ \subseteq (\mathbf{M} \boxminus \mathbf{M}')$ , (d)  $(\mathbf{M} \oplus b) \in \mathcal{M}$ , and (e)  $(\mathbf{M}' \ominus b) \in \mathcal{M}$ . We construct  $b = (\mathbf{b}^+, \mathbf{b}^-)$  by setting  $\mathbf{b}^+ = \mathbf{c}^-$  and  $\mathbf{b}^- = \mathbf{c}^+$ . Notice that, by the construction of  $b$ , we have  $\mathbf{M} \oplus b = \mathbf{M} \ominus c$  and  $\mathbf{M}' \ominus b = \mathbf{M}' \oplus c$ . Therefore, it is clear that  $b$  satisfies the five constraints of the lemma.

Next, let us think when  $\mathbf{M}' = \mathbf{M}^*$  for both cases. Let us consider

$$\boldsymbol{\pi}^1 = \arg \min_{\boldsymbol{\pi} \in \mathcal{M} \setminus \boldsymbol{\pi}^*} \frac{\langle \boldsymbol{\mu}, \boldsymbol{\pi}^* - \boldsymbol{\pi} \rangle}{|\pi_e^* - \pi_e|}. \quad (40)$$

Note that, by the definition of the  $G$ -Gap, we have  $\frac{\langle \boldsymbol{\mu}, \boldsymbol{\pi}^* - \boldsymbol{\pi}^1 \rangle}{|\pi_e^* - \pi_e^1|} = \Delta_e$ . We define  $\boldsymbol{\pi}^0$  such that  $\mathbf{M}(\boldsymbol{\pi}^0) = \mathbf{M}(\boldsymbol{\pi}^*) \ominus b$ . Note that we already have  $\mathbf{M}(\boldsymbol{\pi}^0) \in \mathcal{M}$ . We can see that

$$\frac{\langle \boldsymbol{\mu}, \boldsymbol{\chi}(b) \rangle}{|\chi_e(b)|} = \frac{\langle \boldsymbol{\mu}, \boldsymbol{\pi}^* - \boldsymbol{\pi}^0 \rangle}{|\pi_e^* - \pi_e^0|} \geq \Delta_e, \quad (41)$$

where the inequality follows from the definition of  $G$ -Gap.  $\square$

Next, we establish the confidence bounds used for the analysis of the CSA algorithm.

**Lemma 3.** *Given a phase  $t \in [d]$ , we define random events  $\tau_t$  as follows:*

$$\tau(t) = \left\{ \forall s \in [d] \setminus F(t), |\hat{\mu}_s(t) - \mu_s| < \frac{\Delta_{(d-t+1)}}{(2+L^2)U_{\mathcal{A}}} \right\}, \quad (42)$$

where  $L = \max_{e \in [d], \boldsymbol{\pi}^1, \boldsymbol{\pi}^2, \boldsymbol{\pi}^3 \in \mathcal{A}} \frac{|\pi_e^1 - \pi_e^2|}{|\pi_e^1 - \pi_e^3|}$ . Then, we have

$$\Pr \left[ \bigcap_{t=1}^d \tau(t) \right] \geq 1 - d^2 \exp \left( - \frac{T-d}{2(2+L^2)^2 R^2 \log(d) U_{\mathcal{M}}^2 \mathbf{H}_2} \right) \quad (43)$$

*Proof.* Fix some  $t \in [d]$  and active arm  $s \in [d] \setminus F(t)$  of phase  $t$ . Note that arm  $s$  has been pulled for  $\tilde{T}(t)$  times during the first  $t$  phases. Therefore, by Hoeffding's inequality, we have

$$\Pr \left[ |\hat{\mu}_s(t) - \mu_s| \geq \frac{\Delta_{(d-t+1)}}{(2+L^2)U_{\mathcal{A}}} \right] \leq 2 \exp \left( - \frac{\tilde{T}(t) \Delta_{(d-t+1)}}{2(2+L^2)^2 R^2 U_{\mathcal{A}}^2} \right). \quad (44)$$

By plugging the definition of  $\tilde{T}(t)$ , the quantity  $\tilde{T}(t)\Delta_{(d-t+1)}^2$  on the right hand side of (44) can be further bounded by

$$\begin{aligned}\tilde{T}(t)\Delta_{(d-t+1)}^2 &\geq \frac{T-d}{\tilde{\log}(d)(d-t+1)}\Delta_{(d-t+1)}^2 \\ &\geq \frac{T-d}{\tilde{\log}(d)\mathbf{H}_2},\end{aligned}$$

where the last inequality follows from the definition of  $\mathbf{H}_2 = \max_{s \in [d]} \frac{s}{\Delta_{(s)}^2}$ . By plugging the last equality into (44), we have

$$\Pr \left[ |\hat{\mu}_s(t) - \mu_s| \geq \frac{\Delta_{(d-t+1)}}{(2+L^2)U_{\mathcal{A}}} \right] \leq 2 \exp \left( -\frac{T-d}{2(2+L^2)^2 R^2 \tilde{\log}(d) U_{\mathcal{A}}^2 \mathbf{H}_2} \right). \quad (45)$$

Using (45) and a union bound for all  $t \in [d]$  and all  $s \in [d] \setminus F(t)$ , we have

$$\begin{aligned}\Pr \left[ \bigcap_{t=1}^d \tau(t) \right] &\geq 1 - 2 \sum_{t=1}^d (d-t+1) \exp \left( -\frac{T-d}{2(2+L^2)^2 R^2 \tilde{\log}(d) U_{\mathcal{A}}^2 \mathbf{H}_2} \right) \\ &\geq 1 - d^2 \exp \left( -\frac{T-d}{2(2+L^2)^2 R^2 \tilde{\log}(d) U_{\mathcal{M}}^2 \mathbf{H}_2} \right)\end{aligned}$$

□

The following lemma builds the confidence bound of inner products.

**Lemma 4.** *Fix a phase  $t \in [d]$ . Suppose that random event  $\tau(t)$  occurs. For any vector  $\mathbf{a} \in \mathbb{R}^d$ , we have*

$$|\langle \hat{\boldsymbol{\mu}}(t), \mathbf{a} \rangle - \langle \boldsymbol{\mu}, \mathbf{a} \rangle| < \frac{\Delta_{(d-t+1)}}{(2+L^2)U_{\mathcal{A}}} \|\mathbf{a}\|_1. \quad (46)$$

*Proof.* Suppose that  $\tau_t$  occurs. We have

$$\begin{aligned}|\langle \hat{\boldsymbol{\mu}}(t), \mathbf{a} \rangle - \langle \boldsymbol{\mu}, \mathbf{a} \rangle| &= |\langle \hat{\boldsymbol{\mu}}(t) - \boldsymbol{\mu}, \mathbf{a} \rangle| \\ &= \left| \sum_{s=1}^d (\hat{\mu}_s(t) - \mu_s) a_s \right| \\ &\leq \sum_{s=1}^d |\hat{\mu}_s(t) - \mu_s| |a_s| \\ &< \frac{\Delta_{(d-t+1)}}{(2+L^2)U_{\mathcal{A}}} \sum_{i=1}^d |a_i| \\ &= \frac{\Delta_{(d-t+1)}}{(2+L^2)U_{\mathcal{A}}} \|\mathbf{a}\|_1,\end{aligned} \quad (47)$$

where (47) follows from the definition of  $\tau_t$  in (42). □

### C.3 Main Lemmas

Let  $\mathbf{R}(t) = \{(e, x) \mid \forall e \in [d], \forall x \in \text{POSSIBLE-PI}(\mathbf{e}), (e, x) \notin \mathbf{S}(t)\}$ . We begin with a technical lemma that characterizes several useful lemma properties of  $\mathbf{S}(t)$  and  $\mathbf{R}(t)$ .

**Lemma 5.** *Fix a phase  $t \in [d]$ . Suppose that  $\mathbf{S}(t) \subseteq \mathbf{M}^*$  and  $\mathbf{R}(t) \cap \mathbf{M}^* = \emptyset$ . Let  $\mathbf{M}$  be a set such that  $\mathbf{S}(t) \subseteq \mathbf{M}$  and  $\mathbf{R}(t) \cap \mathbf{M} = \emptyset$ . Let  $\mathbf{a}$  and  $\mathbf{b}$  be two sets satisfying that  $\mathbf{a} \subseteq \mathbf{M} \setminus \mathbf{M}^*$ ,  $\mathbf{b} \subseteq \mathbf{M}^* \setminus \mathbf{M}$  and  $\mathbf{a} \cap \mathbf{b} = \emptyset$ . Then, we have*

$$\mathbf{S}(t) \subseteq (\mathbf{M} \setminus \mathbf{a} \cup \mathbf{b}) \quad \text{and} \quad \mathbf{R}(t) \cap (\mathbf{M} \setminus \mathbf{a} \cup \mathbf{b}) = \emptyset$$

*Proof.* We first prove the first part as follows:

$$\begin{aligned} \mathbf{S}(t) \cap (\mathbf{M} \setminus \mathbf{a} \cup \mathbf{b}) &= (\mathbf{S}(t) \cap (\mathbf{M} \setminus \mathbf{a})) \cup (\mathbf{S}(t) \cap \mathbf{b}) \\ &= \mathbf{S}(t) \cap (\mathbf{M} \setminus \mathbf{a}) \end{aligned} \quad (48)$$

$$\begin{aligned} &= (\mathbf{S}(t) \cap \mathbf{M}) \setminus \mathbf{a} \\ &= \mathbf{S}(t) \setminus \mathbf{a} \end{aligned} \quad (49)$$

$$= \mathbf{S}(t), \quad (50)$$

where (48) holds since we have  $\mathbf{S}(t) \cap \mathbf{b} \subseteq \mathbf{S}(t) \cap (\mathbf{M}^* \setminus \mathbf{M}) \subseteq \mathbf{M} \cap (\mathbf{M}^* \setminus \mathbf{M}) = \emptyset$ ; (49) follows from  $\mathbf{S}(t) \subseteq \mathbf{M}$ ; and (50) follows from  $\mathbf{a} \subseteq \mathbf{M} \setminus \mathbf{M}^*$  and  $\mathbf{S}(t) \subseteq \mathbf{M}^*$  which implies that  $\mathbf{a} \cap \mathbf{S}(t) = \emptyset$ . Notice that  $\mathbf{S}(t) \subseteq (\mathbf{M} \setminus \mathbf{a} \cup \mathbf{b})$ .

Then, we proceed to prove the second part in the following

$$\begin{aligned} \mathbf{R}(t) \cap (\mathbf{M} \setminus \mathbf{a} \cup \mathbf{b}) &= (\mathbf{R}(t) \cap (\mathbf{M} \setminus \mathbf{a})) \cup (\mathbf{R}(t) \cap \mathbf{b}) \\ &= \mathbf{R}(t) \cap (\mathbf{M} \setminus \mathbf{a}) \end{aligned} \quad (51)$$

$$\begin{aligned} &= (\mathbf{R}(t) \cap \mathbf{M}) \setminus \mathbf{a} \\ &= \emptyset \setminus \mathbf{a} = \emptyset, \end{aligned} \quad (52)$$

where (51) follows from the fact that  $\mathbf{R}(t) \cap \mathbf{b} \subseteq \mathbf{R}(t) \cap (\mathbf{M}^* \setminus \mathbf{M}) \subseteq \mathbf{R}(t) \cap \mathbf{M}^* = \emptyset$ ; and (52) follows from the fact that  $\mathbf{R}(t) \cap \mathbf{M} = \emptyset$ . □

Let  $\hat{\mathbf{M}}(t) = \mathbf{M}(\hat{\boldsymbol{\pi}}(t))$ . The next lemma provides an important insight into the correctness of the CSA algorithm. Informally speaking, suppose that the algorithm does not make an error before phase  $t$ . Then, we show that, if arm  $e$  has a gap  $\Delta_e$  larger than the ‘‘reference gap’’  $\Delta_{(d-t+1)}$  of phase  $t$ , then arm  $e$  must be correctly clarified by  $\hat{\mathbf{M}}(t)$ , i.e.,  $M_e(t) = M_e^*$ .

**Lemma 6.** *Fix any phase  $t > 0$ . Suppose that event  $\tau_t$  occurs. Also, assume that  $\mathbf{S}(t) \subseteq \mathbf{M}^*$  and  $\mathbf{R}(t) \cap \mathbf{S}(t) = \emptyset$ . Let  $e \in [d] \setminus \mathbf{F}(t)$  be an active arm. Suppose that  $\Delta_{(d-t+1)} \leq \Delta_e$ . Then, we have  $(e, \pi_e^*) \in \mathbf{M}^* \cap \mathbf{S}(t)$ .*

*Proof.* Suppose that  $(e, \pi_e^*) \notin (\mathbf{M}^* \cap \hat{\mathbf{M}}(t))$ . This is equivalent to the following

$$(e, \pi_e^*) \in (\mathbf{M}^* \cap \neg \hat{\mathbf{M}}(t)) \cup (\neg \mathbf{M}^* \cap \hat{\mathbf{M}}(t)) \quad (53)$$

(53) can be further rewritten as

$$(e, \pi_e^*) \in (\mathbf{M}^* \setminus \hat{\mathbf{M}}(t)) \cup (\hat{\mathbf{M}}(t) \setminus \mathbf{M}^*). \quad (54)$$

From this assumption, it is easy to see that  $\hat{\mathbf{M}}(t) \neq \mathbf{M}^*$ . Therefore, we can apply Lemma 2. We know that there exists  $b = (\mathbf{b}^+, \mathbf{b}^-) \in \mathcal{B}$  such that (a)  $b_e^- = (\hat{\mathbf{M}}(t) \boxplus \mathbf{M}^*)_e$  or  $b_e^+ = (\mathbf{M}^* \boxplus \hat{\mathbf{M}}(t))_e$ , (b)  $\mathbf{b}^- \subseteq (\hat{\mathbf{M}}(t) \boxplus \mathbf{M}^*)$ , (c)  $\mathbf{b}^+ \subseteq \mathbf{M}^* \boxplus \hat{\mathbf{M}}(t)$ , (d)  $(\hat{\mathbf{M}}(t) \oplus b) \in \mathcal{M}$ , (e)  $(\mathbf{M}^* \oplus b) \in \mathcal{M}$ , and  $\frac{\langle \hat{\boldsymbol{\mu}}, \boldsymbol{\chi}(b) \rangle}{\chi_e(b)} \geq \Delta_e > 0$ .

Using Lemma 5, we see that  $(\hat{\mathbf{M}}(t) \oplus b) \cap \mathbf{R}(t) = \emptyset$ ,  $\mathbf{S}(t) \subseteq (\hat{\mathbf{M}}(t) \oplus b)$ , and  $(\mathbf{b}^+ \cup \mathbf{b}^-) \cap (\mathbf{S}(t) \cup \mathbf{R}(t)) = \emptyset$ .

Recall the definition  $\hat{\mathbf{M}}(t) \in \arg \max_{\mathbf{M} \in \mathcal{M}, \mathbf{S}(t) \subseteq \mathbf{M}, \mathbf{R}(t) \cap \mathbf{M} = \emptyset} \langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\mathbf{M}) \rangle$  and also recall that  $\hat{\mathbf{M}}(t) \oplus b \in \mathcal{M}$ . Therefore, we obtain that

$$\frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\hat{\mathbf{M}}(t)) \rangle}{\chi_e(b)} \geq \frac{\langle \hat{\boldsymbol{\mu}}, \boldsymbol{\pi}(\hat{\mathbf{M}}(t) \oplus b) \rangle}{\chi_e(b)}. \quad (55)$$

On the other hand, we have

$$\frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\hat{\mathbf{M}}(t) \oplus b) \rangle}{\chi_e(b)} = \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\hat{\mathbf{M}}(t)) + \boldsymbol{\chi}(b) \rangle}{\chi_e(b)} \quad (56)$$

$$= \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\hat{\mathbf{M}}(t)) \rangle}{\chi_e(b)} + \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\chi}(b) \rangle}{\chi_e(b)} \\ > \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\hat{\mathbf{M}}(t)) \rangle}{\chi_e(b)} + \frac{\langle \boldsymbol{\mu}, \boldsymbol{\chi}(b) \rangle}{\chi_e(b)} - \frac{\Delta_{(d-t+1)}}{(2+L^2)U_{\mathcal{A}}} \frac{\|\boldsymbol{\chi}(b)\|}{\chi_e(b)} \quad (57)$$

$$\geq \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\hat{\mathbf{M}}(t)) \rangle}{\chi_e(b)} + \frac{\langle \boldsymbol{\mu}, \boldsymbol{\chi}(b) \rangle}{\chi_e(b)} - \frac{\Delta_e}{(2+L^2)U_{\mathcal{A}}} \frac{\|\boldsymbol{\chi}(b)\|}{\chi_e(b)} \\ \geq \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\hat{\mathbf{M}}(t)) \rangle}{\chi_e(b)} + \frac{\langle \boldsymbol{\mu}, \boldsymbol{\chi}(b) \rangle}{\chi_e(b)} - \frac{\Delta_e}{(2+L^2)} \quad (58)$$

$$\geq \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\hat{\mathbf{M}}(t)) \rangle}{\chi_e(b)} + \frac{1+L^2}{2+L^2} \Delta_e \quad (59)$$

$$\geq \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\hat{\mathbf{M}}(t)) \rangle}{\chi_e(b)}. \quad (60)$$

This means that  $\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\hat{\mathbf{M}}(t) \oplus b) \rangle > \langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\hat{\mathbf{M}}(t)) \rangle$ . This contradicts the definition of  $\hat{\boldsymbol{\pi}}(t)$ , and therefore, we have  $(e, \pi_e^*) \in (\mathbf{M}^* \cap \hat{\mathbf{M}}(t))$ .  $\square$

The next lemma takes a step further. Hereinafter, we denote  $\tilde{\mathbf{M}}^e(t)$  as  $\mathbf{M}(\tilde{\boldsymbol{\pi}}^e(t))$ .

**Lemma 7.** Fix any phase  $t > 0$ . Suppose that event  $\tau_t$  occurs. Also, assume that  $\mathbf{S}(t) \subseteq \mathbf{M}^*$  and  $\mathbf{R}(t) \cap \mathbf{M}^* = \emptyset$ . Let  $e \in [d] \setminus \mathbf{F}(t)$  satisfy  $\Delta_{(d-t+1)} \leq \Delta_e$ . Then, we have

$$\frac{\langle \hat{\boldsymbol{\mu}}(t), \hat{\boldsymbol{\pi}}(t) - \tilde{\boldsymbol{\pi}}^e(t) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} > \frac{L+1/L}{2+L^2} \Delta_{(d-t+1)}. \quad (61)$$

*Proof.* By Lemma 6, we see that

$$(e, \pi_e^*) \in (\mathbf{M}^* \cap \mathbf{S}(t)). \quad (62)$$

From the definition of  $\tilde{\mathbf{M}}^e(t)$ , which ensures that  $\tilde{\mathbf{M}}^e(t) \neq \mathbf{M}_e^*$ , we have  $(e, \pi_e^*) \in (\mathbf{M}^* \setminus \tilde{\mathbf{M}}^e(t))$ .

Hence, we apply Lemma 2. There exists  $b = (\mathbf{b}^+, \mathbf{b}^-) \in \mathcal{B}$  such that (a)  $b_e^- = \left( \tilde{\mathbf{M}}^e(t) \boxplus \mathbf{M}^* \right)_e$  or  $b_e^+ = \left( \mathbf{M}^* \boxplus \tilde{\mathbf{M}}^e(t) \right)_e$ , (b)  $\mathbf{b}^- \subseteq (\tilde{\mathbf{M}}^e(t) \boxplus \mathbf{M}^*)$ , (c)  $\mathbf{b}^+ \subseteq \mathbf{M}^* \boxplus \tilde{\mathbf{M}}^e(t)$ , (d)  $(\tilde{\mathbf{M}}^e(t) \oplus b) \in \mathcal{M}$ , (e)  $\mathbf{M}^* \oplus b \in \mathcal{M}$ , and  $\frac{\langle \boldsymbol{\mu}, \boldsymbol{\chi}(b) \rangle}{\chi_e(b)} \geq \Delta_e$ .

Define  $\overline{\mathbf{M}}^e(t) \triangleq \tilde{\mathbf{M}}^e(t) \oplus b$ . Using Lemma 5, we have  $\mathbf{S}(t) \subseteq \overline{\mathbf{M}}^e(t)$  and  $\mathbf{R}(t) \cap \overline{\mathbf{M}}^e(t) = \emptyset$ . Since  $\overline{\mathbf{M}}^e(t) \in \mathcal{M}$  and by definition  $\hat{\mathbf{M}}(t) = \arg \max_{\mathbf{M} \in \mathcal{M}, \mathbf{S}_t \subseteq \mathbf{M}, \mathbf{R}_t \cap \mathbf{M} = \emptyset} \langle \hat{\boldsymbol{\mu}}, \boldsymbol{\pi}(\mathbf{M}) \rangle$ , we have

$$\frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\hat{\mathbf{M}}(t)) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \geq \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\overline{\mathbf{M}}^e(t)) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|}. \quad (63)$$

Hence, we have

$$\begin{aligned}
 \frac{\langle \hat{\boldsymbol{\mu}}(t), \hat{\boldsymbol{\pi}}(t) - \tilde{\boldsymbol{\pi}}^e(t) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} &\geq \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\overline{\mathbf{M}}^e(t)) - \tilde{\boldsymbol{\pi}}^e(t) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \\
 &= \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\tilde{\mathbf{M}}^e(t) \oplus b) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} - \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\tilde{\mathbf{M}}^e(t)) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \\
 &= \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\tilde{\mathbf{M}}^e(t) + \boldsymbol{\chi}(b)) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} - \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(\tilde{\mathbf{M}}^e(t)) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \\
 &= \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\chi}(b) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|}
 \end{aligned} \tag{64}$$

$$> \frac{\chi_e(b)}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \left( \frac{\langle \boldsymbol{\mu}, \boldsymbol{\chi}(b) \rangle}{\chi_e(b)} - \frac{\Delta_{(d-t+1)}}{(2+L^2)U_{\mathcal{A}}} \frac{\|\boldsymbol{\chi}(b)\|_1}{\chi_e(b)} \right) \tag{65}$$

$$\geq \frac{\chi_e(b)}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \left( \frac{\langle \boldsymbol{\mu}, \boldsymbol{\chi}(b) \rangle}{\chi_e(b)} - \frac{\Delta_e}{(2+L^2)U_{\mathcal{A}}} \frac{\|\boldsymbol{\chi}(b)\|_1}{\chi_e(b)} \right) \tag{66}$$

$$\geq \frac{\chi_e(b)}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \cdot \frac{1+L^2}{2+L^2} \Delta_e \tag{67}$$

$$\begin{aligned}
 &\geq \frac{L+1/L}{2+L^2} \Delta_e \\
 &\geq \frac{L+1/L}{2+L^2} \Delta_{(d-t+1)},
 \end{aligned} \tag{68}$$

where (65) follows from Lemma 4, the assumption on event  $\tau_t$ ; (66) follows from the assumption that  $\Delta_e \geq \Delta_{(d-t+1)}$ ; (67) holds since  $b \in \mathcal{B}$  and therefore  $\frac{\|\boldsymbol{\chi}(b)\|_1}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \leq U_{\mathcal{A}}$ ; (68) follows from the fact that  $\frac{\langle \boldsymbol{\mu}, \boldsymbol{\chi}(b) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \geq \Delta_e$ .  $\square$

The next lemma shows that, during phase  $t$ , if  $\Delta_e \leq \Delta_{(d-t+1)}$  for some  $e$ , then the empirical gap between  $\boldsymbol{\pi}(t)$  and  $\tilde{\boldsymbol{\pi}}^e(t)$  is smaller than  $\frac{1}{3}\Delta_{(d-t+1)}$ .

**Lemma 8.** Fix any phase  $t > 0$ . Suppose that event  $\tau_t$  occurs. Also, assume that  $\mathbf{S}(t) \subseteq \mathbf{M}^*$  and  $\mathbf{R}(t) \subseteq \mathbf{M}^* = \emptyset$ . Suppose an active arm  $e \in [d] \setminus \mathbf{F}(t)$  satisfies that  $M_e^* \neq M_e(t)$ . Then, we have

$$\frac{\langle \hat{\boldsymbol{\mu}}(t), \hat{\boldsymbol{\pi}}(t) - \tilde{\boldsymbol{\pi}}^e(t) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \leq \frac{L}{2+L^2} \Delta_{(d-t+1)}. \tag{69}$$

*Proof.* Fix any exchange class  $\mathcal{B} = \arg \min_{\mathcal{B}' \in \text{Exchange}(\mathcal{M})} \text{width}(\mathcal{B}')$ .

From the assumption that  $M_e^* \neq \hat{M}_e(t)$ , we can apply Lemma 2, and have (a)  $b_e^- = (\hat{\mathbf{M}}(t) \boxminus \mathbf{M}^*)_e$  or  $b_e^+ = (\mathbf{M}^* \boxminus \hat{\mathbf{M}}(t))_e$ , (b)  $\mathbf{b}^- \subseteq \hat{\mathbf{M}}(t) \boxminus \mathbf{M}^*$ , (c)  $\mathbf{b}^+ \subseteq \mathbf{M}^* \boxminus \hat{\mathbf{M}}(t)$ , (d)  $\hat{\mathbf{M}}(t) \oplus b \in \mathcal{M}$  (e)  $\mathbf{M}^* \oplus b \in \mathcal{M}$ , and  $\frac{\langle \boldsymbol{\mu}, \boldsymbol{\chi}(b) \rangle}{\chi_e(b)} \geq \Delta_e > 0$ .

Define  $\overline{\mathbf{M}}^e(t) \triangleq \hat{\mathbf{M}}(t) \oplus b$ , and let  $\overline{\boldsymbol{\pi}}^e(t) = \boldsymbol{\pi}(\overline{\mathbf{M}}^e(t))$ . We claim that

$$\langle \hat{\boldsymbol{\mu}}(t), \tilde{\boldsymbol{\pi}}^e(t) \rangle \geq \langle \hat{\boldsymbol{\mu}}, \overline{\boldsymbol{\pi}}^e(t) \rangle. \tag{70}$$

From the definition of  $\tilde{\mathbf{M}}^e(t)$  in Algorithm 1, we only need to show that (a):  $\hat{\pi}_e(t) \neq \overline{\pi}_e(t)$  and (b):  $\mathbf{S}(t) \subseteq \overline{\mathbf{M}}^e(t)$  and  $\mathbf{R}(t) \cap \overline{\mathbf{M}}^e(t) = \emptyset$ . Since, either  $\mathbf{b}^+$  or  $\mathbf{b}^-$  has an  $e$ -changer, the  $e$ -th element of  $\overline{\boldsymbol{\pi}}(t)$  is different from that of  $\hat{\boldsymbol{\pi}}(t)$ . Next, we notice that this follows directly from Lemma 5 by setting  $M = \hat{\mathbf{M}}(t)$ . Hence, we have shown that (70) holds.

Therefore, we have

$$\begin{aligned}
 \frac{\langle \hat{\boldsymbol{\mu}}(t), \hat{\boldsymbol{\pi}}(t) - \tilde{\boldsymbol{\pi}}^e(t) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} &\leq \frac{\langle \hat{\boldsymbol{\mu}}(t), \hat{\boldsymbol{\pi}}(t) - \bar{\boldsymbol{\pi}}^e(t) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \\
 &\leq \frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(t) - (\boldsymbol{\pi}(t) + \boldsymbol{\chi}(b)) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \\
 &= -\frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\chi}(b) \rangle}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \\
 &\leq \frac{\chi_e(b)}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \cdot \left( -\frac{\langle \boldsymbol{\mu}, \boldsymbol{\chi}(b) \rangle}{\chi_e(b)} + \frac{\Delta_{(d-t+1)}}{(2+L^2)U_{\mathcal{M}}} \frac{\|\boldsymbol{\chi}(b)\|_1}{\chi_e(b)} \right) \tag{71}
 \end{aligned}$$

$$\begin{aligned}
 &\leq \frac{\chi_e(b)}{|\hat{\pi}_e(t) - \tilde{\pi}_e^e(t)|} \cdot \frac{\Delta_{(d-t+1)}}{(2+L^2)} \tag{72} \\
 &\leq \frac{L}{2+L^2} \Delta_{(d-t+1)}
 \end{aligned}$$

□

#### C.4 Proof of Theorem 2

For the reader's convenience, we first restate Theorem 2 as follows.

**Theorem 2.** *Given any  $T > d$ , action class  $\mathcal{A} \subset \mathbb{R}^d$ , and  $\boldsymbol{\mu} \in \mathbb{R}^d$ , the CSA algorithm uses at most  $T$  samples and outputs a solution  $\boldsymbol{\pi}^{\text{out}} \in \mathcal{A} \cup \{\perp\}$  such that*

$$\begin{aligned}
 &\Pr [\boldsymbol{\pi}^{\text{out}} \neq \boldsymbol{\pi}^*] \\
 &\leq d^2 \exp \left( -\frac{T-d}{2(2+L^2)^2 R^2 \tilde{\log}(d) U_{\mathcal{A}}^2 \mathbf{H}_2} \right), \tag{8}
 \end{aligned}$$

where  $L = \max_{e \in [d], \pi^1, \pi^2, \pi^3 \in \mathcal{A}, \pi_e^1 \neq \pi_e^3} \frac{|\pi_e^1 - \pi_e^2|}{|\pi_e^1 - \pi_e^3|}$ ,  $\tilde{\log}(d) \triangleq \sum_{s=1}^d \frac{1}{s}$ , and  $U_{\mathcal{A}} = \max_{\boldsymbol{\pi}, \boldsymbol{\pi}' \in \mathcal{A}, e \in [d] \mid \pi_s \neq \pi'_s} \frac{\sum_{s=1}^d |\pi_s - \pi'_s|}{|\pi_e - \pi'_e|}$ .

*Proof.* First, we show that the algorithm takes at most  $T$  samples. Note that exactly one arm is pulled for  $\tilde{T}_1$  times, one arm is pulled  $\tilde{T}_2$  times, ..., and one arm is pulled  $\tilde{T}_d$  times. Therefore, the total number of samples used by the algorithm is bounded by

$$\begin{aligned}
 \sum_{t=1}^d \tilde{T}_t &\leq \sum_{t=1}^d \left( \frac{T-d}{\tilde{\log}(d)(d-t+1)} + 1 \right) \\
 &= \frac{T-d}{\tilde{\log}(d)} \tilde{\log}(d) + d = T.
 \end{aligned}$$

By Lemma 3, we know that the event  $\tau = \bigcap_{t=1}^T \tau_t$  occurs with probability at least  $1 - d^2 \exp \left( -\frac{T-d}{R^2 \tilde{\log}(d) U_{\mathcal{A}}^2 \mathbf{H}_2} \right)$ . Therefore, we only need to prove that, under event  $\tau$  the algorithm outputs  $\boldsymbol{M}^*$ . We will assume that event  $\tau$  occurs in the rest of the proof.

We prove the theorem by induction. Fix a phase  $t \in [d]$ . Suppose that the algorithm does not make any error before phase  $t$ , i.e.,  $\boldsymbol{S}(t) \subseteq \boldsymbol{M}^*$  and  $\boldsymbol{R}(t) \cap \boldsymbol{M}^* \neq \emptyset$ . We show that the algorithm does not err at phase  $t$ .

At the beginning of phase  $t$ , there are exactly  $t-1$  inactive arms, i.e.,  $|\boldsymbol{F}(t)| = t-1$ . Therefore, there must exist an active arm  $e(t) \in [d] \setminus \boldsymbol{F}(t)$ , such that  $\Delta_{e(t)} = \Delta_{(d-t+1)}$ . Hence, by Lemma 7, we have

$$\frac{\langle \hat{\boldsymbol{\mu}}(t), \boldsymbol{\pi}(t) - \tilde{\boldsymbol{\pi}}^{e(t)}(t) \rangle}{|\hat{\pi}_{e(t)}(t) - \tilde{\pi}_{e(t)}^{e(t)}(t)|} \geq \frac{L+1/L}{2+L^2} \Delta_{(d-t+1)}, \tag{73}$$

where  $b$  is an exchange set of  $\hat{\boldsymbol{M}}(t)$  and  $\tilde{\boldsymbol{M}}^e(t)$ .

Notice that the algorithm makes an error in phase  $t$  if and only if  $(p(t), \hat{\pi}_{p(t)}(t)) \in (\mathbf{M}^* \cap \neg \hat{\mathbf{M}}(t)) \cup (\neg \mathbf{M}^* \cap \hat{\mathbf{M}}(t))$ .

Suppose that  $(p(t), \hat{\pi}_{p(t)}(t)) \in (\mathbf{M}^* \cap \neg \hat{\mathbf{M}}(t)) \cup (\neg \mathbf{M}^* \cap \hat{\mathbf{M}}(t))$ . From Lemma 8, we have

$$\frac{\langle \hat{\boldsymbol{\mu}}(t), \hat{\boldsymbol{\pi}}(t) - \tilde{\boldsymbol{\pi}}^{p(t)}(t) \rangle}{\left| \hat{\pi}_{p(t)}(t) - \tilde{\pi}_{p(t)}^{p(t)}(t) \right|} \leq \frac{L}{2 + L^2} \Delta_{(d-t+1)}. \quad (74)$$

By combining (73) and (74), we see that

$$\frac{\langle \hat{\boldsymbol{\mu}}(t), \hat{\boldsymbol{\pi}}(t) - \tilde{\boldsymbol{\pi}}^{p(t)}(t) \rangle}{\left| \hat{\pi}_{p(t)}(t) - \tilde{\pi}_{p(t)}^{p(t)}(t) \right|} \leq \frac{L}{2 + L^2} \Delta_{(d-t+1)} < \frac{L + 1/L}{2 + L^2} \Delta_{(d-t+1)} \leq \frac{\langle \hat{\boldsymbol{\mu}}(t), \hat{\boldsymbol{\pi}}(t) - \tilde{\boldsymbol{\pi}}^{e(t)}(t) \rangle}{\left| \hat{\pi}_{e(t)}(t) - \tilde{\pi}_{e(t)}^{e(t)}(t) \right|}. \quad (75)$$

However, (75) is contradictory to the definition of

$$p(t) = \arg \max_{e \in [d] \setminus F(t)} \frac{\langle \hat{\boldsymbol{\mu}}(t), \hat{\boldsymbol{\pi}}(t) - \tilde{\boldsymbol{\pi}}(t) \rangle}{\left| \hat{\pi}_e(t) - \tilde{\pi}_e^e(t) \right|}. \quad (76)$$

Therefore, we have proven that  $(p(t), \hat{\pi}_{p(t)}(t)) \notin (\mathbf{M}^* \cap \neg \hat{\mathbf{M}}(t)) \cup (\neg \mathbf{M}^* \cap \hat{\mathbf{M}}(t))$ . This means that the algorithm does not err at phase  $t$ , or equivalently  $\mathbf{S}(t+1) \subseteq \mathbf{M}^*$  and  $\mathbf{R}(t+1) \cap \mathbf{M}^* = \emptyset$ . By induction, we have proven that the algorithm does not err at any phase  $t \in [d]$ .

Hence, we have  $\mathbf{S}(d+1) \subseteq \mathbf{M}^*$  and  $\mathbf{R}(d+1) \subseteq \neg \mathbf{M}^*$  in the final phase. This means that  $\mathbf{S}(d+1) = \mathbf{M}^*$ , and therefore,  $\boldsymbol{\pi}^{\text{out}} = \boldsymbol{\pi}^*$  after phase  $d$ .  $\square$

## D Proof of Theorem 3

We first introduce some useful lemmas to prove Theorem 3. Lemma 9 shows that Algorithm 2 pulls arms no more than  $T$  times. Recall that  $B = 2^{\lceil \log_2 d \rceil} - 1$  and  $T' = T - \lfloor \frac{T}{d} \beta \rfloor \times d$ .

**Lemma 9.** *Algorithm 2 terminates in phase  $\lceil \log_2 d \rceil$  with no more than a total of  $T$  arm pulls.*

*Proof.* The total number of arm pulls  $T_{\text{total}}$  is bounded as follows.

$$\begin{aligned} T_{\text{total}} &= \left\lfloor \frac{T}{d} \beta \right\rfloor \times d + \sum_{r=1}^{\lceil \log_2 d \rceil} \sum_{s \in [d]} \lceil p_s(r) \cdot m(r) \rceil \\ &\leq \left\lfloor \frac{T}{d} \beta \right\rfloor \times d + \sum_{r=1}^{\lceil \log_2 d \rceil} \left( d + \frac{T' - d \lceil \log_2 d \rceil}{B/2^{r-1}} \right) \\ &\leq \left\lfloor \frac{T}{d} \beta \right\rfloor \times d + d \lceil \log_2 d \rceil + \frac{2^{\lceil \log_2 d \rceil} - 1}{B} (T' - d \lceil \log_2 d \rceil) \\ &= T. \end{aligned}$$

$\square$

Let us write  $\Delta'_{(i)} = \langle \boldsymbol{\mu}, \boldsymbol{\pi}^1 - \boldsymbol{\pi}^i \rangle$ . Note that  $\boldsymbol{\mu}^\top \boldsymbol{\pi}^i \geq \boldsymbol{\mu}^\top \boldsymbol{\pi}^{i+1}$  for all  $i \in [K-1]$ . Lemma 10 bounds the probability that a certain action has its estimate of the expected reward larger than that of the best action at the end of phase  $r$ .

**Lemma 10.** *For a fixed realization of  $\mathbb{A}(r)$  satisfying  $\boldsymbol{\pi}^* \in \mathbb{A}(r)$ , for any action  $\boldsymbol{\pi}^i \in \mathbb{A}(r)$ ,*

$$\Pr \left[ \hat{\boldsymbol{\mu}}^\top(r+1) \boldsymbol{\pi}^1 < \hat{\boldsymbol{\mu}}^\top(r+1) \boldsymbol{\pi}^i \right] \leq \exp \left( - \frac{2 \Delta'_{(i)}{}^2}{\sum_{s=1}^d \frac{(\pi_s^1 - \pi_s^i)^2}{T_s(r+1)} R^2} \right). \quad (77)$$

*Proof.* For any  $i \in \{2, \dots, K\}$ , we have

$$\begin{aligned} \Pr \left[ \hat{\boldsymbol{\mu}}^\top(r+1)\boldsymbol{\pi}^1 < \hat{\boldsymbol{\mu}}^\top(r+1)\boldsymbol{\pi}^i \right] &\leq \Pr \left[ \langle \hat{\boldsymbol{\mu}}(r+1) - \boldsymbol{\mu}, \boldsymbol{\pi}^1 - \boldsymbol{\pi}^i \rangle < -\Delta'_{(i)} \right] \\ &\leq \exp \left( -\frac{2\Delta'_{(i)}{}^2}{\sum_{s=1}^d \frac{(\pi_s^1 - \pi_s^i)^2}{T_s(r+1)} R^2} \right), \end{aligned} \quad (78)$$

where the last inequality follows from Hoeffding's inequality (Hoeffding, 1963).

If we use allocation vector (15), (78) can be upper bounded by

$$\begin{aligned} (78) &\leq \exp \left( -\frac{2\Delta'_{(i)}{}^2}{\sum_{e=1}^d \frac{(\pi_e^1 - \pi_e^i)^2}{\lfloor p_s(r)m(r) \rfloor} R^2} \right) \\ &\leq \exp \left( -\frac{2\Delta'_{(i)}{}^2}{R^2 V'^2} \cdot \frac{T' - \lceil \log_2 d \rceil}{B/2^{r-1}} \right) \\ &\leq \exp \left( -\frac{\Delta'_{(i)}{}^2}{R^2 V'^2} \cdot \frac{T' - \lceil \log_2 d \rceil}{d/2^{r-1}} \right) \end{aligned} \quad (79)$$

where  $V' = \sum_{e=1}^d |\pi_e^1 - \pi_e^i|$ . □

Next, we bound the error probability of a single phase  $r$  in Lemma 11.

**Lemma 11.** *Assume that the best action is not eliminated prior to phase  $r$ , i.e.,  $\boldsymbol{\pi}^1 \notin \mathbb{A}(r)$ . Then, the probability that the best action is eliminated in phase  $r+1$  is bounded as*

$$\Pr \left[ \boldsymbol{\pi}^1 \notin \mathcal{A}_r \mid \boldsymbol{\pi}^1 \in \mathcal{A}_{r-1} \right] \leq \begin{cases} \frac{4K}{d} \exp \left( \frac{T' - \lceil \log_2 d \rceil}{R^2 V'^2} \cdot \frac{\Delta_{(i_r)}^2}{i_r} \right) & (\text{when } r = 1) \\ 3 \exp \left( \frac{T' - \lceil \log_2 d \rceil}{R^2 V'^2} \cdot \frac{\Delta_{(i_r)}^2}{i_r} \right) & (\text{when } r > 1), \end{cases} \quad (80)$$

where  $i_r = \lceil \frac{d}{2^{r+1}} \rceil + 1$  and

$$V = \max_{\boldsymbol{\pi} \in \mathcal{A} \setminus \{\boldsymbol{\pi}^*\}, s \in \{e \in [d] \mid \pi_e^* \neq \pi_e\}} \frac{\sum_{u=1}^d |\pi_u^* - \pi_u|}{|\pi_s^* - \pi_s|}. \quad (81)$$

*Proof.* Define  $\mathbb{B}(r+1)$  as the set of actions in  $\mathbb{A}(r)$  excluding the best action and  $\lceil \frac{d}{2^{r+1}} \rceil - 1$  suboptimal actions with the largest expected rewards. We have  $|\mathbb{B}(r+1)| = |\mathbb{A}(r)| - \lceil \frac{d}{2^{r+1}} \rceil$ .

Let us think of bijective function  $p : [d] \rightarrow [d]$ , which satisfies the following:

$$\forall s \in [d], \Delta_{(p(s))} = \Delta_s. \quad (82)$$

Also, we denote the inverse mapping of  $f$  by  $f^{-1}$ . Next, for any  $i \in [d]$ , we define  $s(i)$  as follows:

$$s(i) = \arg \max_{s \in [d]} |\pi_s^1 - \pi_s^i|.$$

Then, from the definition of  $\{\Delta_s\}_{s=1, \dots, d}$ , we have

$$\frac{\Delta'_{(i)}}{|\pi_{s(i)}^1 - \pi_{s(i)}^i|} \geq \Delta_{s(i)} = \Delta_{(p(s(i)))}. \quad (83)$$

Also, we have

$$\min_{i \in \mathbb{B}(r+1)} \Delta_{(p(s(i)))} \geq \Delta_{(\lceil \frac{d}{2^{r+1}} \rceil + 1)}. \quad (84)$$



If the best action is eliminated in phase  $r$ , then at least  $\lceil \frac{d}{2^r} \rceil - \lceil \frac{d}{2^{r+1}} \rceil + 1$  actions of  $\mathbb{B}(r+1)$  have their estimates of the expected rewards larger than that of the best action.

Let  $N(r)$  denote the number of actions in  $\mathbb{B}(r+1)$  whose estimates of the expected rewards are larger than that of the best action. By Lemma 10, we have

$$\begin{aligned} \mathbb{E}[N_r] &= \sum_{i \in \mathcal{B}_r} \Pr \left[ \hat{\boldsymbol{\mu}}^\top(r+1)\boldsymbol{\pi}^1 < \hat{\boldsymbol{\mu}}^\top(r+1)\boldsymbol{\pi}^i \right] \\ &\leq \sum_{i \in \mathcal{B}_r} \exp \left( -\frac{\Delta'_{(i)}{}^2}{R^2 V'^2} \cdot \frac{T' - \lceil \log_2 d \rceil}{d/2^{r-1}} \right) \\ &\leq |\mathcal{B}_r| \max_{i \in \mathbb{B}(r+1)} \exp \left( -\frac{\Delta'_{(i)}{}^2}{R^2 V'^2} \cdot \frac{T' - \lceil \log_2 d \rceil}{d/2^{r-1}} \right) \\ &= |\mathcal{B}_r| \max_{i \in \mathbb{B}(r+1)} \exp \left( -\frac{|\pi_{s(i)}^1 - \pi_{s(i)}^i|^2 \Delta_{(p(s(i)))}^2}{R^2 V'^2} \cdot \frac{T' - \lceil \log_2 d \rceil}{d/2^{r-1}} \right) \end{aligned} \quad (85)$$

$$\leq \left( |\mathcal{A}_{r-1}| - \left\lceil \frac{d}{2^{r+1}} \right\rceil \right) \exp \left( \frac{T' - \lceil \log_2 d \rceil}{R^2 V^2} \cdot \frac{\Delta(\lceil \frac{d}{2^{r+1}} \rceil + 1)^2}{\left\lceil \frac{d}{2^{r+1}} \right\rceil + 1} \right) \quad (86)$$

where (85) follows from (83), (86) follows from (84), and

$$V = \max_{\boldsymbol{\pi} \in \mathcal{A} \setminus \{\boldsymbol{\pi}^*\}, s \in \{e \in [d] \mid \pi_e^* \neq \pi_e\}} \frac{\sum_{u=1}^d |\pi_u^* - \pi_u|}{|\pi_s^* - \pi_s|}. \quad (87)$$

Then, together with Markov's inequality, we obtain

$$\begin{aligned} \Pr \left[ \boldsymbol{\pi}^1 \notin \mathcal{A}_r \right] &\leq \Pr \left[ N_r \geq \left\lceil \frac{d}{2^r} \right\rceil - \left\lceil \frac{d}{2^{r+1}} \right\rceil + 1 \right] \\ &\leq \frac{\mathbb{E}[N_r]}{\left\lceil \frac{d}{2^r} \right\rceil - \left\lceil \frac{d}{2^{r+1}} \right\rceil + 1} \\ &\leq \frac{|\mathcal{A}_{r-1}| - \left\lceil \frac{d}{2^{r+1}} \right\rceil}{\left\lceil \frac{d}{2^r} \right\rceil - \left\lceil \frac{d}{2^{r+1}} \right\rceil + 1} \exp \left( \frac{T' - \lceil \log_2 d \rceil}{R^2 V^2} \cdot \frac{\Delta(\lceil \frac{d}{2^{r+1}} \rceil + 1)}{\left\lceil \frac{d}{2^{r+1}} \right\rceil + 1} \right) \end{aligned} \quad (88)$$

When  $r = 1$ , we have  $|\mathcal{A}_{r-1}| = K$ . Thus,

$$\begin{aligned} \frac{|\mathcal{A}_{r-1}| - \left\lceil \frac{d}{2^{r+1}} \right\rceil}{\left\lceil \frac{d}{2^r} \right\rceil - \left\lceil \frac{d}{2^{r+1}} \right\rceil + 1} &= \frac{K - \left\lceil \frac{d}{2^{r+1}} \right\rceil}{\left\lceil \frac{d}{2^r} \right\rceil - \left\lceil \frac{d}{2^{r+1}} \right\rceil + 1} \\ &\leq \frac{K}{\frac{d}{2} - \frac{d}{2^2}} \\ &\leq \frac{4K}{d}. \end{aligned} \quad (89)$$

When  $r > 1$ , we have  $|\mathcal{A}_{r-1}| = \left\lceil \frac{d}{2^{r-1}} \right\rceil$ . Thus,

$$\begin{aligned} \frac{|\mathcal{A}_{r-1}| - \left\lceil \frac{d}{2^{r+1}} \right\rceil}{\left\lceil \frac{d}{2^r} \right\rceil - \left\lceil \frac{d}{2^{r+1}} \right\rceil + 1} &= \frac{\left\lceil \frac{d}{2^{r-1}} \right\rceil - \left\lceil \frac{d}{2^{r+1}} \right\rceil}{\left\lceil \frac{d}{2^r} \right\rceil - \left\lceil \frac{d}{2^{r+1}} \right\rceil + 1} \\ &\leq \frac{\frac{d}{2^{r-1}} + 1 - \left\lceil \frac{d}{2^{r+1}} \right\rceil}{\frac{d}{2^r} - \left\lceil \frac{d}{2^{r+1}} \right\rceil + 1} \\ &\leq \frac{3 \cdot \frac{d}{2^{r+1}} + \frac{d}{2^{r+1}} - \left\lceil \frac{d}{2^{r+1}} \right\rceil + 1}{\frac{d}{2^{r+1}} + \frac{d}{2^{r+1}} - \left\lceil \frac{d}{2^{r+1}} \right\rceil + 1} \\ &\leq 3, \end{aligned} \quad (90)$$

where the last inequality results from the fact that for any  $x, y > 0$ ,  $\frac{3x+y}{x+y} \leq 3$ .

Therefore, for this specific realization of  $\mathcal{A}_{r-1}$  satisfying  $1 \in \mathcal{A}_{r-1}$ ,

$$\Pr[\pi^1 \notin \mathcal{A}_r | \pi^1 \in \mathcal{A}_{r-1}] \leq \begin{cases} \frac{4K}{d} \exp\left(\frac{T' - \lceil \log_2 d \rceil}{R^2 V^2} \cdot \frac{\Delta_{i_r}^2}{i_r}\right) & (\text{when } r = 1) \\ 3 \exp\left(\frac{T' - \lceil \log_2 d \rceil}{R^2 V^2} \cdot \frac{\Delta_{i_r}^2}{i_r}\right) & (\text{when } r > 1), \end{cases} \quad (91)$$

where  $i_r = \lceil \frac{d}{2^{r+1}} + 1 \rceil$ . □

Finally, we prove Theorem 3.

*Proof of Theorem 3.* We have

$$\begin{aligned} \Pr[\pi^{\text{out}} \neq \pi^1] &= \Pr[\pi^1 \notin \mathcal{A}_{\lceil \log_2 d \rceil}] \\ &\leq \sum_{r=1}^{\lceil \log_2 d \rceil} \Pr[\pi^1 \notin \mathcal{A}_r | \pi^1 \in \mathcal{A}_{r-1}] \\ &\leq \frac{4K}{d} \exp\left(\frac{T' - \lceil \log_2 d \rceil}{R^2 V^2} \cdot \frac{\Delta_{i_1}^2}{i_1}\right) + \sum_{s=1}^d 3 \exp\left(\frac{T' - \lceil \log_2 d \rceil}{R^2 V^2} \cdot \frac{\Delta_{i_r}^2}{i_r}\right) \\ &\leq \left(\frac{4K}{d} + 3(\lceil \log_2 d \rceil - 1)\right) \exp\left(\frac{T' - \lceil \log_2 d \rceil}{R^2 V^2} \cdot \frac{1}{\max_{1 \leq s \leq d} \frac{s}{\Delta_{(s)}^2}}\right) \\ &< \left(\frac{4K}{d} + 3 \log_2 d\right) \exp\left(\frac{T' - \lceil \log_2 d \rceil}{R^2 V^2} \cdot \frac{1}{\mathbf{H}_2}\right), \end{aligned} \quad (92)$$

where

$$\mathbf{H}_2 = \max_{1 \leq s \leq d} \frac{s}{\Delta_{(s)}^2}.$$

□