
Thompson Sampling for (Combinatorial) Pure Exploration

Siwei Wang¹ Jun Zhu¹

Abstract

Existing methods of combinatorial pure exploration mainly focus on the UCB approach. To make the algorithm efficient, they usually use the sum of upper confidence bounds within arm set S to represent the upper confidence bound of S , which can be much larger than the tight upper confidence bound of S and leads to a much higher complexity than necessary, since the empirical means of different arms in S are independent. To deal with this challenge, we explore the idea of Thompson Sampling (TS) that uses independent random samples instead of the upper confidence bounds, and design the first TS-based algorithm TS-Explore for (combinatorial) pure exploration. In TS-Explore, the sum of independent random samples within arm set S will not exceed the tight upper confidence bound of S with high probability. Hence it solves the above challenge, and achieves a lower complexity upper bound than existing efficient UCB-based algorithms in general combinatorial pure exploration. As for pure exploration of classic multi-armed bandit, we show that TS-Explore achieves an asymptotically optimal complexity upper bound.

1. Introduction

Pure exploration is an important task in online learning, and it tries to find out the target arm as fast as possible. In pure exploration of classic multi-armed bandit (MAB) (Audibert et al., 2010), there are totally m arms, and each arm i is associated with a probability distribution D_i with mean μ_i . Once arm i is pulled, it returns an observation r_i , which is drawn independently from D_i by the environment. At each time step t , the learning policy π either chooses an arm $i(t)$ to pull, or chooses to output an arm $a(t)$. The goal of the learning policy is to pull arms properly, such that

with an error probability at most δ , its output arm $a(t)$ is the optimal arm (i.e., $a(t) = \operatorname{argmax}_{i \in [m]} \mu_i$) with complexity (the total number of observations) as low as possible.

Pure exploration is widely adopted in real applications. For example, in the selling procedure of cosmetics, there is always a testing phase before the commercialization phase (Audibert et al., 2010). The goal of the testing phase is to help to maximize the cumulative reward collected in the commercialization phase. Therefore, instead of regret minimization (Berry & Fristedt, 1985; Auer et al., 2002), the testing phase only needs to do exploration (e.g., to investigate which product is the most popular one), and wants to find out the target with both correctness guarantee and low cost. In the real world, sometimes the system focuses on the best action under a specific combinatorial structure, instead of the best single arm (Chen et al., 2014). For example, a network routing system needs to search for the path with minimum delay between the source and the destination. Since there can be an exponential number of paths, the cost of exploring them separately is unacceptable. Therefore, people choose to find out the best path by exploring single edges, and this is a pure exploration problem instance in combinatorial multi-armed bandit (CMAB). In this setting, we still pull single arms (base arms) at each time step, and there is a super arm set $\mathcal{I} \subseteq 2^{[m]}$. The expected reward of a super arm $S \in \mathcal{I}$ is $\sum_{i \in S} \mu_i$, i.e., the sum of the expected rewards of its contained base arms. And the goal of the player is to find out the optimal super arm with error probability at most δ and complexity as low as possible.

Most of the existing solutions for pure exploration follow the UCB approach (Audibert et al., 2010; Kalyanakrishnan et al., 2012; Chen et al., 2014; Kaufmann & Kalyanakrishnan, 2013). They compute the confidence bounds for all the arms, and claim that one arm is optimal only if its lower confidence bound is larger than the upper confidence bounds of all the other arms. Though this approach is asymptotically optimal in pure exploration of classic MAB problems (Kalyanakrishnan et al., 2012; Kaufmann & Kalyanakrishnan, 2013), it faces some challenges in the CMAB setting (Chen et al., 2014). In the UCB approach, for a super arm S , the algorithm usually uses the sum of upper confidence bounds of all its contained base arms as its upper confidence bound to achieve a low implementation cost. This means that the gap between the empirical mean of super

¹Dept. of Comp. Sci. & Tech., BNRist Center, Tsinghua-Bosch Joint ML Center, Tsinghua University. Correspondence to: Jun Zhu <dcszj@tsinghua.edu.cn>.

arm S and the used upper confidence bound of S is about $\tilde{O}(\sum_{i \in S} \sqrt{1/N_i})$ (here N_i is the number of observations on base arm i). However, since the observations of different base arms are independent, the standard deviation of the empirical mean of super arm S is $\tilde{O}(\sqrt{\sum_{i \in S} 1/N_i})$, which is much smaller than $\tilde{O}(\sum_{i \in S} \sqrt{1/N_i})$. This means that the used upper confidence bound of S is much larger than the tight upper confidence bound of S . Combes et al. (2015) deal with this problem by computing the upper confidence bounds for all the super arms independently, which leads to an exponential time cost and is hard to implement in practice. In fact, existing efficient UCB-based solutions either suffer from a high complexity bound (Chen et al., 2014; Gabillon et al., 2016), or need further assumptions on the combinatorial structure to achieve an optimal complexity bound efficiently (Chen et al., 2017).

Then a natural solution to deal with this challenge is to use random samples of arm i (with mean to be its empirical mean and standard deviation to be $\tilde{O}(\sqrt{1/N_i})$) instead of its upper confidence bound to judge whether a super arm is optimal. If we let the random samples of different base arms be *independent*, then with high probability, the gap between the empirical mean of super arm S and the sum of random samples within S is $\tilde{O}(\sqrt{\sum_{i \in S} 1/N_i})$, which has the same order as the gap between the empirical mean of super arm S and its real mean. Therefore, using independent random samples can behave better than using confidence bounds, and this is the key idea of Thompson Sampling (TS) (Thompson, 1933; Kaufmann et al., 2012; Agrawal & Goyal, 2013). In fact, many prior works show that TS-based algorithms have smaller cumulative regret than UCB-based algorithms in regret minimization of CMAB model (Wang & Chen, 2018; Perrault et al., 2020). However, there still lack studies on adapting the idea of TS, i.e., using random samples to make decisions, to the pure exploration setting.

In this paper, we attempt to fill up this gap, and study using random samples in pure exploration under the frequentist setting for both MAB and CMAB instances. We emphasize that it is non-trivial to design (and analyze) such a TS-based algorithm. The first challenge is that there is a lot more uncertainty in the random samples than in the confidence bounds. In UCB-based algorithms, the upper confidence bound of the optimal arm is larger than its expected reward with high probability. Thus, the arm with the largest upper confidence bound is either an insufficiently learned sub-optimal arm (i.e., the number of its observations is not enough to make sure that it is a sub-optimal arm) or the optimal arm, which means that the number of pulls on sufficiently learned sub-optimal arms is limited. However, for TS-based algorithms that use random samples instead, there is always a constant probability that the random sample of the optimal arm is smaller than its expected reward. In this case, the arm with the largest random sample may be a suffi-

ciently learned sub-optimal arm. Therefore, the mechanism of the TS-based policy should be designed carefully to make sure that we can still obtain an upper bound for the number of pulls on sufficiently learned sub-optimal arms.

Another challenge is that using random samples to make decisions is a kind of Bayesian algorithm and it loses many good properties in the frequentist setting. In the Bayesian setting, at each time step t , the parameters of the game follow a posterior distribution \mathcal{P}_t , and the random samples are drawn from \mathcal{P}_t independently as well. Therefore, using the random samples to output the optimal arm can explicitly ensure the correctness of the TS-based algorithm. However, in the frequentist setting, the parameters of the game are fixed but unknown, and they have no such correlations with the random samples. Because of this, if we still choose to use the random samples to output the optimal arm, then the distributions of random samples in the TS-based algorithm need to be chosen carefully to make sure that we can still obtain the correctness guarantee in the frequentist setting.

Besides, the analysis of the TS-based algorithm in pure exploration is also very different from that in regret minimization. In regret minimization, at each time step, we only need to draw *one* sample for each arm (Thompson, 1933; Agrawal & Goyal, 2013; Wang & Chen, 2018; Perrault et al., 2020). However, in pure exploration, one set of samples at each time step is not enough, since the algorithm needs to i) check whether there is an arm that is the optimal one with high probability; ii) look for an arm that needs exploration. None of these two goals can be achieved by only *one* set of samples. Therefore, we must draw several sets of samples to make decisions, and this may fail the existing analysis of TS in regret minimization.

In this paper, we solve the above challenges, and design a TS-based algorithm TS-Explore for (combinatorial) pure exploration under the frequentist setting with polynomial implementation cost. At each time step t , TS-Explore first draws independent random samples $\theta_i^k(t)$ for all the (base) arms $i \in [m]$ and $k = 1, 2, \dots, M$ (i.e., totally M independent samples for each arm). Then it tries to find out the M best (super) arms \hat{S}_t^k 's under sample sets $\theta^k(t) = [\theta_1^k(t), \theta_2^k(t), \dots, \theta_m^k(t)]$ for $k = 1, 2, \dots, M$. If in all these sample sets, the best (super) arm is the same as the empirically best (super) arm \hat{S}_t (i.e., the best arm under the empirical means), then the algorithm will output that this (super) arm is optimal. Otherwise, for all $k \in [M]$, the algorithm will check the reward gap between \hat{S}_t^k and \hat{S}_t under parameter set $\theta^k(t)$. Then it focuses on the (super) arm $\tilde{S}_t^{k^*}$ with the largest reward gap, i.e., it chooses to pull a (base) arm in the exchange set of \hat{S}_t and $\tilde{S}_t^{k^*}$.

Recording such reward gaps and focusing on $\tilde{S}_t^{k^*}$ is the key mechanism we used to solve the above challenges. On the

one hand, our analysis shows that with a proper value M (the number of sample sets) and proper random sample distributions, $\tilde{S}_t^{k^*}$ has some similar properties with the (super) arm with the largest upper confidence bound. These properties play a critical role in the analysis of UCB approaches. Thus, we can also use them to obtain an upper bound for the number of pulls on sufficiently learned sub-optimal arms in TS-Explore as well as the correctness guarantee of TS-Explore. On the other hand, this novel mechanism saves the key advantage of Thompson Sampling, i.e., the sum of random samples within S will not exceed the tight upper confidence bound of S with high probability. Hence, TS-Explore can achieve a lower complexity upper bound than existing efficient UCB-based solutions. These results indicate that our TS-Explore policy is correct, efficient (with low implementation cost), and effective (with low complexity).

In the general CMAB setting, we show that TS-Explore is near optimal and achieves a lower complexity upper bound than existing efficient UCB-based algorithms (Chen et al., 2014). The optimal algorithm in (Chen et al., 2017) is only efficient when the combinatorial structure satisfies some specific properties (otherwise it suffers from an exponential implementation cost), and is less general than our results. We also conduct experiments to compare the complexity of TS-Explore with existing baselines. The experimental results show that TS-Explore outperforms the efficient baseline CLUCB (Chen et al., 2014), and behaves only a little worse than the optimal but non-efficient baseline NaiveGapElim (Chen et al., 2017). As for the MAB setting, we show that TS-Explore is asymptotically optimal, i.e., it has a comparable complexity to existing optimal algorithms (Kalyanakrishnan et al., 2012; Kaufmann & Kalyanakrishnan, 2013) when $\delta \rightarrow 0$. All these results indicate that our TS-based algorithm is efficient and effective in dealing with pure exploration problems. To the best of our knowledge, this is the first result of using this kind of TS-based algorithm (i.e., always making decisions based on random samples) in pure exploration under the frequentist setting.

2. Related Works

Pure exploration of the classic MAB model is first proposed by Audibert et al. (2010). After that, people have designed lots of learning policies for this problem. The two most representative algorithms are successive-elimination (Even-Dar et al., 2006; Audibert et al., 2010; Kaufmann & Kalyanakrishnan, 2013) and LUCB (Kalyanakrishnan et al., 2012; Kaufmann & Kalyanakrishnan, 2013). Both of them adopt the idea of UCB (Auer et al., 2002) and achieve an asymptotically optimal complexity upper bound (i.e., it matches with the complexity lower bound proposed by Kalyanakrishnan et al. (2012)). Compared to these results, our TS-Explore policy uses a totally different approach, and can achieve an

asymptotically optimal complexity upper bound as well.

Combinatorial pure exploration is first studied by Chen et al. (2014). They propose CLUCB, an LUCB-based algorithm that is efficient as long as there is an offline oracle to output the best super arm under any given parameter set. Chen et al. (2017) then design an asymptotically optimal algorithm for this problem. However, their algorithm can only be implemented efficiently when the combinatorial structure follows some specific constraints. Recently, based on the game approach, Jourdan et al. (2021) provide another optimal learning policy for pure exploration in CMAB. But their algorithm still suffers from an exponential implementation cost. Compared with these UCB-based algorithms, our TS-Explore policy achieves a lower complexity bound than CLUCB (Chen et al., 2014) (with a similar polynomial time cost), and has a much lower implementation cost than the optimal policies (Chen et al., 2017; Jourdan et al., 2021) in the most general combinatorial pure exploration setting.

There also exist some researches about applying Thompson Sampling to pure exploration. For example, Russo (2016) considers a frequentist setting of pure exploration in classic MAB, and proposes algorithms called TTPS, TTVS, and TTTS; Shang et al. (2020) extend the results in (Russo, 2016), design a T3C algorithm, and provide its analysis for Gaussian bandits; Li et al. (2021) study Bayesian contextual pure exploration, propose an algorithm called BRE and obtain its corresponding analysis. However, these results are still very different from ours. The first point is that they mainly use random distributions but not random samples to decide the next chosen arm or when to stop, which may cause a high implementation cost when we extend them to the combinatorial setting, since it is much more complex to deal with random distributions than to deal with random samples. Moreover, our results are still more general even if we only consider pure exploration in classic MAB under the frequentist setting. For example, Russo (2016) does not provide a correct stopping rule, and Shang et al. (2020) only obtain the correctness guarantee for Gaussian bandits. Besides, their complexity bounds are asymptotic ones (which require $\delta \rightarrow 0$), while ours works for any $\delta \in (0, 1)$.

3. Model Setting

3.1. Pure Exploration in Multi-armed Bandit

A pure exploration problem instance of MAB is a tuple $([m], \mathbf{D}, \delta)$. Here $[m] = \{1, 2, \dots, m\}$ is the set of arms, $\mathbf{D} = \{D_1, D_2, \dots, D_m\}$ are the corresponding reward distributions of the arms, and δ is the error constraint. In this paper, we assume that all the distributions D_i 's are supported on $[0, 1]$. Let $\mu_i \triangleq \mathbb{E}_{X \sim D_i}[X]$ denote the expected reward of arm i , and $a^* = \operatorname{argmax}_{i \in [m]} \mu_i$ is the optimal arm with the largest expected reward. Similar to many ex-

isting works (e.g., (Audibert et al., 2010)), we assume that the optimal arm is unique. At each time step t , the learning policy π can either pull an arm $i(t) \in [m]$, or output an arm $a(t) \in [m]$. If it chooses to pull arm $i(t)$, then it will receive an observation $r_{i(t)}(t)$, which is drawn independently from $D_{i(t)}$. The goal of the learning policy is to make sure that with probability at least $1 - \delta$, its output $a(t) = a^*$. Under this constraint, it aims to minimize the complexity

$$Z^\pi \triangleq \sum_{i=1}^m N_i(T^\pi),$$

where T^π denotes the time step t that policy π chooses to output $a(t)$, and $N_i(t)$ denotes the number of observations on arm i until time step t .

Let $\Delta_{i,m} \triangleq \mu_{a^*} - \mu_i$ denote the expected reward gap between the optimal arm a^* and any other arm $i \neq a^*$. For the optimal arm a^* , its $\Delta_{a^*,m}$ is defined as $\mu_{a^*} - \max_{i \neq a^*} \mu_i$. We also define $H_m \triangleq \sum_{i \in [m]} \frac{1}{\Delta_{i,m}^2}$, and existing works (Kalyanakrishnan et al., 2012) show that the complexity lower bound of any pure exploration algorithm is $\Omega(H_m \log \frac{1}{\delta})$.

3.2. Pure Exploration in Combinatorial Multi-armed Bandit

A pure exploration problem instance of CMAB is an extension of the MAB setting. The arms $i \in [m]$ are called base arms, and there is also a super arm set $\mathcal{I} \subseteq 2^{[m]}$. For each super arm $S \in \mathcal{I}$, its expected reward is $\sum_{i \in S} \mu_i$. Let $S^* = \operatorname{argmax}_{S \in \mathcal{I}} \sum_{i \in S} \mu_i$ denote the optimal super arm with the largest expected reward, and we assume that the optimal super arm is unique (as in (Chen et al., 2014)). At each time step t , the learning policy π can either pull a base arm $i(t) \in [m]$, or output a super arm $S(t) \in \mathcal{I}$. The goal of the learning policy is to make sure that with probability at least $1 - \delta$, its output $S(t) = S^*$. Under this constraint, it also wants to minimize its complexity Z^π .

As in many existing works (Chen et al., 2013; 2014; Wang & Chen, 2018), we also assume that there exists an offline Oracle, which takes a set of parameters $\theta = [\theta_1, \dots, \theta_m]$ as input, and outputs the best super arm under this parameter set, i.e., $\text{Oracle}(\theta) = \operatorname{argmax}_{S \in \mathcal{I}} \sum_{i \in S} \theta_i$.

In this paper, for $i \notin S^*$, we use $\Delta_{i,c} \triangleq \sum_{j \in S^*} \mu_j - \max_{S \in \mathcal{I}: i \in S} \sum_{j \in S} \mu_j$ to denote the expected reward gap between the optimal super arm S^* and the best super arm that contains i . As for $i \in S^*$, its $\Delta_{i,c}$ is defined as $\Delta_{i,c} \triangleq \sum_{j \in S^*} \mu_j - \max_{S \in \mathcal{I}: i \notin S} \sum_{j \in S} \mu_j$, i.e., the expected reward gap between S^* and the best super arm that does not contain i . We also define $S \oplus S' = (S \setminus S') \cup (S' \setminus S)$ and $\text{width} \triangleq \max_{S \neq S'} |S \oplus S'|$, and let $H_{1,c} \triangleq \text{width} \sum_{i \in [m]} \frac{1}{\Delta_{i,c}^2}$, $H_{2,c} \triangleq \text{width}^2 \sum_{i \in [m]} \frac{1}{\Delta_{i,c}^2}$.

Chen et al. (2017) prove that the complexity lower bound for combinatorial pure exploration is $\Omega(H_{0,c} \log \frac{1}{\delta})$, where $H_{0,c}$ is the optimal value of the following convex program (here $\Delta_{S^*,S} = \sum_{i \in S^*} \mu_i - \sum_{i \in S} \mu_i$):

$$\begin{aligned} \min \quad & \sum_{i \in [m]} N_i \\ \text{s.t.} \quad & \sum_{i \in S^* \oplus S} \frac{1}{N_i} \leq \Delta_{S^*,S}^2, \quad \forall S \in \mathcal{I}, S \neq S^* \end{aligned}$$

The following result shows the relationships between $H_{0,c}$, $H_{1,c}$ and $H_{2,c}$.

Proposition 3.1. *For any combinatorial pure exploration instance, $H_{0,c} \leq H_{1,c} \leq H_{2,c}$.*

Proof. Since $\text{width} \geq 1$, we have that $H_{1,c} \leq H_{2,c}$. As for the first inequality, note that $\forall i \in [m]$, $N_i = \frac{\text{width}}{\Delta_{i,c}^2}$ is a feasible solution of the above convex program. Hence we have that $H_{0,c} \leq H_{1,c}$. \square

4. Thompson Sampling-based Pure Exploration Algorithm

Note that pure exploration of classic multi-armed bandit is a special case of combinatorial multi-armed bandit (i.e., $\mathcal{I} = \{\{1\}, \{2\}, \dots, \{m\}\}$). Therefore, in this section, we mainly focus on the TS-based pure exploration algorithm in the combinatorial multi-armed bandit setting. Its framework and analysis can directly lead to results in the classic multi-armed bandit setting.

In the following, we let $\Phi(x, \mu, \sigma^2) \triangleq \Pr_{X \sim \mathcal{N}(\mu, \sigma^2)}[X \geq x]$. For any $x \in (0, 0.5)$, we also define $\phi(x)$ as a function of x such that $\Phi(\phi(x), 0, 1) = x$.

4.1. Algorithm Framework

Our Thompson Sampling-based algorithm (TS-Explore) is described in Algorithm 1. We use N_i to denote the number of observations on arm i , R_i to denote the sum of all the observations from arm i , and $N_i(t), R_i(t)$ to denote the value of N_i, R_i at the beginning of time step t .

At each time step t , for any $i \in [m]$, TS-Explore first draws $M(\delta, q, t) \triangleq \frac{1}{q} \log(12|\mathcal{I}|^2 t^2 / \delta)$ random samples $\{\theta_i^k(t)\}_{k=1}^{M(\delta, q, t)}$ independently from a Gaussian distribution with mean $\hat{\mu}_i(t) \triangleq \frac{R_i(t)}{N_i(t)}$ (i.e., the empirical mean of arm i) and variance $\frac{C(\delta, q, t)}{N_i(t)}$ (i.e., inversely proportional to $N_i(t)$), where $C(\delta, q, t) \triangleq \frac{\log(12|\mathcal{I}|^2 t^2 / \delta)}{\phi^2(q)}$. Then it checks which super arm is optimal in the empirical mean parameter set $\hat{\mu}(t) = [\hat{\mu}_1(t), \hat{\mu}_2(t), \dots, \hat{\mu}_m(t)]$ and the k -th sample set $\theta^k(t) = [\theta_1^k(t), \theta_2^k(t), \dots, \theta_m^k(t)]$ for all k by using the offline Oracle, i.e., $\hat{S}_t = \text{Oracle}(\hat{\mu}(t))$ and

Algorithm 1 TS-Explore

```

1: Input: Error constraint  $\delta, q \in [\delta, 0.1], t \leftarrow m, N_i \leftarrow 0, R_i \leftarrow 0$  for all  $i \in [m]$ .
2: Pull each arm once, update their number of pulls  $N_i$ 's, the sum of their observations  $R_i$ 's.
3: while true do
4:    $t \leftarrow t + 1$ .
5:   For all base arm  $i \in [m], \hat{\mu}_i(t) \leftarrow \frac{R_i}{N_i}$ .
6:    $\hat{\boldsymbol{\mu}}(t) \leftarrow [\hat{\mu}_1(t), \hat{\mu}_2(t), \dots, \hat{\mu}_m(t)]$ .
7:    $\hat{S}_t \leftarrow \text{Oracle}(\hat{\boldsymbol{\mu}}(t))$ .
8:   for  $k = 1, 2, \dots, M(\delta, q, t)$  do
9:     For each arm  $i$ , draw  $\theta_i^k(t)$  independently from distribution  $\mathcal{N}(\hat{\mu}_i(t), \frac{C(\delta, q, t)}{N_i})$ .
10:     $\boldsymbol{\theta}^k(t) \leftarrow [\theta_1^k(t), \theta_2^k(t), \dots, \theta_m^k(t)]$ .
11:     $\tilde{S}_t^k \leftarrow \text{Oracle}(\boldsymbol{\theta}^k(t))$ .
12:     $\tilde{\Delta}_t^k \leftarrow \sum_{i \in \tilde{S}_t^k} \theta_i^k(t) - \sum_{i \in \hat{S}_t} \theta_i^k(t)$ .
13:  end for
14:  if  $\forall 1 \leq k \leq M(\delta, q, t), \tilde{S}_t^k = \hat{S}_t$  then
15:    Return:  $\hat{S}_t$ .
16:  else
17:     $k_t^* \leftarrow \operatorname{argmax}_k \tilde{\Delta}_t^k, \tilde{S}_t \leftarrow \tilde{S}_t^{k_t^*}$ .
18:    Pull arm  $i(t) \leftarrow \operatorname{argmin}_{i \in \tilde{S}_t \oplus \tilde{S}_t} N_i$ , update its number of pulls  $N_{i(t)}$  and the sum of its observations  $R_{i(t)}$ .
19:  end if
20: end while

```

$\tilde{S}_t^k = \text{Oracle}(\boldsymbol{\theta}^k(t))$. If all the best super arms \tilde{S}_t^k 's are the same as \hat{S}_t , then TS-Explore outputs that this super arm is the optimal one. Otherwise, for all $1 \leq k \leq M(\delta, q, t)$, it will compute the reward gap between \tilde{S}_t^k and \hat{S}_t under the k -th sample set $\boldsymbol{\theta}^k(t)$, and focus on k_t^* with the largest reward gap, i.e., TS-Explore will choose to pull the base arm with the least number of observations in $\hat{S}_t \oplus \tilde{S}_t^{k_t^*}$ (in the following, we use \tilde{S}_t to represent $\tilde{S}_t^{k_t^*}$ to simplify notations). Note that $\hat{S}_t \neq \tilde{S}_t$ (otherwise we will output \hat{S}_t as the optimal super arm), thus the rule of choosing arms to pull is well defined.

4.2. Analysis of TS-Explore

Theorem 4.1. *In the CMAB setting, for $q \in [\delta, 0.1]$, with probability at least $1 - \delta$, TS-Explore will output the optimal super arm S^* with complexity upper bounded by $O(H_{1,c}(\log \frac{1}{\delta} + \log(|\mathcal{I}|H_{1,c}))) \frac{\log \frac{1}{\delta}}{\log \frac{1}{q}} + H_{1,c} \frac{\log^2(|\mathcal{I}|H_{1,c})}{\log \frac{1}{q}}$. Specifically, if we choose $q = \delta$, then the complexity upper bound is $O(H_{1,c} \log \frac{1}{\delta} + H_{1,c} \log^2(|\mathcal{I}|H_{1,c}))$.*

Remark 4.2. When the error constraint $\delta \in (0.1, 1)$, we can still let the parameters (δ, q) in TS-Explore be $(\delta_0, q_0) = (0.1, 0.1)$. In this case: i) its error probability is upper bounded by $\delta_0 = 0.1 < \delta$; and ii) since $\frac{\delta}{\delta_0} < \frac{1}{0.1} = 10$, the complexity of TS-Explore is still

upper bounded by $O(H_{1,c} \log \frac{1}{\delta_0} + H_{1,c} \log^2(|\mathcal{I}|H_{1,c})) = O(H_{1,c} \log \frac{1}{\delta} + H_{1,c} \log^2(|\mathcal{I}|H_{1,c}))$.

By Theorem 4.1, we can directly obtain the correctness guarantee and the complexity upper bound for applying TS-Explore in pure exploration of classic multi-armed bandit.

Corollary 4.3. *In the MAB setting, for $q \in [\delta, 0.1]$, with probability at least $1 - \delta$, TS-Explore will output the optimal arm a^* with complexity upper bounded by $O(H_m(\log \frac{1}{\delta} + \log(mH_m))) \frac{\log \frac{1}{\delta}}{\log \frac{1}{q}} + H_m \frac{\log^2(mH_m)}{\log \frac{1}{q}}$. Specifically, if we choose $q = \delta$, then the complexity upper bound is $O(H_m \log \frac{1}{\delta} + H_m \log^2(mH_m))$.*

Remark 4.4. The value q in TS-Explore is used to control the number of times that we draw random samples at each time step. Note that $M(\delta, q, t) = \frac{1}{q} \log(12|\mathcal{I}|^2 t^2 / \delta)$. Hence when q becomes larger, we need fewer samples, but the complexity bound becomes worse. Here is a trade-off between the algorithm's complexity and the number of random samples it needs to draw. Our analysis shows that using $q = \delta^{\frac{1}{\beta}}$ for some constant $\beta \geq 1$ can make sure that the complexity upper bound remains the same order, and reduce the number of random samples significantly.

Remark 4.5. If the value $|\mathcal{I}|$ is unknown (which is common in real applications), we can use 2^m instead (in $M(\delta, q, t)$ and $C(\delta, q, t)$). This only increases the constant term in the complexity upper bound and does not influence the major term $O(H_{1,c} \log \frac{1}{\delta})$.

Theorem 4.1 shows that the complexity upper bound of TS-Explore is width lower than the CLUCB policy in (Chen et al., 2014). To the best of our knowledge, TS-Explore is the first algorithm that efficiently achieves an $O(H_{1,c} \log \frac{1}{\delta})$ complexity upper bound in the most general combinatorial pure exploration setting. Besides, this is also the first theoretical analysis of using a TS-based algorithm (i.e., using random samples to make decisions) to deal with combinatorial pure exploration problems under the frequentist setting.

Though the complexity bound of TS-Explore still has some gap with the optimal one $O(H_{0,c} \log \frac{1}{\delta})$, we emphasize that this is because we only use the simple offline oracle and do not seek more detailed information about the combinatorial structure. This makes our policy more efficient. As a comparison, the existing optimal policies (Chen et al., 2017; Jourdan et al., 2021) either suffer from an exponential time cost, or require the combinatorial structure to satisfy some specific constraints so that they can adopt more powerful offline oracles that explore detailed information about the combinatorial structure efficiently (please see Appendix C for discussions about this). Therefore, our algorithm is more efficient in the most general CMAB setting, and can be attractive in real applications with large scale and complex combinatorial structures.

On the other hand, the gap between the complexity upper bound and the complexity lower bound does not exist in the MAB setting. Corollary 4.3 shows that the complexity upper bound of applying TS-Explore in the classic MAB setting matches the complexity lower bound $\Omega(H_m \log \frac{1}{\delta})$ (Kalyanakrishnan et al., 2012) when $\delta \rightarrow 0$, i.e., TS-Explore is asymptotically optimal in the classic MAB setting.

Now we provide the proof of Theorem 4.1.

In this proof, we denote $\mathcal{J} = \{U : \exists S, S' \in \mathcal{I}, U = S \setminus S'\}$. Note that $|\mathcal{J}| \leq |\mathcal{I}|^2$ and for any $U \in \mathcal{J}$, $|U| \leq \text{width}$. We also denote $L_1(t) = \log(12|\mathcal{I}|^2 t^2 / \delta)$ and $L_2(t) = \log(12|\mathcal{I}|^2 t^2 M(\delta, q, t) / \delta)$ to simplify notations.

We first define three events as follows: $\mathcal{E}_{0,c}$ is the event that for all $t > 0$, $U \in \mathcal{J}$,

$$\left| \sum_{i \in U} (\hat{\mu}_i(t) - \mu_i) \right| \leq \sqrt{\sum_{i \in U} \frac{1}{2N_i(t)} L_1(t)};$$

$\mathcal{E}_{1,c}$ is the event that for all $t > 0$, $1 \leq k \leq M(\delta, q, t)$, $U \in \mathcal{J}$,

$$\left| \sum_{i \in U} (\theta_i^k(t) - \hat{\mu}_i(t)) \right| \leq \sqrt{\sum_{i \in U} \frac{2C(\delta, q, t)}{N_i(t)} L_2(t)};$$

and $\mathcal{E}_{2,c}$ is the event that for all $t > 0$, $S, S' \in \mathcal{I}$. there exists $1 \leq k \leq M(\delta, q, t)$ such that

$$\sum_{i \in S} \theta_i^k(t) - \sum_{i \in S'} \theta_i^k(t) \geq \sum_{i \in S} \mu_i - \sum_{i \in S'} \mu_i.$$

Roughly speaking, $\mathcal{E}_{0,c}$ means that for any $U \in \mathcal{J}$, the gap between its real mean and its empirical mean lies in the corresponding confidence radius; $\mathcal{E}_{1,c}$ means that for any $U \in \mathcal{J}$, the gap between its empirical mean and the sum of its random samples lies in the corresponding confidence radius; and $\mathcal{E}_{2,c}$ means that for any two super arms $S \neq S'$, at each time step, there is at least one sample set such that the gap between the sum of their random samples is larger than the gap between their real means.

We will first prove that $\mathcal{E}_{0,c} \wedge \mathcal{E}_{1,c} \wedge \mathcal{E}_{2,c}$ happens with high probability, and then show the correctness guarantee and complexity upper bound under $\mathcal{E}_{0,c} \wedge \mathcal{E}_{1,c} \wedge \mathcal{E}_{2,c}$.

Lemma 4.6. *In Algorithm 1, we have that*

$$\Pr[\mathcal{E}_{0,c} \wedge \mathcal{E}_{1,c} \wedge \mathcal{E}_{2,c}] \geq 1 - \delta.$$

Proof. Note that the random variable $(\hat{\mu}_i(t) - \mu_i)$ is zero-mean and $\frac{1}{4N_i(t)}$ sub-Gaussian, and for different i , the random variables $(\hat{\mu}_i(t) - \mu_i)$'s are independent. Therefore, $\sum_{i \in U} (\hat{\mu}_i(t) - \mu_i)$ is zero-mean and $\sum_{i \in U} \frac{1}{4N_i(t)}$ sub-Gaussian. Then by concentration inequality of sub-Gaussian

random variables (see details in Appendix D),

$$\begin{aligned} & \Pr \left[\left| \sum_{i \in U} (\hat{\mu}_i(t) - \mu_i) \right| > \sqrt{\sum_{i \in U} \frac{1}{2N_i(t)} L_1(t)} \right] \\ & \leq 2 \exp(-L_1(t)) \\ & = \frac{\delta}{6|\mathcal{I}|^2 t^2}. \end{aligned}$$

This implies that

$$\Pr[-\mathcal{E}_{0,c}] \leq \sum_{U,t} \frac{\delta}{6|\mathcal{I}|^2 t^2} \leq \sum_t \frac{\delta}{6t^2} \leq \frac{\delta}{3},$$

where the second inequality is because that $|\mathcal{J}| \leq |\mathcal{I}|^2$, and the third inequality is because that $\sum_t \frac{1}{t^2} \leq 2$.

Similarly, the random variable $(\theta_i^k(t) - \hat{\mu}_i(t))$ is a zero-mean Gaussian random variable with variance $\frac{C(\delta, q, t)}{N_i(t)}$, and for different i , the random variables $(\theta_i^k(t) - \hat{\mu}_i(t))$'s are also independent. Then by concentration inequality,

$$\begin{aligned} & \Pr \left[\left| \sum_{i \in U} (\theta_i^k(t) - \hat{\mu}_i(t)) \right| > \sqrt{\sum_{i \in U} \frac{2C(\delta, q, t)}{N_i(t)} L_2(t)} \right] \\ & \leq 2 \exp(-L_2(t)) \\ & = \frac{\delta}{6|\mathcal{I}|^2 t^2 M(\delta, q, t)}. \end{aligned}$$

This implies that

$$\Pr[-\mathcal{E}_{1,c}] \leq \sum_{U,t,k} \frac{\delta}{6|\mathcal{I}|^2 t^2 M(\delta, q, t)} \leq \sum_{U,t} \frac{\delta}{6|\mathcal{I}|^2 t^2} \leq \frac{\delta}{3},$$

where the second inequality is because that there are totally $M(\delta, q, t)$ sample sets at time step t .

Finally we consider the probability $\Pr[-\mathcal{E}_{2,c} | \mathcal{E}_{0,c}]$. In the following, we denote $\Delta_{S,S'} \triangleq \sum_{i \in S} \mu_i - \sum_{i \in S'} \mu_i = \sum_{i \in S \setminus S'} \mu_i - \sum_{i \in S' \setminus S} \mu_i$ as the reward gap (under the real means) between S and S' .

For any fixed $S \neq S'$, we denote $A(t) = \sum_{i \in S \setminus S'} \frac{1}{2N_i(t)}$, $B(t) = \sum_{i \in S' \setminus S} \frac{1}{2N_i(t)}$ and $C(t) = C(\delta, q, t)$ to simplify notations. Then under event $\mathcal{E}_{0,c}$, we must have that $\sum_{i \in S \setminus S'} \hat{\mu}_i(t) - \sum_{i \in S' \setminus S} \hat{\mu}_i(t) \geq (\sum_{i \in S \setminus S'} \mu_i - \sqrt{A(t)L_1(t)}) - (\sum_{i \in S' \setminus S} \mu_i + \sqrt{B(t)L_1(t)}) \geq \Delta_{S,S'} - \sqrt{A(t)L_1(t)} - \sqrt{B(t)L_1(t)}$.

Since $\sum_{i \in S \setminus S'} \theta_i^k(t) - \sum_{i \in S' \setminus S} \theta_i^k(t)$ is a Gaussian random variable with mean $\sum_{i \in S \setminus S'} \hat{\mu}_i(t) - \sum_{i \in S' \setminus S} \hat{\mu}_i(t)$ and variance $2A(t)C(t) + 2B(t)C(t)$, then under event $\mathcal{E}_{0,c}$, (recall that $\Phi(x, \mu, \sigma^2) = \Pr_{X \sim \mathcal{N}(\mu, \sigma^2)}[X \geq x]$):

$$\Pr \left[\sum_{i \in S} \theta_i^k(t) - \sum_{i \in S'} \theta_i^k(t) \geq \Delta_{S,S'} \right]$$

$$\begin{aligned}
 &= \Pr \left[\sum_{i \in S \setminus S'} \theta_i^k(t) - \sum_{i \in S' \setminus S} \theta_i^k(t) \geq \Delta_{S, S'} \right] \\
 &= \Phi \left(\Delta_{S, S'}, \sum_{i \in S \setminus S'} \hat{\mu}_i(t) - \sum_{i \in S' \setminus S} \hat{\mu}_i(t), 2A(t)C(t) + 2B(t)C(t) \right) \\
 &\geq \Phi(\Delta_{S, S'}, \Delta_{S, S'} - \sqrt{A(t)L_1(t)} - \sqrt{B(t)L_1(t)}, \\
 &\quad 2A(t)C(t) + 2B(t)C(t)) \\
 &= \Phi(\sqrt{A(t)L_1(t)} + \sqrt{B(t)L_1(t)}, 0, 2A(t)C(t) + 2B(t)C(t)) \\
 &= \Phi \left(\sqrt{\frac{L_1(t)}{C(t)}} \cdot \frac{\sqrt{A(t)} + \sqrt{B(t)}}{\sqrt{2A(t) + 2B(t)}}, 0, 1 \right) \\
 &\geq \Phi \left(\sqrt{\frac{L_1(t)}{C(t)}}, 0, 1 \right) \\
 &= q,
 \end{aligned}$$

where the last equation is because that we choose $C(t) = \frac{L_1(t)}{\phi^2(q)}$ and $\Phi(\phi(q), 0, 1) = q$ (by definition of ϕ).

Note that the parameter sets $\{\theta^k(t)\}_{k=1}^{M(\delta, q, t)}$ are chosen independently, therefore under event $\mathcal{E}_{0,c}$, we have that

$$\begin{aligned}
 &\Pr \left[\forall k, \sum_{i \in S} \theta_i^k(t) - \sum_{i \in S'} \theta_i^k(t) < \Delta_{S, S'} \right] \\
 &\leq (1 - q)^{M(\delta, q, t)} \\
 &\leq \frac{\delta}{12|\mathcal{I}|^2 t^2},
 \end{aligned}$$

where that last inequality is because that we choose $M(\delta, q, t) = \frac{1}{q} \log(12|\mathcal{I}|^2 t^2 / \delta)$.

This implies that

$$\Pr[-\mathcal{E}_{2,c} | \mathcal{E}_{0,c}] \leq \sum_{t, S, S'} \frac{\delta}{12|\mathcal{I}|^2 t^2} \leq \sum_t \frac{\delta}{12t^2} \leq \frac{\delta}{3}.$$

All these show that $\Pr[\mathcal{E}_{0,c} \wedge \mathcal{E}_{1,c} \wedge \mathcal{E}_{2,c}] \geq 1 - \delta$. \square

Then it is sufficient to prove that under event $\mathcal{E}_{0,c} \wedge \mathcal{E}_{1,c} \wedge \mathcal{E}_{2,c}$, TS-Explore works correctly with complexity upper bound shown in Theorem 4.1.

Firstly, we prove that TS-Explore will output S^* . The proof is quite straightforward: if $\hat{S}_t \neq S^*$, then under event $\mathcal{E}_{2,c}$, there exists k such that $\sum_{i \in S^*} \theta_i^k(t) - \sum_{i \in \hat{S}_t} \theta_i^k(t) \geq \Delta_{S^*, \hat{S}_t} > 0$. Therefore $\tilde{S}_t^k \neq \hat{S}_t$ and we will not output \hat{S}_t . Because of this, TS-Explore can only return S^* under event $\mathcal{E}_{0,c} \wedge \mathcal{E}_{1,c} \wedge \mathcal{E}_{2,c}$, and we finish the proof of its correctness.

Then we come to bound the complexity of TS-Explore, and we will use the following lemma in our analysis.

Lemma 4.7. *Under event $\mathcal{E}_{0,c} \wedge \mathcal{E}_{1,c} \wedge \mathcal{E}_{2,c}$, a base arm i will not be pulled if $N_i(t) \geq \frac{98 \text{ width } C(t) L_2(t)}{\Delta_{i,c}^2}$.*

Due to space limit, we only prove Lemma 4.7 for the case that $(\hat{S}_t = S^*) \vee (\tilde{S}_t = S^*)$ in our main text, and defer the complete proof to Appendix A.

Proof. We will prove this lemma by contradiction.

First we consider the case that $(\hat{S}_t = S^*) \wedge (\tilde{S}_t \neq S^*)$. In this case, $i \in S^* \oplus \tilde{S}_t$, which implies that $\Delta_{i,c} \leq \Delta_{S^*, \tilde{S}_t}$. If we choose a base arm i with $N_i(t) \geq \frac{98 \text{ width } C(t) L_2(t)}{\Delta_{i,c}^2} \geq \frac{98 \text{ width } C(t) L_2(t)}{\Delta_{S^*, \tilde{S}_t}^2}$ to pull,

we know that $\forall j \in S^* \oplus \tilde{S}_t$, $N_j(t) \geq \frac{98 \text{ width } C(t) L_2(t)}{\Delta_{S^*, \tilde{S}_t}^2}$.

This implies that $\sqrt{\sum_{j \in S^* \setminus \tilde{S}_t} \frac{2}{N_j(t)}} \leq \frac{\Delta_{S^*, \tilde{S}_t}}{7\sqrt{C(t)L_2(t)}}$ and $\sqrt{\sum_{j \in \tilde{S}_t \setminus S^*} \frac{2}{N_j(t)}} \leq \frac{\Delta_{S^*, \tilde{S}_t}}{7\sqrt{C(t)L_2(t)}}$.

By $\mathcal{E}_{2,c}$, there exists k such that $\sum_{j \in S^* \setminus \tilde{S}_t} \theta_j^k(t) - \sum_{j \in \tilde{S}_t \setminus S^*} \theta_j^k(t) \geq \Delta_{S^*, \tilde{S}_t}$. By $\mathcal{E}_{1,c}$, $|\sum_{j \in S^* \setminus \tilde{S}_t} (\theta_j^k(t) - \hat{\mu}_j(t))| \leq \sqrt{\sum_{j \in S^* \setminus \tilde{S}_t} \frac{2C(t)}{N_j(t)} L_2(t)} \leq \frac{\Delta_{S^*, \tilde{S}_t}}{7}$ and similarly $|\sum_{j \in \tilde{S}_t \setminus S^*} (\theta_j^k(t) - \hat{\mu}_j(t))| \leq \frac{\Delta_{S^*, \tilde{S}_t}}{7}$. Hence

$$\begin{aligned}
 &\sum_{j \in S^* \setminus \tilde{S}_t} \hat{\mu}_j(t) - \sum_{j \in \tilde{S}_t \setminus S^*} \hat{\mu}_j(t) \\
 &\geq \sum_{j \in S^* \setminus \tilde{S}_t} \theta_j^k(t) - \sum_{j \in \tilde{S}_t \setminus S^*} \theta_j^k(t) \\
 &\quad - |\sum_{j \in S^* \setminus \tilde{S}_t} (\theta_j^k(t) - \hat{\mu}_j(t))| - |\sum_{j \in \tilde{S}_t \setminus S^*} (\theta_j^k(t) - \hat{\mu}_j(t))| \\
 &\geq \frac{5\Delta_{S^*, \tilde{S}_t}}{7}.
 \end{aligned}$$

$\mathcal{E}_{1,c}$ also means $|\sum_{j \in S^* \setminus \tilde{S}_t} (\theta_j^{k^*}(t) - \hat{\mu}_j(t))| \leq \frac{\Delta_{S^*, \tilde{S}_t}}{7}$ and $|\sum_{j \in \tilde{S}_t \setminus S^*} (\theta_j^{k^*}(t) - \hat{\mu}_j(t))| \leq \frac{\Delta_{S^*, \tilde{S}_t}}{7}$. Thus we know that

$$\begin{aligned}
 &\sum_{j \in S^* \setminus \tilde{S}_t} \theta_j^{k^*}(t) - \sum_{j \in \tilde{S}_t \setminus S^*} \theta_j^{k^*}(t) \\
 &\geq \sum_{j \in S^* \setminus \tilde{S}_t} \hat{\mu}_j(t) - \sum_{j \in \tilde{S}_t \setminus S^*} \hat{\mu}_j(t) \\
 &\quad - |\sum_{j \in S^* \setminus \tilde{S}_t} (\theta_j^{k^*}(t) - \hat{\mu}_j(t))| - |\sum_{j \in \tilde{S}_t \setminus S^*} (\theta_j^{k^*}(t) - \hat{\mu}_j(t))| \\
 &\geq \frac{3\Delta_{S^*, \tilde{S}_t}}{7}.
 \end{aligned}$$

This contradicts with the fact that \tilde{S}_t is the optimal super arm under the k_t^* -th sample set $\theta^{k_t^*}(t)$.

Then we come to the case that $(\hat{S}_t \neq S^*) \wedge (\tilde{S}_t = S^*)$. In this case, $i \in S^* \oplus \hat{S}_t$, which implies that $\Delta_{i,c} \leq \Delta_{S^*, \hat{S}_t}$. If we choose a base arm i with $N_i(t) \geq \frac{98 \text{ width } C(t) L_2(t)}{\Delta_{i,c}^2} \geq$

$\frac{98 \text{ width } C(t) L_2(t)}{\Delta_{S^*, \hat{S}_t}^2}$ to pull, we have that $\sqrt{\sum_{j \in S^* \setminus \hat{S}_t} \frac{2}{N_j(t)}} \leq \frac{\Delta_{S^*, \hat{S}_t}}{7\sqrt{C(t) L_2(t)}}$ and $\sqrt{\sum_{j \in \hat{S}_t \setminus S^*} \frac{2}{N_j(t)}} \leq \frac{\Delta_{S^*, \hat{S}_t}}{7\sqrt{C(t) L_2(t)}}$.

By $\mathcal{E}_{2,c}$, there exists k such that $\sum_{j \in S^* \setminus \hat{S}_t} \theta_j^k(t) - \sum_{j \in \hat{S}_t \setminus S^*} \theta_j^k(t) \geq \Delta_{S^*, \hat{S}_t}$. By $\mathcal{E}_{1,c}$, $|\sum_{j \in S^* \setminus \hat{S}_t} (\theta_j^k(t) - \hat{\mu}_j(t))| \leq \frac{\Delta_{S^*, \hat{S}_t}}{7}$ and $|\sum_{j \in \hat{S}_t \setminus S^*} (\theta_j^k(t) - \hat{\mu}_j(t))| \leq \frac{\Delta_{S^*, \hat{S}_t}}{7}$. All these imply

$$\begin{aligned} & \sum_{j \in S^* \setminus \hat{S}_t} \hat{\mu}_j(t) - \sum_{j \in \hat{S}_t \setminus S^*} \hat{\mu}_j(t) \\ \geq & \sum_{j \in S^* \setminus \hat{S}_t} \theta_j^k(t) - \sum_{j \in \hat{S}_t \setminus S^*} \theta_j^k(t) \\ & - \left| \sum_{j \in S^* \setminus \hat{S}_t} (\theta_j^k(t) - \hat{\mu}_j(t)) \right| - \left| \sum_{j \in \hat{S}_t \setminus S^*} (\theta_j^k(t) - \hat{\mu}_j(t)) \right| \\ \geq & \frac{5\Delta_{S^*, \hat{S}_t}}{7}. \end{aligned}$$

This contradicts with the fact that \hat{S}_t is the empirically optimal super arm. \square

Lemma 4.7 is similar to Lemma 10 in (Chen et al., 2014). Both of them give an upper bound for the number of pulls on base arm i . The key difference is that in (Chen et al., 2014), for an arm set $U \in \mathcal{J}$, the gap between its real mean and its upper confidence bound is $\tilde{O}(\sum_{i \in U} \sqrt{\frac{1}{N_i}})$, which means that we require all the N_i 's to be $\tilde{\Theta}(\frac{\text{width}^2}{\Delta_{i,c}^2})$ to make sure that this gap is less than $\Delta_{i,c}$. In our paper, based on event $\mathcal{E}_{0,c} \wedge \mathcal{E}_{1,c}$, the gap between U 's real mean and the sum of random samples in U is $\tilde{O}(\sqrt{\sum_{i \in U} \frac{C(t)}{N_i}})$. Therefore, we only require all the N_i 's to be $\tilde{\Theta}(\frac{\text{width } C(t)}{\Delta_{i,c}^2})$ to make sure that this gap is less than $\Delta_{i,c}$, and this reduces a factor of width in the number of pulls on base arm i (our analysis shows that $C(t)$ is approximately a constant). In fact, reducing a width factor in Lemma 4.7 is the key reason that the complexity upper bound of TS-Explore is width lower than the CLUCB policy in (Chen et al., 2014).

The novelty of our proof for Lemma 4.7 mainly lies in the event $\mathcal{E}_{2,c}$ (as well as the mechanism of recording all the $\tilde{\Delta}_t^k$'s and focusing on the largest one), i.e., we show that under event $\mathcal{E}_{2,c}$, \tilde{S}_t has some similar properties as the super arm with the largest upper confidence bound (which is used in LUCB-based policies such as CLUCB). For example, when \hat{S}_t does not equal the optimal super arm S^* , $\mathcal{E}_{2,c}$ tells us that there must exist k such that $\sum_{i \in S^*} \theta_i^k(t) - \sum_{i \in \hat{S}_t} \theta_i^k(t) \geq \Delta_{S^*, \hat{S}_t}$. Along with the fact $k_t^* = \text{argmax}_k \tilde{\Delta}_t^k$, we know $\tilde{\Delta}_t^{k_t^*} \geq \Delta_{S^*, \hat{S}_t}$. This means that the reward gap between \hat{S}_t and \tilde{S}_t ($\tilde{S}_t^{k_t^*}$) could

be larger than Δ_{S^*, \hat{S}_t} , which implies that \tilde{S}_t is either an insufficiently learned sub-optimal arm or the optimal arm (see details in the complete proof in Appendix A). This method solves the challenge of the uncertainty in random samples, and allows us to use similar analysis techniques (e.g., (Chen et al., 2014)) to prove Lemma 4.7.

By Lemma 4.7, if $\forall i, N_i(t) \geq \frac{98 \text{ width } C(t) L_2(t)}{\Delta_{i,c}^2}$, then TS-Explore must terminate (and output the correct answer). Thus, the complexity Z satisfies

$$Z \leq \sum_{i \in [m]} \frac{98 \text{ width } C(Z) L_2(Z)}{\Delta_{i,c}^2} = 98 H_{1,c} C(Z) L_2(Z).$$

For $q \leq 0.1$, $\frac{1}{\phi^2(q)} = O(\frac{1}{\log \frac{1}{q}})$. Then with $C(Z) = \frac{\log(12|\mathcal{I}|^2 Z^2 / \delta)}{\phi^2(q)}$, $L_2(Z) = \log(12|\mathcal{I}|^2 Z^2 M(Z) / \delta)$ and $M(Z) = \frac{\log(12|\mathcal{I}|^2 Z^2 / \delta)}{q}$, we have that (note that $q \geq \delta$):

$$Z \leq O\left(H_{1,c} \frac{(\log(|\mathcal{I}|Z) + \log \frac{1}{\delta})^2}{\log \frac{1}{q}}\right). \quad (1)$$

Therefore, after some basic calculations (the details are deferred to Appendix B), we know that

$$Z = O\left(H_{1,c} \frac{(\log(|\mathcal{I}|H_{1,c}) + \log \frac{1}{\delta})^2}{\log \frac{1}{q}}\right).$$

5. Experiments

In this section, we conduct some experiments to compare the complexity performances of efficient learning algorithms (i.e., TS-Explore and CLUCB (Chen et al., 2014)) and the optimal but non-efficient algorithm NaiveGapElim (Chen et al., 2017) in the CMAB setting. We consider the following combinatorial pure exploration problem instance with parameter $n \in \mathbb{N}^+$, and always choose $q = \delta$ in TS-Explore. All the results (average complexities and their standard deviations) take an average of 100 independent runs.

Problem 5.1. For fixed value n , there are totally $2n$ base arms. For the first n base arms, their expected rewards equal 0.1, and for the last n base arms, their expected rewards equal 0.9. There are only two super arms: S_1 contains the first n base arms and S_2 contains the last n base arms.

We first fix $\delta = 10^{-3}$, and compare the complexity of the above algorithms under different n 's (Fig. 1(a)). We can see that when n increases, the complexities of TS-Explore and NaiveGapElim do not increase a lot, while the complexity of CLUCB increases linearly. This accords with our analysis, since $H_{0,c}(n) = H_{1,c}(n) = 2n \cdot \frac{2n}{(0.8n)^2} = 6.25$ is a constant but $H_{2,c}(n) = 2n \cdot \frac{(2n)^2}{(0.8n)^2} = 12.5n$ is linear

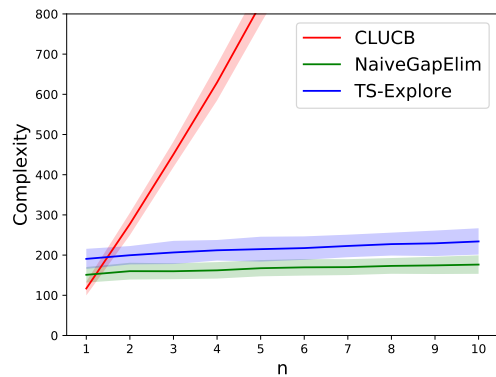
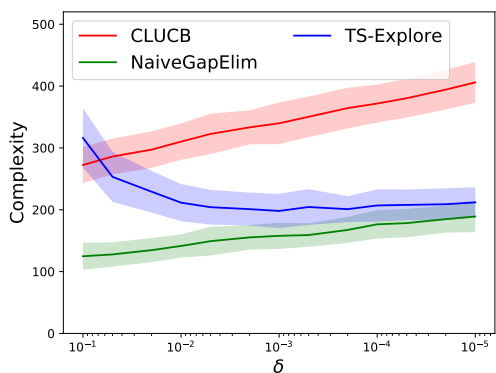

 (a) $\delta = 10^{-3}$

 (b) $n = 2$

Figure 1. Comparison of TS-Explore, CLUCB and NaiveGapElim.

with n (here $H_{0,c}(n)$, $H_{1,c}(n)$, $H_{2,c}(n)$ are the values of $H_{0,c}$, $H_{1,c}$, $H_{2,c}$ under the problem instance with parameter n , respectively).

Then we fix $n = 2$, and compare the complexity of the above algorithms under different δ 's (Fig. 1(b)). We can see that when δ is large, the complexity of TS-Explore decreases as δ decreases, and when δ is small, the complexity of TS-Explore increases as δ decreases. Moreover, the complexities of TS-Explore and NaiveGapElim increase much slower than CLUCB (when δ decreases). This also accords with our analysis. Note that there is a term $O(H_{1,c} \frac{\log^2(|\mathcal{I}|H_{1,c})}{\log \frac{1}{q}})$ in our complexity bound. Since we choose $q = \delta$, this term decreases as δ decreases. When $\delta = 10^{-1}$, this term is very large and becomes the majority term in complexity, and therefore the complexity decreases when δ decreases from 10^{-1} to 10^{-3} . When $\delta = 10^{-3}$, the term $O(H_{1,c} \log \frac{1}{\delta})$ becomes the majority term in complexity, therefore the complexity increases when δ decreases from 10^{-3} to 10^{-5} .

All the experimental results indicate that TS-Explore outper-

forms CLUCB (especially when the size of the problem is large or the error constraint δ is small). On the other hand, the complexity of our efficient algorithm TS-Explore is only a little higher than the optimal but non-efficient algorithm NaiveGapElim. These results demonstrate the effectiveness of our algorithm.

6. Conclusions

In this paper, we explore the idea of Thompson Sampling to solve the pure exploration problems under the frequentist setting. We first propose TS-Explore, an efficient policy that uses random samples to make decisions, and then show that: i) in the combinatorial multi-armed bandit setting, our policy can achieve a lower complexity bound than existing efficient policies; and ii) in the classic multi-armed bandit setting, our policy can achieve an asymptotically optimal complexity bound.

There remain many interesting topics to be further studied, e.g., how to achieve the optimal complexity bound in the combinatorial multi-armed bandit setting by seeking detailed information about the combinatorial structure in TS-Explore; and what are the complexity bounds of using our TS-Explore framework in other pure exploration problems (e.g., dueling bandit and linear bandit). It is also worth studying how to design TS-based pure exploration algorithms for the fixed budget setting.

Acknowledgement

This work was supported by the National Key Research and Development Program of China (2017YFA0700904, 2020AAA0104304), NSFC Projects (Nos. 62061136001, 62106122, 62076147, U19B2034, U1811461, U19A2081), Beijing NSF Project (No. JQ19016), Tsinghua-Huawei Joint Research Program, and a grant from Tsinghua Institute for Guo Qiang.

References

- Agrawal, S. and Goyal, N. Further optimal regret bounds for thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, 2013.
- Audibert, J. Y., Bubeck, S., and Munos, R. Best arm identification in multi-armed bandits. In *COLT 2010 - the Conference on Learning Theory, Haifa, Israel, June*, pp. 41–53, 2010.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- Berry, D. A. and Fristedt, B. *Bandit problems: sequential*

- allocation of experiments (Monographs on statistics and applied probability)*. Springer, 1985.
- Chen, L., Gupta, A., Li, J., Qiao, M., and Wang, R. Nearly optimal sampling algorithms for combinatorial pure exploration. In *Conference on Learning Theory*, pp. 482–534. PMLR, 2017.
- Chen, S., Lin, T., King, I., Lyu, M. R., and Chen, W. Combinatorial pure exploration of multi-armed bandits. In *International Conference on Neural Information Processing Systems*, pp. 379–387, 2014.
- Chen, W., Wang, Y., and Yuan, Y. Combinatorial multi-armed bandit: General framework, results and applications. In *Proceedings of the 30th International Conference on Machine Learning*, pp. 151–159, 2013.
- Combes, R., Talebi, S., Proutière, A., and Lelarge, M. Combinatorial bandits revisited. In *Advances in Neural Information Processing Systems*, 2015.
- Even-Dar, E., Mannor, S., Mansour, Y., and Mahadevan, S. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of machine learning research*, 7(6), 2006.
- Gabillon, V., Lazaric, A., Ghavamzadeh, M., Ortner, R., and Bartlett, P. Improved learning complexity in combinatorial pure exploration bandits. In *Artificial Intelligence and Statistics*, pp. 1004–1012. PMLR, 2016.
- Jourdan, M., Mutn y, M., Kirschner, J., and Krause, A. Efficient pure exploration for combinatorial bandits with semi-bandit feedback. In *Algorithmic Learning Theory*, pp. 805–849. PMLR, 2021.
- Kalyanakrishnan, S., Tewari, A., Auer, P., and Stone, P. Pac subset selection in stochastic multi-armed bandits. In *International Conference on Machine Learning*, 2012.
- Kaufmann, E. and Kalyanakrishnan, S. Information complexity in bandit subset selection. In *Conference on Learning Theory*, pp. 228–251, 2013.
- Kaufmann, E., Korda, N., and Munos, R. Thompson sampling: an asymptotically optimal finite-time analysis. In *International Conference on Algorithmic Learning Theory*, 2012.
- Li, J., Du, C., and Zhu, J. A bayesian approach for subset selection in contextual bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 8384–8391, 2021.
- Perrault, P., Boursier, E., Perchet, V., and Valko, M. Statistical efficiency of thompson sampling for combinatorial semi-bandits. In *Neural Information Processing Systems*, 2020.
- Russo, D. Simple bayesian algorithms for best arm identification. In *Conference on Learning Theory*, pp. 1417–1418. PMLR, 2016.
- Shang, X., Heide, R., Menard, P., Kaufmann, E., and Valko, M. Fixed-confidence guarantees for bayesian best-arm identification. In *International Conference on Artificial Intelligence and Statistics*, pp. 1823–1832. PMLR, 2020.
- Thompson, W. R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- Wang, S. and Chen, W. Thompson sampling for combinatorial semi-bandit. In *International Conference on Machine Learning*, 2018.

A. The Complete Proof of Lemma 4.7

Lemma 4.7 (restated) Under event $\mathcal{E}_{0,c} \wedge \mathcal{E}_{1,c} \wedge \mathcal{E}_{2,c}$, a base arm i will not be pulled if $N_i(t) \geq \frac{98 \text{ width } C(t)L_2(t)}{\Delta_{i,c}^2}$.

Proof. We still prove this lemma by contradiction.

Assume that arm i is pulled with $N_i(t) \geq \frac{98 \text{ width } C(t)L_2(t)}{\Delta_{i,c}^2}$, then there are four probabilities: $(i \in \hat{S}_t) \wedge (i \in S^*)$, $(i \in \hat{S}_t) \wedge (i \notin S^*)$, $(i \in \tilde{S}_t) \wedge (i \in S^*)$, $(i \in \tilde{S}_t) \wedge (i \notin S^*)$.

Case i): $(i \in \hat{S}_t) \wedge (i \in S^*)$ or $(i \in \tilde{S}_t) \wedge (i \notin S^*)$. In this case, $i \in S^* \oplus \tilde{S}_t$ and therefore $\Delta_{i,c} \leq \Delta_{S^*, \tilde{S}_t}$.

Since $\hat{S}_t = \text{Oracle}(\hat{\mu}(t))$, we have that $\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \hat{\mu}_j(t) \leq \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \hat{\mu}_j(t)$. By event $\mathcal{E}_{1,c}$, we also have that for any k , $|\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} (\theta_j^k(t) - \hat{\mu}_j(t))| \leq \sqrt{C(t) \sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \frac{2}{N_j(t)} L_2(t)} \leq \frac{\Delta_{i,c}}{7}$, and similarly $|\sum_{j \in \hat{S}_t \setminus \tilde{S}_t} (\theta_j^k(t) - \hat{\mu}_j(t))| \leq \sqrt{C(t) \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \frac{2}{N_j(t)} L_2(t)} \leq \frac{\Delta_{i,c}}{7}$ (recall that $N_i(t) \geq \frac{98 \text{ width } C(t)L_2(t)}{\Delta_{i,c}^2}$ and $i = \text{argmin}_{j \in \hat{S}_t \oplus \tilde{S}_t} N_j(t)$).

Therefore, for any k , we have that

$$\begin{aligned}
 & \sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \theta_j^k(t) - \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \theta_j^k(t) \\
 &= \left(\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \hat{\mu}_j(t) - \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \hat{\mu}_j(t) \right) + \left(\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} (\theta_j^k(t) - \hat{\mu}_j(t)) \right) - \left(\sum_{j \in \hat{S}_t \setminus \tilde{S}_t} (\theta_j^k(t) - \hat{\mu}_j(t)) \right) \\
 &\leq \left(\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \hat{\mu}_j(t) - \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \hat{\mu}_j(t) \right) + \left(\left| \sum_{j \in \tilde{S}_t \setminus \hat{S}_t} (\theta_j^k(t) - \hat{\mu}_j(t)) \right| \right) + \left(\left| \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} (\theta_j^k(t) - \hat{\mu}_j(t)) \right| \right) \\
 &\leq 0 + \frac{2}{7} \Delta_{i,c} \\
 &\leq \frac{2}{7} \Delta_{i,c}.
 \end{aligned}$$

This means that $\tilde{\Delta}_t^{k_i^*} = \sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \theta_j^{k_i^*}(t) - \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \theta_j^{k_i^*}(t) \leq \frac{2}{7} \Delta_{i,c}$.

Moreover, since $\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \theta_j^{k_i^*}(t) \geq \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \theta_j^{k_i^*}(t)$, we have that

$$\begin{aligned}
 & \sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \hat{\mu}_j(t) - \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \hat{\mu}_j(t) \\
 &= \left(\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \theta_j^{k_i^*}(t) - \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \theta_j^{k_i^*}(t) \right) - \left(\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} (\theta_j^{k_i^*}(t) - \hat{\mu}_j(t)) \right) + \left(\sum_{j \in \hat{S}_t \setminus \tilde{S}_t} (\theta_j^{k_i^*}(t) - \hat{\mu}_j(t)) \right) \\
 &\geq \left(\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \theta_j^{k_i^*}(t) - \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \theta_j^{k_i^*}(t) \right) - \left(\left| \sum_{j \in \tilde{S}_t \setminus \hat{S}_t} (\theta_j^{k_i^*}(t) - \hat{\mu}_j(t)) \right| \right) - \left(\left| \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} (\theta_j^{k_i^*}(t) - \hat{\mu}_j(t)) \right| \right) \\
 &\geq 0 - \frac{2}{7} \Delta_{i,c} \\
 &\geq -\frac{2}{7} \Delta_{i,c}.
 \end{aligned} \tag{2}$$

On the other hand, by event $\mathcal{E}_{2,c}$, we know that $\exists k'$ such that $\sum_{j \in S^*} \theta_j^{k'}(t) - \sum_{j \in \tilde{S}_t} \theta_j^{k'}(t) \geq \Delta_{S^*, \tilde{S}_t}$. Then

$$\sum_{j \in S^*} \theta_j^{k'}(t) - \sum_{j \in \tilde{S}_t} \theta_j^{k'}(t)$$

$$\begin{aligned}
 &= \left(\sum_{j \in S^*} \theta_j^{k'}(t) - \sum_{j \in \tilde{S}_t} \theta_j^{k'}(t) \right) + \left(\sum_{j \in \tilde{S}_t} \theta_j^{k'}(t) - \sum_{j \in \hat{S}_t} \theta_j^{k'}(t) \right) \\
 &\geq \Delta_{S^*, \tilde{S}_t} + \left(\sum_{j \in \tilde{S}_t} \theta_j^{k'}(t) - \sum_{j \in \hat{S}_t} \theta_j^{k'}(t) \right) \\
 &= \Delta_{S^*, \tilde{S}_t} + \left(\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \theta_j^{k'}(t) - \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \theta_j^{k'}(t) \right) \\
 &= \Delta_{S^*, \tilde{S}_t} + \left(\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \hat{\mu}_j(t) - \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \hat{\mu}_j(t) \right) + \left(\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} (\theta_j^{k'}(t) - \hat{\mu}_j(t)) \right) - \left(\sum_{j \in \hat{S}_t \setminus \tilde{S}_t} (\theta_j^{k'}(t) - \hat{\mu}_j(t)) \right) \\
 &\geq \Delta_{S^*, \tilde{S}_t} - \frac{2}{7} \Delta_{i,c} - \left(\left| \sum_{j \in \tilde{S}_t \setminus \hat{S}_t} (\theta_j^{k'}(t) - \hat{\mu}_j(t)) \right| \right) - \left(\left| \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} (\theta_j^{k'}(t) - \hat{\mu}_j(t)) \right| \right) \quad (3) \\
 &\geq \Delta_{S^*, \tilde{S}_t} - \frac{4}{7} \Delta_{i,c} \\
 &> \frac{2}{7} \Delta_{i,c},
 \end{aligned}$$

where Eq. (3) is because of Eq. (2). This means that $\tilde{\Delta}_t^{k'} > \frac{2}{7} \Delta_{i,c}$, which contradicts with $\tilde{\Delta}_t^{k_t^*} \leq \frac{2}{7} \Delta_{i,c}$ (since $k_t^* = \operatorname{argmax}_k \tilde{\Delta}_t^k$).

Case ii): $(i \in \hat{S}_t) \wedge (i \notin S^*)$ or $(i \in \tilde{S}_t) \wedge (i \in S^*)$. In this case, $i \in S^* \oplus \hat{S}_t$ and therefore $\Delta_{i,c} \leq \Delta_{S^*, \hat{S}_t}$.

By event $\mathcal{E}_{2,c}$, we know that $\exists k, \sum_{j \in S^*} \theta_j^k(t) - \sum_{j \in \hat{S}_t} \theta_j^k(t) \geq \Delta_{S^*, \hat{S}_t}$. Hence $\tilde{\Delta}_t^k \geq \Delta_{S^*, \hat{S}_t}$. Moreover, since $k_t^* = \operatorname{argmax}_k \tilde{\Delta}_t^k$, we have that $\sum_{j \in \tilde{S}_t} \theta_j^{k_t^*}(t) - \sum_{j \in \hat{S}_t} \theta_j^{k_t^*}(t) \geq \Delta_{S^*, \hat{S}_t}$, which is the same as $\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \theta_j^{k_t^*}(t) - \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \theta_j^{k_t^*}(t) \geq \Delta_{S^*, \hat{S}_t}$. On the other hand, by event $\mathcal{E}_{1,c}$, we have that $|\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} (\theta_j^{k_t^*}(t) - \hat{\mu}_j(t))| \leq \sqrt{C(t) \sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \frac{2}{N_j(t)} L_2(t)} \leq \frac{\Delta_{i,c}}{7}$, and similarly $|\sum_{j \in \hat{S}_t \setminus \tilde{S}_t} (\theta_j^{k_t^*}(t) - \hat{\mu}_j(t))| \leq \sqrt{C(t) \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \frac{2}{N_j(t)} L_2(t)} \leq \frac{\Delta_{i,c}}{7}$ (recall that $N_i(t) \geq \frac{98 \text{ width } C(t) L_2(t)}{\Delta_{i,c}^2}$ and $i = \operatorname{argmin}_{j \in \hat{S}_t \oplus \tilde{S}_t} N_j(t)$).

Therefore,

$$\begin{aligned}
 &\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \hat{\mu}_j(t) - \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \hat{\mu}_j(t) \\
 &= \left(\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \theta_j^{k_t^*}(t) - \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \theta_j^{k_t^*}(t) \right) - \left(\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} (\theta_j^{k_t^*}(t) - \hat{\mu}_j(t)) \right) + \left(\sum_{j \in \hat{S}_t \setminus \tilde{S}_t} (\theta_j^{k_t^*}(t) - \hat{\mu}_j(t)) \right) \\
 &\geq \left(\sum_{j \in \tilde{S}_t \setminus \hat{S}_t} \theta_j^{k_t^*}(t) - \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} \theta_j^{k_t^*}(t) \right) - \left(\left| \sum_{j \in \tilde{S}_t \setminus \hat{S}_t} (\theta_j^{k_t^*}(t) - \hat{\mu}_j(t)) \right| \right) - \left(\left| \sum_{j \in \hat{S}_t \setminus \tilde{S}_t} (\theta_j^{k_t^*}(t) - \hat{\mu}_j(t)) \right| \right) \\
 &\geq \Delta_{S^*, \hat{S}_t} - \frac{2}{7} \Delta_{i,c} \\
 &> 0.
 \end{aligned}$$

This means that $\sum_{j \in \tilde{S}_t} \hat{\mu}_j(t) - \sum_{j \in \hat{S}_t} \hat{\mu}_j(t) > 0$, which contradicts with the fact that $\hat{S}_t = \operatorname{Oracle}(\hat{\boldsymbol{\mu}}(t))$. \square

B. How to Obtain Complexity Upper Bound by Eq. (1)

Eq. (1) is restated below:

$$Z \leq O \left(H_{1,c} \frac{(\log(|\mathcal{I}|Z) + \log \frac{1}{\delta})^2}{\log \frac{1}{q}} \right).$$

We can then use the following lemma to find an upper bound for complexity Z .

Lemma B.1. *Given K functions $f_1(x), \dots, f_K(x)$ and K positive values X_1, \dots, X_K , if $\forall x \geq X_k, K f_k(x) < x$ holds for all $1 \leq k \leq K$, then for any $x \geq \sum_k X_k, \sum_k f_k(x) < x$.*

Proof. Since X_1, \dots, X_K are positive values, for any $x \geq \sum_k X_k$, we must have that $x \geq X_k$. Therefore $K f_k(x) < x$, which implies that

$$\sum_k K f_k(x) < \sum_k x.$$

This is the same as $\sum_k f_k(x) < x$. □

To apply Lemma B.1, we set $f_1(Z) = H_{1,c} \frac{\log^2(|\mathcal{I}|Z)}{\log \frac{1}{q}}$, $f_2(Z) = H_{1,c} \frac{\log(|\mathcal{I}|Z) \log \frac{1}{\delta}}{\log \frac{1}{q}}$, and $f_3(Z) = H_{1,c} \frac{\log^2 \frac{1}{\delta}}{\log \frac{1}{q}}$. After some basic calculations, we get that $X_1 \leq c_1 H_{1,c} \frac{\log^2(|\mathcal{I}|H_{1,c})}{\log \frac{1}{q}}$, $X_2 \leq c_{2,1} H_{1,c} \frac{\log(|\mathcal{I}|H_{1,c}) \log \frac{1}{\delta}}{\log \frac{1}{q}} + c_{2,2} H_{1,c} \frac{\log^2 \frac{1}{\delta}}{\log \frac{1}{q}}$ and $X_3 \leq c_3 H_{1,c} \frac{\log^2 \frac{1}{\delta}}{\log \frac{1}{q}}$. Here $c_1, c_{2,1}, c_{2,2}, c_3$ are universal constants.

Then we know that for $Z \geq \Theta(H_{1,c} (\log \frac{1}{\delta} + \log(|\mathcal{I}|H_{1,c})) \frac{\log \frac{1}{\delta}}{\log \frac{1}{q}} + H_{1,c} \frac{\log^2(|\mathcal{I}|H_{1,c})}{\log \frac{1}{q}})$, $f_1(Z) + f_2(Z) + f_3(Z) < Z$ (by Lemma B.1). This contradicts with Eq. (1). Therefore, we know that

$$Z = O \left(H_{1,c} \frac{(\log(|\mathcal{I}|H_{1,c}) + \log \frac{1}{\delta})^2}{\log \frac{1}{q}} \right) = O \left(H_{1,c} \left(\log \frac{1}{\delta} + \log(|\mathcal{I}|H_{1,c}) \right) \frac{\log \frac{1}{\delta}}{\log \frac{1}{q}} + H_{1,c} \frac{\log^2(|\mathcal{I}|H_{1,c})}{\log \frac{1}{q}} \right),$$

and this is the complexity upper bound in Theorem 4.1.

C. Discussions about the Optimal Algorithms for Combinatorial Pure Exploration

Chen et al. (2017) prove that the complexity lower bound for combinatorial pure exploration is $\Omega(H_{0,c} \log \frac{1}{\delta})$, where $H_{0,c}$ is the optimal value of the following convex program (here $\Delta_{S^*, S} = \sum_{i \in S^*} \mu_i - \sum_{i \in S} \mu_i$):

$$\begin{aligned} \min \quad & \sum_{i \in [m]} N_m \\ \text{s.t.} \quad & \sum_{i \in S^* \oplus S} \frac{1}{N_i} \leq \Delta_{S^*, S}^2, \quad \forall S \in \mathcal{I}, S \neq S^* \end{aligned}$$

In other words, for any correct combinatorial pure exploration algorithm, $[\frac{\mathbb{E}[N_1]}{\log \frac{1}{\delta}}, \frac{\mathbb{E}[N_2]}{\log \frac{1}{\delta}}, \dots, \frac{\mathbb{E}[N_m]}{\log \frac{1}{\delta}}]$ must be a feasible solution for the above convex program, where $\mathbb{E}[N_i]$ represents the expected number of pulls on base arm i . We say base arm i needs exploration the most at time t if $\alpha N_i^* \log \frac{1}{\delta} - N_i(t)$ is positive, where α is some universal constant and N_i^* is the value of N_i in the optimal solution $H_{0,c}$ (note that there may be several base arms that need exploration the most in one time step). By this definition, if we always pull a base arm that needs exploration the most, then the frequency of pulling each base arm converges to the optimal solution of $H_{0,c}$, which leads to an optimal complexity upper bound.

However, the simple offline oracle used in TS-Explore (as described in Section 3.2) is not enough to look for a base arm that needs exploration the most. In fact, both the existing optimal policies (Chen et al., 2017; Jourdan et al., 2021) not only need to use this simple offline oracle, but also require some other mechanisms to explore detailed information about the combinatorial structure of the problem instance to look for a base arm that needs exploration the most. The algorithms in (Chen et al., 2017) need to record all the super arms in \mathcal{I} with an empirical mean larger than some threshold η . This is one kind of information about the combinatorial structure that can help to find out a base arm that needs exploration the most.

Nevertheless, the authors only provide an efficient way to do this when the combinatorial structure satisfies some specific constraints. In the most general setting, the algorithms in (Chen et al., 2017) must pay an exponential time cost for collecting the detailed information. As for (Jourdan et al., 2021), the best-response oracle used by the λ -player needs to return a super arm within set $N(S_t)$ that has the shortest distance to a fixed target. Here S_t is a super arm, $N(S_t)$ represents the set of super arms whose cells' boundaries intersect the boundary of the cell of S_t , and the cell of a super arm S_t is defined as all the possible parameter sets $[\theta_1, \theta_2, \dots, \theta_m]$ in which S_t is the optimal super arm. This is another kind of information about the combinatorial structure that can help to find out a base arm that needs exploration the most. Nevertheless, this best-response oracle also has an exponential time cost (which is scaled with $|N(S_t)|$). By using these exponential time cost mechanisms, the optimal algorithms (Chen et al., 2017; Jourdan et al., 2021) can find out a base arm that needs exploration the most, which is critical to achieving the optimal complexity upper bound.

In this paper, to make TS-Explore efficient in the most general setting, we only use the simple offline oracle in our algorithm and our mechanism can only inform us of one of the violated constraints in the optimization problem (if all the constraints are not violated, TS-Explore will output the correct optimal arm). This means that we know nothing about the combinatorial structure of the problem instance. Therefore, the best thing we can do is to treat all the base arms in the violated constraint equally, i.e., we choose to pull the base arm (in the violated constraint) with the smallest number of pulls. This leads to a complexity upper bound of $O(H_{1,c} \log \frac{1}{\delta})$. If we want TS-Explore to achieve an optimal complexity upper bound $O(H_{0,c} \log \frac{1}{\delta})$, then we need to know which base arm needs exploration the most, e.g., by applying a powerful offline oracle that takes the empirical means and random samples as input and outputs a base arm that needs exploration the most. How to design such offline oracles and how to implement them efficiently is one of our future research topics.

D. Concentration Inequality of Sub-Gaussian Random Variables

Fact D.1. If X is zero-mean and σ^2 sub-Gaussian, then

$$\Pr[|X| > \epsilon] \leq 2 \exp\left(-\frac{\epsilon^2}{2\sigma^2}\right).$$