

Collaborative Top Distribution Identifications with Limited Interaction*

Nikolai Karpov
Computer Science Department
Indiana University
nkarpov@iu.edu

Qin Zhang
Computer Science Department
Indiana University
qzhangcs@indiana.edu

Yuan Zhou
Department of ISE
University of Illinois at Urbana-Champaign
yuanz@illinois.edu

Abstract

We consider the following problem in this paper: given a set of n distributions, find the top- m ones with the largest means. This problem is also called *top- m arm identifications* in the literature of reinforcement learning, and has numerous applications. We study the problem in the collaborative learning model where we have multiple agents who can draw samples from the n distributions in parallel. Our goal is to characterize the tradeoffs between the running time of learning process and the number of rounds of interaction between agents, which is very expensive in various scenarios. We give optimal time-round tradeoffs, as well as demonstrate complexity separations between top-1 arm identification and top- m arm identifications for general m and between fixed-time and fixed-confidence variants. As a byproduct, we also give an algorithm for selecting the distribution with the m -th largest mean in the collaborative learning model.

*Accepted for presentation at the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS 2020). Nikolai Karpov and Qin Zhang are supported in part by NSF IIS-1633215, CCF-1844234 and CCF-2006591. Yuan Zhou is supported in part by NSF CCF-2006526 and a JPMorgan Chase AI Research Faculty Research Award.

1 Introduction

In this paper we study the following problem: given a set of n distributions, try to find the m ones with the largest means via sampling. We study the problem in the multi-agent setting where we have K agents, who try to identify the top- m distributions collaboratively via communication. Suppose sampling from each distribution takes a unit time, our goal is to minimize both the running time and the number of rounds of communication of the collaborative learning process.

The problem of top- m distribution identifications originates from the literature of *multi-armed bandits* (MAB) [53], where each distribution is called an *arm*, and each sampling from a distribution is called an *arm pull*. When $m = 1$, the problem is called *best arm identification*, and has been studied extensively in the centralized setting where there is only one agent [5, 11, 24, 27, 46, 40, 34, 41, 20, 13, 29]. Some of these algorithms can be easily modified to handle top- m arm identification (e.g., [5, 12]). The problem of best arm identification has also been studied in the multi-agent collaborative learning model [31, 55]. Surprisingly, we found that in the multi-agent setting, the tasks of identifying the best arm and the top- m arms look to be very different in terms of problem complexities; the algorithm design and lower bound proof for the top- m case require significantly new ideas, and need to address some fundamental challenges in collaborative learning.

Collaborative Learning with Limited Interaction. A natural way to speed up machine learning tasks is to introduce multiple agents, and let them learn the target function collaboratively. In recent years some works have been done to address the power of parallelism (under the name of *concurrent learning*, e.g., [51, 30, 23, 22]). Most of these works assume that agents have the full ability of communication. That is, they can send/receive messages to/from each other at *any* time step. This assumption, unfortunately, is unrealistic in real-world applications, as it would be very expensive to implement unrestricted communication, which is usually the biggest drain of time, data, energy and network bandwidth. For example, once we deploy sensors/robots to unknown environment such as deep sea and outer space, it would be almost impossible to recharge them; when we train a model in a central server by interacting with hundreds of thousands of mobile devices, the communication cost will directly contribute to our data bills, not mentioning the excessive energy and bandwidth consumption.

In this paper we consider the model of *collaborative learning with limited interaction*, where the learning process is partitioned into rounds of predefined time intervals. In each round, each of the K agents takes a series of actions individually like in the centralized model, and they can only communicate at the end of each round. At the end of the last round before any communication, all agents should agree on the same output; otherwise we say the algorithm fails. Our goal is to minimize both the number of rounds of computation R and the running time T (assuming each action takes a unit time step).¹

Naturally, there is a tradeoff between R and T : If $R = 1$, that is, no communication is allowed, then $T \geq T_C$ where T_C is the running time of the best centralized algorithm. When R increases, T may decrease. On the other hand we always have $T \geq T_C/K$ even when $R = T$. We are mostly interested in understanding the number of rounds needed to achieve almost *full speedup*, that is, when $T = \tilde{O}(T_C/K)$ where $\tilde{O}(\cdot)$ hides logarithmic factors.

We do not put any constraints on the lengths of the messages that each agent can send at the end of each round, but in the MAB setting they will not be very large – the information that each agent collects can

¹We note that our model is a simplified version of the one formulated in [55]. The model defined in [55] allows each agent to perform different numbers of actions in each round, and the length of each round can be determined adaptively by the agents. However, we noticed that all the existing algorithms for collaborative learning in the literature have predefined round lengths, under which there is no point for an agent to stop early in a round.

always be compressed to an array of n pairs in the form of $(x_i, \tilde{\theta}_i)$, where x_i is the number of arm pulls on the i -th arm, and $\tilde{\theta}_i$ is the empirical mean of the x_i arm pull.

Top- m Arm Identification. To be consistent with the MAB literature, we will use the term *arm* instead of *distribution* throughout this paper. The top- m arm identification problem is motivated by a variety of applications ranging from industrial engineering [42] to medical tests [56], and from evolutionary computation [50] to crowdsourcing [1]. The readers may refer to [5, 37, 21, 18, 19] and references therein for the state-of-the-art results on the top- m arm identification in the centralized model.

In this paper we mainly focus on the *fixed-time* case, where given a fixed time horizon T , the task is to identify the set of m arms with the largest means with the smallest error probability. We will also discuss the *fixed-confidence* case, where given a fixed error probability δ , the task is to identify the top- m arms with error δ using the smallest amount of time.

Without loss of generality, we assume that each of the underlying distributions has support on $(0, 1)$. In the centralized setting, Bubeck et al. [12] introduced the following complexity to characterize the hardness of an input instance V for the top- m arm identification problem. Let θ_i be the mean of the i -th arm. Let $[j]$ be the index of the arm in V with the j -th largest mean, and let $\theta_{[j]}(V)$ be the corresponding mean. Given an input instance I of n arms, let $\Delta_i^{(m)}(I)$ be the gap between the mean of the i -th arm and that of the $[m]$ -th arm or the $[m + 1]$ -th arm, whichever is larger. In other words,

$$\Delta_i^{(m)}(I) \triangleq \begin{cases} \theta_i - \theta_{[m+1]}(I), & \text{if } \theta_i \geq \theta_{[m]}(I), \\ \theta_{[m]}(I) - \theta_i, & \text{if } \theta_i \leq \theta_{[m+1]}(I). \end{cases} \quad (1)$$

Definition 1 (Instance Complexity). Given an input instance I of n arms and a parameter m (call it the pivot), we define the following quantity which characterizes the complexity of I .

$$H^{(m)}(I) \triangleq \sum_{i \in I} \left(\Delta_i^{(m)}(I) \right)^{-2}.$$

We also define a related quantity which we call the ϵ -truncated instance complexity.

$$H_\epsilon^{(m)}(I) \triangleq \sum_{i \in I} \max \left\{ \Delta_i^{(m)}(I), \epsilon \right\}^{-2}.$$

To see why $H^{(m)}(I)$ is the right measure for the instance complexity, note that if the mean of an arm is either $(\theta + \Delta)$ or $(\theta - \Delta)$ where θ is a known threshold, it takes $\Omega(\Delta^{-2})$ samples to decide whether the mean is above or below the threshold θ (as long as $\theta \pm \Delta$ are bounded away from 0 and 1). Therefore, suppose all the means are bounded away from 0 and 1, even if we are given the means of the $[m]$ -th and the $[m + 1]$ -th arms, it still takes $\Omega(H^{(m)}(I))$ samples to decide for each arm whether it is one of the top- m arms or not. Such intuition can be formalized to show that, in the fixed-confidence case, $\Omega(H^{(m)}(I) \log(1/\delta))$ samples are needed to identify the top- m arms with success probability $(1 - \delta)$ [52, 19]. On the other hand, there are centralized algorithms to achieve $O(H^{(m)}(I) \log(1/\delta) + H^{(m)}(I) \log H^{(m)}(I))$ (see, e.g., [37]), almost matching the lower bound (up to logarithmic factors).

For the fixed-time case, in [12] it was shown that there is a centralized algorithm that identifies the top- m arms with probability at least

$$1 - \exp \left(-\tilde{\Omega} \left(\frac{T}{H^{(m)}(I)} \right) \right) \quad (2)$$

using at most T time steps, where $\tilde{\Omega}(\cdot)$ hides logarithmic factors in n . This upper bound can also be shown to be tight up to logarithmic factors [41, 13, 52, 19]. In the collaborative learning setting, our goal is to replace the T factor in (2) with KT where K is the number of agents, so as to achieve a full speedup.

Our Contributions. We summarize our main results and their implications.

1. We give an algorithm for the fixed-time top- m arm identification problem in the collaborative learning model with K agents and a set I of n arms. For any choice of r , the algorithm uses T time steps and $O(\log \frac{\log m}{\log K} + r)$ rounds of communication, and successfully computes the set of top- m arms with probability at least $1 - \exp\left(-\tilde{\Omega}\left(\frac{K^{(r-1)/r} \cdot T}{H^{(m)}(I)}\right)\right)$. In particular, when $r = \log K$, the algorithm uses T time steps and $O(\log \frac{\log m}{\log K} + \log K)$ rounds of communication to compute the set of top- m arms with probability at least $1 - \exp\left(-\tilde{\Omega}\left(\frac{KT}{H^{(m)}(I)}\right)\right)$, achieving a full speedup. See Section 3.
2. We prove that under the same setting, any collaborative algorithm that uses $T = \frac{1}{\sqrt{K}} \cdot H^{(m)}(I)$ time steps and aims to achieve success probability 0.99 needs at least $\Omega(\log \frac{\log m}{\log K})$ rounds of communication. By leveraging a result in [55], we can also show that any collaborative algorithm that uses $T = \frac{\alpha}{K} \cdot H^{(m)}(I)$ time steps and aims to achieve success probability 0.99 needs at least $\Omega(\log K / (\log \log K + \log \alpha))$ rounds of communication. These indicate that our upper bound is almost the best possible. See Section 4.
3. Our lower bound gives a strong separation between the best arm identification and top- m identifications: there is a collaborative algorithm for best arm identification (i.e., when $m = 1$) that uses $T = \tilde{O}\left(\frac{1}{\sqrt{K}} \cdot H^{(1)}(I)\right)$ time and 2 rounds of communication (see [55, 31]), while Item 2 states that for general m , to achieve the same time bound we need $\Omega(\log K / (\log \log K + \log \alpha))$ rounds of communication.
4. We give an algorithm for the fixed-confidence top- m identification problem in the collaborative model with K agents and a set of n arms; the algorithm uses $O\left(\frac{H^{(m)}(I)}{K} \log\left(\frac{n}{\delta} \log H^{(m)}\right)\right)$ time steps and $O\left(\log(1/\Delta_{[m]}^{(m)})\right)$ rounds of communication, and successfully computes the set of top- m arms with probability at least $1 - \delta$. This is almost tight by a previous result in [55]. See Section 5.
5. Combining Items 1, 2, and 4, we have given a separation between fixed-time and fixed-confidence top- m arm identification. We note that a similar separation result is also proved for the best arm identification problem [55], although the round complexities for top- m identification are quite different from the $m = 1$ special case (i.e., best arm identification).

Speedup. In [55] the authors introduced a concept called *speedup* for presenting the power of collaborative learning algorithms. The precise definition of speedup is rather complicated due to the definition of the instance complexity of MAB. Roughly, the speedup is defined to be the ratio between the *best* running time of centralized algorithm and that of a collaborative algorithm (given a predefined round budget R) under the condition that the two algorithms achieve the same success probability. In this paper we simply focus on a fixed success probability 0.99, and define the speedup of a collaborative algorithm which identifies the top- m arms on input instance I with accuracy 0.99 using time $T_{\mathcal{A}}(I)$ to be $T_{\mathcal{A}}(I)/H^{(m)}(I)$, since the best centralized algorithm achieving success probability 0.99 has running time $\tilde{\Theta}(H^{(m)}(I))$ [12]. Interpreting our results in terms of speedup, we have the following remarks:

1. Our algorithm for fixed-time top- m arm identification achieves a speedup of $\tilde{O}(K^{\frac{r-1}{r}})$ and uses $O(\log \frac{\log m}{\log K} + r)$ rounds.
2. Our lower bound shows that in order to achieve even an $\tilde{\Omega}(\sqrt{K})$ speedup, any algorithm for top- m arm identification needs at least $\Omega(\log \frac{\log m}{\log K})$ rounds.
3. Compared with the main result for the best arm identification in [55], which states that there is a R -round algorithm achieving a speedup of $\tilde{O}(K^{\frac{R-1}{R}})$, we have shown a separation between the complexities of the two problems (e.g., when $R = 2$).

Selection under Uncertainty. As a byproduct, we also get almost tight bounds for a closely related problem we call *selection under uncertainty*. This problem is similar to the classic *selection* problem where given a set of n numbers, one needs to find the m -th largest number. The difference is that now instead of having n (deterministic) numbers, we have n distributions/arms, and our goal is to find the one with the m -th largest mean via sampling. It is easy to see that this problem can be solved by first identifying the top- m arms, and then finding the worst arm in these top- m arms, which can be done in the same way as identifying the best arm.

For convenience, let us introduce a new (but very similar) definition of instance complexity for the selection under uncertainty problem:

$$\bar{H}^{(m)}(I) \triangleq \sum_{i \neq [m]} (\theta_i - \theta_{[m]})^{-2}.$$

With $\bar{H}^{(m)}$ we have the following immediate result:

- There exists an algorithm for the fixed-time m -th arm selection problem in the collaborative learning model with K agents and a set I of n arms; the algorithm uses T time steps and $O(\log \frac{\log m}{\log K} + r)$ rounds of communication, and successfully identifies the m -th arm with probability at least $1 - \exp\left(-\tilde{\Omega}\left(\frac{K^{(r-1)/r} \cdot T}{\bar{H}^{(m)}(I)}\right)\right)$.

Why Top- m Arm Identification is Difficult in the Collaborative Learning Model? Before presenting our results, let us first try to give some intuition on why top- m arm identification is difficult in the collaborative learning setting, as one may think that the top- m arm identification is a natural generalization of best arm identification (when $m = 1$), and the algorithm for the latter in [55] may be adapted to the former.

The key procedure used in previous collaborative algorithms for best arm identification [31, 55] is that in the first round, we *randomly* partition the set I of n arms into K groups, and feed each group to one agent as a subproblem. Now if each of the K agents computes the best arm in its subproblem, then we can reduce the number of best arm candidates from n to K after the first round, which is critical for us to achieve $\log K$ communication rounds. The question now is whether each subproblem can be solved time-efficiently (more precisely, in $\tilde{O}(H^{(1)}(I)/K)$ time steps if we target a $\tilde{\Omega}(K)$ speedup) at each agent in the first round.

A nice property for the best arm identification is that if we randomly partition the set I of n arms to the K groups, then the group (denoted by G) containing the global best arm has a subproblem complexity $H' = \sum_{i=2}^{|G|} (\Delta'_i)^{-2}$, where Δ'_i is the difference between the mean of the best arm and that of the i -th best arm in group G . It is easy to show that

$$\mathbb{E}[H'] = \Theta\left(H^{(1)}(I)/K\right). \quad (3)$$

Therefore, even though we cannot guarantee that each of the K subproblems can be solved successfully under time budget $\tilde{O}(H/K)$, we still know that the global best arm will advance to the next round with a good probability, which is enough for the algorithm to succeed.

Unfortunately, the above property does not hold in the top- m setting due to its “multi-objective” goal. First, the global m -th arm will only be assigned to one agent, and thus others do not know what pivot to use for defining its subproblem complexity. Second, even for the agent who gets the m -th arm j , it does not know what is the local rank of j , and, thus, still does not know when to stop the local pruning. Third, even if the agents know the local ranks of the m -th arm, it may not have enough time budget to solve the sub-problem; note that this is an issue only for the top- m case but *not* for the best arm case, since in the top- m case each subproblem may contain some top- m arms.

We will design an algorithm which addresses all of these challenges, and then complement it with an almost tight lower bound. Looking back, we feel that in the best arm case it was just lucky for us to have Equation (3), while in the general top- m case we have to deal with some inherent challenges in collaborative learning, which, unfortunately, also make our algorithm for top- m much more complicated than that for best arm identification. We will give a technical overview for both the algorithm and lower bound proof in Section 2.

Related Work. To the best of our knowledge, the collaborative learning model studied in this paper was first proposed in [31], where the authors studied the best arm identification problem in MAB. The model was recently formalized in [55], where almost tight time-round tradeoffs for best arm identification are given.

A number of works studied *regret minimization*, which is another important problem in MAB, in various distributed models, most of which are different from the collaborative learning model considered in this paper. For example, several works [45, 49, 9] studied regret minimization in the setting of cognitive radio network, where radio channels are models as arms, and the rewards by pulling each arm depend on the number of simultaneous pulls by the K agents (i.e., penalty is introduced for collisions). In [16] the authors considered a model where at each time step each agent can choose either to pull an arm, or broadcast a message to other agents, but cannot do both. Authors of [54, 43, 58] considered regret minimization in communication networks. Distributed regret minimization has also been studied in the non-stochastic setting [6, 38, 15].

The collaborative learning model is closely related to the *batched* model (or, *learning with limited adaptivity*), where one wants to minimize the number of policy switches in the learning process. In the batched model we want to minimize the number of policy switches when trying to achieve our learning goal. Algorithms designed in the batched model can naturally be translated to a *restricted* version of the collaborative model in which at each time step, the action taken by each agent is determined by the information (historical actions and outcomes, messages received from other agents, and the randomness of the algorithm) the agent has at the beginning of the round, and the agents cannot change their policies in the middle of the a round. A number of problems have been studied in the batched model in recent years, including best arm identification [36, 2, 35], regret minimization in MAB [48, 28, 26], Q -learning [7], convex optimization [25], online learning [14]. We note that our collaboratively learning algorithm for top- m arm identification in the fixed-confidence case also works in the batched model, and improves the algorithm in [35].

Finally, we note that there is also a large body of work on sample/communication-efficient distributed algorithms for various learning-related tasks such as classification [8, 32, 39], convex optimization [60, 59, 3], linear programming [4, 57]. Sample-efficient PAC learning in the collaborative setting is recently studied by [10, 17, 47]. However, the models considered in the papers mentioned above mainly focus on reducing the sample/communication cost, and are all different from the collaborative learning with limited interaction

n	number of arms in the input instance.
K	number of agents.
T	running time.
θ_i	mean of the i -th arm.
$\theta_{[i]}(V)$	the i -th largest mean among arms in V .
$Top_m(V)$	indices of the m arms with the largest means in V .
$Top_1(V)$	index of the best arm in V .
$\Delta_i^{(m)}(V)$	mean gap of the i -th arm; defined in (1).
$H^{(m)}(V)$	instance complexity; see Definition 1.
$H_\epsilon^{(m)}(V)$	ϵ -truncated instance complexity; see Definition 1.

Table 1: Summary of Notations

model we study in this paper.

Notations and Conventions. Let $Top_m(V)$ be the indices of m arms in V with the largest means, and $Top_1(V)$ be the index of the best arm in V .

We say the i -th arm is (ϵ, j) -top in V if and only if $\theta_i \geq \theta_{[j]}(V) - \epsilon$. Similarly, the i -th arm is (ϵ, j) -bottom in V if and only if $\theta_i \leq \theta_{[|V|+1-j]}(V) + \epsilon$.

In this paper we focus on the case when $\theta_{[m]}(I) > \theta_{[m+1]}(I)$, since otherwise the instance complexity of I will be infinity.

For simplicity, we will write $Top_m(V)$, $Top_1(V)$, $\theta_{[i]}(V)$, $\Delta_i^{(m)}(V)$, $H^{(m)}(V)$, and $H_\epsilon^{(m)}(V)$ as Top_m , Top_1 , $\theta_{[i]}$, $\Delta_i^{(m)}$, $H^{(m)}$, and $H_\epsilon^{(m)}$, when $V = I$ (I is the input instance) or it is clear from the context.

We include a list of frequently used (global) notations in Table 1.

Roadmap. In the rest of the paper, we first give a technical overview of our main results in Section 2. We next present our algorithm for the fixed-time case in Section 3, and then complement it with a matching lower bound in Section 4. Finally in Section 5, we give an algorithm for the fixed-confidence case and discuss the corresponding lower bound.

2 Technical Overview

In this section we give a technical overview for our upper and lower bounds for fixed-time top- m arm identification.

2.1 Upper Bounds for the Fixed-Time Setting

For simplicity we consider the full speedup setting (i.e., we target a speedup of $\tilde{\Omega}(K)$); the general speedup is an easy extension. We achieve our upper bound result for fixed-time top- m arm identification in three stages. We first design an algorithm for a special time horizon $T = \tilde{\Theta}(H^{(m)}/K)$ which uses $O(\log \frac{\log n}{\log K} + \log K)$ rounds of communication and has an error probability 0.01. We next consider general time horizon T , and target an error probability that is exponentially small in T . Finally, we try to improve the round complexity to $O(\log \frac{\log m}{\log K} + \log K)$. In each stage we face new challenges which stem from the collaborative learning model, each of which demands novel ideas.

Stage 1: A Basic Algorithm. We start with our basic algorithm. A natural idea for achieving the $T = \tilde{O}(H^{(m)}/K)$ running time is to *randomly* partition the n arms to K agents, and then ask each agent to solve a top- η arms identification (for some value η) on its sub-instance. At the end we try to aggregate the K outputs. As briefly mentioned in the introduction, there are multiple hurdles associated with this approach. First, it is not clear how to set the value η , since we do not know how many global top- m arms will be distributed to each agent. Second, even if we know the number of global top- m arms assigned to each agent, there are cases in which the global instance complexity is rarely distributed evenly across the K agents. In other words, we cannot guarantee that each agent can solve the subproblem within our time budget $\tilde{O}(H^{(m)}/K)$.

We resolve these issues using the following ideas: we take a *conservative* approach by setting $\eta \approx (m/K - \sqrt{n})$, and ask each agent to adopt a PAC algorithm for multiple arm identification and compute an *approximate* set of top- η arms on its sub-instance using $\tilde{O}(H^{(m)}/K)$ time steps. The approximation error is a *random variable* depending on the random partition process. We then show that with a good probability this error is smaller than the gap between the smallest mean of the outputted arms and that of the *global* m -th arm. In this way we can guarantee that the approximate top- η arms outputted by each agent are indeed in the set of global top- m arms. Using the same idea we try to prune a set of “bottom” arms of size $\approx ((n - m)/K - \sqrt{n})$. After these operations we recurse on the rest $O(K\sqrt{n})$ arms. We continue the recursion for $O(\log \frac{\log n}{\log K})$ rounds until the number of arms is reduced to K^{10} , and then use a simple $O(\log n)$ -round collaborative algorithm which is modified from an existing centralized algorithm. Note that for $n' = K^{10}$ we have $O(\log n') = O(\log K)$, and thus overall we have used $O(\log \frac{\log n}{\log K} + \log K)$ rounds.

Stage 2: General Time Horizon. The basic algorithm only guarantees that the set of top- m arms are correctly identified with probability 0.99. Our next goal is to make the error probability exponentially small in T , which is achievable in the centralized setting. The standard technique to achieve this is to perform parallel repetition and then take the majority. That is, we guess the instance complexity to be $H = 1, 2, 4, \dots$, and for each guess we run the basic algorithm with time horizon H for T/H times. Finally, we take the majority of the output. In the case that the budget T is larger than the actual instance complexity, at each run with probability 0.99 we are guaranteed to obtain the correct answer. Unfortunately, when T is smaller than the actual instance complexity, not much can be guaranteed. For some bad input instances, the output of the basic algorithm can be *consistently* wrong, resulting in a wrong majority.

We resolve this difficulty by introducing a notion we call *top- m certificate*, which takes form of a pair $(S, \{\tilde{\theta}_i\}_{i \in I})$, with the property that $S = \text{Top}_m$ and for each $i \in I$, it holds that $|\tilde{\theta}_i - \theta_i| < \Delta_i^{(m)}/4$. We can augment our basic algorithm to output a $(S, \{\tilde{\theta}_i\}_{i \in I})$ pair (instead of simply a set of top- m arms). We then design a verification algorithm which is able to check for each $(S, \{\tilde{\theta}_i\}_{i \in I})$ pair whether it is indeed a top- m certificate. Our verification step can be fully parallelized and can finish within our guessed instance complexity H . With such a verification step at hand, the situation that we take a wrong majority will *not* happen with high probability.

Stage 3: Better Round Complexity. Our ultimate goal is to achieve an $O(\log \frac{\log m}{\log K} + \log K)$ round complexity, instead of $O(\log \frac{\log n}{\log K} + \log K)$ in the basic algorithm. We approach this by first reducing the number of arms in the input instance to $\tilde{O}(m)$, and then applying the basic algorithm. Such a reduction, however, is highly non-trivial, especially when we require the error probability introduced by the reduction to again be exponentially small in T .

Our basic idea for performing the reduction is the following: we construct a random sub-instance V by

sampling each of the n arms with probability $1/m$. We can show that with constant probability, V contains exactly one global top- m arm, and $H^{(1)}(V) = O(H^{(m)}/m)$. Therefore we have enough time budget to compute the best arm of V and include it into set S as a top- m candidate. We perform this subsampling procedure for $\tilde{O}(m)$ times, getting $\tilde{O}(m)$ sub-instances. By the Coupon Collector’s problem we know that all global top- m arms will be included in S with a good probability.

The challenging part is to reduce the error probability of this reduction to a value that is exponentially small in T . Unfortunately, the idea of “guess-then-verify” that we have used previously does not apply here – there is simply no $(S, \{\tilde{\theta}_i\}_{i \in I})$ pair for us to verify in the reduction process.

We take the following new approach. We try to make sure that for each randomly sampled sub-instance on which we try to compute the best arm, the probability of outputting any arm in Top_m is at least *half* of that of any arm outside Top_m . This turns out to be enough for us to guarantee that the set S contains all top- m arms. We comment that the relaxation “half” is necessary here for a technical reason which we will elaborate next.

Our key observation is that if we provide sufficient time budget, say, $T \geq \lambda H^{(1)}(V)$ where λ is a polylogarithmic factor, for solving a randomly sampled sub-instance V , then provided that there is only one arm $a \in Top_m$ in V , we will output a correctly with a good probability. Now for any two arms $a \in Top_m$ and $b \notin Top_m$, by the uniformity of the sampling they will be in the sub-instance with equal probability. We are thus able to conclude that the probability of outputting a is at least as large as that of outputting b . On the other hand, if $T \leq H^{(1)}(V)$, then we can use our verification step to detect this event. The subtle part is the middle case when $H^{(1)}(V) \leq T \leq \lambda H^{(1)}(V)$, to handle which we perturb our time budget T such that it takes values T/λ or λT with equal probability. Using this trick we are able to “reduce” the third case to the first two cases with probability at least $1/2$, which leads to our desired property. The actual implementation of this idea is more involved, and we refer the readers to Section 3.4 for details.

2.2 Lower Bounds for the Fixed-Time Setting

In the lower bound part, we present two results. The first result is that $\Omega(\log K / (\log \log K + \log \alpha))$ communication rounds are needed for any algorithm with (K/α) speedup to identify the top- m arms for any m . This matches (up to logarithmic factors) the R term in the $O(\log \frac{\log m}{\log K} + R)$ rounds vs $\tilde{O}(K^{(R-1)/R})$ speedup trade-off in our upper bound result. This lower bound theorem is derived via a simple reduction together with the similar type of lower bound proved in [55] for the $m = 1$ special case.

Our main contribution in the lower bound part is the second theorem. The theorem states that even if the goal is an $O(\sqrt{K})$ speedup, the $\log \frac{\log m}{\log K}$ term in the round-speedup trade-off is necessary. (In fact, the $\log \frac{\log m}{\log K}$ can be shown to be necessary for any K^ζ speedup where ζ is a positive constant.) This marks a completely different phenomenon from the $m = 1$ special case where only 2 rounds of communication are needed to achieve an $\tilde{O}(\sqrt{K})$ speedup [55, 31]. Below we sketch the proof idea for this lower bound theorem.

The need for the $\log \frac{\log m}{\log K}$ term in the round complexity stems from the hardness of collaboratively learning the splitting position (i.e., where the m -th largest arm locates), which turns out to be substantially more difficult than estimating the best arm (the $m = 1$ special case). We start from the fact that any (possibly randomized) algorithm cannot identify the number of 1’s in the n -bit binary vector with success probability $\omega(n^{-1/2})$, if the algorithm is allowed to probe only $o(n)$ entries in the vector. A strengthened statement we will prove as the building block is the following lower bound for the “learning the bias” problem: given n Bernoulli arms (i.e., the stochastic reward of the arm is either 0 or 1), each of which has mean reward $(\mu + \epsilon)$ or $(\mu - \epsilon)$, then any algorithm using $o(n\epsilon^{-2} / \log(n/\epsilon))$ samples will not be able to identify the number of

two types of arms with probability $\omega(n^{-1/2})$.²

Now we explain the connection between the learning the bias problem and the top- m arm identification problem by sketching the plan of constructing the hard instances as follows. Suppose that we set all but $n^{1/2}$ arms in the hard instance to be Bernoulli with mean reward either $(\mu + \epsilon)$ (namely “the top arms”) or $(\mu - \epsilon)$ (namely “the bottom arms”). We denote the set of the rest $n^{1/2}$ arms by M , and their mean rewards are sandwiched between $(\mu + \epsilon)$ and $(\mu - \epsilon)$. We will set $m = n/2$, i.e., the goal is to identify the top half of the arms. Now, as long as the number of top arms, denoted by X , is bounded between $\frac{n}{2} - \sqrt{n}$ and $\frac{n}{2} + \sqrt{n}$, the goal is equivalent to identify the X top arms and the top- $(\frac{n}{2} - X)$ arms in M . We then vary the number of the top arms and consequently the number of the bottom arms (say, let X be uniformly randomly chosen from the range), and will argue that each agent will not be able to identify X much better than a random guess without communication, and therefore must perform one round of communication to learn X in order to identify the top- $(\frac{n}{2} - X)$ arms in M . Here, the need for communication is due to the lower bound for learning the bias and the fact that any agent in a \sqrt{K} -speedup algorithm is allowed to make only $O(n\epsilon^{-2}/\sqrt{K}) = o(n\epsilon^{-2})$ samples (where we make a crucial assumption that the $H^{(m)}$ complexity of the constructed hard instance is $O(n\epsilon^{-2})$). The last piece of plan is to argue that since X is not known before the first round of communication, each agent cannot make much progress before the communication towards identifying the top- $(\frac{n}{2} - X)$ arms in M , which is a necessary sub-task. We will finally inductively prove a communication lower bound for this sub-task. Note that the number of arms in M is $n^{1/2}$, and this plan will lead to a $\log \log n$ -style round complexity lower bound.

There are several challenges for the plan above. Note that in the sub-task for M , the goal is no longer to identify the top half arms, which is not well aligned with the (planned) induction hypothesis. Moreover, to make the induction work, M would naturally have the similar structure as the n -arm instance, i.e., with many top and bottom arms (possibly with different μ and ϵ parameters). However, such a construction would hardly ensure that the $H^{(m')}$ complexity is still $O(n\epsilon^{-2})$. Indeed, if the goal of the sub-task is to identify, for example, the top $|M|/4$ arms, since most of the top half arms are the same, the corresponding the $H^{(m')}$ complexity would become infinitely large. Finally, it is not clear how to make sure that any agent will not gain much information about M before the first round of communication so as to quickly identify the top $(\frac{n}{2} - X)$ arms in M whenever X is learned.

To address these challenges, we craft a more complex distribution of hierarchical instances. The main highlight is that we let M consist of multiple blocks I_1, I_2, \dots, I_k , where each block has the same number of arms and is independently sampled from a recursively defined hard distribution. We restrict the possible values of $(\frac{n}{2} - X)$ to be the half multiples of the block size so that the sub-task always becomes to identify the top half arms in I_ξ for some $\xi \in \{1, 2, \dots, k\}$. We will make careful selection of the block parameters so that the $H^{(m)}$ complexity for any instance in the support of the distribution, and the $H^{(m')}$ complexity of any sub-task, are all $\Theta(n\epsilon^{-2})$, where both upper and lower bounds are crucial to the proof.

The detailed construction and proof idea are presented in Section 4.2.1. For the analysis of the learning the bias problem and the ultimate round complexity lower bound, please refer to the subsequent subsections (Section 4.2.2 and Section 4.2.3).

²The sample complexity lower bound for a similar problem is proved in a recent work [44]. Our lower bound is different from theirs in two aspects. First, in their setting, the number of arms is not bounded and the goal is to estimate the fraction of the two types of arms up to an additive error, while in our setting, the number of arms is n , and the goal is to find out the exact numbers of arms for the two types. Second, their lower bound is for algorithms with constant success probability, while our lower bound is for algorithms with only $\omega(n^{-1/2})$ success probability.

3 A Collaborative Algorithm for the Fixed-Time Case

3.1 Preparation

In this section we give a few auxiliary lemmas to be used in our algorithm and analysis.

The following lemma gives connections between instance complexities and sub-instance complexities.

Lemma 1. *Let $V \subseteq I$ ($|I| = n$) be a subset of arms. Let $j \in \{1, \dots, n-1\}$ and $k \in \{1, \dots, |V|-1\}$ be two indices such that $\theta_{[k]}(V) \geq \theta_{[j]} \geq \theta_{[j+1]} \geq \theta_{[k+1]}(V)$. We have*

1. $H^{(k)}(V) \leq \sum_{i \in V} \left(\Delta_i^{(j)} \right)^{-2} \leq H^{(j)}$.
2. $H_\epsilon^{(k)}(V) \leq \sum_{i \in V} \max\{\Delta_i^{(j)}, \epsilon\}^{-2} \leq H_\epsilon^{(j)}$.

Proof. We start with the first item. The second inequality is straightforward by the definition of $H^{(j)}$. For the first inequality, for each arm $i \in V$, we consider two cases.

1. $\theta_i \geq \theta_{[k]}(V)$: It holds that $\Delta_i^{(k)}(V) = \theta_i - \theta_{[k+1]}(V) \geq \theta_i - \theta_{[j+1]} = \Delta_i^{(j)}$.
2. $\theta_i \leq \theta_{[k+1]}(V)$: It holds that $\Delta_i^{(k)}(V) = \theta_{[k]}(V) - \theta_i \geq \theta_{[j]} - \theta_i = \Delta_i^{(j)}$.

We thus have

$$H^{(k)}(V) = \sum_{i \in V} \left(\Delta_i^{(k)}(V) \right)^{-2} \leq \sum_{i \in V} \left(\Delta_i^{(j)} \right)^{-2}.$$

The second item follows from the same line of arguments. □

The next lemma gives a connection between instance complexities under different pivots.

Lemma 2. *For any $t \in \{1, \dots, n\}$, $H_{\Delta_{[t]}^{(t)}}^{(t)} \leq 4H^{(m)}$.*

Proof. We consider $t \leq m$ and $t > m$ separately. In the case that $t \leq m$, for any $i \in I$ we consider three cases.

1. $\theta_i < \theta_{[m]}$: In this case we have

$$\max\{\Delta_i^{(t)}, \Delta_{[t]}^{(m)}\} \geq \Delta_i^{(t)} = \theta_{[t]} - \theta_i \geq \theta_{[m]} - \theta_i = \Delta_i^{(m)}.$$

2. $\theta_{[m]} \leq \theta_i < \theta_{[m+1]} + 2\Delta_{[t]}^{(m)}$: In this case we have

$$\max\{\Delta_i^{(t)}, \Delta_{[t]}^{(m)}\} \geq \Delta_{[t]}^{(m)} \geq \frac{\theta_i - \theta_{[m+1]}}{2} \geq \frac{\Delta_i^{(m)}}{2}.$$

3. $\theta_{[m+1]} + 2\Delta_{[t]}^{(m)} \leq \theta_i$: In this case we have

$$\begin{aligned} \max\{\Delta_i^{(t)}, \Delta_{[t]}^{(m)}\} \geq \Delta_i^{(t)} &= \theta_i - \theta_{[t+1]} = \theta_i - \theta_{[m+1]} + \theta_{[m+1]} - \theta_{[t+1]} \\ &\geq \frac{\theta_i - \theta_{[m+1]}}{2} + \frac{\theta_i - \theta_{[m+1]}}{2} - (\theta_{[t+1]} - \theta_{[m+1]}) \\ &\geq \frac{\Delta_i^{(m)}}{2} + \Delta_{[t]}^{(m)} - (\theta_{[t+1]} - \theta_{[m+1]}) \geq \frac{\Delta_i^{(m)}}{2}. \end{aligned}$$

Therefore,

$$H_{\Delta_{[t]}^{(m)}}^{(t)} = \sum_{i \in I} \max \left\{ \Delta_i^{(t)}, \Delta_{[t]}^{(m)} \right\}^{-2} \leq \sum_{i \in I} 4 \left(\Delta_i^{(m)} \right)^{-2} \leq 4H^{(m)}.$$

In the case when $t > m$, the proof is symmetric by considering the following three cases: (1') $\theta_i > \theta_{[m+1]}$, (2') $\theta_{[m]} \geq \theta_i > \theta_{[m+1]} + 2\Delta_{[t]}^{(m)}$, and (3') $\theta_{[m+1]} + 2\Delta_{[t]}^{(m)} \geq \theta_i$. \square

The following simple fact gives an upper bound of the contribution (to the instance complexity) of an arm that is not very close to the pivot.

Lemma 3. *For any $t \in \{1, \dots, n\}$, if $t \leq m - z$ or $t \geq m + z$, then*

$$\left(\Delta_{[t]}^{(m)} \right)^{-2} \leq H^{(m)} / z.$$

Proof. In the case that $t \leq m - z$, there are at least z arms i such that $\theta_{[t]} \geq \theta_i \geq \theta_{[m]}$, or $\Delta_i^{[m]} \leq \Delta_{[t]}^{(m)}$. Consequently we have

$$H^{(m)} = \sum_{i \in I} \left(\Delta_i^{(m)} \right)^{-2} \geq z \cdot \left(\Delta_{[t]}^{(m)} \right)^{-2}.$$

The case $t \geq m + z$ can be proved in the same way. \square

We need two centralized algorithms `CentralApproxTop` and `CentralApproxBtm` for computing (ϵ, m) -top/bottom arms. We leave their detailed description to Section 3.5.1. The following lemma summarizes the guarantees of these two algorithms; it is a direct consequence of Lemma 26 and Lemma 27, which will be presented and proved in Section 3.5.1.

Lemma 4. *Let I be a set of n arms, $m \in \{1, \dots, n - 1\}$, and $\epsilon \in (0, 1)$ be an approximation parameter. Let*

$$T_1(I, a, \epsilon, \delta) = c_1 H_{\epsilon/2}^{(a)}(I) \cdot \log \left(H_{\epsilon/2}^{(a)}(I) / \delta \right)$$

for a universal constant c_1 . We have that

- If $T \geq T_1(I, m, \epsilon, \delta)$ then `CentralApproxTop`(I, m, T, δ) with probability at least $1 - \delta$, returns m arms each of which is (ϵ, m) -top in I using at most T time steps.
- If $T \geq T_1(I, n - m, \epsilon, \delta)$ then `CentralApproxBtm`(I, m, T, δ) with probability at least $1 - \delta$, returns m arms each of which is (ϵ, m) -bottom in I using at most T time steps.

The following lemma says that there is a simple collaborative algorithm `CollabTopMSimple` for top- m arm identification that uses $O(\log n)$ rounds of communication. Note that this bound is still much larger than our final target $O(\log \log m + \log K)$ rounds. `CollabTopMSimple` is a simple modification of a centralized algorithm in [12], and will be described in details in Section 3.5.2. Lemma 5 is a direct consequence of Lemma 28, which will be presented and proved in Section 3.5.2.

Lemma 5. *Let I be a set of n arms, and $m \in \{1, \dots, n - 1\}$. Let*

$$T_2(I, m, \delta) = c_2 \cdot \frac{H^{(m)}(I)}{K} \cdot \log n \cdot \log \frac{n}{\delta} \quad (4)$$

for a universal constant c_2 . There is a collaborative algorithm `CollabTopMSimple`(I, m, T) such that if $T \geq T_2(I, m, \delta)$ then with probability at least $1 - \delta$, one computes the set of top- m arms of I using at most T time steps and $O(\log n)$ rounds.

Algorithm 1: CollabTopM(I, m, T)

Input: a set of n arms I , parameter m , and time horizon T .
Output: the set of top- m arms of I .

- 1 Let R be the global upper bound on the number of rounds
- 2 and δ be also the global parameter equal to $1/(100R)$;
- 3 $q \leftarrow 4K\sqrt{n \log(nR)}$;
- 4 **if** $n > K^{10}$ **then**
- 5 $Acc \leftarrow \emptyset, Rej \leftarrow \emptyset$;
- 6 randomly assign each arm in I to one of the K agents, and let I_i be the set of arms assigned to i -th agent ³;
- 7 **if** $m > q$ **then**
- 8 $\ell \leftarrow (m - q)/K$;
- 9 **for agent** $i = 1$ **to** K **do**
- 10 $Acc_i \leftarrow \text{CentralApproxTop}(I_i, \ell, \frac{T}{4R}, \frac{\delta}{2K})$;
- 11 $Acc \leftarrow \bigcup_{i=1}^K Acc_i$;
- 12 **if** $n - m > q$ **then**
- 13 $r \leftarrow (n - m - q)/K$;
- 14 **for agent** $i = 1$ **to** K **do**
- 15 $Rej_i \leftarrow \text{CentralApproxBtm}(I_i, r, \frac{T}{4R}, \frac{\delta}{2K})$;
- 16 $Rej \leftarrow \bigcup_{i=1}^K Rej_i$;
- 17 **return** $Acc \cup \text{CollabTopM}(I \setminus (Acc \cup Rej), m - |Acc|, T)$;
- 18 **else**
- 19 **return** $\text{CollabTopMSimple}(I, m, T/2)$.

3.2 Special Time Horizon T

In this section we prove the following theorem concerning a special time horizon T .

Theorem 6. *Let I be a set of n arms, and $m \in \{1, \dots, n - 1\}$. Let*

$$T_0 = c_0 \cdot \frac{H^{(m)}}{K} \cdot \left(\log(H^{(m)} \cdot K) + \log^2 n \right) \cdot \log \log n \quad (5)$$

for a large enough constant c_0 . There exists a collaborative algorithm $\text{CollabTopM}(I, m, T)$ that computes the set of top- m arms of I with probability at least 0.99 when $T \geq T_0$, and uses at most T time steps and $O(\log \frac{\log n}{\log K} + \log K)$ rounds of communication.

Algorithm and Intuition. Our algorithm is described in Algorithm 1. Note that we have used *recursion* instead of *iteration* to omit a superscript r . But we still call each recursive step a *round*.

Let us briefly describe Algorithm 1 in words. At the beginning of each round we first randomly partition the set of arms to the K agents. Then each agent tries to identify a subset of arms Acc_i of size $\ell \approx (m/K - \sqrt{n})$ to be included to Top_m , and a subset of arms Rej_i of size $r \approx ((n - m)/K - \sqrt{n})$ to be pruned. The intuition to introduce the additive \sqrt{n} term is that by a concentration bound, we have with a

good probability that at least ℓ true top- m arms will be assigned to each agent, and similarly at least r non-top- m arms will be assigned to each agent. However, even with this fact, we still cannot guarantee that each agent can identify the top and bottom arms successfully given its limited budget, which is approximately $H^{(m)}/K$. Such a budget in some sense demands that the global instance complexity is *evenly* divided into the K agents, which is not necessary true. We thus adopt a PAC algorithm for top- m arm identification which returns a set of ℓ (ϵ, ℓ) -top arms at each agent A_i , where ϵ is a random variable which, with a high probability, is smaller than the gap between the ℓ -th top arm *locally* at A_i and that of the m -th *global* top arm. In this way we can guarantee that it is safe to include each Acc_i that A_i computes into Top_m . By essentially the same arguments, we can show that it is safe to prune the set of bottom arms Rej_i .

The following lemma is critical for the correctness of the algorithm.

Lemma 7. *if $T \geq T_0$ (defined in 5). When running Algorithm 1 $\text{CollabTopM}(I, m, T)$, we have that at each round,*

$$\Pr[(Acc \subseteq Top_m) \wedge (Rej \subseteq I \setminus Top_m)] \geq 1 - 1/(200R), \quad (6)$$

where Acc and Rej are defined in Algorithm 1.

Before proving Lemma 7, we first show that Lemma 7 implies Theorem 6.

Proof of Theorem 6. W.l.o.g. we assume that $K = \omega(\log \log n)$. Note that when $K = O(\log \log n)$, we have $T_0 \geq K \cdot T_2(I, m, 1/100)$, and thus each agent can simply solve the problem independently using the centralized algorithm in [12].

Let $m' = m - |Acc|$. If $Acc \subseteq Top_m$ and $Rej \subseteq I \setminus Top_m$, then we have

1. $Top_m = Acc \cup Top_{m'}(I \setminus (Acc \cup Rej))$.
2. $H^{(m')}(I \setminus (Acc \cup Rej)) \leq H^{(m)}$.

The first item is obvious. The second item is due to Lemma 1. The second item ensures that the recursion goes through under the same time horizon.

By the first item, Lemma 5 (note that $T/2 \geq T_2(I, m, 1/200)$), and a union bound, we have that with error at most $(1/(200R) \cdot R + 1/200) = 1/100$, Algorithm 1 computes Top_m .

Now we analyze the running time. Under the condition that at each round we have $Acc \subseteq Top_m$ and $Rej \subseteq I \setminus Top_m$, it follows that when $n > K^8$ and $K = \omega(\log \log n)$,

$$n - |Acc| - |Rej| \leq 2 \cdot 4K \sqrt{n \log(nR)} \leq n^{7/8}.$$

Therefore after $R = 10 \log \left(\frac{\log n}{10 \log K} \right)$ rounds, we have $n^{(7/8)^R} \leq K^{10}$. Consequently Algorithm 1 must have already reached Line 19. The algorithm CollabTopMSimple in Line 19 takes $O(\log K^{10}) = O(\log K)$ rounds by Lemma 28 (setting $R = \log n$ where $n = K^{10}$ here). Thus the total number of rounds is bounded by $O \left(\log \left(\frac{\log n}{\log K} \right) + \log K \right)$. \square

In the rest of this section we will prove Lemma 7.

Let $q = 4K \sqrt{n \log(nR)}$, $\ell = (m - q)/K$ and $r = ((n - m) - q)/K$ be defined in Algorithm 1. Further, define $a = m - q/2$ and $b = (n - m) - q/2$.

The following lemma concerns properties of the random partition in Line 6 in Algorithm 1.

³This randomness can be precomputed and stored at each agent.

Lemma 8. Let V be a random subset of I ($n > K^{10}$) by taking each arm independently with probability $1/K$. We have

1. If $m > q$, then $\Pr \left[(|V| \geq \ell) \wedge (\theta_{[\ell]}(V) \geq \theta_{[a]}) \right] \geq 1 - \frac{1}{1600KR}$.
2. If $n - m > q$, then $\Pr \left[(|V| \geq r) \wedge (\theta_{[|V|-r+1]}(V) \leq \theta_{[n-b+1]}) \right] \geq 1 - \frac{1}{1600KR}$.

Proof. We focus on the first item; the second item is symmetric and can be proved by similar arguments.

For each $i \in \{1, \dots, a\}$, we define a random variable X_i which is 1 if the arm with mean $\theta_{[i]}$ lies in the set V , and 0 otherwise. Let $X = \sum_{i=1}^a X_i$. Thus $\mathbb{E}[X] = a/K$. By Chernoff-Hoeffding (Lemma 41) we have

$$\Pr[X < \ell] = \Pr \left[X < \frac{a}{K} - \frac{q}{2K} \right] \leq \exp \left(\frac{-8n \log(nR)}{a} \right) \leq \frac{1}{1600KR}.$$

Thus with probability at least $(1 - \frac{1}{1600KR})$, V contains at least ℓ arms with mean at least $\theta_{[a]}$. \square

The next lemma connects the global instance complexity with the local instance complexity.

Lemma 9. Let V be a random subset of I ($n > K^{10}$) by taking each arm independently with probability $1/K$. Let P be a random variable such that $\theta_{[P]} = \theta_{[\ell]}(V)$, and Q be a random variable such that $\theta_{[Q]} = \theta_{[|V|-r+1]}(V)$. We have

1. If $m > q$, then $\Pr \left[(|V| \geq \ell) \wedge \left(H_{\Delta_{[P]}^{(m)}}^{(\ell)}(V) \leq 5H^{(m)}/K \right) \right] \geq 1 - \frac{1}{800KR}$.
2. If $n - m > q$, then $\Pr \left[(|V| \geq r) \wedge \left(H_{\Delta_{[Q]}^{(m)}}^{(r)}(V) \leq 5H^{(m)}/K \right) \right] \geq 1 - \frac{1}{800KR}$.

Proof. We only need to prove the first item. The second item follows by symmetry.

When $m > q$, let χ denote the event $(|V| < \ell) \vee (\theta_{[\ell]}(V) < \theta_{[a]})$. By Lemma 8 we have that χ happens with probability at most $1/(1600KR)$. We have

$$\begin{aligned} \Pr \left[(|V| < \ell) \vee \left(H_{\Delta_{[P]}^{(m)}}^{(\ell)}(V) > \frac{5H^{(m)}}{K} \right) \right] &\leq \Pr \left[\bar{\chi} \wedge \left(H_{\Delta_{[P]}^{(m)}}^{(\ell)}(V) > \frac{5H^{(m)}}{K} \right) \right] + \Pr[\chi] \\ &\leq \Pr \left[\bar{\chi} \wedge \left(H_{\Delta_{[P]}^{(m)}}^{(\ell)}(V) > \frac{5H^{(m)}}{K} \right) \right] + \frac{1}{1600KR}. \end{aligned} \quad (7)$$

Apply the law of total probability to the first term:

$$\Pr \left[\bar{\chi} \wedge \left(H_{\Delta_{[P]}^{(m)}}^{(\ell)}(V) > \frac{5H^{(m)}}{K} \right) \right] = \sum_{p=1}^a \Pr \left[\left(H_{\Delta_{[p]}^{(m)}}^{(\ell)}(V) > \frac{5H^{(m)}}{K} \right) \wedge (P = p) \right] \quad (8)$$

$$\begin{aligned} &\leq \sum_{p=1}^a \Pr \left[\left(\sum_{i \in V} \max \{ \Delta_i^{(p)}, \Delta_{[p]}^{(m)} \}^{-2} > \frac{5H^{(m)}}{K} \right) \wedge (P = p) \right] \\ &= \sum_{p=1}^a \Pr \left[\sum_{i \in V} \max \{ \Delta_i^{(p)}, \Delta_{[p]}^{(m)} \}^{-2} > \frac{5H^{(m)}}{K} \right], \end{aligned} \quad (9)$$

where the inequality is due to Lemma 1.

For each $i \in I$, define $X_i = \max \left\{ \Delta_i^{(p)}, \Delta_{[p]}^{(m)} \right\}^{-2}$ if $i \in V$, and $X_i = 0$ otherwise. Let $X = \sum_{i \in I} X_i$. We thus have $\mathbb{E}[X] = H_{\Delta_{[p]}^{(m)}}^{(p)}/K$. By Lemma 2 we have

$$\mathbb{E}[X] \leq \frac{4H^{(m)}}{K}. \quad (10)$$

By Lemma 3 and $p \leq a = m - q/2$, we have

$$X_i = \left[0, H^{(m)} / \frac{q}{2} \right]. \quad (11)$$

By (10), (11), and Chernoff-Hoeffding we have

$$\begin{aligned} \Pr \left[X > 5H^{(m)}/K \right] &= \Pr \left[X > \mathbb{E}[X] + H^{(m)}/K \right] \\ &\leq \exp \left(\frac{-2 \left(H^{(m)}/K \right)^2}{n \left(2H^{(m)}/q \right)^2} \right) = \exp \left(\frac{-q^2}{2K^2n} \right) \\ &\leq \frac{1}{1600nKR}. \end{aligned}$$

We thus have (9) $\leq a \cdot 1/(1600nKR) \leq 1/(1600KR)$, which, together with (7), gives the first item of the lemma. \square

Now we are ready to prove Lemma 7.

Proof of Lemma 7. We first analyze the probability that $Acc \subseteq Top_m$.

Again let P be the random variable such that $\theta_{[P]} = \theta_{[\ell]}(I_i)$ for a partition I_i . Since at Line 10 of Algorithm 1 we call `CentralApproxTop` with time budget $T/4R$, we have that if $H_{\Delta_{[P]}^{(m)}}^{(\ell)}(V) \leq 5H^{(m)}/K$ then

$$\frac{T}{4R} \geq T_1 \left(I_i, \ell, \Delta_{[P]}^{(m)}/4, 1/(800KR) \right),$$

and with probability at least $1/(800KR) \cdot K = 1/(800R)$, it holds that `CentralApproxTop` $(I_i, \ell, \frac{T}{4R})$ succeeds for all $i \in \{1, \dots, K\}$, which, together with the first item of Lemma 9, gives that

$$\Pr[Acc \subseteq Top_m] \geq 1 - \frac{1}{400R}.$$

By symmetric arguments we can also show that $\Pr[Rej \subseteq I \setminus Top_m] \geq 1 - 1/(400R)$. \square

3.3 General Time Horizon T

Theorem 6 only achieves a constant error probability for a special case of the time horizon $T = \tilde{\Theta}(T_0)$ where $T_0 = H^{(m)}/K$. Our next goal is to consider general time horizon $T \geq T_0$, and try to make the error probability decrease exponentially with respect to T/T_0 . More precisely, we have the following theorem.

Theorem 10. *Let I be a set of n arms, and $m \in \{1, \dots, n-1\}$. Let T be a time horizon. There exists a collaborative algorithm `CollabTopMGeneral` that computes the set of top- m arms of I with probability at least*

$$1 - n \cdot \exp \left(-\Omega \left(\frac{TK}{H^{(m)} \cdot (\log(H^{(m)}K) + \log^2 n) \cdot \log \log n \cdot \log^2 (TK/H^{(m)})} \right) \right)$$

using at most T time steps and $O \left(\log \frac{\log n}{\log K} + \log K \right)$ rounds.

High Level Idea. A standard technique to achieve an error probability that is exponentially small in terms of T/T_0 is to perform parallel repetition and then take the majority. This is straightforward if we know the value T_0 . Unfortunately, T_0 depends on the instance complexity which we do not know in advance. A standard trick to handle this issue is to use the doubling method. That is, we guess $T_0 = 1, 2, 4, \dots$, and for each value we repeat T/T_0 times (ignoring logarithmic factors). We know that one of these values is very close to the actual T_0 . We hope that this value is the *first* value in $\{1, 2, 4, \dots\}$ for which the T/T_0 runs of `CollabTopM`(I, m, T) contain a majority output.

The main issue in this approach is that when $T \leq T_0$, the output of the algorithm can be *consistently* wrong, which leads to a wrong majority. Note that we do not have much control on the output of the algorithm when the time horizon is very small.

We handle this issue by introducing a concept called *top- m certificate*. We require each algorithm for top- m arm identification to output a pair $(S, \{\tilde{\theta}_i\}_{i \in I})$, where S is a subset of I of size m and $\{\tilde{\theta}_i\}_{i \in I}$ are the estimated means for *all* arms in I (not just those in S). We say a pair $(S, \{\tilde{\theta}_i\}_{i \in I})$ is a top- m certificate if it can pass an additional verification step which checks whether S is indeed the set of top- m arms of I given the estimated means $\{\tilde{\theta}_i\}_{i \in I}$. With such a verification step at hand, we do not need to worry about the case that `CollabTopM` will output a wrong answer when T is too small, since a wrong output will simply *not* pass the verification step. Finally, we make sure that this verification step is perfectly parallelizable and thus fit in our time budget.

In the rest of this section we will first give the definition of the top- m certificate and describe the verification algorithm, and then give the collaborative algorithm for general time horizon T .

3.3.1 Top- m Certificate

Definition 2 (Top- m Certificate). *Let $S \subseteq I$. We say $(S, \{\tilde{\theta}_i\}_{i \in I})$ is a top- m certificate of I if*

$$(S = \text{Top}_m) \wedge \left(\forall i \in I : \left| \tilde{\theta}_i - \theta_i \right| < \Delta_i^{(m)}/4 \right).$$

Observation 11. *If $(S, \{\tilde{\theta}_i\}_{i \in I})$ is a top- m certificate, then for any $i \in S$ and $j \in I \setminus S$, we have $\tilde{\theta}_i > \tilde{\theta}_j$.*

Given an arbitrary pair $(S, \{\tilde{\theta}_i\}_{i \in I})$, we can design an algorithm to verify whether $(S, \{\tilde{\theta}_i\}_{i \in I})$ is a top- m certificate of I ; see Algorithm 2 `VerifyTopM`. We note that Algorithm 2 can be easily implemented in $O(1)$ rounds since the number of pulls on each arm is determined in advance and can thus be fully parallelised.

The following lemma shows that if $(S, \{\tilde{\theta}_i\}_{i \in I})$ is indeed a certificate and T is large enough, then `VerifyTopM` returns the set S with a good probability.

Algorithm 2: $\text{VerifyTopM}(I, m, S, \{\tilde{\theta}_i\}_{i \in I}, \gamma, T)$

Input: a set of n arms I , parameter m , pair $(S, \{\tilde{\theta}_i\}_{i \in I})$, parameter γ , and time horizon T .

Output: the set of top- m arms, or \perp .

- 1 Set $\ell \leftarrow \arg \min_{i \in S} \tilde{\theta}_i$ and $r \leftarrow \arg \max_{i \in I \setminus S} \tilde{\theta}_i$;
 - 2 for each $i \in S$ set $\Delta_i \leftarrow \tilde{\theta}_i - \tilde{\theta}_r$, and for $i \in I \setminus S$ set $\Delta_i \leftarrow \tilde{\theta}_\ell - \tilde{\theta}_i$;
 - 3 **if** $\exists i$ s.t. $\Delta_i \leq 0$ or $|S| \neq m$ **then return** \perp ;
 - 4 **if** $\sum_{i \in I} (64\gamma\Delta_i^{-2}) \leq KT$ **then**
 - 5 pull i -th arm for $64\gamma\Delta_i^{-2}$ times and let $\hat{\theta}_i$ be the empirical mean;
 - 6 **if** $\min_{i \in S} \{\hat{\theta}_i - \Delta_i/4\} > \max_{i \in I \setminus S} \{\hat{\theta}_i + \Delta_i/4\}$ **then return** S ;
 - 7 **return** \perp .
-

Lemma 12. For any $I, m, S, \{\tilde{\theta}_i\}_{i \in I}, \gamma, T$, we have

$$\Pr[\text{VerifyTopM}(I, m, S, \{\tilde{\theta}_i\}_{i \in I}, \gamma, T) \notin \{Top_m, \perp\}] \leq 2n \cdot e^{-\gamma}. \quad (12)$$

Moreover, if $(S, \{\tilde{\theta}_i\}_{i \in I})$ is a top- m certificate of I , and $T \geq 200\gamma H^{(m)}/K$ then

$$\Pr[\text{VerifyTopM}(I, m, S, \{\tilde{\theta}_i\}_{i \in I}, \gamma, T) \neq Top_m] \leq 2n \cdot e^{-\gamma}. \quad (13)$$

Proof. By Chernoff-Hoeffding, for any $i \in I$, after $64\gamma\Delta_i^{-2}$ pulls, we have $\Pr[|\hat{\theta}_i - \theta_i| > \Delta_i/8] \leq 2e^{-\gamma}$. By a union bound we have

$$\Pr[\exists i : |\hat{\theta}_i - \theta_i| > \Delta_i/8] \leq 2n \cdot e^{-\gamma}. \quad (14)$$

We now show that if

$$\forall i : |\hat{\theta}_i - \theta_i| \leq \Delta_i/8, \quad (15)$$

then Algorithm 2 can only output S when $S = Top_m$.

We prove by contradiction. Suppose Algorithm 2 outputs S when $S \neq Top_m$, then there must exist a pair (i, j) such that $i \in S \setminus Top_m$ and $j \in Top_m \setminus S$, and consequently $\theta_i < \theta_j$. Meanwhile, by Line 6 of Algorithm 2 we have

$$\hat{\theta}_i - \Delta_i/4 > \hat{\theta}_j + \Delta_j/4. \quad (16)$$

Combining (15) and (16) we have

$$\theta_i > \theta_i - \Delta_i/8 > \theta_j + \Delta_j/8 > \theta_j,$$

which contradicts to the choices of (i, j) . This proves (12).

We next prove (13). If $(S, \{\tilde{\theta}_i\}_{i \in I})$ is a top- m certificate, then by (15) we have

$$\Delta_i^{(m)}/2 \leq \Delta_i \leq 3\Delta_i^{(m)}/2.$$

Thus if $T \geq 200\gamma H^{(m)}/K$, then

$$\sum_{i \in I} (64\gamma\Delta_i^{-2}) \leq 200\gamma \sum_{i \in I} (\Delta_i^{(m)})^{-2} \leq KT.$$

We thus only need to show that

$$\min_{i \in Top_m} \{\hat{\theta}_i - \Delta_i/4\} > \max_{i \notin Top_m} \{\hat{\theta}_i + \Delta_i/4\}. \quad (17)$$

We again prove by contradiction. Suppose that there exists a pair (i, j) such that $i \in Top_m$, $j \notin Top_m$, and $\hat{\theta}_i - \Delta_i/4 \leq \hat{\theta}_j + \Delta_j/4$, then by (15) we have

$$\theta_i - \theta_j \leq \frac{3}{8}(\Delta_i + \Delta_j) \leq \frac{3}{8}(\Delta_i^{(m)} + \Delta_j^{(m)}), \quad (18)$$

which contradicts to the fact that $\frac{1}{2}(\Delta_i^{(m)} + \Delta_j^{(m)}) \leq \theta_i - \theta_j$. \square

For technical reasons we need the following lemma, which says that `VerifyTopM` is very likely to output \perp when the time horizon T is small.

Lemma 13. *For any $I, m, S, \{\tilde{\theta}_i\}_{i \in I}$, and γ , if $T < \frac{\gamma}{16}H^{(m)}/K$, then*

$$\Pr[\text{VerifyTopM}(I, m, S, \{\tilde{\theta}_i\}_{i \in I}, \gamma, T) = \perp] \geq 1 - 2n \cdot \exp(-\gamma).$$

Proof. We can assume that $Top_m = S$, and for any $i \in S$ and $j \notin S$ we have $\tilde{\theta}_i > \tilde{\theta}_j$, since otherwise `VerifyTopM` will simply output \perp .

If for all i we have $\Delta_i \leq 32\Delta_i^{(m)}$, then $\sum_{i \in I} \Delta_i^{-2} \geq H^{(m)}/1024$, and consequently

$$\sum_{i \in I} (64\gamma\Delta_i^{-2}) \geq \frac{\gamma}{16}H^{(m)} > KT.$$

In this case, according to Line 4 of Algorithm 2, `VerifyTopM` will output \perp .

Now consider the case that there exists i such that $\Delta_i > 32\Delta_i^{(m)}$. We first consider the case $i \in S$. By (14) we have that with probability $1 - 2n \cdot \exp(-\gamma)$, the event holds: $\forall i \in I : |\hat{\theta}_i - \theta_i| < \Delta_i/8$; denote this event by \mathcal{E}_1 . Consequently,

$$\hat{\theta}_i - \Delta_i/4 \leq \theta_i - \Delta_i/8 \leq \theta_i - 4\Delta_i^{(m)}. \quad (19)$$

Consider j such that $\theta_j = \theta_{[m+1]}$. When \mathcal{E}_1 holds, we have

$$\hat{\theta}_j + \Delta_j/4 \geq \theta_j. \quad (20)$$

Note that at Line 4, Algorithm 2 returns a set only if $\hat{\theta}_i - \Delta_i/4 > \hat{\theta}_j + \Delta_j/4$, which, by (19) and (20), is equivalent to $\theta_i - 4\Delta_i^{(m)} > \theta_j$. We now have

$$\Delta_i^{(m)} = \theta_i - \theta_j > 4\Delta_i^{(m)}.$$

A contradiction.

The case that $i \in I \setminus S$ can be proved by essentially the same arguments. \square

Algorithm 3: CollabTopMGeneral(I, m, T)

Input: a set of n arms I , parameter m , and time horizon T .

Output: the set of top- m arms of I .

```
1 for  $s = 1, 2, \dots$  do
2   run  $4^s$  copies of CollabTopM( $I, m, 3T/(\pi^2 s^2 4^s)$ ), and record the returned pair  $(S, \{\tilde{\theta}_i\}_{i \in I})$ ;
3   let  $S^{(s)}$  be the most frequent answer of the  $4^s$  returned sets of top- $m$  arms, and for each  $i \in [n]$ ,
   let  $\tilde{\theta}_i^{(s)}$  be the median of the  $4^s$  estimated means for the  $i$ -th arm;
4   run VerifyTopM( $I, m, S^{(s)}, \tilde{\theta}^{(s)}, 4^s, 3T/(\pi^2 s^2)$ ) and let  $A_s$  be the output ;
5 let  $s^*$  be the maximum  $s$  such that  $A_s \neq \perp$ ;
6 if we cannot find such an  $s^*$  after  $T$  time steps then
7   return an arbitrary set of  $m$  arms;
8 else
9   return  $A_{s^*}$ .
```

3.3.2 Algorithm for General Time Horizon

In this section we present an algorithm for general time horizon. We first slightly augment Algorithm 1 CollabTopM so that it also outputs an estimate of the mean (i.e., the empirical mean) of each of the n arms. Now the output of CollabTopM is a top- m certificate $(S, \{\tilde{\theta}_i\}_{i \in I})$.

We have the following lemma regarding CollabTopM. The proof is straightforward based on the properties of CentralApproxTop and CentralApproxBtm, which can be found in Section 3.5.

Lemma 14. *If $T \geq T_0$, then CollabTopM(I, m, T) is able to output a top- m certificate of I with probability at least 0.99.*

Our final algorithm for the general time horizon is described in Algorithm 3. It follows the guess-and-verify framework mentioned earlier. Now we are ready to prove the main theorem of this section.

Proof of Theorem 10. Let s' be the largest $s \geq 1$ such that $3T/(\pi^2 s^2 4^s) \geq T_0$, and consequently $4^{s'} = \Omega\left(\frac{T/T_0}{\log^2(T/T_0)}\right)$. If no such s exists then Theorem 10 holds trivially. By Lemma 14 and the standard median trick we have that with probability at least $1 - \exp(-4^{s'})$, $(S^{(s')}, \tilde{\theta}^{(s')})$ is a top- m certificate. We also have

$$3T/(\pi^2 (s')^2) \geq 200 \cdot 4^{s'} \cdot \frac{H^{(m)}}{K},$$

which, combined with Lemma 12, guarantees that the call of VerifyTopM at Line 4 in Algorithm 3 returns Top_m with probability at least $1 - 2n \cdot \exp(-4^{s'})$. By a union bound, we have that

$$\Pr[A_{s'} = Top_m] \geq 1 - 4n \cdot \exp(-4^{s'}). \quad (21)$$

On the other hand, for each $s = s' + j$ ($j \geq 1$), by Lemma 12 we have

$$\Pr[A_{s'} \notin \{Top_m, \perp\}] \leq 2n \cdot \exp(-4^{s'+j}). \quad (22)$$

Combining (21) and (22), we have

$$\begin{aligned} \Pr[A_{s^*} = \text{Top}_m] &\geq \Pr[A_{s'} = \text{Top}_m] - \sum_{j=s'+1}^{\infty} \Pr[A_j \notin \{\text{Top}_m, \perp\}] \\ &\geq 1 - 4n \sum_{s \geq s'} \exp(-4^s) \geq 1 - 8n \cdot \exp(-4^{s'}). \end{aligned}$$

Plugging the fact that $4^{s'} = \Omega\left(\frac{T/T_0}{\log^2(T/T_0)}\right)$, we have $\Pr[A_{s^*} = \text{Top}_m] \geq 1 - n \cdot \exp\left(-\Omega\left(\frac{T/T_0}{\log^2(T/T_0)}\right)\right)$.

Finally, it is easy to see that `CollabTopMGeneral` can be implemented in $O(\log \frac{\log n}{\log K} + \log K)$ rounds of communication by Theorem 6, because all runs of `CollabTopM` can be done in parallel and `VerifyTopM` requires only the constant number of rounds of communication. \square

3.4 An Improved Algorithm

In this section we further improve the round complexity of Algorithm 1 to $O(\log \frac{\log m}{\log K} + \log K)$.

High Level Idea. The general idea to improve the $\log \log n$ term to $\log \log m$ in the round complexity is to first reduce the number of arms from n to $\tilde{O}(m)$. This idea is relatively easy to implement if we only target for a constant error probability: We sample each of the n arms with probability $1/m$, getting a subset V . By an easy calculation, the event that V contains *exactly* one top- m arms of I happens with a constant probability. Conditioned on this event, the expected sub-instance complexity $H^{(1)}(V)$ is upper bounded by $O(H^{(m)}(I)/m)$. Thus if we sample $\tilde{O}(m)$ sub-instances, and compute the best arm in each sub-instance, then the collection of the $\tilde{O}(m)$ best arms will be a *superset* of Top_m with a good probability.

As before, our ultimate goal is to make the error probability exponentially small in terms of T . To this end we also need to amplify the success probability of the aforementioned reduction. Unfortunately, using a verification step as that in Section 3.3 is not sufficient here, and once again we need new ideas. We note that this section is the most technically challenging part in the whole algorithm design of this paper.

To start with, we again try to guess the sub-instance complexity using a geometric sequence, and use parallel repetition to amplify the success probability. Similar as Section 3.3.2, the key is to avoid outputting a wrong majority when the guess is too small. This is possible if the following statement holds:

For a randomly sampled sub-instance on which we compute the best arm, the probability of outputting any arm inside global Top_m is larger than that of any arm outside Top_m .

Unfortunately, this statement does not always hold. We thus try to prove a slightly weaker version of the statement. We show that the probability of outputting any arm inside global Top_m is *not much smaller* than that of any arm outside Top_m .

To show this, we consider three cases. Let V be the randomly sampled sub-instance. Let T_0 be our guessed complexity of V , and $H = H^{(1)}(V)$ be the real expected complexity of V .

1. If $T_0 \geq \lambda H$ where $\lambda = \log^{\Theta(1)}(T_0 K n)$, then if we apply the best arm identification algorithm from [55] on V with time budget T , with a good probability we can successfully identify the best arm. Note that since our subsampling is uniform, for any pair of two arms a, b where $a \in \text{Top}_m, b \notin \text{Top}_m$, they will be included in the sub-instance with equal probability. Therefore the probability of outputting a must be at least that of outputting b .

2. If $T_0 \leq H$, then by applying [55] on V with time budget T_0 , together with a verification step (Algorithm 2), we can detect with a good probability that our guess T_0 is too small compared with H .
3. To handle the gap $H < T_0 < \lambda H$, we apply the following trick: We replace T_0 by T_0 or T_0/λ with equal probability, so as to “reduce” this case to either the first or the second case, for which we know how to handle. In this way we can show that for $a \in \text{Top}_m$, $b \notin \text{Top}_m$, the probability of outputting a is at least *half* that of outputting b . We note that this intuitive description is not entirely precise, but it conveys the idea.

In this section we prove the following theorem.

Theorem 15. *Let I be a set of n arms, $m \in \{1, \dots, n-1\}$, and T be the time horizon. There exists a collaborative algorithm that computes the set of top- m arms of I with probability at least*

$$1 - n \cdot \exp \left(-\Omega \left(\frac{KT}{H^{(m)} \log^6(KT) \log^2(KT/H^{(m)}) \log n} \right) \right),$$

using at most T time steps and $O(\log \frac{\log m}{\log K} + \log K)$ rounds.

3.4.1 Subsampling and Its Properties

We need the following technical lemma.

Lemma 16. *Let V be a set of n arms, and η be the time horizon. There exists a collaborative algorithm $\mathcal{A}(V, \delta, \eta)$, and functions*

$$f(V, \delta, \eta) = C_f \cdot \frac{H^{(1)}(V)}{K} \cdot \log^3 \eta K \cdot \log \frac{|V|}{\delta}, \quad (23)$$

and

$$g(V, \delta) = C_g \cdot \frac{H^{(1)}(V)}{K} \cdot \log \frac{|V|}{\delta}, \quad (24)$$

where C_f and C_g are universal constants such that

1. \mathcal{A} returns an arm from V or \perp using at most η time steps and $O(\log K)$ rounds of communication;
2. $\Pr[\mathcal{A}(V, \delta, \eta) \notin \{\text{Top}_1(V), \perp\}] \leq \delta$;
3. If $\eta \geq f(V, \delta, \eta)$, then $\Pr[\mathcal{A}(V, \delta, \eta) = \text{Top}_1(V)] \geq 1 - \delta$;
4. If $\eta \leq g(V, \delta)$, then $\Pr[\mathcal{A}(V, \delta, \eta) = \perp] \geq 1 - \delta$.

We will show that the algorithm for fixed-time best arm identification in [55] (denoted by $\mathcal{B}(V, T)$, where V is the input and T is the time horizon), combined with Algorithm 2 (setting $m = 1$), satisfies Lemma 16.

We first recall the following result from [55].

Lemma 17 ([55]). *For any V and T and fixed R , $\mathcal{B}(V, T)$ uses in at most T time steps and R rounds of communication, and returns the best arm with probability at least*

$$1 - |V| \exp \left(-\Omega \left(\frac{TK^{(R-1)/R}}{H^{(1)}(V) \cdot \log^3(TK)} \right) \right).$$

Algorithm 4: $\mathcal{A}(V, \delta, T)$

- 1 Run $\mathcal{B}(V, T/2)$, and let the pair $(\{a\}, \tilde{\theta})$ be the certificate produced by \mathcal{B} ;
 - 2 **return** $\text{VerifyTopM}(V, 1, \{a\}, \tilde{\theta}, \log(|V|/\delta), T/2)$.
-

Algorithm 5: $\mathcal{A}'(I, m, \delta, T)$

Input: a set of n arms I , parameters m and δ , and time horizon T .

Output: the top-1 arm in a subset of I obtained by randomly sampling each arm in I with probability $1/m$.

- 1 Sample each element from I independently with probability $1/m$; let V be the sampled subset ;
 - 2 choose $\tau \in \{T, T/\beta\}$ uniformly at random;
 - 3 **return** $\mathcal{A}(V, \delta, \tau)$.
-

In Algorithm 4 we describe how to augment $\mathcal{B}(V, T)$ with a verification step to construct $\mathcal{A}(V, \delta, T)$. By Lemma 17, Lemma 12, Lemma 13, and setting $R = \log K$, $\mathcal{A}(V, \delta, T)$ satisfies all the four items of Lemma 16.

By (23) and (24), we have

$$\frac{f(V, \delta, \eta)}{g(V, \delta)} \leq \frac{C_f}{C_g} \cdot \log^3 \eta K .$$

Define

$$\beta \triangleq \frac{C_f}{C_g} \cdot \log^3 TK .$$

We now design another algorithm \mathcal{A}' for finding the best arm in a random subset of I , using Algorithm 4 as a subroutine. \mathcal{A}' is described in Algorithm 5. The following lemma says that any arm in the top- m arms of I will be returned with a good probability by \mathcal{A}' .

Lemma 18. For any I , $\delta \in (0, \frac{1}{2})$ and arm $a \in \text{Top}_m$, if

$$T \geq c_{\mathcal{A}} \cdot \frac{H^{(m)}}{\delta K m} \cdot \log^6(TK) \cdot \log \frac{n}{\delta} \quad (25)$$

for a universal constant $c_{\mathcal{A}}$, then we have

$$\Pr [\mathcal{A}'(I, m, \delta, T) = a] \geq \frac{(1 - \delta)^2}{em} . \quad (26)$$

Proof. We first note that for each $a \in \text{Top}_m$ and $V \subseteq I$ sampled at Line 1 in Algorithm 5, if $f(V, \delta, \tau) \leq T/\beta$ and $\text{Top}_m \cap V = \{a\}$, then by Item 3 of Lemma 16 we have $\Pr[\mathcal{A}'(I, m, \delta, T) = a] \geq 1 - \delta$.

According to the uniform subsampling, we have

$$\Pr[\text{Top}_m \cap V = \{a\}] = \frac{1}{m} \cdot \left(1 - \frac{1}{m}\right)^{m-1} \geq \frac{1}{em} . \quad (27)$$

For every $i \in I$, let $X_i = 1$ if $i \in V$ and $X_i = 0$ otherwise. For any V such that $V \cap \text{Top}_m = \{a\}$, we have

$$H^{(1)}(V) \leq 2 \sum_{i \in I \setminus \text{Top}_m} (\theta_a - \theta_i)^{-2} X_i \leq 2 \sum_{i \in I \setminus \text{Top}_m} X_i \left(\Delta_i^{(m)}\right)^{-2} .$$

We thus have

$$\mathbb{E}[H^{(1)}(V) \mid \text{Top}_m \cap V = \{a\}] \leq \frac{2}{m} H^{(m)}.$$

By a Markov inequality, we have that conditioned on $\text{Top}_m \cap V = \{a\}$, with probability $(1 - \delta)$,

$$H^{(1)}(V) \leq \frac{2}{\delta m} H^{(m)} \quad (28)$$

and so on

$$f(V, \delta, \tau) \leq C_f \frac{H^{(1)}(V)}{K} \cdot \log^3 TK \cdot \log \frac{n}{\delta} \leq C_f \frac{2H^{(m)}}{\delta Km} \cdot \log^3 TK \cdot \log \frac{n}{\delta} \leq \frac{T}{\beta} \leq \tau, \quad (29)$$

under which

$$\mathcal{A}(V, \delta, \tau) \text{ outputs } a \text{ with probability } (1 - \delta) \quad (30)$$

according to Item 3 of Lemma 16.

Finally, we have

$$\begin{aligned} \Pr[\mathcal{A}'(I, m, \delta, T) = a] &= \Pr[\mathcal{A}(V, \delta, \tau) = a] \\ &\geq \Pr[\mathcal{A}(V, \delta, \tau) = a \mid H^{(1)}(V) \leq \tau, \text{Top}_m \cap V = \{a\}] \\ &\quad \cdot \Pr[H^{(1)}(V) \leq \tau \mid \text{Top}_m \cap V = \{a\}] \cdot \Pr[\text{Top}_m \cap V = \{a\}] \\ &\geq (1 - \delta) \cdot (1 - \delta) \cdot 1/(em) = (1 - \delta)^2/(em), \end{aligned}$$

where the last inequality is due to (27), (28) and (30). \square

The next lemma is critical. It says that the probability that Algorithm 5 outputs any arm from Top_m cannot be significantly smaller than that of outputting any arm from $I \setminus \text{Top}_m$.

Lemma 19. For any arms $a \notin \text{Top}_m$, $b \in \text{Top}_m$ and $\delta \in (0, \frac{1}{2})$ we have

$$\Pr[\mathcal{A}'(I, m, \delta, T) = a] \leq \Pr[\mathcal{A}'(I, m, \delta, T) = b] + \left(\frac{1}{2em} + \frac{2\delta}{m} \right).$$

Proof. Let us consider a pair of random sets (U, V) , where V is formed by picking each of the n arms of I with probability $1/m$, and U is the set of V after exchanging the assignments of a and b . By this construction we have that

$$\text{Top}_1(V) = a \Rightarrow \text{Top}_1(U) = b. \quad (31)$$

Moreover, it is easy to see that the marginal distribution of U is *identical* to that of V .

Note that $\mathcal{A}'(I, m, \delta, T) = \mathcal{A}(V, \delta, \tau)$. We thus only need to prove

$$\Pr[\mathcal{A}(V, \delta, \tau) = a] \leq \Pr[\mathcal{A}(V, \delta, \tau) = b] + \left(\frac{1}{2em} + \frac{2\delta}{m} \right). \quad (32)$$

We start from the left hand side.

$$\begin{aligned} \Pr[\mathcal{A}(V, \delta, \tau) = a] &= \Pr[(\mathcal{A}(V, \delta, \tau) = a) \wedge (\text{Top}_1(V) = a)] \\ &\quad + \Pr[(\mathcal{A}(V, \delta, \tau) = a) \wedge (\text{Top}_1(V) \neq a)] \\ &\leq \Pr[(\mathcal{A}(V, \delta, \tau) = a) \wedge (\text{Top}_1(V) = a)] \\ &\quad + \Pr[(a \in V) \wedge \mathcal{A}(V, \delta, \tau) \notin \{\text{Top}_1(V), \perp\}] \\ &\leq \Pr[(\mathcal{A}(V, \delta, \tau) = a) \wedge (\text{Top}_1(V) = a)] + \frac{\delta}{m}, \end{aligned} \quad (33)$$

where the last inequality follows from the fact that $\Pr[a \in V] = 1/m$ and Item 2 of Lemma 16.

By (31) we have

$$\Pr[(\mathcal{A}(U, \delta, \tau) = b) \wedge (\text{Top}_1(V) = a)] \leq \Pr[(\mathcal{A}(U, \delta, \tau) = b) \wedge (\text{Top}_1(U) = b)]. \quad (34)$$

Since the marginal distribution of U and V are identical, we have

$$\Pr[A(U, \delta, \tau) = b \wedge \text{Top}_1(U) = b] \leq \Pr[A(U, \delta, \tau) = b] = \Pr[A(V, \delta, \tau) = b]. \quad (35)$$

The following claim enables us to connect $\Pr[\mathcal{A}(V, \delta, \tau) = a]$ and $\Pr[\mathcal{A}(V, \delta, \tau) = b]$.

Claim 20. For any arm $a \notin \text{Top}_m$, $b \in \text{Top}_m$, and $\delta \in (0, \frac{1}{2})$, we have

$$\Pr[(\mathcal{A}(V, \delta, \tau) = a) \wedge (\text{Top}_1(V) = a)] \leq \Pr[(\mathcal{A}(U, \delta, \tau) = b) \wedge (\text{Top}_1(V) = a)] + \frac{1 + \delta}{2em}. \quad (36)$$

We will prove Claim 20 shortly. By (33), (36), (34), (35) (one for each inequality below; in order),

$$\begin{aligned} \Pr[\mathcal{A}(V, \delta, \tau) = a] &\leq \Pr[(\mathcal{A}(V, \delta, \tau) = a) \wedge (\text{Top}_1(V) = a)] + \frac{\delta}{m} \\ &\leq \Pr[(\mathcal{A}(U, \delta, \tau) = b) \wedge (\text{Top}_1(V) = a)] + \frac{1 + \delta}{2em} + \frac{\delta}{m} \\ &\leq \Pr[(\mathcal{A}(U, \delta, \tau) = b) \wedge (\text{Top}_1(U) = b)] + \left(\frac{1}{2em} + \frac{2\delta}{m} \right) \\ &\leq \Pr[\mathcal{A}(V, \delta, \tau) = b] + \left(\frac{1}{2em} + \frac{2\delta}{m} \right). \end{aligned}$$

This proves Inequality (32) and gives the lemma. \square

Proof of Claim 20. We first note that if $\text{Top}_1(V) = a$, then by our construction of U , the fact that $a \notin \text{Top}_m$ and $b \in \text{Top}_m$, and the monotonicity of functions f and g , we have

$$f(U, \delta, \tau) \leq f(V, \delta, \tau) \text{ and } g(U, \delta) \leq g(V, \delta). \quad (37)$$

We expand the left hand side of (36) by running over all subsets $S \subseteq I$.

$$\Pr[(\mathcal{A}(V, \delta, \tau) = a) \wedge (\text{Top}_1(V) = a)] = \sum_{S: \text{Top}_1(S)=a} \Pr[(\mathcal{A}(V, \delta, \tau) = a) \wedge (V = S)]. \quad (38)$$

For each $V = S$ where $\text{Top}_1(S) = a$, we analyze the quantity

$$\Psi \triangleq \Pr[(\mathcal{A}(V, \delta, \tau) = a) \wedge (V = S)] - \Pr[(\mathcal{A}(U, \delta, \tau) = b) \wedge (V = S)]$$

in two cases which cover all possible T , in both cases we use properties of \mathcal{A} from Lemma 16.

1. $f(U, \delta, T) \leq T$:

$$\begin{aligned} \Psi &= \frac{1}{2} \Pr[(\mathcal{A}(V, \delta, T) = a) \wedge (V = S)] + \frac{1}{2} \Pr[(\mathcal{A}(V, \delta, T/\beta) = a) \wedge (V = S)] \\ &\quad - \frac{1}{2} \Pr[(\mathcal{A}(U, \delta, T) = b) \wedge (V = S)] - \frac{1}{2} \Pr[(\mathcal{A}(U, \delta, T/\beta) = b) \wedge (V = S)] \\ &\leq \frac{1}{2} + \frac{1}{2} - \frac{1 - \delta}{2} - \frac{0}{2} \leq \frac{1 + \delta}{2} \end{aligned}$$

Algorithm 6: $\text{Reduction}(I, m, \delta, \gamma, T)$

Input: a set of n arms I , parameters m, δ , and γ , and time horizon T .

Output: a super set of Top_m arms or \perp .

- 1 Run independently $z = 25m \cdot 4^\gamma / \delta^2$ copies of $\mathcal{A}'(I, m, \delta, T/z)$;
 - 2 let $\hat{\alpha}_i$ be the frequency of i -th arm among the z returned values by \mathcal{A}' ;
 - 3 **if** (m -th largest $\hat{\alpha}_i$) $< 3/(4em)$ **then**
 - 4 **return** \perp ;
 - 5 **else**
 - 6 **return** all arms i for which $\hat{\alpha}_i \geq 1/(16em)$.
-

2. $\frac{T}{\beta} \leq g(U, \delta)$:

$$\begin{aligned} \Psi &= \frac{1}{2} \Pr[(\mathcal{A}(V, \delta, T) = a) \wedge (V = S)] + \frac{1}{2} \Pr[(\mathcal{A}(V, \delta, T/\beta) = a) \wedge (V = S)] \\ &\quad - \frac{1}{2} \Pr[(\mathcal{A}(U, \delta, T) = b) \wedge (V = S)] - \frac{1}{2} \Pr[(\mathcal{A}(U, \delta, T/\beta) = b) \wedge (V = S)] \\ &\leq \frac{1}{2} + \frac{\delta}{2} - \frac{0}{2} - \frac{0}{2} \leq \frac{1+\delta}{2} \end{aligned}$$

Combining the two cases, we have that

$$\begin{aligned} &\Pr[(\mathcal{A}(V, \delta, \tau) = a) \wedge (\text{Top}_1(V) = a)] - \Pr[(\mathcal{A}(U, \delta, \tau) = b) \wedge (\text{Top}_1(V) = a)] \\ &\leq \frac{1+\delta}{2} \sum_{S: \text{Top}_1(S)=a} \Pr[V = S] \leq \frac{1+\delta}{2} \left(1 - \frac{1}{m}\right)^m \frac{1}{m} \\ &\leq \frac{1+\delta}{2} \cdot \frac{1}{em}. \end{aligned}$$

□

3.4.2 Reduction to $O(m)$ Arms

With Lemma 19 we are able to design an algorithm such that given a time horizon T , it either returns a superset of Top_m of size $O(m)$, or \perp . The algorithm is described in Algorithm 6. The following lemma characterizes the property of Algorithm 6.

Lemma 21. For any $I, m, \delta \in (0, \frac{1}{24})$, and γ , we have

$$\Pr[(\text{Top}_m \subseteq \text{Reduction}(I, m, \delta, \gamma, T)) \vee (\text{Reduction}(I, m, \delta, \gamma, T) = \perp)] \geq 1 - n \cdot \exp(-4^\gamma).$$

Moreover, If

$$T \geq c_R \cdot 4^\gamma \cdot \frac{H^{(m)}}{\delta^3 K} \cdot \log^6 TK \cdot \log \frac{n}{\delta}$$

for a universal constant c_R , then

$$\Pr[\text{Top}_m \subseteq \text{Reduction}(I, m, \delta, \gamma, T)] \geq 1 - n \cdot \exp(-4^\gamma).$$

Finally, if $\text{Reduction}(I, m, \delta, \gamma, T) \neq \perp$, then the number of returned arms is bounded by $O(m)$.

Proof. Let α_i be the probability that \mathcal{A}' returns the i -th arm, that is,

$$\alpha_i \triangleq \Pr \left[\mathcal{A}' \left(I, m, \delta, \frac{T}{z} \right) = i \right].$$

Let $H = \{i \in I \mid \alpha_i \geq \frac{5}{8em}\}$. By Chernoff-Hoeffding we have that

$$\forall i \in I \setminus H : \Pr \left[\hat{\alpha}_i \geq (1 + \delta) \cdot \frac{5}{8em} \right] \leq \exp \left(-\frac{5\delta^2 z}{24em} \right) \leq \exp(-4^\gamma).$$

Since $\frac{5(1+\delta)}{8em} < \frac{3}{4em}$, it holds that

$$\Pr \left[\forall i \in I \setminus H : \hat{\alpha}_i < \frac{3}{4em} \right] \geq 1 - n \cdot \exp(-4^\gamma). \quad (39)$$

For the first part of the lemma, we consider two cases regarding H .

1. $|H| < m$: By (39), Algorithm 6 returns \perp with probability at least $1 - n \cdot \exp(-4^\gamma)$.
2. $|H| \geq m$: We must have $Top_m \subseteq H$ or $H \setminus Top_m \neq \emptyset$. In the first case, by the definition of H we have that $\forall i \in Top_m : \alpha_i \geq \frac{5}{8em} \geq \frac{1}{8em}$. In the second case, by Lemma 19 we have

$$\forall i \in Top_m : \alpha_i \geq \frac{3}{4em} - \left(\frac{1}{2em} + \frac{2\delta}{m} \right) \geq \frac{1}{8em}.$$

By Chernoff-Hoeffding we have

$$\forall i \in Top_m : \Pr \left[\hat{\alpha}_i < \frac{1}{16em} \right] \leq \Pr \left[\hat{\alpha}_i < (1 - \delta) \cdot \frac{1}{8em} \right] \leq \exp \left(-\frac{\delta^2 z}{8em} \right) \leq \exp(-4^\gamma).$$

Therefore with probability at least $1 - n \cdot \exp(-4^\gamma)$, Algorithm 6 outputs a super set of Top_m .

Combining the two cases, with probability at least $1 - n \cdot \exp(-4^\gamma)$, Algorithm 6 outputs either \perp or a super set of Top_m .

For the second part, since

$$\frac{T}{z} \geq c_T \cdot \frac{H^{(m)}}{25\delta Km} \log^6(TK) \log \frac{n}{\delta} \geq c_A \cdot \frac{H^{(m)}}{\delta Km} \log^6(TK) \log \frac{n}{\delta}$$

for a large enough constant c_T , by Lemma 18 we have $\alpha_i \geq \frac{(1-\delta)^2}{em}$ for any $i \in Top_m$. By Chernoff-Hoeffding we have

$$\forall i \in Top_m : \Pr \left[\hat{\alpha}_i \leq \frac{(1-\delta)^3}{em} \right] \leq \exp \left(-\frac{\delta^2(1-\delta)^2 z}{em} \right) \leq \exp(-4^\gamma). \quad (40)$$

Since $\frac{(1-\delta)^3}{em} > \frac{3}{4em}$, Algorithm 6 outputs a superset of Top_m with probability at least $1 - n \cdot \exp(-4^\gamma)$.

Finally, since $\sum_{i \in I} \hat{\alpha}_i \leq 1$, there are at most $O(m)$ arms i with $\hat{\alpha}_i \geq 1/(16em)$. \square

We now use Algorithm 6 as a building block for designing the reduction algorithm for general time horizon. The final reduction is described in Algorithm 7.

The following lemma summarizes the property of Algorithm 7.

Algorithm 7: ReductionGeneral(I, m, T)

Input: a set of n arms I , parameter m , and time horizon T .

Output: a super set of top- m arms.

- 1 **for** $s = 1, 2, \dots$ **do**
 - 2 $B_s \leftarrow \text{Reduction}(I, m, \frac{1}{2^s}, s, \frac{6T}{\pi^2 s^2})$;
 - 3 let s^* be the largest s such that $B_s \neq \perp$;
 - 4 **if** there is no such s **then**
 - 5 \perp **return** an arbitrary set of m arms;
 - 6 **else**
 - 7 \perp **return** B_{s^*} .
-

Lemma 22. For any I, m and T , the collaborative algorithm ReductionGeneral(I, m, T) returns a super set of Top_m of size $O(m)$ with probability at least

$$1 - 2n \cdot \exp\left(-\Omega\left(\frac{KT}{H^{(m)} \log^6(TK) \log^2(TK/H^{(m)}) \log n}\right)\right).$$

Proof. Let s' be the largest s such that $\frac{6T}{\pi^2 s^2} \geq 4^s \cdot c_R \frac{H^{(m)}}{\delta^3 K} \log^6 TK \log \frac{n}{\delta}$ with $\delta = \frac{1}{2^s}$. If there is no such s then the lemma follows trivially. Otherwise we have

$$4^{s'} = \Omega\left(\frac{KT}{H^{(m)} \log^6(KT) \log^2(TK/H^{(m)}) \log n}\right). \quad (41)$$

By Lemma 21 we have

$$\Pr[\text{Top}_m \subseteq B_{s'}] \geq 1 - n \cdot \exp(-4^{s'}). \quad (42)$$

For any $s = s' + j$ ($j \geq 1$), by Lemma 21 we have

$$\Pr[(\text{Top}_m \subseteq B_s) \vee (B_s = \perp)] \geq 1 - n \cdot \exp(-4^s). \quad (43)$$

By (42) and (43) we have

$$\begin{aligned} \Pr[\text{Top}_m \subseteq B_{s^*}] &\geq \Pr[\text{Top}_m \subseteq B_{s'}] - \sum_{s=s'+1}^{\infty} \Pr[(\text{Top}_m \not\subseteq B_s) \wedge (B_s \neq \perp)] \\ &\geq 1 - n \cdot \exp(-4^{s'}) - n \cdot \sum_{s=s'+1}^{\infty} \exp(-4^s) \\ &\geq 1 - 2n \cdot \exp(-4^{s'}), \end{aligned}$$

which, combined with (41), concludes the lemma. \square

Now we are ready to prove Theorem 15.

Proof of Theorem 15. We describe the algorithm in Algorithm 8. We first use Algorithm 7 to reduce the number of arms to $O(m)$, and then call Algorithm 3 CollabTopMGeneral to compute the set of top- m arms.

Algorithm 8: CollabTopMImproved(I, m, T)

Input: a set of n arms I , parameter m , and time horizon T .

Output: the set of top- m arms of I .

- 1 $B \leftarrow \text{ReductionGeneral}(I, m, T/2)$;
 - 2 **return** CollabTopMGeneral($B, m, T/2$).
-

The success probability of Theorem 15 follows directly from Lemma 22 and Theorem 10.

The round complexity follows from Lemma 16, Theorem 10, and the ways Algorithm 7, Algorithm 6, and Algorithm 5 are designed (they can be perfectly parallelized). Basically, the first term $O\left(\log \frac{\log m}{\log K}\right)$ comes from Algorithm 6 in which we reduce the number of arms from $O(m)$ to $O(K^{10})$. The second term $O(\log K)$ comes from two sources. One is from Line 2 of Algorithm 8 where we have used the best arm identification algorithm in [55] (setting $R = \log K$), and the other is due to the run of CollabTopMSimple on the remaining K^{10} arms (setting $R = \log n = \log K^{10}$). The round complexities introduced by other steps are negligible. \square

We comment that we can make the the second term in the round complexity of Theorem 15 to be an arbitrary number R , at the cost of slightly reducing the speedup. These can be accomplished by using the general round-speedup of Lemma 17 and Lemma 28, by setting round complexity to R and $10R$ respectively. On the other hand, the first term $O\left(\log \frac{\log m}{\log K}\right)$ remain the same even we target a $\tilde{\Theta}(\sqrt{K})$ speedup. We will show in Section 4 that this is inevitable.

Theorem 23. *Let I be a set of n arms, and $m \in \{1, \dots, n-1\}$. Let T be a time horizon and R be a parameter ($1 \leq R \leq \log K$). There exists a collaborative algorithm that computes the set of top- m arms of I with probability at least*

$$1 - n \cdot \exp\left(-\Omega\left(\frac{TK^{(R-1)/R}}{H^{(m)} \cdot (\log(H^{(m)}K) + R \log n) \cdot \log \log n \cdot \log^2(TK/H^{(m)})}\right)\right)$$

using at most T time steps and $O\left(\log \frac{\log m}{\log K} + R\right)$ rounds.

If we want to present Theorem 23 in the form of round-speedup tradeoff, we have the following corollary.

Corollary 24. *There is a collaborative algorithm for the top- m arms identification problem that achieves $\tilde{\Omega}(K^{(R-1)/R})$ speedup using at most $O\left(\log \frac{\log m}{\log K} + R\right)$ rounds of communication.*

3.5 Auxiliary Algorithms

3.5.1 PAC Top- m Arm Identification

Let $\beta(u, t) := \sqrt{\frac{1}{2u} \log\left(\frac{5nt^4}{4\delta}\right)}$. We first recall the LUCB algorithm from [37], which is described in Algorithm 9.

The following lemma summarizes the properties of Algorithm 9.

Lemma 25 ([37]). *Algorithm 9 LUCB(I, m, ϵ, δ) returns the set of (ϵ, m) -top arms with probability at least $(1 - \delta)$ using at most $O\left(H_\epsilon^{(m)} \log\left(H_\epsilon^{(m)}/\delta\right)\right)$ time steps. Moreover,*

$$\Pr[\forall i \in I : |\hat{\theta}_i - \theta_i| \leq \max\{\Delta_i^{(m)}/4, \epsilon/4\}] \geq 1 - \delta.$$

Algorithm 9: LUCB(I, m, ϵ, δ)

Input: a set of n arms I , parameter m , parameters $\epsilon, \delta \in (0, 1)$
Output: the set of m arms such that each arm is (ϵ, m) -top

- 1 Pull each arm once; $\forall i, p_i \leftarrow 1$;
- 2 **for** $t = |I| + 1, |I| + 2, \dots$ **do**
- 3 Let H be the set of m arms with highest estimated mean $\hat{\theta}_i$, and $L \leftarrow I \setminus H$;
- 4 $h^* \leftarrow \arg \min_{i \in H} \{\hat{\theta}_i - \beta(p_i, t)\}$, and $l^* \leftarrow \arg \max_{i \in L} \{\hat{\theta}_i + \beta(p_i, t)\}$;
- 5 **if** $(\hat{\theta}_{l^*} + \beta(p_{l^*}, t)) - (\hat{\theta}_{h^*} - \beta(p_{h^*}, t)) < \epsilon/2$ **then**
- 6 **return** H ;
- 7 pull h^* and l^* , and increment p_{l^*} and p_{h^*} .

Algorithm 10: CentralApproxTop(I, m, T, δ)

Input: a set of n arms I , parameter m , time horizon T , and parameter δ .
Output: the set of m arms such that each arm is (ϵ, m) -top.

- 1 Pull each arm once; $\forall i, p_i \leftarrow 1$ and set $t \leftarrow n$;
- 2 **while** $t \leq T - 2$ **do**
- 3 let H be the set of m arms with highest estimated mean $\hat{\theta}_i$, and $L \leftarrow I \setminus H$;
- 4 $h^* \leftarrow \arg \min_{i \in H} \{\hat{\theta}_i - \beta(p_i, t)\}$, and $l^* \leftarrow \arg \max_{i \in L} \{\hat{\theta}_i + \beta(p_i, t)\}$;
- 5 pull h^* and l^* , update $\hat{\theta}_{l^*}$ and $\hat{\theta}_{h^*}$, set $t \leftarrow t + 2$, and increment p_{l^*} and p_{h^*} ;
- 6 **return** H .

Note that Algorithm 9 is used to minimize the time (number of pulls) given PAC parameters ϵ and δ . In our task we need to (implicitly) minimize ϵ given T and δ . For this purpose we need to modify Algorithm 9 a bit; we described our new algorithm in Algorithm 10.

The following lemma summarizes the properties of Algorithm 10. It can be proved in essentially the same way as that for Lemma 25 in [37].

Lemma 26. *If $T \geq c_1 H_\epsilon^{(m)} \log(H_\epsilon^{(m)}/\delta)$ for a large enough constant c_1 , then with probability at least $(1 - \delta)$, CentralApproxTop(I, m, T, δ) returns a set of (ϵ, m) -top arms. Moreover,*

$$\Pr[\forall i \in I : |\hat{\theta}_i - \theta_i| \leq \max\{\Delta_i^{(m)}/4, \epsilon/4\}] \geq 1 - \delta.$$

The algorithm CentralApproxBtm is almost identical to CentralApproxTop: we can just follow CentralApproxTop, but replace all the sample values x with $1 - x$.

Lemma 27. *If $T \geq c_1 H_\epsilon^{(n-m)} \log(H_\epsilon^{(n-m)}/\delta)$ for a large enough constant c_1 , then with probability at least $(1 - \delta)$, CentralApproxBtm(I, m, T, δ) returns a set of (ϵ, m) -bottom arms and*

$$\Pr[\forall i \in I : |\hat{\theta}_i - \theta_i| \leq \max\{\Delta_i^{(n-m)}/4, \epsilon/4\}] \geq 1 - \delta.$$

3.5.2 A Simple Collaborative Algorithm for Top- m Identification

The CollabTopMSimple algorithm, described in Algorithm 11, is a slightly modified version of the *successive accepts and rejects* (SAR) algorithm in [5]. The goal of the modification is to achieve a small

Algorithm 11: CollabTopMSimple(I, m, T)

Input: a set of arms I , parameter m , and time horizon T .

Output: the set of top- m arms.

- 1 $I_0 \leftarrow I, m_0 \leftarrow m, Acc_1 \leftarrow \emptyset, \Theta \leftarrow \emptyset$;
 - 2 $T_0 \leftarrow 0, T_r \leftarrow \lfloor \frac{n^{r/R} T}{n^{1+1/R}(R+1)} \rfloor$ for $r = 1, \dots, R+1$;
 - 3 $n_r \leftarrow \lfloor \frac{n}{n^{r/R}} \rfloor$ for $r = 0, \dots, R+1$;
 - 4 **for** $r = 0, \dots, R$ **do**
 - 5 each agent pulls each arm $i \in I_r$ for $(T_{r+1} - T_r)$ times;
 - 6 let $\hat{\theta}_i^{(r)}$ for $i \in I_r$ be the aggregated mean of i -th arm after KT_r pulls;
 - 7 let $\sigma_r : \{1, \dots, |I_r|\} \rightarrow I_r$ be the bijection such that $\hat{\theta}_{\sigma_r(1)}^{(r)} \geq \hat{\theta}_{\sigma_r(2)}^{(r)} \geq \dots \geq \hat{\theta}_{\sigma_r(|I_r|)}^{(r)}$;
 - 8 for $i \in I_r$ define empirical gaps $\Delta_i^{(r)} = \begin{cases} \hat{\theta}_i^{(r)} - \hat{\theta}_{\sigma_r(m_r+1)}^{(r)}, & \text{if } \hat{\theta}_i^{(r)} \geq \hat{\theta}_{\sigma_r(m_r)}^{(r)}, \\ \hat{\theta}_{\sigma_r(m_r)}^{(r)} - \hat{\theta}_i^{(r)}, & \text{if } \hat{\theta}_i^{(r)} \leq \hat{\theta}_{\sigma_r(m_r+1)}^{(r)} \end{cases}$;
 - 9 let E_r be the set of $(n_r - n_{r+1})$ arms from I_r with the largest gaps $\Delta_i^{(r)}$;
 - 10 $Acc_{r+1} \leftarrow Acc_r \cup \{i \in E_r \mid \hat{\theta}_i^{(r)} \geq \hat{\theta}_{\sigma_r(m_r)}^{(r)}\}$;
 - 11 for $i \in E_r$ add $\hat{\theta}_i^{(r)}$ to Θ ;
 - 12 set $I_{r+1} \leftarrow I_r \setminus E_r$ and $m_{r+1} \leftarrow m - |Acc_{r+1}|$;
 - 13 **return** (Acc_{R+1}, Θ) .
-

number of rounds of communication in the collaborative setting.

Lemma 28. For any fixed R , Algorithm 11 CollabTopMSimple(I, m, T) uses at most T time steps and $R+1$ rounds of communication, and returns a top- m certificate (S, Θ) with probability at least

$$1 - 2n(R+1) \cdot \exp\left(-\frac{KT}{256 \cdot n^{1/R} H^{(m)} \cdot (R+1) \log 2n}\right).$$

Proof. The $R+1$ round complexity is clear from the description of the algorithm. The running time can be upper bounded by

$$\sum_{r=0}^R n_r \cdot T_{r+1} \leq \sum_{r=0}^R \frac{n}{n^{r/R}} \cdot \frac{n^{r/R} T}{n(R+1)} = T.$$

We next bound the error probability. The argument is very similar to the one in [5], and we include it here for completeness.

Let $\pi : \{1, \dots, n\} \rightarrow I$ be the bijection such that

$$\Delta_{\pi(1)}^{(m)} \leq \Delta_{\pi(2)}^{(m)} \leq \dots \leq \Delta_{\pi(n)}^{(m)}.$$

Consider the event

$$\mathcal{E}_2 : \left\{ \forall i \in I, \forall r = 0, \dots, R : |\hat{\theta}_i^{(r)} - \theta_i| \leq \Delta_{\pi(n_{r+1})}^{(m)} / 8 \right\}.$$

We have

$$\Pr [\bar{\mathcal{E}}_2] \leq \sum_{i \in I} \sum_{r=0}^R \Pr \left[|\hat{\theta}_i^{(r)} - \theta_i| > \Delta_{\pi(n_{r+1})}^{(m)}/8 \right] \quad (44)$$

$$\leq \sum_{i \in I} \sum_{r=0}^R 2 \exp \left(-KT_{r+1} \cdot \left(\Delta_{\pi(n_{r+1})}^{(m)}/8 \right)^2 \right) \quad (45)$$

$$\leq 2n(R+1) \cdot \exp \left(-\frac{KT}{256 \cdot n^{1/R} H^{(m)} \cdot (R+1) \log 2n} \right), \quad (46)$$

where (44) \rightarrow (45) is due to Chernoff-Hoeffding, and (45) \rightarrow (46) we have used an inequality from [5]:

$$\max_{i \in I} \left\{ i / \left(\Delta_{\pi(i)}^{(m)} \right)^2 \right\} \leq \log 2n \cdot H^{(m)}.$$

Once \mathcal{E}_2 holds, the proof for that fact that the returned pair (S, Θ) is a top- m certificate is straightforward. \square

4 Lower Bounds for the Fixed-Time Case

In this section, we prove the following lower bound theorems for the fixed-time setting.

Theorem 29. *For every K , m ($m \leq K$), and α ($\alpha \in [1, K^{0.1}]$), if a fixed-time collaborative algorithm \mathcal{A} with K agents returns the top- m arms for every instance J with probability at least 0.99, when given time budget $\frac{\alpha}{17K} \cdot H^{(m)}(J)$, then there exists an instance J' such that \mathcal{A} uses $\Omega(\log K / (\log \log K + \log \alpha))$ rounds of communication in expectation given instance J' and time budget $\frac{\alpha}{17K} \cdot H^{(m)}(J')$.*

In other words, to achieve (K/α) speedup for identifying the top m arms, the collaborative algorithm needs $\Omega(\log K / (\log \log K + \log \alpha))$ communication rounds.

We will prove Theorem 29 in Section 4.1. It is relatively easy and resembles the round complexity lower bound $\Omega(\log K / (\log \log K + \log \alpha))$ for top arm identification in the fixed-time setting [55].

Theorem 30. *For every large enough K and m such that $K \geq \Omega(\log^4 m)$, if a fixed-time collaborative algorithm \mathcal{A} with K agents returns the top- m arms for every instance J with probability at least 0.99, when given time budget $\frac{1}{\sqrt{K}} \cdot H^{(m)}(J)$, then there exists an instance J' such that \mathcal{A} uses $\Omega(\log(\log m / \log K))$ rounds of communication given instance J' and time budget $\frac{1}{\sqrt{K}} \cdot H^{(m)}(J')$.*

In other words, even if one only aims at \sqrt{K} speedup, the collaborative algorithm needs

$$\Omega(\log(\log m / \log K))$$

rounds of communication.

Theorem 30 will be proved in Section 4.2. It marks the different round complexity requirement for collaborative multiple arm identification compared to the best arm identification problem. It is known that only constant number of round is needed to achieve 0.99 success probability using $\tilde{O}(K^{-\zeta} \cdot H^{(m)}(J))$ time budget (i.e., $\tilde{O}(K^\zeta)$ speedup) for every constant $\zeta \in (0, 1)$ [31, 55]. However, Theorem 30 rules out such possibility for the multiple arm identification problem, proving it much harder than best arm identification in the collaborative setting. We note that we only prove the lower bound for $\zeta = 1/2$, for the simplicity of the exposition. However, the proof can be easily extended to any constant $\zeta > 0$. The only differences are that, in the theorem statement, the constraint $K \geq \Omega(\log^4 m)$ will become $K \geq \log^{f(\zeta)} m$, and the round complexity lower bound will become $\frac{1}{f(\zeta)} \cdot \log(\log m / \log K)$ where $f(\zeta) > 0$ increases as ζ approaches 0.

4.1 Proof of Theorem 29

The proof of Theorem 29 is via a simple reduction from the following lower bound for the best arm identification problem.

Theorem 31 (Theorem 10 in [55]). *For every K and α ($\alpha \in [1, K^{0.1}]$), if a fixed-time collaborative algorithm \mathcal{B} with K agents returns the best arm for every instance I with probability at least 0.99, when given time budget $\frac{\alpha}{K} \cdot H^{(1)}(I)$, then there exists an instance I' with $H^{(1)}(I') \geq K$ such that 1) the mean reward of the best arm in I' is less or equal to $1/2$ and 2) \mathcal{B} uses $\Omega(\log K / (\log \log K + \log \alpha))$ rounds of communication in expectation given instance I' and time budget $\frac{\alpha}{K} \cdot H^{(1)}(I')$.*

Proof of Theorem 29. Given an algorithm \mathcal{A} to identify the top- m arms, we construct an algorithm \mathcal{B} to identify the best arm as follows.

For every instance I where the mean reward of the best arm is at most $1/2$, we create $(m - 1)$ artificial arms with mean reward $3/4$, and together with I we have an instance J . We have that $H^{(m)}(J) \leq 16(m - 1) + H^{(1)}(I)$. When $H^{(1)}(I) \geq K \geq m$, we further have $H^{(m)}(J) \leq 17H^{(1)}(I)$. If \mathcal{A} is given time budget T , the algorithm \mathcal{B} will simulate \mathcal{A} with instance J and time budget T . Whenever an artificial arm is queried by \mathcal{A} , \mathcal{B} will generate a sample from a Bernoulli variable with mean $3/4$; otherwise, \mathcal{B} queries the real arm in I and feed the observation back to \mathcal{A} . The total time used by \mathcal{B} (only counting queries to the real arms) is at most T , satisfying the time budget constraint. When \mathcal{A} returns the identified arms, \mathcal{B} will return any real arm from the set (or declare failure if no such arm exists). If \mathcal{A} returns the correct top m arms in J (which happens with probability at least 0.99), \mathcal{B} will also return the correct best arm of I .

Now invoking Theorem 31, we know that there exists an instance I' with $H^{(1)}(I') \geq K$ such that \mathcal{B} uses $\Omega(\log K / (\log \log K + \log \alpha))$ rounds of communication in expectation for instance I' and time budget $T = \frac{\alpha}{K} \cdot H^{(1)}(I')$. Therefore, we have that \mathcal{A} uses $\Omega(\log K / (\log \log K + \log \alpha))$ rounds of communication in expectation for the corresponding instance J' (constructed from I') and time budget T . Also note that $T \geq \frac{\alpha}{17K} H^{(m)}(J')$, and we complete the proof. \square

Remark 32. *Given the above proof, one may observe that the $m \leq K$ constraint in Theorem 29 can be relaxed to admit much bigger m if Theorem 31 can be strengthened to provide lower bound instances I' such that $H^{(1)}(I') \gg K$. We highly believe this is possible, and we do have a proof sketch in the core of which is an improved argument of Section 3.2 in [55]. Since the proof is quite involved and not directly related to this paper, we leave the formal proof as a future work.*

4.2 Proof of Theorem 30

4.2.1 The Hard Instances and the Proof Intuition

For any fixed K (the number of agents), we define a distribution of instances $\mathcal{I}(C, \mu, n)$ with n Bernoulli arms, where we assume n is an odd integer, the parameter $C \in (0, 1/4)$ denotes the gap between the mean rewards of the top arm and the bottom arm, the parameter $\mu \in (3/8, 5/8)$. We always set $m = (n - 1)/2$, i.e., the goal is to identify the top half arms (not including the median arm).

When $n \leq K^{10}$, we set $\mathcal{I}(C, \mu, n)$ to be a deterministic instance where the $(n - 1)/2$ top arms have mean reward $(\mu + C/2)$, and the $(n - 1)/2$ bottom arms have mean reward $(\mu - C/2)$. There is 1 middle arm with mean reward μ sandwiched between the top and the bottom arms.

When $n > K^{10}$, we define a random sample $I \sim \mathcal{I}(C, \mu, n)$ in a recursive fashion as follows (and illustrated in Figure 1). Let η be the smallest odd integer that is greater than $\sqrt[4]{n}$. There are $((n - \eta(2\eta + 1))/2 + \xi\eta)$ top arms with mean reward $(\mu + C/2)$, and there are $((n - \eta(2\eta + 1))/2 - \xi\eta)$ bottom arms




Figure 1: Illustration of the mean rewards of the arms in $I \sim \mathcal{I}(C, \mu, n)$ for $n > K^{10}$.

with mean reward $(\mu - C/2)$, where the *bias* ξ is an integer independently and uniformly sampled from $[-\eta, \eta]$. For the rest $\eta(2\eta + 1)$ arms in the middle, we make $(2\eta + 1)$ independent samples $I_1, I_2, \dots, I_{2\eta+1}$ (each of which has η arms), such that for each $j \in [2\eta + 1]$, we have

$$I_j \sim \mathcal{I}\left(C\sqrt{\eta/n}, \mu + \frac{j - \eta - 1}{8} \cdot Cn^{-1/4}, \eta\right).$$

The final instance I consists of the union of the arms in I_j ($j \in [2\eta + 1]$) together with the top and the bottom arms. We also say that the arm is in the j -th *block* if it is an arm in I_j .

Below we will claim a few properties about our constructed hard instances. First, it is straightforward to verify the lemma.

Lemma 33. *For sufficiently large n , we have the following claims.*

1. *For any arm in the j -th block, its mean reward $\theta \in \frac{1}{2} + Cn^{-1/4} \cdot (\frac{j - \eta - 1}{8} \pm \frac{1}{100})$. Therefore, the mean rewards of all middle arms are sandwiched between the top and the bottom arms, and any two distinct blocks do not overlap.*
2. *The median arm of I is the median arm of $I_{\xi + \eta + 1}$.*

In the following lemma, we show the order of the complexity measure of the constructed instances.

Lemma 34. *For each instance I in the support of $\mathcal{I}(C, \mu, n)$, we have $H^{(m)}(I) = \Theta(C^{-2}n \cdot \log(e + \frac{\log n}{\log K}))$.*

Proof. We prove this lemma via induction, where the base case $n \leq K^{10}$ is straightforward to verify.

When $n > K^{10}$, let η be defined in the construction of the instances. Let $I_1, I_2, \dots, I_{2\eta+1}$ be any instances such that I_j is in the support of $\mathcal{I}(C\sqrt{\eta/n}, \mu + \frac{j - \eta - 1}{8} \cdot Cn^{-1/4}, \eta)$ for each $j \in [2\eta + 1]$. Let $\xi \in [-\eta, \eta]$ be any integer, and let I be the instance constructed using the parameters above. Let $\theta_{[m]}$ and $\theta_{[m+1]}$ be the mean rewards of the m -th and the $(m + 1)$ -th best arm, respectively. We have

$$\begin{aligned} H^{(m)}(I) = & \left(\frac{n - \eta(2\eta + 1)}{2} + \xi\eta\right) \cdot \left(\frac{1 + C}{2} - \theta_{[m+1]}\right)^{-2} + \left(\frac{n - \eta(2\eta + 1)}{2} - \xi\eta\right) \cdot \left(\frac{1 - C}{2} - \theta_{[m]}\right)^{-2} \\ & + \sum_{j=1}^{\xi+\eta} \sum_{\text{arm } i \text{ in block } j} (\theta_i - \theta_{[m]})^{-2} + \sum_{j=\xi+\eta+2}^{2\eta+1} \sum_{\text{arm } i \text{ in block } j} (\theta_i - \theta_{[m+1]})^{-2} + H^{((\eta-1)/2)}(I_{\xi+\eta+1}). \quad (47) \end{aligned}$$

By Lemma 33, we have $\theta_{[m]}, \theta_{[m+1]} \in \frac{1}{2} + Cn^{-1/4} \cdot (\frac{\xi}{8} \pm \frac{1}{100})$. Therefore, we have

$$\begin{aligned} \left(\frac{n - \eta(2\eta + 1)}{2} + \xi\eta \right) \cdot \left(\frac{1 + C}{2} - \theta_{[m+1]} \right)^{-2} + \left(\frac{n - \eta(2\eta + 1)}{2} - \xi\eta \right) \cdot \left(\frac{1 - C}{2} - \theta_{[m]} \right)^{-2} \\ \in \left[\frac{C^{-2}n}{4}, 16C^{-2}n \right]. \end{aligned} \quad (48)$$

and

$$\begin{aligned} \sum_{j=1}^{\xi+\eta} \sum_{i \text{ in block } j} (\theta_i - \theta_{[m]})^{-2} + \sum_{j=\xi+\eta+2}^{2\eta+1} \sum_{i \text{ in block } j} (\theta_i - \theta_{[m+1]})^{-2} \\ \leq \sum_{k=1}^{2\eta} \eta \cdot \left(Cn^{-1/4} \cdot (k/16) \right)^{-2} \leq 256\eta C^{-2} n^{1/2} \cdot \frac{\pi^2}{6} \leq 512C^{-2} n^{3/4}. \end{aligned} \quad (49)$$

Combining (47), (48), and (49), for sufficiently large n , we have

$$H^{(m)}(I) \in \left[\frac{C^{-2}n}{4}, 17C^{-2}n \right] + H^{((\eta-1)/2)}(I_{\xi+\eta+1}).$$

Apply induction hypothesis to $H^{((\eta-1)/2)}(I_{\xi+\eta+1})$ and we prove the lemma. \square

Proof Intuition. The intuition about our lower bound instance distribution is as follows. As pointed out in Lemma 33 (Item 2), the top $m = (n - 1)/2$ arms consists of the top arms, the blocks from $I_{\xi+\eta+2}$ to $I_{2\eta+1}$, and finally the top half (excluding the median) arms in block $I_{\xi+\eta+1}$. Therefore, two necessary tasks are i) to complete is to identify the value of ξ , and ii) to identify the top half arms in $I_{\xi+\eta+1}$.

For the first task, in Section 4.2.2, we will introduce a sub-problem named ‘‘learning the bias’’. Via studying this problem, we will show that, any agent, if using at most $H^{(m)}(I)/\sqrt{K}$ queries, cannot learn the correct value of ξ with probability $\omega(n^{-1/4})$. Note that since there are only $2\eta + 1 = O(n^{1/4})$ possible values for ξ , this means that the agent cannot do much better than random guessing. Also, we note that we prove the impossibility statement for agents with even $\Theta(nC^{-2}/\log(n/C))$ queries, which is a stronger statement as Lemma 34 shows that $H^{(m)}(I)$ is always $\tilde{\Theta}(nC^{-2})$.

The above discussion suggests that a communication step is needed for the agents to collectively decide the exact value of ξ . It also suggests that not too many queries are made to the block $I_{\xi+\eta+1}$ before the first communication step (more specifically, the amount is at most $O(n^{-1/4})$ fraction of the total number of queries before the communication step, see Item 2 of Lemma 37 for detailed justification), which is negligible for further identifying the top half arms in $I_{\xi+\eta+1}$ (the second necessary task). On the other hand, by Lemma 34, $H^{((\eta-1)/2)}(I_{\xi+\eta+1})$ is still $\tilde{\Theta}(nC^{-2})$. Therefore, we can recursively apply the similar argument to $I_{\xi+\eta+1}$, yielding a communication round lower bound that is proportional to the number of hierarchies in the definition of $\mathcal{I}(C, \mu, n)$, which is $\Theta(\log(\log n / \log k))$.

This recursive (or inductive) argument is presented in Section 4.2.3. Note that in the simplified explanation above, we neglected the extra queries made to the $I_{\xi+\eta+1}$ before the first communication step. To formally deal with these extra queries, we need to strengthen the inductive hypothesis, and introduce the definition of *augmented algorithms* where the agents enjoy a small number of free (and shared) queries before the very first communication round. We will show that the communication round lower bound still holds even for augmented algorithms.

4.2.2 The ‘‘Learning the Bias’’ Sub-problem and Its Analysis

In this subsection, we identify a critical sub-problem for identifying the top m arms in our constructed hard instances. We then prove the sample complexity lower bounds for the sub-problem, which will be a crucial building block for the ultimate lower bound theorem for identifying the top m arms.

We first define the sub-problem as follows.

Problem Definition (Learning the Bias). *There are n Bernoulli arms. Given the parameters $\epsilon \in (0, 1/8)$, $\mu \in (3/8, 5/8)$, and a distribution \mathcal{D} supported on $\{\pm 1\}^n$ (which are publicly known), a hidden vector (b_1, b_2, \dots, b_n) is sampled from \mathcal{D} , and the mean reward of the i -th arm is set to be $\theta_i = \mu + b_i\epsilon$ (also hidden from the algorithm). The algorithm has a budget of T adaptive samples from the arms, and the goal is to decide the bias $B = b_1 + b_2 + \dots + b_n$.*

The following lemma shows the sample complexity lower bound for learning the bias when \mathcal{D} is the uniform distribution.

Lemma 35. *Assume that \mathcal{D} is the uniform distribution of $\{\pm 1\}^n$. For sufficiently large n , if*

$$T \leq n\epsilon^{-2}/(20000 \log(n/\epsilon)),$$

then the probability that the player correctly identifies B is at most $O(n^{-1/2})$.

Proof. Without loss of generality we can assume that the player makes the guess about B after using all T samples. Let $\tau = (i_1, y_1, i_2, y_2, \dots, i_T, y_T)$ be the *transcript* of all samples, where $i_t \in [n]$ denotes the arm sampled from at time t , and $y_t \in \{0, 1\}$ denotes the observation at time t . The player will finally uses an algorithm $\mathcal{A}(\tau)$ to decide the guess about the bias. The probability that the player makes a correct guess is

$$\begin{aligned} \Pr[\text{correct guess}] &= \mathbb{E}_{\tau} \Pr_{(b_1, \dots, b_n) \sim \mathcal{D}} [\mathcal{A}(\tau) = B | \tau] \\ &\leq \mathbb{E}_{\tau} \max_{\beta} \left\{ \Pr_{(b_1, \dots, b_n) \sim \mathcal{D}} [B = \beta | \tau] \right\} = \mathbb{E}_{\tau} \max_{\beta} \left\{ \Pr_{(b_1, \dots, b_n) \sim \mathcal{D}(\tau)} [B = \beta] \right\}, \end{aligned} \quad (50)$$

where we let $\mathcal{D}(\tau)$ be the posterior distribution of (b_1, b_2, \dots, b_n) given τ .

For fixed τ , let $r_{i,t}$ be the number of 1’s the player observes among the first t samples made from the i -th arm. Also let T_i be the total number of samples made from the i -th arm, we have $T = T_1 + T_2 + \dots + T_n$. Since \mathcal{D} is a product distribution, we have

$$\mathcal{D}(\tau) = \otimes_{i=1}^n \mathcal{D}_i(r_{i,T_i}, T_i), \quad (51)$$

where $\mathcal{D}_i(r_{i,T_i}, T_i)$ is the posterior of b_i given that r_{i,T_i} 1’s are observed from T_i samples from arm i .

Note that the posterior distribution $\mathcal{D}_i(r_{i,T_i}, T_i)$ is completely determined by the posterior probability $p_i \triangleq \Pr[b_i = +1 | r_{i,T_i}, T_i]$. Let $\sigma_i^2 \triangleq \text{Var}[\mathcal{D}_i(r_{i,T_i}, T_i)] = 4p_i(1 - p_i)$ for every arm i . Let $\sigma^2 \triangleq \sigma_1^2 + \sigma_2^2 + \dots + \sigma_n^2$ (where $\sigma \geq 0$). By (51) and invoking the Berry-Esseen theorem (Theorem 42) with $\rho = 8$, we have

$$\forall x \in \mathbb{R}, \Pr_{(b_1, \dots, b_n) \sim \mathcal{D}(\tau)} \left[\sigma^{-1} \sum_{i=1}^n (b_i - (1 - 2p_i)) \leq x \right] \in \Phi(x) \pm 4.5n\sigma^{-3/2}.$$

Since $\Phi(x)$ is a continuous function, we have

$$\forall \beta \in \mathbb{R}, \Pr_{(b_1, \dots, b_n) \sim \mathcal{D}(\tau)} [B = \beta] \leq 9n\sigma^{-3/2}. \quad (52)$$

Now we will estimate the posterior probability p_i and give a lower bound on σ^2 to upper bound the probability in (52).

Via standard concentration inequalities (e.g., Hoeffding's inequality), we have

$$\forall i, t, \Pr_{\tau} \left[|r_{i,t} - t\theta_i| \leq 2\sqrt{t \log(n/\epsilon)} \right] \geq 1 - 2(\epsilon/n)^8.$$

Therefore, if we define \mathcal{E}_4 be the event

$$\mathcal{E}_4 \triangleq \{\forall i, t, |r_{i,t} - t\theta_i| \leq 2\sqrt{t \log(n/\epsilon)}\},$$

we have $\Pr_{\tau}[\mathcal{E}_4] \geq 1 - 1/n^6$.

Say an arm i is *sufficiently explored* if $T_i \geq \epsilon^{-2}/(10000 \log(n/\epsilon))$. By Markov's inequality, there are at most $n/2$ sufficiently explored arms. Fix an arm i , let $r = r_{i,T_i}$ for notational convenience. Conditioned on the event \mathcal{E}_4 and that it is insufficiently explored, we have $|r - T_i\theta_i| \leq 1/(50\epsilon)$. Since $\theta_i = \mu \pm \epsilon$, we further have

$$|r - \mu T_i| \leq \frac{1}{50\epsilon} + \epsilon T_i \leq \frac{1}{40\epsilon}. \quad (53)$$

By the definition of p_i , we have

$$p_i = \frac{(\mu + \epsilon)^r (1 - \mu - \epsilon)^{T_i - r}}{(\mu + \epsilon)^r (1 - \mu - \epsilon)^{T_i - r} + (\mu - \epsilon)^r (1 - \mu + \epsilon)^{T_i - r}} = \frac{1}{1 + \left(\frac{\mu - \epsilon}{\mu + \epsilon}\right)^r \left(\frac{1 - \mu + \epsilon}{1 - \mu - \epsilon}\right)^{T_i - r}}. \quad (54)$$

By (53), for $\epsilon \in (0, 1/8)$ and $\mu \in (3/8, 5/8)$, we have that

Note that

$$\left(\frac{\mu - \epsilon}{\mu + \epsilon}\right)^r \left(\frac{1 - \mu + \epsilon}{1 - \mu - \epsilon}\right)^{T_i - r} = \left(\frac{\mu - \epsilon}{\mu + \epsilon}\right)^{\mu T_i + (r - \mu T_i)} \left(\frac{1 - \mu + \epsilon}{1 - \mu - \epsilon}\right)^{(1 - \mu)T_i + (\mu T_i - r)}, \quad (55)$$

and for $\epsilon \in (0, 1/8)$, $\mu \in (3/8, 5/8)$, and $T_i \leq \epsilon^{-2}/10000$, it holds that

$$\left(\frac{\mu - \epsilon}{\mu + \epsilon}\right)^{\mu T_i} \left(\frac{1 - \mu + \epsilon}{1 - \mu - \epsilon}\right)^{(1 - \mu)T_i} \in [0.99, 1.01]. \quad (56)$$

Also, by (53) and the ranges for ϵ and μ , we have

$$\begin{aligned} & \left(\frac{\mu - \epsilon}{\mu + \epsilon}\right)^{r - \mu T_i} \left(\frac{1 - \mu + \epsilon}{1 - \mu - \epsilon}\right)^{\mu T_i - r} \\ & \in \left[\left(\frac{(\mu - \epsilon)(1 - \mu - \epsilon)}{(\mu + \epsilon)(1 - \mu + \epsilon)}\right)^{1/(40\epsilon)}, \left(\frac{(\mu + \epsilon)(1 - \mu + \epsilon)}{(\mu - \epsilon)(1 - \mu - \epsilon)}\right)^{1/(40\epsilon)} \right] \subseteq [0.75, 1.3]. \end{aligned} \quad (57)$$

Combining (54), (55), (56), and (57), and conditioned on \mathcal{E}_4 and that arm i is insufficiently explored, we have that $p_i \in [0.4, 0.6]$, meaning that $\sigma_i^2 \geq 0.96$. Conditioned on \mathcal{E}_4 , since there are at most $n/2$ sufficiently explored arms, we have that $\sigma^2 \geq 0.48n$. Using (52), we have

$$\begin{aligned} \mathbb{E}_{\tau} \max_{\beta} \left\{ \Pr_{(b_1, \dots, b_n) \sim \mathcal{D}(\tau)} [B = \beta] \right\} & \leq \mathbb{E}_{\tau} \left[\max_{\beta} \left\{ \Pr_{(b_1, \dots, b_n) \sim \mathcal{D}(\tau)} [B = \beta] \right\} \middle| \mathcal{E}_4 \right] + \Pr[\bar{\mathcal{E}}_4] \\ & \leq 9n(0.48n)^{-3/2} + n^{-6} \leq 30n^{-1/2}. \end{aligned}$$

Together with (50), we prove the lemma. \square

To analyze the lower bound for our hard instances $\mathcal{I}(C, \mu, n)$, we need to adapt Lemma 35 to a different distribution \mathcal{D} , as shown in the following corollary.

Corollary 36. *For any $S \subseteq [n]$ ($S \neq \emptyset$), let $\mathcal{D} = \mathcal{D}(S)$ be the following distribution supported on $\{\pm 1\}^n$: first sample a uniformly random integer s from S , and then sample a uniformly random vector from $\{\pm 1\}^n$ such that the number of $+1$'s in the vector is exactly s . Let $q = \min_{s \in S} \{2^{-n} \binom{n}{s}\}$. For sufficiently large n , if $T \leq n\epsilon^{-2}/(20000 \log(n/\epsilon))$, then in the learning the bias problem, the probability that the player correctly identifies B is at most $O(n^{-1/2}q^{-1} \cdot |S|^{-1})$*

Proof. Construct the joint distribution with probability mass function p for the random variables

$$(b_1, b_2, \dots, b_n, Y) \in \{\pm 1\}^n \times \{0, 1\}$$

as follows (where we let $\beta = \sum_{i=1}^n (b_i + 1)/2$ be the number of $+1$'s in the vector (b_1, b_2, \dots, b_n)),

$$p(b_1, b_2, \dots, b_n, Y) = \begin{cases} 0 & \text{when } \beta \notin S \text{ and } Y = 1 \\ 2^{-n} \binom{n}{\beta} & \text{when } \beta \notin S \text{ and } Y = 0 \\ q \binom{n}{\beta}^{-1} & \text{when } \beta \in S \text{ and } Y = 1 \\ 2^{-n} \binom{n}{\beta} - q \binom{n}{\beta}^{-1} & \text{when } \beta \in S \text{ and } Y = 0 \end{cases}.$$

It is clear that the marginal distribution on (b_1, b_2, \dots, b_n) is uniform over $\{\pm 1\}^n$ and the conditional distribution on (b_1, b_2, \dots, b_n) given that $Y = 1$ is $\mathcal{D}(S)$. Therefore, the probability that the player correctly guesses B given $\mathcal{D} = \mathcal{D}(S)$ is

$$\begin{aligned} \Pr_{(b_1, b_2, \dots, b_n) \sim \mathcal{D}(S)} [\text{correct guess}] &= \Pr_{(b_1, b_2, \dots, b_n, Y) \sim p} [\text{correct guess} | Y = 1] \\ &\leq \frac{\Pr_{(b_1, b_2, \dots, b_n, Y) \sim p} [\text{correct guess}]}{\Pr_{(b_1, b_2, \dots, b_n, Y) \sim p} [Y = 1]} \leq O(n^{-1/2}) \cdot \frac{1}{q|S|}, \end{aligned}$$

where in the last inequality, we invoked Lemma 35. □

4.2.3 The Lower Bound Theorems for Communication Rounds and Concluding the Proof

The following lemma helps to relate the lower bound results derived for the learning the bias problem in Section 4.2.2 to the form of top m arm identification.

Lemma 37. *For sufficiently large K , any odd integer n such that $n > K^{10}$, any $C \in (0, 1/8)$, and any $\mu \in (3/8, 5/8)$, consider a random instance from $\mathcal{I}(C, \mu, n)$. For any player that makes at most T sequential samples, when $T \leq nC^{-2}/(40000 \log(n/C))$, we have the following claims.*

1. *The probability that the player correctly identifies the top m arms is at most $O(n^{-1/4})$ (recall that $m = (n - 1)/2$).*
2. *The expected number of samples made to any arm in the block that contains the median arm is at most $O(n^{-1/4}T)$.*

Proof. We prove the lemma by reducing the learning of the bias problem to the top- m -arm identification problem for the instance distribution $\mathcal{I}(C, \mu, n)$

Let us consider the learning the bias problem with $n' = n - \eta(2\eta + 1)$ arms (note that $n' \geq n/2$), $\epsilon = C$, the same μ parameter, and the distribution $\mathcal{D} = \mathcal{D}(S)$ where $S = \{n'/2 + z\eta : z \in \{-\eta, \eta + 1, \dots, \eta - 1, \eta\}\}$. Once the expected rewards of the n' arms are determined by a sample from \mathcal{D} , let J be set of the arms. To construct a top- m -arm identification problem instance, we independently sample smaller problem instances $I_1, I_2, \dots, I_{2\eta+1}$ such that $I_j \sim \mathcal{I}(C\sqrt{\eta/n}, \mu + \frac{j-\eta-1}{8} \cdot Cn^{-1/4}, \eta)$ for each $j \in [2\eta + 1]$. Let $I = J \cup I_1 \cup I_2 \cup \dots \cup I_{2\eta+1}$. One can verify that I follows the distribution $\mathcal{I}(C, \mu, n)$.

Now we prove the first claim. Suppose that the player correctly identifies the top m arms with probability p . According to Lemma 33, only the arms in block $(z + \eta + 1)$ have non-empty intersection with both the set of top m arms and the set of the remaining arms. Therefore, when conditioned on that the top m arms are correctly identified, by checking the identified set of arms, the player can find out the value of z , and deduce that the bias $B = 2z\eta$. Therefore, there exists an algorithm correctly identifying the bias with probability at least p . Invoking Corollary 36, and noting that $|S| = \Theta(n^{-1/4})$, $q = \Theta(1)$, we have that

$$p \leq O\left((n')^{-1/2} q^{-1} \cdot |S|^{-1}\right) \leq O(n^{-1/4}).$$

Regarding the second claim, let us consider the player, after making at most T samples, guessing $z = \tilde{z}$ with probability $t_{\tilde{z}+\eta+1}/T$, where t_j is the number of samples made to the j -th block, and finally guessing $B = 2\tilde{z}\eta$. The probability that the player successfully identifies the bias B is $\mathbb{E}[t_{\tilde{z}+\eta+1}/T]$. Invoking Corollary 36, we have that this value is upper bounded by $O(n^{-1/4})$. Therefore, we have $\mathbb{E}[t_{\tilde{z}+\eta+1}] \leq O(n^{-1/4}T)$, which proves the claim by noting that block $(z + \eta + 1)$ contains the median arm in I , due to Lemma 33. \square

Definition of Augmented Algorithms and Uniform Upper Bound p. For fixed and sufficiently large K , let $\mathcal{A}_{R,T}^{(\alpha)}$ be the set of R -round K -agent *augmented* algorithms defined as follows. Any algorithm in $\mathcal{A}_{R,T}^{(\alpha)}$ has R rounds of communication, where during each round, the time budget is T/\sqrt{K} . Before the first round, there is an *augmented round* where a single thread is allowed to make αT sequential samples and broadcast the observations to all agents. Let

$$\mathfrak{p}_{R,T}^{(\alpha)}(C, n) \triangleq \sup_{\mathbb{A} \in \mathcal{A}_{R,T}^{(\alpha)}} \sup_{\mu: |\mu-1/2| < 1/8 - C} \Pr_{I \sim \mathcal{I}(C, \mu, n), \mathbb{A}} [\mathbb{A} \text{ identifies the top } m = (n-1)/2 \text{ arms}] \quad (58)$$

be the best success probability of augmented algorithms in $\mathcal{A}_{R,T}^{(\alpha)}$ when the input instance follows $\mathcal{I}(C, \mu, n)$ for any $\mu : |\mu - 1/2| < 1/8 - C$, where the subscript of \Pr specifies that the probability is taken over both I and the randomness of algorithm \mathbb{A} . Clearly, if we can prove that $\mathfrak{p}_{R,T}^{(\alpha)}(C, n) < 0.1$ for any α , we obtain the round complexity lower bound R for fixed-time algorithms with time budget T . In what follows, we will prove upper bounds for $\mathfrak{p}_{R,T}^{(\alpha)}(C, n)$ via induction.

Lemma 38. *For any positive $C \in (0, 1/8)$, any constant $\iota \geq 0$, suppose $n > K^{10}$, $R \geq 2$, and $\sqrt{K} \geq 80000 \log^{\iota+1}(nC^{-2})$. Let $T = (nC^{-2}) \log^{\iota}(nC^{-2})$, and let η be the smallest odd integer that is greater than $\sqrt[4]{n}$. It holds that*

$$\mathfrak{p}_{R,T}^{(n^{-1/8})}(C, n) \leq O\left(n^{-3/32}\right) + \mathfrak{p}_{R-1,T}^{(\eta^{-1/8})}(C\sqrt{\eta/n}, \eta).$$

Proof. Fix any algorithm $\mathbb{A} \in \mathcal{A}_{R,T}^{(\alpha)}$ and any $\mu \in (3/8 + C, 5/8 - C)$, we will upper bound the success probability of \mathbb{A} given the input instance $I \sim \mathcal{I}(C, \mu, n)$.

Recall in the construction of the instance $I \sim \mathcal{I}(C, \mu, n)$, $(2\eta + 1)$ blocks are independently sampled. Let $\zeta \in [2\eta + 1]$ be the block which the median arm is in. Let n_0 be the number of samples made in the augmented round to arms in block ζ , and let n_i be the number of samples made in the first round to arms in block ζ by agent i . For every agent i , by the second claim of Lemma 37 (and observing that the total number of samples made in the augmented round and the first round by agent i is at most $T/\sqrt{K} + n^{-1/8}T \leq T(K^{-1/2} + K^{-5/4}) \leq nC^{-2}/(40000 \log(n/C))$), we have that

$$\mathbb{E}_{I \sim \mathcal{I}(C, \mu, n), \mathbb{A}} [n_0 + n_i] \leq O\left(nC^{-2} \cdot n^{-1/4}\right).$$

Therefore,

$$\mathbb{E}_{I \sim \mathcal{I}(C, \mu, n), \mathbb{A}} [n_0 + n_1 + n_2 + \dots + n_K] \leq O\left(n^{3/4}C^{-2} \cdot K\right) \leq O\left(n^{7/8}C^{-2}\right).$$

Let \mathcal{E}_5 be the event that $n_0 + n_1 + n_2 + \dots + n_K \leq nC^{-2} \cdot \eta^{-1/8}$. By Markov's inequality, we have

$$\Pr_{I \sim \mathcal{I}(C, \mu, n), \mathbb{A}} [\mathcal{E}_5] \geq 1 - \frac{O(n^{7/8}C^{-2})}{nC^{-2} \cdot \eta^{-1/8}} = 1 - O\left(n^{-3/32}\right). \quad (59)$$

Our next goal is to establish (60) for every $j \in [2\eta + 1]$. Fix such j , consider the following algorithm \mathbb{B} that works for an instance $I_j \sim \mathcal{I}(C\sqrt{\eta/n}, \mu + \frac{j-\eta-1}{8} \cdot Cn^{-1/4}, \eta)$. The algorithm first samples $I_{j'} \sim \mathcal{I}(C\sqrt{\eta/n}, \mu + \frac{j'-\eta-1}{8} \cdot Cn^{-1/4}, \eta)$ for all $j' \neq j$. The algorithm \mathbb{B} also creates $((n - \eta(2\eta + 1))/2 + (j - \eta - 1)\eta)$ Bernoulli arms with mean reward $(\mu + C/2)$, and $((n - \eta(2\eta + 1))/2 - (j - \eta - 1)\eta)$ Bernoulli arms with mean reward $(\mu - C/2)$. Combining all the arms (including those in $I_1, I_2, \dots, I_{2\eta+1}$), we have an instance I^b of n arms. The algorithm \mathbb{B} simulates algorithm \mathbb{A} with input instance I^b in the following manner. Whenever \mathbb{A} is to sample an arm in I_j , \mathbb{B} queries the real arm in I_j , otherwise \mathbb{B} simulates a sample to the artificial arm, and feed the observation to \mathbb{A} . More importantly, only a single thread is used to simulate \mathbb{A} during the augmented round and the first round. Then, if event \mathcal{E}_5 holds, all K agents are used to simulate the corresponding agents in \mathbb{A} from the second round and reports I_j intersecting the set of top arms returned by \mathbb{A} ; otherwise, \mathbb{B} reports failure and terminates.

Note that $\mathbb{B} \in \mathcal{A}_{R-1, T}^{(\eta^{-1/8})}$. By the definition in (58), we have

$$\Pr_{I_j, \mathbb{B}} [\mathbb{B} \text{ identifies the top } (\eta - 1)/2 \text{ arms in } I_j] \leq \mathfrak{p}_{R-1, T}^{(\eta^{-1/8})}(C\sqrt{\eta/n}, \eta).$$

Note that I^b constructed above follows the conditional distribution $\mathcal{I}(C, \mu, n)$ given that the median arm is in the j -th block. Also note that when \mathcal{E}_5 holds and \mathbb{A} is correct, \mathbb{B} is also correct. Therefore, we have

$$\begin{aligned} & \Pr_{I \sim \mathcal{I}(C, \mu, n), \mathbb{A}} [\mathbb{A} \text{ identifies the top } (n - 1)/2 \text{ arms in } I | \zeta = j] \\ & \leq \Pr_{I \sim \mathcal{I}(C, \mu, n), \mathbb{A}} [\mathbb{A} \text{ identifies the top } (n - 1)/2 \text{ arms in } I \wedge \mathcal{E}_5 | \zeta = j] + \Pr_{I \sim \mathcal{I}(C, \mu, n), \mathbb{A}} [\overline{\mathcal{E}_5} | \zeta = j] \\ & \leq \Pr_{I_j, \mathbb{B}} [\mathbb{B} \text{ identifies the top } (\eta - 1)/2 \text{ arms in } I_j] + \Pr_{I \sim \mathcal{I}(C, \mu, n), \mathbb{A}} [\overline{\mathcal{E}_5} | \zeta = j] \\ & \leq \mathfrak{p}_{R-1, T}^{(\eta^{-1/8})}(C\sqrt{\eta/n}, \eta) + \Pr_{I \sim \mathcal{I}(C, \mu, n), \mathbb{A}} [\overline{\mathcal{E}_5} | \zeta = j], \end{aligned} \quad (60)$$

where ζ is the block which the median arm is in.

Since (60) holds for all $j \in [2\eta + 1]$. We have

$$\Pr_{I \sim \mathcal{I}(C, \mu, n), \mathbb{A}} [\mathbb{A} \text{ identifies the top } (n-1)/2 \text{ arms in } I] \leq \mathfrak{p}_{R-1, T}^{(\eta^{-1/8})}(C\sqrt{\eta/n}, \eta) + \Pr_{I \sim \mathcal{I}(C, \mu, n), \mathbb{A}} [\overline{\mathcal{E}}_5]. \quad (61)$$

Combining (59) and (61), we conclude the proof of the lemma. \square

When $R = 1$, the agents do not communicate except for the shared observation from the augmented round. Therefore, Lemma 37 implies that for $C \in (0, 1/8)$, any constant $\iota \geq 0$, $n > K^{10}$, $\sqrt{K} \geq 80000 \log^{\iota+1}(nC^{-2})$, and $T = (nC^{-2}) \log^{\iota}(nC^{-2})$, it holds that

$$\mathfrak{p}_{1, T}^{(n^{-1/8})}(C, n) \leq O(n^{-1/4}). \quad (62)$$

Combining Lemma 38 and (62), we have the following lemma.

Lemma 39. *For $C \in (0, 1/8)$, any constant $\iota \geq 0$, $n > K^{10}$, $\sqrt{K} \geq 80000 \log^{\iota+1}(nC^{-2})$, $T = (nC^{-2}) \log^{\iota}(nC^{-2})$, and $R = \lceil \log_4 \frac{\log n}{10 \log K} \rceil$, it holds that*

$$\mathfrak{p}_{R, T}^{(n^{-1/8})}(C, n) \leq \sum_{r=0}^{R-1} O(n^{-4^{-r} \cdot 3/32}) \leq O(K^{-1/5}).$$

Theorem 30 is proved by setting $\iota = 1$ and $C = 1/10$ in Lemma 39, together with Lemma 34 and the discussion below the definition of \mathfrak{p} (Equation (58)).

5 The Fixed-Confidence Case

In this section we discuss the fixed-confidence case. We first present a collaborative algorithm for the fixed-confidence case. The algorithm is inspired by [31] and [12], and described in Algorithm 12.

Theorem 40. *There is an algorithm (Algorithm 12) that solves top- m arm identification with probability at least $1 - \delta$, using $O\left(\log\left(1/\Delta_{[m]}^{(m)}\right)\right)$ rounds of communication and $O\left(\frac{H^{(m)}}{K} \log\left(\frac{n}{\delta} \log H^{(m)}\right)\right)$ time.*

Proof. First, by Chernoff-Hoeffding we have that for any $r \geq 0$ and $i \in I_r$,

$$\Pr \left[\left| \hat{\theta}_i^{(r)} - \theta_i \right| \geq \frac{\epsilon_r}{4} \right] \leq 2 \exp\left(-\frac{\epsilon_r^2}{8} K T_r\right) \leq \frac{\delta}{2n(r+1)^2}.$$

By a union bound, the event $\mathcal{E}_3 : \forall i, r, \left| \hat{\theta}_i^{(r)} - \theta_i \right| \leq \epsilon_r/4$ holds with probability at least $1 - \delta$.

It suffices to show that conditioned on event \mathcal{E}_3 , the algorithm does not make any error and terminates using the stated time and rounds. We prove this by induction on r . We have the following induction hypothesis:

1. $m_r = |\text{Top}_m \cap I_r|$,
2. $\text{Acc}_r \subseteq \text{Top}_m$ and $\text{Rej}_r \subseteq I \setminus \text{Top}_m$,
3. $\{i \in I \mid \Delta_i^{(m)} \geq 4\epsilon_r\} \cap I_{r+1} = \emptyset$.

Algorithm 12: Collaborative algorithm for fixed-confidence setting.

Input: a set of arms I , parameter m , and a confidence parameter δ .

Output: a set of top- m arms of I .

- 1 Initialize $I_0 \leftarrow I$, $m_0 \leftarrow m$, $Acc_0 \leftarrow \emptyset$, $Rej_0 \leftarrow \emptyset$, $r \leftarrow 0$, $T_{-1} \leftarrow 0$;
 - 2 for $r = 0, 1, \dots$, let $\epsilon_r = 2^{-(r+1)}$ and $T_r = 8 \log(4n(r+1)^2 \delta^{-1}) / (K \epsilon_r^2)$;
 - 3 **while** $I_r \neq \emptyset$ **do**
 - 4 each agent pulls each arm in I_r for $T_r - T_{r-1}$ times;
 - 5 for each $i \in I_r$, let $\hat{\theta}_i^{(r)}$ be the estimated mean of the i -th arm in I_r after KT_r pulls (over all rounds and agents so far);
 - 6 let $\pi_r : \{1, \dots, |I_r|\} \rightarrow I_r$ be the bijection such that $\hat{\theta}_{\pi_r(1)}^{(r)} \geq \hat{\theta}_{\pi_r(2)}^{(r)} \geq \dots \geq \hat{\theta}_{\pi_r(|I_r|)}^{(r)}$;
 - 7 $Acc_{r+1} \leftarrow Acc_r \cup \{i \in I_r : \hat{\theta}_i^{(r)} > \hat{\theta}_{\pi_r(m_{r+1})}^{(r)} + \epsilon_r\}$;
 - 8 $Rej_{r+1} \leftarrow Rej_r \cup \{i \in I_r : \hat{\theta}_i^{(r)} < \hat{\theta}_{\pi_r(m_r)}^{(r)} - \epsilon_r\}$;
 - 9 $m_{r+1} \leftarrow m - |Acc_{r+1}|$;
 - 10 $I_{r+1} \leftarrow I_r \setminus (Acc_{r+1} \cup Rej_{r+1})$;
 - 11 $r \leftarrow r + 1$;
 - 12 **return** Acc_r .
-

It is easy to see that the base case ($r = 0$) holds trivially. Let us assume that the hypothesis holds for round $(r - 1)$, and consider round r . By event \mathcal{E}_3 we have

$$\forall i \in I_r : \left| \hat{\theta}_i^{(r)} - \theta_i \right| \leq \epsilon_r/4. \quad (63)$$

And for any $a \in \{1, \dots, |I_r|\}$ we have

$$\theta_{[a]}(I_r) - \epsilon_r/4 \leq \hat{\theta}_{\pi_r(a)}^{(r)} \leq \theta_{[a]}^{(r)}(I_r) + \epsilon_r/4. \quad (64)$$

If $\hat{\theta}_i^{(r)} > \hat{\theta}_{\pi_r(m_{r+1})}^{(r)} + \epsilon_r$, in which case the algorithm adds i to Acc_{r+1} , then by (63) and (64) we have

$$\theta_i + \epsilon_r/4 \geq \hat{\theta}_i^{(r)} > \hat{\theta}_{\pi_r(m_{r+1})}^{(r)} + \epsilon_r \geq \theta_{[m_{r+1}]}(I_r) + 3\epsilon_r/4.$$

We thus have $\theta_i - \theta_{[m_{r+1}]}(I_r) > \epsilon_r/2$, which implies that $\theta_i \geq \theta_{[m_r]}(I_r)$. By the first and second items of the induction hypothesis, we have $i \in Top_m$, which implies $m_{r+1} = |Top_m \cap I_{r+1}|$ and $Acc_{r+1} \subseteq Top_m$. Similarly we can also show $Rej_{r+1} \subseteq I \setminus Top_m$.

We next consider the third item of the induction hypothesis. For an arm $i \in I_r \supseteq I_{r+1}$ such that $\Delta_i^{(m)} \geq 4\epsilon_r$ and $\theta_i \geq \theta_{[m]}$, we have

$$\begin{aligned} \hat{\theta}_i^{(r)} &\geq \theta_i - \epsilon_r/4 = \Delta_i^{(m)} + \theta_{[m+1]} - \epsilon_r/4 \\ &\geq \theta_{[m_{r+1}]}(I_r) + \Delta_i^{(m)} - \epsilon_r/4 \\ &\geq \hat{\theta}_{\pi_r(m_{r+1})}^{(r)} + \Delta_i^{(m)} - \epsilon_r/2 \\ &\geq \hat{\theta}_{\pi_r(m_{r+1})}^{(r)} + 2\epsilon_r. \end{aligned}$$

Thus the i -th item in I will be added into Acc_{r+1} , and thus will not appear in I_{r+1} . By the same line of arguments, we can show that for any arm $i \in I_r$, if $\Delta_i^{(m)} \geq 4\epsilon_r$ and $\theta_i \leq \theta_{[m+1]}$, then $i \in Rej_{r+1}$ and will not appear in I_{r+1} .

With the three items in the induction hypothesis, we prove the correctness of the algorithm and analyze its time and round complexities. By the definition of ϵ_r , when $r \geq r_0 = \log(4/\Delta_i^{(m)})$, we have $\{i \in I \mid \Delta_i^{(m)} \geq 4\epsilon_r\} = I$. We have the followings:

1. By the third item of the induction hypothesis, the algorithm will terminate in r_0 rounds.
2. By the second item of the induction hypothesis, we have $Acc_{r_0} = Top_m$.
3. Note that if $\epsilon_r \leq \Delta_i^{(m)}/4$, then each agent pulls the i -th arm for at most $T_r = 8 \log(4n(r+1)^2 \delta^{-1}) / (K \epsilon_r^2)$ times. Let $r(i) = \min_r \{\epsilon_r \leq \Delta_i^{(m)}/4\}$; we thus have $\Delta_i^{(m)}/8 \leq \epsilon_{r(i)} \leq \Delta_i^{(m)}/4$. By the third item of the induction hypothesis, each agent pulls the i -th arm for at most

$$T_{r(i)} \leq \frac{512}{K (\Delta_i^{(m)})^2} \log \left(\frac{16n}{\delta} \log^2 \left(4/\Delta_i^{(m)} \right) \right)$$

times. Therefore, the total running time is bounded by $\sum_{i \in I} T_{r(i)} = O \left(\frac{H^{(m)}}{K} \log \left(\frac{n}{\delta} \log H^{(m)} \right) \right)$.

□

Finally we comment on the lower bound. In [55] it was shown that for the special case when $m = 1$, to achieve a running time of $\tilde{O}(H^{(1)}/K)$ with success probability 0.99 one needs at least $\log(1/\Delta_{[1]}^{(1)})$ rounds. Therefore the upper bound in Theorem 40 is tight up to logarithmic factors.

References

- [1] Ittai Abraham, Omar Alonso, Vasilis Kandylas, and Aleksandrs Slivkins. Adaptive crowdsourcing algorithms for the bandit survey problem. In *COLT*, pages 882–910, 2013.
- [2] Arpit Agarwal, Shivani Agarwal, Sepehr Assadi, and Sanjeev Khanna. Learning with limited rounds of adaptivity: Coin tossing, multi-armed bandits, and ranking from pairwise comparisons. In *COLT*, pages 39–75, 2017.
- [3] Yossi Arjevani and Ohad Shamir. Communication complexity of distributed convex learning and optimization. In *NIPS*, pages 1756–1764, 2015.
- [4] Sepehr Assadi, Nikolai Karpov, and Qin Zhang. Distributed and streaming linear programming in low dimensions. In *PODS*, pages 236–253. ACM, 2019.
- [5] Jean-Yves Audibert, Sébastien Bubeck, and Rémi Munos. Best arm identification in multi-armed bandits. In *COLT*, pages 41–53, 2010.
- [6] Baruch Awerbuch and Robert D. Kleinberg. Competitive collaborative learning. In *COLT*, pages 233–248, 2005.

- [7] Yu Bai, Tengyang Xie, Nan Jiang, and Yu-Xiang Wang. Provably efficient q-learning with low switching cost. In *NeurIPS*, 2019.
- [8] Maria-Florina Balcan, Avrim Blum, Shai Fine, and Yishay Mansour. Distributed learning, communication complexity and privacy. In *COLT*, pages 26.1–26.22, 2012.
- [9] Ilai Bistriz and Amir Leshem. Distributed multi-player bandits - a game of thrones approach. In *NeurIPS*, pages 7222–7232, 2018.
- [10] Avrim Blum, Nika Haghtalab, Ariel D. Procaccia, and Mingda Qiao. Collaborative PAC learning. In *NIPS*, pages 2392–2401, 2017.
- [11] Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure exploration in multi-armed bandits problems. In *ALT*, pages 23–37, 2009.
- [12] Sébastien Bubeck, Tengyao Wang, and Nitin Viswanathan. Multiple identifications in multi-armed bandits. In *ICML*, pages 258–265, 2013.
- [13] Alexandra Carpentier and Andrea Locatelli. Tight (lower) bounds for the fixed budget best arm identification bandit problem. In *COLT*, pages 590–604, 2016.
- [14] Nicolò Cesa-Bianchi, Ofer Dekel, and Ohad Shamir. Online learning with switching costs and other adaptive adversaries. In *NIPS*, pages 1160–1168, 2013.
- [15] Nicolò Cesa-Bianchi, Claudio Gentile, Yishay Mansour, and Alberto Minora. Delay and cooperation in nonstochastic bandits. In *COLT*, pages 605–622, 2016.
- [16] Mithun Chakraborty, Kai Yee Phoebe Chua, Sanmay Das, and Brendan Juba. Coordinated versus decentralized exploration in multi-agent multi-armed bandits. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 164–170, 2017.
- [17] Jiecao Chen, Qin Zhang, and Yuan Zhou. Tight bounds for collaborative pac learning via multiplicative weights. In *NIPS*, pages 3602–3611, 2018.
- [18] Lijie Chen, Anupam Gupta, and Jian Li. Pure exploration of multi-armed bandit under matroid constraints. In *COLT*, pages 647–669, 2016.
- [19] Lijie Chen, Jian Li, and Mingda Qiao. Nearly instance optimal sample complexity bounds for top-k arm selection. In *AISTATS*, pages 101–110, 2017.
- [20] Lijie Chen, Jian Li, and Mingda Qiao. Towards instance optimal bounds for best arm identification. In *COLT*, pages 535–592, 2017.
- [21] Shouyuan Chen, Tian Lin, Irwin King, Michael R. Lyu, and Wei Chen. Combinatorial pure exploration of multi-armed bandits. In *NIPS*, pages 379–387, 2014.
- [22] Maria Dimakopoulou, Ian Osband, and Benjamin Van Roy. Scalable coordinated exploration in concurrent reinforcement learning. In *NeurIPS*, pages 4223–4232, 2018.
- [23] Maria Dimakopoulou and Benjamin Van Roy. Coordinated exploration in concurrent reinforcement learning. In *ICML*, pages 1270–1278, 2018.

- [24] Carlos Domingo, Ricard Gavaldà, and Osamu Watanabe. Adaptive sampling methods for scaling up knowledge discovery algorithms. *Data Min. Knowl. Discov.*, 6(2):131–152, 2002.
- [25] John C. Duchi, Feng Ruan, and Chulhee Yun. Minimax bounds on stochastic batched convex optimization. In *COLT*, pages 3065–3162, 2018.
- [26] Hossein Esfandiari, Amin Karbasi, Abbas Mehrabian, and Vahab S. Mirrokni. Batched multi-armed bandits with optimal regret. *CoRR*, abs/1910.04959, 2019.
- [27] Eyal Even-Dar, Shie Mannor, and Yishay Mansour. PAC bounds for multi-armed bandit and markov decision processes. In *COLT*, pages 255–270, 2002.
- [28] Zijun Gao, Yanjun Han, Zhimei Ren, and Zhengqing Zhou. Batched multi-armed bandits problem. In *NeurIPS*, 2019.
- [29] Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *COLT*, pages 998–1027, 2016.
- [30] Zhaohan Guo and Emma Brunskill. Concurrent PAC RL. In *AAAI*, pages 2624–2630, 2015.
- [31] Eshcar Hillel, Zohar Shay Karnin, Tomer Koren, Ronny Lempel, and Oren Somekh. Distributed exploration in multi-armed bandits. In *NIPS*, pages 854–862, 2013.
- [32] Hal Daumé III, Jeff M. Phillips, Avishek Saha, and Suresh Venkatasubramanian. Efficient protocols for distributed classification and optimization. In *ALT*, pages 154–168, 2012.
- [33] Shevtsova Irina. An improvement of convergence rate estimates in the lyapunov theorem. In *Doklady Mathematics*, volume 82, pages 862–864. Springer, 2010.
- [34] Kevin Jamieson, Matthew Malloy, Robert Nowak, and Sébastien Bubeck. lil^2 : An optimal exploration algorithm for multi-armed bandits. In *COLT*, pages 423–439, 2014.
- [35] Tianyuan Jin, Jieming Shi, Xiaokui Xiao, and Enhong Chen. Efficient pure exploration in adaptive round model. In *NeurIPS*, pages 6605–6614, 2019.
- [36] Kwang-Sung Jun, Kevin G. Jamieson, Robert D. Nowak, and Xiaojin Zhu. Top arm identification in multi-armed bandits with batch arm pulls. In *AISTATS*, pages 139–148, 2016.
- [37] Shivaram Kalyan Krishnan, Ambuj Tewari, Peter Auer, and Peter Stone. Pac subset selection in stochastic multi-armed bandits. In *ICML*, pages 227–234, 2012.
- [38] Varun Kanade, Zhenming Liu, and Bozidar Radunovic. Distributed non-stochastic experts. In *NIPS*, pages 260–268, 2012.
- [39] Daniel Kane, Roi Livni, Shay Moran, and Amir Yehudayoff. On communication complexity of classification problems. In *COLT*, pages 1903–1943, 2019.
- [40] Zohar Karnin, Tomer Koren, and Oren Somekh. Almost optimal exploration in multi-armed bandits. In *ICML*, pages 1238–1246, 2013.
- [41] Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *J. Mach. Learn. Res.*, 17:1:1–1:42, 2016.

- [42] Lloyd W Koenig and Averill M Law. A procedure for selecting a subset of size m containing the l best of k independent normal populations, with applications to simulation. *Communications in Statistics-Simulation and Computation*, 14(3):719–734, 1985.
- [43] Peter Landgren, Vaibhav Srivastava, and Naomi Ehrich Leonard. On distributed cooperative decision-making in multiarmed bandits. In *ECC*, pages 243–248, 2016.
- [44] Jasper CH Lee and Paul Valiant. Uncertainty about uncertainty: Near-optimal adaptive algorithms for estimating binary mixtures of unknown coins. *arXiv preprint arXiv:1904.09228*, 2019.
- [45] Keqin Liu and Qing Zhao. Distributed learning in multi-armed bandit with multiple players. *IEEE Trans. Signal Processing*, 58(11):5667–5681, 2010.
- [46] Shie Mannor and John N. Tsitsiklis. The sample complexity of exploration in the multi-armed bandit problem. *J. Mach. Learn. Res.*, 5:623–648, 2004.
- [47] Huy L. Nguyen and Lydia Zakyntinou. Improved algorithms for collaborative PAC learning. In *NeurIPS*, pages 7642–7650, 2018.
- [48] Vianney Perchet, Philippe Rigollet, Sylvain Chassang, and Erik Snowberg. Batched bandit problems. In *COLT*, page 1456, 2015.
- [49] Jonathan Rosenski, Ohad Shamir, and Liran Szlak. Multi-player bandits - a musical chairs approach. In *ICML*, pages 155–163, 2016.
- [50] Christian Schmidt, Jürgen Branke, and Stephen E. Chick. Integrating techniques from statistical ranking into evolutionary algorithms. In *Applications of Evolutionary Computing, EvoWorkshops 2006: EvoBIO, EvoCOMNET, EvoHOT, EvoIASP, EvoINTERACTION, EvoMUSART, and EvoSTOC*, pages 752–763, 2006.
- [51] David Silver, Leonard Newnham, David Barker, Suzanne Weller, and Jason McFall. Concurrent reinforcement learning from customer interactions. In *ICML*, pages 924–932, 2013.
- [52] Max Simchowitz, Kevin G. Jamieson, and Benjamin Recht. The simulator: Understanding adaptive sampling in the moderate-confidence regime. In *COLT*, pages 1794–1834, 2017.
- [53] Richard S. Sutton and Andrew G. Barto. *Reinforcement learning - an introduction*. Adaptive computation and machine learning. MIT Press, 1998.
- [54] Balázs Szörényi, Róbert Busa-Fekete, István Hegedűs, Róbert Ormándi, Márk Jelasity, and Balázs Kégl. Gossip-based distributed stochastic bandit algorithms. In *ICML*, pages 19–27, 2013.
- [55] Chao Tao, Qin Zhang, and Yuan Zhou. Collaborative learning with limited interaction: Tight bounds for distributed exploration in multi-armed bandits. In *FOCS*, pages 126–146, 2019.
- [56] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- [57] Santosh S. Vempala, Ruosong Wang, and David P. Woodruff. The communication complexity of optimization. *CoRR*, abs/1906.05832, 2019.

- [58] Jie Xu, Cem Tekin, Simpson Zhang, and Mihaela Van Der Schaar. Distributed multi-agent online learning based on global feedback. *IEEE Transactions on Signal Processing*, 63(9):2225–2238, 2015.
- [59] Yuchen Zhang, John C. Duchi, and Martin J. Wainwright. Communication-efficient algorithms for statistical optimization. In *NIPS*, pages 1511–1519, 2012.
- [60] Martin Zinkevich, Markus Weimer, Alexander J. Smola, and Lihong Li. Parallelized stochastic gradient descent. In *NIPS*, pages 2595–2603, 2010.

A Probability Tools

Lemma 41. Let $X_1, \dots, X_n \in [0, d]$ be independent random variables and $X = \sum_{i=1}^n X_i$. Then

$$\Pr[X > \mathbb{E}[X] + t] \leq \exp\left(-\frac{2t^2}{nd^2}\right) \quad \text{and} \quad \Pr[X < \mathbb{E}[X] - t] \leq \exp\left(-\frac{2t^2}{nd^2}\right).$$

Moreover, if $X_1, \dots, X_n \in [0, 1]$ and $\mu_L \leq \mathbb{E}[X] \leq \mu_H$, then we also have for every $\delta \in [0, 1]$,

$$\Pr[X \geq (1 + \delta)\mu_H] \leq \exp\left(-\frac{\delta^2\mu_H}{3}\right) \quad \text{and} \quad \Pr[X \leq (1 - \delta)\mu_L] \leq \exp\left(-\frac{\delta^2\mu_L}{3}\right).$$

Theorem 42 (Berry-Esseen, [33]). Let X_1, X_2, \dots, X_n be independent random variables with $\mathbb{E}[X_i] = 0$, $\mathbb{E}[X_i^2] = \sigma_i^2$ and $\mathbb{E}[|X_i|^3] \leq \rho$ for all $i \in [n]$. Let

$$S = \frac{X_1 + X_2 + \dots + X_n}{\sqrt{\sigma_1^2 + \sigma_2^2 + \dots + \sigma_n^2}}.$$

Let F be the cumulative distribution function of S , and Φ be the cumulative distribution function of the standard normal distribution. It holds that

$$\sup_{x \in \mathbb{R}} |F(x) - \Phi(x)| \leq 0.5601 \cdot \rho n \left(\sum_{i=1}^n \sigma_i^2 \right)^{-3/2}.$$