
On Top-k Selection in Multi-Armed Bandits and Hidden Bipartite Graphs

Wei Cao¹

Jian Li¹

Yufei Tao²

Zhize Li¹

¹Tsinghua University ²Chinese University of Hong Kong

¹fcao-w13@mails, lijian83@mail, zz-li14@mailsg.tsinghua.edu.cn ²taoyf@cse.cuhk.edu.hk

Abstract

This paper discusses how to efficiently choose from n unknown distributions the k ones whose means are the greatest by a certain metric, up to a small relative error. We study the topic under two standard settings—*multi-armed bandits* and *hidden bipartite graphs*—which differ in the nature of the input distributions. In the former setting, each distribution can be sampled (in the i.i.d. manner) an arbitrary number of times, whereas in the latter, each distribution is defined on a population of a finite size m (and hence, is fully revealed after m samples). For both settings, we prove lower bounds on the total number of samples needed, and propose optimal algorithms whose sample complexities match those lower bounds.

1 Introduction

This paper studies a class of problems that share a common high-level objective: from a number n of probabilistic distributions, find the k ones whose means are the greatest by a certain metric.

Crowdsourcing. A crowdsourcing algorithm (see recent works [1, 13] and the references therein) summons a certain number, say k , of individuals, called *workers*, to collaboratively accomplish a complex task. Typically, the algorithm breaks the task into a potentially very large number of micro-tasks, each of which makes a binary decision (*yes* or *no*) by taking the majority vote from the participating workers. Each worker is given an (often monetary) reward for every micro-task that s/he participates in. It is therefore crucial to identify the most reliable workers that have the highest rates of making correct decisions. Because of this, a crowdsourcing algorithm should ideally be preceded by an *exploration phase*, which selects the best k workers from n candidates by a series of “control questions”. Every control-question must be paid for in the same way as a micro-task. The challenge is to find the best workers with the least amount of money.

Frequent Pattern Discovery. Let B and W be two relations. Given a join predicate $Q(b; w)$, the *joining power* of a tuple $b \in B$ equals the number of tuples $w \in W$ such that b and w satisfy Q . A *top-k semi-join* [14, 17] returns the k tuples in B with the greatest joining power. This type of semi-joins is notoriously difficult to process when the evaluation of Q is complicated, and thus unfriendly to tailored-made optimization. A well-known example from graph databases is the discovery of frequent patterns [14], where B is a set of graph patterns, W a set of data graphs, and $Q(b; w)$ decides if a pattern b is a subgraph of a data graph w . In this case, top-k semi-join essentially returns the set of k graph patterns most frequently found in the data graphs. Given a black box for resolving subgraph isomorphism $Q(b; w)$, the challenge is to minimize the number of calls to the black box.

We refer to the reader to [14, 15] for more examples of difficult top-k semi-joins of this sort.

1.1 Problem Formulation

We first define four problems that capture the essence of the above applications.

Multi-Armed Bandit. We consider a standard setting of stochastic multi-armed bandit selection. Specifically, there is a bandit with a set B of n arms, where the i -th arm is associated with a Bernoulli

distribution with an unknown mean $\mu_i \in (0; 1]$. In each round, we choose an arm, pull it, and then collect a reward, which is an i.i.d. sample from the arm's reward distribution.

Given a subset $V \subseteq B$ of arms, we denote by $a_i(V)$ the arm with the i -th largest mean in V , and by $\mu_i(V)$ the mean of $a_i(V)$. Define $\mu_{\text{avg}}(V) = \frac{1}{k} \sum_{i=1}^k \mu_i(V)$, namely, the average of the means of the top- k arms in V .

Our first two problems aim to identify k arms whose means are the greatest either individually or aggregatively:

Problem 1 [Top- k Arm Selection (k -AS)] *Given parameters $\epsilon \in (0; \frac{1}{4})$; $\delta \in (0; \frac{1}{48})$, and $k, n=2$, we want to select a k -sized subset V of B such that, with probability at least $1 - \epsilon - \delta$, it holds that*

$$\mu_i(V) \geq (1 - \epsilon) \mu_i(B); \quad 8i \leq k:$$

Furthermore, we study a variation of k -AS where we change the multiplicative guarantee $\mu_i(V) \geq (1 - \epsilon) \mu_i(B)$ to an additive guarantee $\mu_i(V) \geq \mu_i(B) - \epsilon$ and keep all the other settings unchanged. We refer to the modified problem as Top- k_{add} Arm Selection (k_{add} -AS). Note that k_{add} -AS achieves a strictly stronger guarantee than EXPLORE- k (see Footnote 2). Since we have $\mu_i \in (0; 1]$, the multiplicative guarantee in k -AS is always stronger than the additive guarantee in k_{add} -AS.

Problem 2 [Top- k_{avg} Arm Selection (k_{avg} -AS)] *Given the same parameters as in k -AS, we want to select a k -sized subset V of B such that, with probability at least $1 - \epsilon - \delta$, it holds that*

$$\mu_{\text{avg}}(V) \geq (1 - \epsilon) \mu_{\text{avg}}(B):$$

For both problems, the *cost* of an algorithm is the total number of arms pulled, or equivalently, the total number of samples drawn from the arms' distributions. For this reason, we refer to the cost as the algorithm's *sample complexity*. It is easy to see that k -AS is more stringent than k_{avg} -AS; hence, a feasible solution to the former is also a feasible solution to the latter, but not the vice versa.

Hidden Bipartite Graph. The second main focus of the paper is the exploration of *hidden bipartite graphs*. Let $G = (B; W; E)$ be a bipartite graph, where the nodes in B are colored black, and those in W colored white. Set $n = |B|$ and $m = |W|$. The edge set E is *hidden* in the sense that an algorithm does not see any edge at the beginning. To find out whether an edge exists between a black vertex b and a white vertex w , the algorithm must perform a *probe* operation. The *cost* of the algorithm equals the number of such operations performed.

If an edge exists between b and w , we say that there is a *solid edge* between them; otherwise, we say that they have an *empty edge*. Let $\deg(b)$ be the degree of a black vertex b , namely, the number of solid edges of b . Given a subset of black vertices $V \subseteq B$, we denote by $b_i(V)$ the black vertex with i -th largest degree in V , and by $\deg_i(V)$ the degree of $b_i(V)$. Furthermore, define $\deg_{\text{avg}}(V) = \frac{1}{k} \sum_{i=1}^k \deg_i(V)$.

We now state the other two problems studied in this work, which aim to identify k black vertices whose degrees are the greatest either individually or aggregatively:

Problem 3 [k -Most Connected Vertex [14] (k -MCV)] *Given parameters $\epsilon \in (0; \frac{1}{4})$; $\delta \in (0; \frac{1}{48})$, and $k, n=2$, we want to select a k -sized subset V of B such that, with probability at least $1 - \epsilon - \delta$, it holds that*

$$\deg_i(V) \geq (1 - \epsilon) \deg_i(B); \quad 8i \leq k:$$

Problem 4 [k_{avg} -Most Connected Vertex (k_{avg} -MCV)] *Given the same parameters as in k -MCV, we want to select a k -sized subset V of B such that, with probability at least $1 - \epsilon - \delta$, it holds that*

$$\deg_{\text{avg}}(V) \geq (1 - \epsilon) \deg_{\text{avg}}(B):$$

A feasible solution to k -MCV is also feasible for k_{avg} -MCV, but not the vice versa. We will refer to the cost of an algorithm also as its *sample complexity*, by regarding a probe operation as "sampling" the edge probed. Note that there is no randomness in the given instance of k -MCV or k_{avg} -MCV. However, for any deterministic algorithm, the adversary can force the algorithm to always probe $\Omega(mn)$ edges. Hence, we only consider randomized algorithms.

Relevance and Discrepancy. One notable distinction of our formulations from most previous work (e.g., [6, 8, 11, 16]) is that we require a small multiplicative error instead an additive one. While this seems to be a technical nuance, we have several motivating reasons. First, an multiplicative

error is a more stringent requirement, in particular when the optimal solution is very small (this scenario is especially relevant in the hidden graph problems where the maximum degree can be much smaller than the total number of nodes). It may appear that we can always set the additive error $\epsilon' = \epsilon \cdot \mu_{\text{avg}}(B)$ ($\mu_{\text{avg}}(B)$ is the optimal value of $k_{\text{avg}}\text{-AS}$) and use the previous algorithms developed for additive errors. However, a priori estimate of the optimal value may not be available at all (e.g., in the semi-join application mentioned before) and a wrong estimate may lead to completely meaningless result. Lastly, even if a good estimation of the optimal value is known magically in advance, directly applying the previous algorithms would lead to suboptimal sample complexity (see Section 1.3 for details).

$k\text{-MCV}$ can be reduced to $k\text{-AS}$. Given a hidden bipartite graph $(B; W; E)$, we can treat every black vertex $b \in B$ as an “arm” associated with a Bernoulli reward distribution: the reward is 1 with probability $\deg(b)/m$ (recall $m = |W|$), and 0 with probability $1 - \deg(b)/m$. Any algorithm A for $k\text{-AS}$ can be deployed to solve $k\text{-MCV}$ as follows. Whenever A samples from arm b , we randomly choose a white vertex $w \in W$, and probe the edge between b and w . A reward of 1 is returned to A if and only if the edge exists.

$k\text{-AS}$ and $k\text{-MCV}$ differ, however, in the size of the population that a reward distribution is defined on. For $k\text{-AS}$, the reward of each arm is sampled from a population of an *indefinite* size, which can even be infinite. Consequently, $k\text{-AS}$ nicely models situations such as the crowdsourcing application mentioned earlier.

For $k\text{-MCV}$, the reward distribution of each “arm” (i.e., a black vertex b) is defined on a population of size $m = |W|$ (i.e., the edges of b). This has three implications. First, $k\text{-MCV}$ is a better modeling of applications like top- k semi-join (where an edge exists between $b \in B$ and $w \in W$ if and only if $Q(b; w)$ is true). Second, the problem admits an obvious algorithm with cost $O(nm)$ (recall $n = |B|$): simply probe all the hidden edges. Third, an algorithm never needs to probe the same edge between b and w twice—once probed, whether the edge is solid or empty is perpetually revealed. We refer to the last implication as the *history-awareness property*.

Remark. The above discussion on $k\text{-AS}$ and $k\text{-MCV}$ also applies to $k_{\text{avg}}\text{-AS}$ and $k_{\text{avg}}\text{-MCV}$. For each of above problems, we refer to an algorithm which achieves the precision and failure requirements prescribed by ϵ and δ as an $(\epsilon; \delta)$ -*approximate* algorithm.

1.2 Previous Results

Problem 1. Sheng et al. [14] presented an algorithm¹ that solves $k\text{-AS}$ with expected cost $O(\frac{n}{\epsilon^2} \frac{1}{\theta_k(B)} \log \frac{n}{\delta})$.

No lower bound is known on the sample complexity of $k\text{-AS}$. The closest work is due to Kalyanakrishnan et al. [11]. They considered the $\text{EXPLORE-}k$ problem, where the goal is to return a set V of k arms such that, with probability at least $1 - \delta$, the mean of each arm in V is at least $\mu_k(B) - \epsilon'$. They showed an optimal algorithm with sample complexity $\Theta(\frac{n}{\epsilon'^2} \log \frac{k}{\delta})$ in expectation by establishing a matching lower bound. Note that $\text{EXPLORE-}k$ ensures an *absolute-error* guarantee, which is weaker than the *individually relative-error* guarantee of $k\text{-AS}$. Therefore, the same $\text{EXPLORE-}k$ lower bound also applies to $k\text{-AS}$.

The readers may be tempted to set $\epsilon' = \epsilon \cdot \mu_k(B)$ to derive a “lower bound” of $\Omega(\frac{n}{\epsilon^2} \frac{1}{(\theta_k(B))^2} \log \frac{k}{\delta})$ for $k\text{-AS}$. This, however, is clearly wrong because when $\mu_k(B) = o(1)$ (a typical case in practice) this “lower bound” may be even higher than the upper bound of [14] mentioned earlier. The cause of the error lies in that the hard instance constructed in [11] requires $\mu_k(B) = \Omega(1)$.

Problem 2. The $O(\frac{n}{\epsilon^2} \frac{1}{\theta_k(B)} \log \frac{n}{\delta})$ upper bound of [14] on $k\text{-AS}$ carries over to $k_{\text{avg}}\text{-AS}$ (which, as mentioned before, can be solved by any $k\text{-AS}$ algorithm).

Zhou et al. [16] considered an OPTMAI problem whose goal is to find a k -sized subset V such that $\mu_{\text{avg}}(V) \geq \mu_{\text{avg}}(B) - \epsilon'$ holds with probability at least $1 - \delta$. Note, once again, that this is an absolute-error guarantee, as opposed to the relative-error guarantee of $k_{\text{avg}}\text{-AS}$. For OPTMAI , Zhou et al. presented an algorithm with sample complexity $O(\frac{n}{\epsilon'^2} (1 + \frac{\log(1/\delta)}{k}))$ in expectation. Observe that if $\mu_{\text{avg}}(B)$ is available *magically* in advance, we can immediately apply the OPTMAI algorithm of [16] to settle $k_{\text{avg}}\text{-AS}$ by setting $\epsilon' = \epsilon \cdot \mu_{\text{avg}}(B)$. The expected cost of the algorithm becomes

¹The algorithm was designed for $k\text{-MCV}$, but it can be adapted to $k\text{-AS}$ as well.

Table 1: Comparison of our and previous results (all bounds are in expectation)

problem		sample complexity	source
k -AS	upper bound	$O\left(\frac{n}{\epsilon^2} \frac{1}{\theta_k(B)} \log \frac{n}{\delta}\right)$ $O\left(\frac{n}{\epsilon^2} \frac{1}{\theta_k(B)} \log \frac{k}{\delta}\right)$	[14] new
	lower bound	$\Omega\left(\frac{n}{\epsilon^2} \log \frac{k}{\delta}\right)$ $\Omega\left(\frac{n}{\epsilon^2} \frac{1}{\theta_k(B)} \log \frac{k}{\delta}\right)$	[11] new
k_{avg} -AS	upper bound	$O\left(\frac{n}{\epsilon^2} \frac{1}{\theta_k(B)} \log \frac{n}{\delta}\right)$	[14]
		$O\left(\frac{n}{\epsilon^2} \frac{1}{(\theta_{\text{avg}}(B))^2} \left(1 + \frac{\log(1/\delta)}{k}\right)\right)$ $O\left(\frac{n}{\epsilon^2} \frac{1}{\theta_{\text{avg}}(B)} \left(1 + \frac{\log(1/\delta)}{k}\right)\right)$	[16] new
	lower bound	$\Omega\left(\frac{n}{\epsilon^2} \left(1 + \frac{\log(1/\delta)}{k}\right)\right)$ $\Omega\left(\frac{n}{\epsilon^2} \frac{1}{\theta_{\text{avg}}(B)} \left(1 + \frac{\log(1/\delta)}{k}\right)\right)$	[16] new
k -MCV	upper bound	$O\left(\min\left\{\frac{n}{\epsilon^2} \frac{m}{\deg_k(B)} \log \frac{n}{\delta}; nm\right\}\right)$ $O\left(\min\left\{\frac{n}{\epsilon^2} \frac{m}{\deg_k(B)} \log \frac{k}{\delta}; nm\right\}\right)$	[14] new
	lower bound	$\begin{cases} \Omega\left(\frac{n}{\epsilon^2} \frac{m}{\deg_k(B)} \log \frac{k}{\delta}\right) & \text{if } \deg_k(B) \geq \Omega\left(\frac{1}{\epsilon^2} \log \frac{n}{\delta}\right) \\ \Omega(nm) & \text{if } \deg_k(B) < O\left(\frac{1}{\epsilon}\right) \end{cases}$	new
k_{avg} -MCV	upper bound	$O\left(\min\left\{\frac{n}{\epsilon^2} \frac{m}{\deg_k(B)} \log \frac{n}{\delta}; nm\right\}\right)$ $O\left(\min\left\{\frac{n}{\epsilon^2} \frac{m^2}{(\deg_{\text{avg}}(B))^2} \left(1 + \frac{\log(1/\delta)}{k}\right); nm\right\}\right)$ $O\left(\min\left\{\frac{n}{\epsilon^2} \frac{m}{\deg_{\text{avg}}(B)} \left(1 + \frac{\log(1/\delta)}{k}\right); nm\right\}\right)$	[14] [16] new
	lower bound	$\begin{cases} \Omega\left(\frac{n}{\epsilon^2} \frac{m}{\deg_{\text{avg}}(B)} \left(1 + \frac{\log(1/\delta)}{k}\right)\right) & \text{if } \deg_{\text{avg}}(B) \geq \Omega\left(\frac{1}{\epsilon^2} \log \frac{n}{\delta}\right) \\ \Omega(nm) & \text{if } \deg_{\text{avg}}(B) < O\left(\frac{1}{\epsilon}\right) \end{cases}$	new

$O\left(\frac{n}{\epsilon^2} \frac{1}{(\theta_{\text{avg}}(B))^2} \left(1 + \frac{\log(1/\delta)}{k}\right)\right)$ (which is suboptimal. See the table).

No lower bound is known on the sample complexity of k_{avg} -AS. For OPTMAI, Zhou et al. [16] proved a lower bound of $\Omega\left(\frac{n}{\epsilon^2} \left(1 + \frac{\log(1/\delta)}{k}\right)\right)$, which directly applies to k_{avg} -AS due to its stronger quality guarantee.

Problems 3 and 4. Both problems can be trivially solved with cost $O(nm)$. Furthermore, as explained in Section 1.1, k -MCV and k_{avg} -MCV can be reduced to k -AS and k_{avg} -AS respectively. Indeed, the best existing k -AS and k_{avg} -AS algorithms (surveyed in the above) serve as the state of the art for k -MCV and k_{avg} -MCV, respectively.

Prior to this work, no lower bound results were known for k -MCV and k_{avg} -MCV. Note that *none* of the lower bounds for k -AS (or k_{avg} -AS) is applicable to k -MCV (or k_{avg} -MCV, resp.), because there is no reduction from the former problem to the latter.

1.3 Our Results

We obtain tight upper and lower bounds for all of the problems defined in Section 1.1. Our main results are summarized in Table 1 (all bounds are in expectation). Next, we explain several highlights, and provide an overview into our techniques.

k -AS. Our algorithm improves the $\log n$ factor of [14] to $\log k$ (in practice $k \ll n$), thereby achieving the optimal sample complexity (Theorem 1).

For k_{add} -AS, we can solve it with the sample complexity $\Omega\left(\frac{n}{\epsilon^2} \log \frac{k}{\delta}\right)$ by slightly modifying our algorithm for k -AS (Theorem 3). Note that the sample complexity of the modified algorithm is exactly the same as the algorithm in [10] for EXPLORE- k , however our algorithm achieves a strictly stronger guarantee than EXPLORE- k^2 .

² Our problem guarantees that $\theta_i(V) \geq \theta_i(B) - \epsilon'$ for all $i \in [k]$ while EXPLORE- k guarantees that $\theta_i(V) \geq \theta_k(B) - \epsilon'$ for all $i \in [k]$.

Our analysis for k -AS is inspired by [8, 10, 11] (in particular the *median elimination* technique in [8]). However, the details are very different and more involved than the previous ones (the application of median elimination of [8] was in a much simpler context where the analysis was considerably easier). On the lower bound side, ours is based on and very similar to the argument of [11]. We need to get rid of the $\deg_k(B) = \Omega(1)$ assumption (as explained in Section 1.2), which requires several changes in the analysis (Theorem 2).

k_{avg} -AS. Our algorithm improves both existing solutions in [14, 16] significantly, noticing that both $\deg_k(B)$ and $(\deg_{\text{avg}}(B))^2$ are never larger, but can be far smaller, than $\deg_{\text{avg}}(B)$. This improvement results from an enhanced version of median elimination, and once again, requires a non-trivial analysis specific to our context (Theorem 5). Our lower bound is established with a novel reduction from the 1-AS problem (Theorem 6). It is worth nothing that the reduction can be used to simplify the proof of the lower bound in [16, Theorem 5.5].

k -MCV and k_{avg} -MCV. The stated upper bounds for k -MCV and k_{avg} -MCV in Table 1 can be obtained directly from our k -AS and k_{avg} -AS algorithms. In contrast, all the lower-bound arguments for k -AS and k_{avg} -AS—which crucially rely on the samples being i.i.d.—break down for the two MCV problems, due to the history-awareness property explained in Section 1.1.

For k -MCV, we remedy the issue by (i) (when $\deg_k(B)$ is large) a reduction from k -AS, and (ii) (when $\deg_k(B)$ is small) a reduction from a sampling lower bound for distinguishing two extremely similar distributions (Theorem 4). Analogous ideas are deployed for k_{avg} -MCV (Theorem 7). Note that for a small range of $\deg_k(B)$ (i.e., $\Omega(\frac{1}{\epsilon}) < \deg_k(B) < O(\frac{1}{\epsilon^2} \log \frac{n}{\delta})$), we do not have the optimal lower bounds yet for k -MCV and k_{avg} -MCV. Closing the gap is left as an interesting open problem.

2 Top- k Arm Selection

In this section, we describe a new algorithm for the k -AS problem and its variation k_{add} -AS. We present the detailed analysis in Appendix B.

Our k -AS algorithm consists of three components: ME-AS, *Median Elimination* (ME), and *Uniform Sampling* (US), as shown in Algorithms 1, 2, and 3, respectively.

Given parameters B , ϵ , δ , k (as in Problem 1), ME-AS takes a “guess” $\hat{a}_k$ (Line 2) on the value of $\deg_k(B)$, and then applies ME (Line 3) to prune B down to a set S of at most $4k$ arms. Then, at Line 4, US is invoked to process S . At Line 5, (as will be clear shortly) the value of $\hat{a}_k^{\text{US}}$ is what ME-AS thinks should be the value of $\deg_k(B)$; thus, the algorithm performs a quality check to see whether $\hat{a}_k^{\text{US}}$ is larger than but close to $\hat{a}_k$. If the check fails, ME-AS halves its guess $\hat{a}_k$ (Line 2), and repeats the above steps; otherwise, the output of US from Line 4 is returned as the final result.

ME runs in rounds. Round ℓ ($\ell = 1; 2; \dots$) is controlled by parameters S_ℓ , Q_ℓ , ϵ_ℓ , and δ_ℓ (their values for Round 1 are given at Line 1). In general, S_ℓ is the set of arms from which we still want to sample. For each arm $a \in S_\ell$, ME takes Q_ℓ (Line 4) samples from a , and calculates its *empirical value* $\hat{a}_a^{\text{ME}}$ (Lines 5 and 6). ME drops (at Lines 7 and 8) half of the arms in S_ℓ with the smallest empirical values, and then (at Line 9) sets the parameters of the next round. ME terminates by returning S_ℓ as soon as $|S_\ell|$ is at most $4k$ (Lines 3 and 10).

US simply takes Q samples from each arm $a \in S$ (Line 2), and calculates its *US-empirical value* $\hat{a}_a^{\text{US}}$ (Lines 3 and 4). Finally, US returns the k arms in S with the largest US-empirical values (Lines 5 and 6).

Remark. If we ignore Line 3 of Algorithm 1 and simply set $S = B$, then ME-AS degenerates into the algorithm in [14].

Theorem 1 *ME-AS solves the k -AS problem with expected cost $O\left(\frac{n}{\epsilon^2} \frac{1}{\deg_k(B)} \log \frac{k}{\delta}\right)$.*

Furthermore, we extend the proof in [11] and establish the lower bound for k -AS problem as shown in Theorem 2.

Theorem 2 *For any $\epsilon \in (0; \frac{1}{4})$ and $\delta \in (0; \frac{1}{48})$, given any algorithm, there is an instance of the k -AS problem on which the algorithm must entail $\Omega\left(\frac{n}{\epsilon^2} \frac{1}{\deg_k(B)} \log \frac{k}{\delta}\right)$ cost in expectation.*

Algorithm 1: ME-AS

```
1 input:  $B, \epsilon, \delta, k$ 
2 for  $\ell = 1; 2; 4; \dots; k$  do
3    $S = \text{ME}(B; \epsilon; \delta; k)$ ;
4    $f(a_i; \hat{US}(a_i))_{j=1}^i \quad kg = \text{US}(S; \epsilon; (1-\epsilon/2); k)$ ;
5   if  $\hat{US}(a_k) \geq 2$  then
6     return  $f(a_1; \dots; a_k)g$ ;
```

Algorithm 2: Median Elimination (ME)

```
1 input:  $B, \epsilon, \delta, k$ 
2  $S_1 = B, \ell_1 = \lceil 16/\epsilon \rceil, \ell_1 = \lceil 8/\delta \rceil$ , and  $\ell = 1$ ;
3 while  $jS_\ell \geq 4k$  do
4   sample every arm  $a \in S_\ell$  for  $Q_\ell = (12/\epsilon^2)(1-\epsilon)\log(6k/\epsilon)$  times;
5   for each arm  $a \in S_\ell$  do
6     its empirical value  $\hat{a} =$  the average of the  $Q_\ell$  samples from  $a$ ;
7    $a_1; \dots; a_{|S_\ell|}$  = the arms sorted in non-increasing order of their empirical values;
8    $S_{\ell+1} = f(a_1; \dots; a_{|S_\ell|/2})g$ ;
9    $\ell_{\ell+1} = 3\ell/4, \ell_{\ell+1} = \ell/2, \ell_{\ell+1} = (1-\epsilon)\ell$ , and  $\ell = \ell + 1$ ;
10 return  $S_\ell$ ;
```

Algorithm 3: Uniform Sampling (US)

```
1 input:  $S; \epsilon; \delta; k$ 
2 sample every arm  $a \in S$  for  $Q = (96/\epsilon^2)(1-\epsilon)\log(4jS/\epsilon)$  times;
3 for each arm  $a \in S$  do
4   its US-empirical value  $\hat{US}(a)$  = the average of the  $Q$  samples from  $a$ ;
5    $a_1; \dots; a_{|S|}$  = the arms sorted in non-increasing order of their US-empirical values;
6 return  $f(a_1; \hat{US}(a_1)); \dots; (a_k; \hat{US}(a_k))g$ 
```

For $k_{\text{add-AS}}$, we simply drop the “guess” stage of ME-AS and adjust the parameters appropriately. We call such the algorithm Additive Median Elimination (AME). We present the detailed algorithm and analysis for AME in Appendix B.4.

Theorem 3 *AME solves the $k_{\text{add-AS}}$ problem with expected cost $O\left(\frac{n}{\epsilon^2} \log \frac{k}{\delta}\right)$.*

3 k -MOST CONNECTED VERTEX

This section is devoted to the k -MCV problem (Problem 3). We will focus on lower bounds because our k -AS algorithm in the previous section also settles k -MCV with the cost claimed in Table 1 by applying the reduction described in Section 1.1. We establish matching lower bounds below:

Theorem 4 *For any $\epsilon \in (0; \frac{1}{12})$ and $\delta \in (0; \frac{1}{48})$, the following statements are true about any k -MCV algorithm:*

when $\deg_k(B) \geq \Omega\left(\frac{1}{\epsilon^2} \log \frac{n}{\delta}\right)$, there is an instance on which the algorithm must probe $\Omega\left(\frac{n}{\epsilon^2} \frac{m}{\deg_k(B)} \log \frac{k}{\delta}\right)$ edges in expectation.

when $\deg_k(B) < O\left(\frac{1}{\epsilon}\right)$, there is an instance on which the algorithm must probe $\Omega(nm)$ edges in expectation.

For large $\deg_k(B)$ in Theorem 4, we utilize an instance for k -AS to construct a random hidden bipartite graph and feed it to any algorithm solves k -MCV. By doing this, we reduce k -AS to k -MCV and thus, establish our first lower bound.

For small $\deg_k(B)$, we define the *single-vertex* problem where the goal is to distinguish two extremely distributions. We prove the lower bound of single-vertex problem and reduce it to k -MCV. Thus, we establish our second lower bound. The details are presented in Appendix C.

4 Top- k_{avg} Arm Selection

Our k_{avg} -AS algorithm QE-AS is similar to ME-AS described in Section 2, except that the parameters are adjusted appropriately. We present the details in Appendix D.

Theorem 5 *QE-AS solves the k_{avg} -AS problem with expected cost $O\left(\frac{n}{\epsilon^2} \frac{1}{\theta_{\text{avg}}(B)} \left(1 + \frac{\log(1/\delta)}{k}\right)\right)$.*

We establish the lower bound for k_{avg} -AS as shown in Theorem 6.

Theorem 6 *For any $\epsilon \in (0; \frac{1}{12})$ and $\delta \in (0; \frac{1}{48})$, given any $(\epsilon; \delta)$ -approximate algorithm, there is an instance of the k_{avg} -AS problem on which the algorithm must entail $\Omega\left(\frac{n}{\epsilon^2} \frac{1}{\theta_{\text{avg}}(B)} \left(1 + \frac{\log(1/\delta)}{k}\right)\right)$ cost in expectation.*

We show that the lower bound of k_{avg} -AS is the maximum of $\Omega\left(\frac{n}{\epsilon^2} \frac{1}{\theta_{\text{avg}}(B)} \frac{\log(1/\delta)}{k}\right)$ and $\Omega\left(\frac{n}{\epsilon^2} \frac{1}{\theta_{\text{avg}}(B)}\right)$. Our proof of the first lower bound is based on a novel reduction from 1-AS. We stress that our reduction can be used to simplify the proof of the lower bound in [16, Theorem 5.5].

5 k_{avg} -MOST CONNECTED VERTEX

Our k_{avg} -AS algorithm, combined with the reduction described in Section 1.1, already settles k_{avg} -MCV with the sample complexity given in Table 1. We establish the following lower bound and prove it in Appendix E.

Theorem 7 *For any $\epsilon \in (0; \frac{1}{12})$ and $\delta \in (0; \frac{1}{48})$, the following statements are true about any k_{avg} -MCV algorithm:*

when $\deg_{\text{avg}}(B) \geq \Omega\left(\frac{1}{\epsilon^2} \log \frac{n}{\delta}\right)$, there is an instance on which the algorithm must probe

$$\Omega\left(\frac{n}{\epsilon^2} \frac{m}{\deg_{\text{avg}}(B)} \left(1 + \frac{\log(1/\delta)}{k}\right)\right)$$

edges in expectation.

when $\deg_k(B) < O(\frac{1}{\epsilon})$, there is an instance on which the algorithm must probe $\Omega(nm)$ edges in expectation.

6 Experiment Evaluation

We first compare ME-AS and AMCV [14] for k -MCV problem and we provide the additional experimental results in Appendix F. We use two synthetic data sets and one real world data set to evaluate the algorithms. Each data is represented as a bipartite graph with $n = m = 5000$. For the synthetic data, the degrees of the black vertices follow a power law distribution. For each black vertex $b \in B$, its degree equals d with probability $c(d+1)^{-\tau}$ where τ is the parameter to be set and c is the normalizing factor. We build two bipartite graphs by setting the proper parameters in order to control the average degrees equal to 50 and 3000 correspondingly. For the real world data, we crawl 5000 active users from twitter with their corresponding relationships. We construct a bipartite graph $G = (B; W; E)$ where each of B and W represent all the users and E represents the 2-hop relationships. We say two users $b \in B$ and $w \in W$ has a 2-hop relationship if they share at least one common friend.

As the theoretical analysis is rather pessimistic due to the extensive usage of the union bound, to make a fair comparison, we adopt the same strategy as in [14], i.e., to divide the sample cost in theory by a heuristic constant β . We use the same parameter $\beta = 2000$ for AMCV as in [14]. For ME-AS, we first take $\beta = 10^7$ for each round of the median elimination step and then we use the previous sample cost dividing 250 as the samples of the uniform sampling step. Notice that it does not conflict the theoretical sample complexity since the median elimination step dominates the sample complexity of the algorithm.

We fix the parameters $\epsilon = 0.1; k = 20$ and enumerate δ from 0.01 to 0.1. We then calculate the actual failure probability by counting the successful runs in 100 repeats. Recall that due to the heuristic nature, the algorithm may not achieve the theoretical guarantees prescribed by $(\epsilon; \delta)$. Whenever this happens, we label the percentage of actual error ϵ_a it achieves according to the failure




Figure 1: Performance comparison for k -MCV vs.

probability ϵ . For example 2:9 means the algorithm actually achieves an error $\epsilon_a = 0:029$ with failure probability ϵ . The experiment result is shown in Fig 1.

As we can see, ME-AS outperforms AMCV in both sample complexity and the actual error in all data sets. We stress that in the worst case, it seems ME-AS only shows a different when $n = k$. However for the most of the real world data, the degrees of the vertices usually follow a power law distribution or a Gaussian distribution. For such cases, our algorithm only need to take a few samples in each round of the elimination step and drop half of vertices with high confidence. Therefore, the experimental result shows that the sample cost of ME-AS is much less than AMCV.

7 Related Work

Multi-armed bandit problems are classical decision problems with exploration-exploitation trade-offs, and have been extensively studied for several decades (dating back to 1930s). In this line of research, k -AS and k_{avg} -AS fit into the *pure exploration* category, which has attracted significant attentions in recent years due to its abundant applications such as online advertisement placement [6], channel allocation for mobile communications [2], crowdsourcing [16], etc. We mention some closely related work below, and refer the interested readers to a recent survey [4].

Even-Dar et al. [8] proposed an optimal algorithm for selecting a single arm which approximates the best arm with an additive error at most ϵ (a matching lower bound was established by Mannor et al. [12]). Kalyanakrishnan et al. [10, 11] considered the EXPLORE- k problem which we mentioned in Section 1.2. They provided an algorithm with the sample complexity $O(\frac{n}{\epsilon^2} \log \frac{k}{\delta})$. Similarly, Zhou et al. [16] studied the OPTMAI problem which, again as mentioned in Section 1.2, is the absolute-error version of k_{avg} -AS.

Audibert et al. [2] and Bubeck et al. [4] investigated the *fixed budget* setting where, given a fixed number of samples, we want to minimize the so-called *misidentification probability* (informally, the probability that the solution is not optimal). Bubeck et al. [5] also showed the links between the simple regret (the gap between the arm we obtain and the best arm) and the cumulative regret (the gap between the reward we obtained and the expected reward of the best arm). Gabillon et al. [9] provide a unified approach UGapE for EXPLORE- k in both the fixed budget and the fixed confidence settings. They derived the algorithms based on “lower and upper confidence bound” (LUCB) where the time complexity depends on the gap between $k(B)$ and the other arms. Note that each time LUCB samples the two arms that are most difficult to distinguish. Since our problem ensures an individually guarantee, it is unclear whether only sampling the most difficult-to-distinguish arms would be enough. We leave it as an intriguing direction for future work. Chen et al. [6] studied how to select the best arms under various combinatorial constraints.

Acknowledgements. Jian Li, Wei Cao, Zizhe Li were supported in part by the National Basic Research Program of China grants 2015CB358700, 2011CBA00300, 2011CBA00301, and the National NSFC grants 61202009, 61033001, 61361136003. The research of Yufei Tao was supported in part by projects GRF 4168/13 and GRF 142072/14 from HKRGC.

References

- [1] Y. Amsterdamer, S. B. Davidson, T. Milo, S. Novgorodov, and A. Somech. OASSIS: query driven crowd mining. In *SIGMOD*, pages 589–600, 2014.
- [2] J.-Y. Audibert, S. Bubeck, et al. Best arm identification in multi-armed bandits. *COLT*, 2010.
- [3] Z. Bar-Yossef. *The complexity of massive data set computations*. PhD thesis, University of California, 2002.
- [4] S. Bubeck, N. Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and trends in machine learning*, 5(1):1–122, 2012.
- [5] S. Bubeck, R. Munos, and G. Stoltz. Pure exploration in finitely-armed and continuous-armed bandits. *Theoretical Computer Science*, 412(19):1832–1852, 2011.
- [6] S. Chen, T. Lin, I. King, M. R. Lyu, and W. Chen. Combinatorial pure exploration of multi-armed bandits. In *Advances in Neural Information Processing Systems*, pages 379–387, 2014.
- [7] D. P. Dubhashi and A. Panconesi. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.
- [8] E. Even-Dar, S. Mannor, and Y. Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *The Journal of Machine Learning Research*, 7:1079–1105, 2006.
- [9] V. Gabillon, M. Ghavamzadeh, and A. Lazaric. Best arm identification: A unified approach to fixed budget and fixed confidence. In *Advances in Neural Information Processing Systems*, pages 3212–3220, 2012.
- [10] S. Kalyanakrishnan and P. Stone. Efficient selection of multiple bandit arms: Theory and practice. In *ICML*, pages 511–518, 2010.
- [11] S. Kalyanakrishnan, A. Tewari, P. Auer, and P. Stone. PAC subset selection in stochastic multi-armed bandits. In *ICML*, pages 655–662, 2012.
- [12] S. Mannor and J. N. Tsitsiklis. The sample complexity of exploration in the multi-armed bandit problem. *The Journal of Machine Learning Research*, 5:623–648, 2004.
- [13] A. G. Parameswaran, S. Boyd, H. Garcia-Molina, A. Gupta, N. Polyzotis, and J. Widom. Optimal crowd-powered rating and filtering algorithms. *PVLDB*, 7(9):685–696, 2014.
- [14] C. Sheng, Y. Tao, and J. Li. Exact and approximate algorithms for the most connected vertex problem. *TODS*, 37(2):12, 2012.
- [15] J. Wang, E. Lo, and M. L. Yiu. Identifying the most connected vertices in hidden bipartite graphs using group testing. *TKDE*, 25(10):2245–2256, 2013.
- [16] Y. Zhou, X. Chen, and J. Li. Optimal PAC multiple arm identification with applications to crowdsourcing. In *ICML*, pages 217–225, 2014.
- [17] M. Zhu, D. Papadias, J. Zhang, and D. L. Lee. Top-k spatial joins. *TKDE*, 17(4):567–579, 2005.

Organization

The appendix is organized as follows. We first review several useful forms of Chernoff bounds in Appendix A. In Appendix B, we analyze our algorithm for k -AS and the corresponding lower bound (Theorem 1 and 2). We also provide the algorithm and the analysis for k_{add} -AS in this section (Theorem 3). In Appendix C, we provide the deferred proof of our lower bound for k -MCV. Next, we provide the deferred analysis of algorithm and the proof of the lower bound for k_{avg} -AS (Theorem 5 and 6) in Appendix D. In Appendix E, we prove our lower bound for k_{avg} -MCV (Theorem 7). Finally, we present the additional experiment evaluation in Appendix F.

A Chernoff Bounds

Proposition 1 ([7]) *Let $X_1, \dots, X_n$ be independent random variables with $\Pr[X_i = 1] = p_i$, and $\Pr[X_i = 0] = 1 - p_i$. Define $X = \sum_{i=1}^n X_i$. Clearly, $\mathbb{E}[X] = \sum_{i=1}^n p_i$. Then, for $0 < \epsilon < 1$:*

$$\Pr[X \geq (1 + \epsilon)\mathbb{E}(X)] \leq \exp\left(-\frac{\epsilon^2 \mathbb{E}(X)}{3}\right); \quad (1)$$

$$\Pr[X \leq (1 - \epsilon)\mathbb{E}(X)] \leq \exp\left(-\frac{\epsilon^2 \mathbb{E}(X)}{2}\right); \quad (2)$$

For $\epsilon \leq 1$:

$$\Pr[X \geq (1 + \epsilon)\mathbb{E}(X)] \leq \exp\left(-\frac{(1 + \epsilon)\mathbb{E}(X)}{6}\right); \quad (3)$$

For $\epsilon > 0$:

$$\Pr[X \geq (1 + \epsilon)\mathbb{E}(X)] \leq \left(\frac{e}{1 + \epsilon}\right)^{(1 + \epsilon)\mathbb{E}(X)}; \quad (4)$$

$$\Pr[X \leq \mathbb{E}(X) - \epsilon] \leq \exp\left(-\frac{\epsilon^2}{2\mathbb{E}(X)}\right); \quad (5)$$

$$\Pr[jX \leq \mathbb{E}(X)j - \epsilon] \leq 2\exp(-\epsilon^2/2n); \quad (6)$$

Moreover, suppose $\mathbb{E}[X_i] < a$ for some real $0 < a < 1$. For every $t > 0$, we have that

$$\Pr[X = n > a + t] < \left(\left(\frac{a}{a + t}\right)^{a + t} \left(\frac{1}{1 - \frac{a}{a + t}}\right)^{1 - \frac{a}{a + t}}\right)^n \quad (7)$$

B Analysis of Top- k Arm Selection

In this section we analyze our k -AS algorithm. We first analyze the Median Elimination in Appendix B.1. In Appendix B.2, we analyze the correctness and the sample complexity which prove Theorem 1. Then, we prove Theorem 2 in Appendix B.3. Finally, we present the algorithm and the analysis of k_{add} -AS in Appendix B.4.

B.1 Analysis of Median Elimination

This subsection serves as a proof for the following lemma, which says that if the guess (passed in by ME-AS) is low enough, with a large probability, ME returns a good arm set S that preserves approximately the k best arms in B :

Lemma 1 *In Median Elimination, if $\hat{a}_k(B) \leq \frac{\delta}{4}$, with probability at least $1 - \frac{\delta}{4}$, ME returns an arm set S satisfying $\hat{a}_i(S) \leq (1 - \frac{\epsilon}{2}) \hat{a}_i(B)$ for all $i \leq k$.*

Let us first focus on the ℓ -th round for a specific ℓ . Recall that, in this round, ME calculates an empirical value for each arm in S_ℓ . Henceforth, for any subset $X \subseteq S_\ell$, define $\hat{a}_i(X)$ to be the empirical value of $a_i(X)$.

Let E_ℓ be the event where the following holds for all $i \leq k$:

$$\hat{a}_i(S_\ell) \leq 2 \left[\left(1 - \frac{\ell}{2}\right) \hat{a}_i(S_\ell); \left(1 + \frac{\ell}{2}\right) \hat{a}_i(S_\ell) \right];$$

Moreover, we say that an arm $a \in S_\ell$ is a *bad arm* if there exists an $i \leq k$ such that $\hat{a}_i(a) < (1 - \frac{\ell}{2}) \hat{a}_i(S_\ell)$ whereas $\hat{a}_i(a) > (1 + \frac{\ell}{2}) \hat{a}_i(S_\ell)$.

Lemma 2 In Round ℓ , if E_ℓ occurs and the number of bad arms is no more than $\frac{|S_\ell|}{2} - k$, then $\hat{\mu}_i(S_{\ell+1}) \leq (1 - \frac{\epsilon_\ell}{2}) \hat{\mu}_i(S_\ell)$ holds for all $i \leq k$.

Proof: Under E_ℓ , no arm $a_i(S_\ell)$ of any $i \leq k$ can be bad. To see this, suppose that there exists $j < i$ satisfying $\hat{\mu}_i(S_\ell) < (1 - \frac{\epsilon_\ell}{2}) \hat{\mu}_j(S_\ell)$ but $\hat{\mu}_i(S_\ell) \leq (1 - \frac{\epsilon_\ell}{2}) \hat{\mu}_j(S_\ell)$. From E_ℓ , we have $\hat{\mu}_i(S_\ell) \leq (1 + \frac{\epsilon_\ell}{2}) \hat{\mu}_i(S_\ell) < (1 + \frac{\epsilon_\ell}{2})(1 - \frac{\epsilon_\ell}{2}) \hat{\mu}_j(S_\ell) < (1 - \frac{\epsilon_\ell^2}{4}) \hat{\mu}_j(S_\ell)$, which is a contradiction.

When the number of bad arms is at most $\frac{|S_\ell|}{2} - k$, at least k arms in $S_{\ell+1}$ are not bad. Denote the set of these arms as T . The rest of the proof will establish the claim:

$$\hat{\mu}_i(T) \leq (1 - \frac{\epsilon_\ell}{2}) \hat{\mu}_i(S_\ell), \quad 8i \leq k$$

which is sufficient for the lemma to hold.

Suppose that the claim does not hold. Consider the smallest $i \leq k$ such that $\hat{\mu}_i(T) < (1 - \frac{\epsilon_\ell}{2}) \hat{\mu}_i(S_\ell)$. Since $a_i(T)$ is not bad, it holds that $\hat{\mu}_i(T) < (1 - \frac{\epsilon_\ell}{2}) \hat{\mu}_i(S_\ell)$. Event E_ℓ suggests that $\hat{\mu}_i(S_\ell) \leq (1 + \frac{\epsilon_\ell}{2}) \hat{\mu}_i(S_\ell) > \hat{\mu}_i(T)$ which implies that $a_i(S_\ell)$ has not been eliminated in this round. This, together with the fact (shown earlier) that $a_i(S_\ell)$ is not bad, indicates that $a_i(S_\ell)$ must match an arm $a_{i'}(T)$ with $i' \leq i - 1$.

For any $j < i$, since $\hat{\mu}_j(S_\ell) \leq \hat{\mu}_i(S_\ell)$, we know that $\hat{\mu}_j(T) < (1 - \frac{\epsilon_\ell}{2}) \hat{\mu}_j(S_\ell)$. A similar argument shows that $a_j(S_\ell)$ matches an arm $a_{j'}(T)$ with $j' \leq i - 1$ and yet $j' \neq i'$. This means that the first i elements of S_ℓ form a one-to-one mapping to the first $i - 1$ elements of T , which is a contradiction. $\square$

Lemma 3 Under the condition $\hat{\mu}_k(B) \leq \frac{1}{2}$, if for all $j < k$, Median Elimination ensures

$$\hat{\mu}_i(S_{j+1}) \leq (1 - \frac{\epsilon_j}{2}) \hat{\mu}_i(S_j); \quad 8i \leq k \quad (8)$$

then, with probability at least $1 - \frac{\epsilon_\ell}{2}$, $\hat{\mu}_i(S_{\ell+1}) \leq (1 - \frac{\epsilon_\ell}{2}) \hat{\mu}_i(S_\ell)$ holds for all $i \leq k$.

Proof: By Lemma 2, it suffices to show that with probability no less than $1 - \frac{\epsilon_\ell}{2}$, E_ℓ occurs and the number of bad arms does not exceed $\frac{|S_\ell|}{2} - k$.

Recall that in ME $\hat{\mu}_k(S_1) = \hat{\mu}_k(B) \leq \frac{1}{2}$. Therefore, from (8), we know $\hat{\mu}_k(S_\ell) \leq \prod_{j=1}^{\ell-1} (1 - \frac{\epsilon_j}{2}) \hat{\mu}_k(S_1) \leq \frac{1}{2}$. Thus, by Chernoff bounds (1) and (2), we have that for any $i \leq k$:

$$\Pr \left[\hat{\mu}_i(S_\ell) \leq \left(1 - \frac{\epsilon_\ell}{2}\right) \hat{\mu}_i(S_\ell) \right] \geq \exp \left(- \frac{(\frac{\epsilon_\ell}{2})^2 Q_\ell \hat{\mu}_i(S_\ell)}{2} \right) \geq \frac{\ell}{6k};$$

and also

$$\Pr \left[\hat{\mu}_i(S_\ell) \geq \left(1 + \frac{\epsilon_\ell}{2}\right) \hat{\mu}_i(S_\ell) \right] \leq \frac{\ell}{6k};$$

By the union bound, the failure probability of E_ℓ is at most $\frac{\ell}{3}$.

Next, consider the probability of an arm a being bad. For any $i \leq k$ which satisfies $\hat{\mu}_i(a) < (1 - \frac{\epsilon_\ell}{2}) \hat{\mu}_i(S_\ell)$, define $\theta = \frac{\theta_i(S_\ell)}{\theta(a)} (1 - \frac{\epsilon_\ell}{2}) \leq 1$. We distinguish two cases:

Case 0 < $\theta < 1$. From Chernoff bound (1),

$$\begin{aligned} & \Pr \left[\hat{\mu}_i(a) \leq \left(1 - \frac{\epsilon_\ell}{2}\right) \hat{\mu}_i(S_\ell) \right] \\ &= \Pr \left[\hat{\mu}_i(a) \leq (1 + \theta) \hat{\mu}_i(a) \right] \\ & \leq \exp \left(- \frac{(1 - \frac{\epsilon_\ell}{2})^2 Q_\ell \hat{\mu}_i(S_\ell)}{3} \frac{1}{1 + \theta} \right); \end{aligned}$$

$$\Pr \left[\hat{a} = (1 + \epsilon) a \right] \\ \exp \left(- \frac{(1 - \epsilon)^2}{3} Q_{\ell}(S_{\ell}) \right) \frac{\ell^2}{(2 - \epsilon)(2 + \epsilon)} \\ \exp \left(- \frac{(1 - \epsilon)^2}{3} Q_{\ell}(S_{\ell}) \right) \frac{\ell}{6k}:$$
$$\Pr \left[\hat{a} = (1 + \frac{\epsilon_\ell}{2}) a \right] \\ \exp \left(- \frac{(1 - \frac{\epsilon_\ell}{2})^2}{6} Q_{\ell} (S_\ell) \right) \\ \exp \left(- \frac{(\ell - 2)^2}{2} Q_{\ell} (S_\ell) \right) \frac{\ell}{6k};$$
$$\Pr \left[\sum_{i=1}^{|S_\ell|} Z_i > \frac{jS_\ell j}{2} \mid k \right] \leq \frac{\mathbb{E} \left[\sum_{i=1}^{|S_\ell|} Z_i \right]}{jS_\ell j} \leq \frac{jS_\ell j}{jS_\ell j} = \frac{2}{3},$$
$$1 \sum_{i=1}^L i = 1 \sum_{i=1}^L \frac{1}{2^{i-1}} \quad 1 \quad 2 \quad 1 = 1 \quad = 4;$$
$$\begin{aligned} \prod_{j=1}^L (1 - \frac{1}{2^j}) &= \exp \left(\sum_{j=1}^L \log \left(1 - \frac{1}{2^j} \right) \right) \\ &= \exp \left(- \sum_{j=1}^{\infty} \frac{1}{2^j} \right) = \exp \left(- \frac{1}{2} \right) = \frac{1}{\sqrt{e}} \end{aligned}$$

12

B.2 Correctness and Sample Complexity

Analysis of Uniform Sampling. As mentioned in Section 2, US was already applied in the algorithm developed in by Sheng et al. [14]. We can re-use much of their analysis in our context. First, they proved:

Lemma 4 ([14]) *In Uniform Sampling, if $s_k(S) \leq \frac{\delta}{4}$, then with probability at least $1 - \frac{\delta}{4}$, US returns an arm set V which satisfies $i(V) \leq (1 - \frac{\epsilon}{2}) i(S)$ for any $i \leq k$.*

Proof: By the proof of Lemma 8 in [14] and adjusting constants appropriately. $\square$

Next, we show that with a large probability ME-AS terminates at $2 \lceil \frac{1}{8} k(B); k(B) \rceil$. If this is not what happens, one of the following two events must have happened:

- (1) *Premature:* ME-AS terminates when $s_k(B) > \frac{\delta}{4}$.
- (2) *Overdue:* ME-AS does not terminate after invoking US with $s_k(B) < \frac{1}{4} k(B)$.

We say that ME *succeeds* if the arm set S it returns satisfies $i(S) \leq (1 - \frac{\epsilon}{2}) i(B)$ for all $i \leq k$.

Lemma 5 *Both of the following statements are true:*

(Due to [14]) The premature event happens with probability at most $\delta/4$.

If ME succeeds, the overdue event happens with probability at most $\delta/4$.

Proof: The first statement has been proved in [14] (see Lemma 9 therein and adjusting constants appropriately). For the overdue event, when $s_k(B) < \frac{1}{4} k(B)$, given that ME has succeeded, we have $s_k(S) = (1 - \frac{\epsilon}{2}) s_k(B) < \frac{1}{4} k(S)$. Define $\hat{a}_k^{US}(S)$ to be the US-empirical value of $a_k(S)$. We have:

$$\begin{aligned} \Pr \left[\hat{a}_k^{US}(S) \geq \frac{1}{2} \right] &= \Pr \left[\hat{a}_k^{US}(S) \geq \frac{1}{2} \mid s_k(B) \leq \frac{1}{4} k(B) \right] \\ &= \Pr \left[\hat{a}_k^{US}(S) \geq \frac{s_k(S)}{2(1 - \frac{\epsilon}{2})} \right] \\ &= \Pr \left[\hat{a}_k^{US}(S) \geq (1 - \frac{\epsilon}{2}) \frac{s_k(S)}{k(S)} \right] \\ &\stackrel{\text{(by (2))}}{=} \delta/4; \end{aligned}$$

where the third inequality used the fact that $\frac{1}{2(1 - \epsilon/2)} \geq (1 - \frac{\epsilon}{2})$ holds for any $\epsilon \in (0, \frac{1}{4})$. $\square$

Correctness of ME-AS. Combining Lemmas 1, 4, and 5, we now prove:

Lemma 6 *With probability at least $1 - \delta$, ME-AS returns an arm set V which satisfies $i(V) \leq (1 - \frac{\epsilon}{2}) i(B)$ for all $i \leq k$.*

Proof: If the premature event does not occur, we have that $s_k(B) \leq \frac{\delta}{4}$. Provided that ME has succeeded, Lemma 5 has showed that, with probability at least $1 - \frac{\delta}{4}$, ME-AS terminates at $2 \lceil \frac{1}{8} k(B); k(B) \rceil$. ME having succeeded ensures $i(S) \leq (1 - \frac{\epsilon}{2}) i(B)$ for all $i \leq k$. Hence, $s_k(S) \leq (1 - \frac{\epsilon}{2}) s_k(B) \leq \frac{\delta}{4}$. Then, by Lemma 4, we know that with probability at least $1 - \frac{\delta}{4}$, the following holds for all $i \leq k$:

$$\begin{aligned} i(V) &\leq (1 - \frac{\epsilon}{2}) i(S) \\ &\leq (1 - \frac{\epsilon}{2})^2 i(B) \leq (1 - \frac{\epsilon}{2}) i(B); \end{aligned}$$

The premature event occurs with probability at most $\frac{\delta}{4}$. By Lemma 1, ME fails with probability only at most $\frac{\delta}{4}$ when $s_k(B) > \frac{\delta}{4}$. Lemma 6 thus follows from the union bound. $\square$

Sample Complexity of ME-AS. We first analyze the cost of ME and US under a specific δ . This is trivial for US, for which the answer is obviously $O(\frac{k}{\epsilon^2} \frac{1}{\mu} \log \frac{k}{\delta})$ (recall that $jS_j \leq 4k$).

Regarding the ℓ -th round of ME, we know $|S_\ell| = \frac{n}{2^{\ell-1}}$, $\ell = (3=4)^{\ell-1} = 16$, $\ell = (1=2)^{\ell-1} = 8$, and $\ell = \sum_{j=1}^{\ell} (1 - \frac{1}{j}) u_1 = (1 - \frac{1}{2}) = \frac{1}{2}$. Hence, the cost of ME is bounded by:

$$\begin{aligned} & \sum_{\ell=1}^{\ell} \frac{12}{\ell} \frac{1}{2^{\ell-1}} \log \frac{6k}{\ell} |S_\ell| \\ &= O \sum_{\ell=1}^{\ell} \frac{12}{\ell} \frac{1}{2^{\ell-1}} \log \frac{k}{\ell} + \frac{1}{2^{\ell-1}} \frac{n}{2^{\ell-1}} \\ &= O \sum_{\ell=1}^{\ell} \frac{8}{9} \frac{1}{2^{\ell-1}} \log \frac{k}{\ell} + \frac{n}{2^{\ell-1}} \\ &= O \frac{n}{2} \log \frac{k}{\ell} : \end{aligned}$$

We have proved that one for-iteration of ME-AS (i.e., Lines 3-6 of Algorithm 1) under a specific ℓ has cost $O(\frac{k}{2^{\ell-1}} \log \frac{k}{\ell})$. The cost doubles each time when we halve the parameter ℓ until it drops below $k(B)=4$. Hence, the total complexity so far is $O(\frac{n}{2} \frac{1}{k(B)} \log k)$.

It remains to show that if ME-AS terminates at $\ell = k(B)=8$, the extra cost is still bounded by $O(\frac{n}{2} \frac{1}{k(B)} \log k)$ in expectation. From Lemmas 4 and 5, we know that, at every $\ell < k(B)=8$, ME-AS fails to terminate with probability at most $\frac{1}{2}$. Thus, its cost for all $\ell < k(B)=8$ is bounded by

$$\begin{aligned} & O \sum_{i=1}^{\ell} \frac{n}{2} \frac{1}{k(B)} \log \frac{k}{\ell} 2^{i-1} \frac{1}{2^i} \\ &= O \frac{n}{2} \frac{1}{k(B)} \log \frac{k}{\ell} : \end{aligned}$$

We thus have proved Theorem 1.

Theorem 1 ME-AS solves the k -AS problem with expected cost $O(\frac{n}{2} \frac{1}{k(B)} \log k)$.

B.3 Lower Bound

This subsection serves as a proof for:

Theorem 2 For any $\epsilon \in (0, \frac{1}{4}]$ and $\delta \in (0, \frac{1}{48}]$, given any algorithm, there is an instance of the k -AS problem on which the algorithm must entail $\Omega(\frac{n}{2} \frac{1}{k(B)} \log k)$ cost in expectation.

Hard Instances. We use $a_1, \dots, a_n$ to represent the arms in B . Let I_x denote the collection of all the x -sized subset $I = \{a_2, \dots, a_n\}$. For each $1 \leq \ell \leq k-1$, we create an instance with:

$$(a_1) = (1 + 4^{-\ell}) \quad (9)$$

$$(a_i) = (1 + 8^{-\ell}) ; 8a_i \geq 1 \quad (10)$$

$$(a_i) = \epsilon ; 8a_i \geq 2 \cup n \quad (11)$$

where ϵ is an arbitrary real value in $(0, \frac{1}{12})$.

We use I_ℓ to denote the instance associated with ℓ . Notice for $\ell \in (0, \frac{1}{4}]$, $(1 - \epsilon)(1 + 4^{-\ell}) > (1 - \epsilon)(1 + 8^{-\ell}) > (1 + 4^{-\ell})$. Thus, by the problem definition of k -AS, any algorithm, with probability at least $1 - \delta$, must return a_1 for $\ell \geq k-1$ and return I_ℓ for $\ell \leq k-1$.

Deterministic Algorithms. We will first consider an algorithm A deterministic in the following sense. At the beginning, the first arm that samples from is fixed. From then on, iteratively, the next arm to sample from is always a function of the results of the previous samples (i.e., which arms have been sampled, in what order, and what are the results). In this way, be uniquely

described by history W , which is the sequence of arms chosen at each step, and the result of each sample. Finally, the output of A is a function of W .

It is important to note that A is a random variable, due to the randomness from the arms, namely, the samples taken from them are random. Even on a specific problem instance, various histories occur with different probabilities. Therefore, the cost of A (i.e., the length of W) is still a random variable; and also A may choose to fail on some I , provided that the overall failure probability on the instance is at most ϵ .

The crux of our proof is to show that, for any deterministic algorithm A , there exists $I \in \mathcal{I}_{k-1}$ such that A must take at least $\frac{n-1}{2} \ln \frac{k}{\epsilon}$ samples in expectation to solve the AS problem on instance I . Assume, on the contrary, that for any $I \in \mathcal{I}_{k-1}$, A takes at most $\frac{n-1}{36864} \ln \frac{k}{\epsilon}$ samples in expectation. We will prove that A must fail with a probability greater than ϵ on some instance $I \in \mathcal{I}_{k-1}$.

A Deterministic Lower Bound. Set $T = \frac{1}{2304} \ln \frac{k}{\epsilon}$. Define random variable T_j be the number of samples that A takes from arm a_j . It is easy to see that $T_j > 0$ (i.e., A must sample at least once from a_j).

Given a specific $I \in \mathcal{I}_{k-1}$, denote by $E_I[T_j]$ the expectation of T_j when A runs on I . At most $n=4$ arms $a_j \in \mathcal{I}_{k-1}$ satisfy

$$E_I[T_j] > \frac{1}{9216} \ln \frac{k}{\epsilon}.$$

Since $k \geq n=2$, at least $k - \frac{n}{4} = \frac{n-k}{2}$ arms $a_j \in \mathcal{I}_{k-1}$ satisfy $E_I[T_j] \leq \frac{1}{9216} \ln \frac{k}{\epsilon}$. We refer to these arms as ordinary arms. For any such an arm a_j , by Markov's inequality, we know:

$$\Pr_I[T_j > T] < \frac{E_I[T_j]}{T} \leq \frac{1}{4}; \quad (12)$$

where the subscript of $\Pr_I$ (just like the subscript in E_I) indicates "conditioned on the specific instance I ".

Now let us fix an ordinary arm a_j . Recall that each of the T_j samples that A draws from a_j is either 1 or 0. Let X be the sum of all those T_j samples. Set $\tau = \frac{1}{T} \ln(k/\epsilon)$ and define E_A be the event where both $T_j \leq T$ and $X > \tau$ are true.

Lemma 7 E_A happens with probability at least $\frac{1}{2}$.

Proof: By substituting $(a_j) = \dots$, we can upper bound the probability of not happening by:

$$\begin{aligned} & \Pr_I[T_j > T] + \Pr_I[T_j \leq T; X \leq \tau] \\ & \leq \frac{1}{4} + \Pr_I[T_j \leq T; X \leq \tau] \\ & = \frac{1}{4} + \sum_{t=1}^T \Pr[T_j = t] \Pr[X \leq \tau | T_j = t] \\ & \leq \frac{1}{4} + \sum_{t=1}^T \Pr[T_j = t] \exp\left(-\frac{1}{2} \ln \frac{k}{\epsilon}\right) \\ & = \frac{1}{4} + \exp\left(-\frac{1}{2} \ln \frac{k}{\epsilon}\right) = \frac{1}{4} + \frac{1}{k} \leq \frac{1}{2} \end{aligned}$$

where the first inequality used (12), and the second inequality is due to Chernoff bound (5).

Let V be the arm set returned by A and E_B be the event that $V = I \setminus \{a_1\}$. Define event $E_S = E_A \setminus E_B$. As A (when executed on I) returns $I \setminus \{a_1\}$ with probability at least $\frac{1}{2}$, we know that

$$\Pr_I[E_S] \geq \frac{1}{2} \quad (13)$$

Let us focus on one specific history w . Consider the instance $[f, a_j, g, 2, l_k]$, which can be conveniently understood as raising the mean of a_j from μ to $(1+8\mu)$ in the instance $\mathcal{I}$. Denote by $\Pr_{I[f, a_j, g]}[W]$ the probability of W when A runs on $I[f, a_j, g]$, and similarly, by $\Pr_I[W]$ the probability of W on I .

Lemma 8 For any $\mu \in (0; \frac{1}{4})$; $\mu \in (0; \frac{1}{48})$, when conditioned on $\mathcal{E}_S$, it holds that

$$\frac{\Pr_{I[f, a_j, g]}[W]}{\Pr_I[W]} \leq \frac{5}{k}.$$

Proof: As the mean of a_j is larger under $I[f, a_j, g]$, the ratio of the two probabilities is minimized when the fewest samples a_j turn out to be 1 in W . Recall that the event $\mathcal{E}_S$ ensures $X > T_j$. With this, we can derive

$$\begin{aligned} & \frac{\Pr_{I[f, a_j, g]}[W]}{\Pr_I[W]} \\ &= \frac{[(1+8\mu)]^{T_j}}{(1+8\mu)^{T_j}} \cdot \frac{[1 - (1+8\mu)]^{(1-\mu)T_j + \frac{1}{2} \#T_j}}{(1 - (1+8\mu))^{(1-\mu)T_j + \frac{1}{2} \#T_j}} \\ &= (1+8\mu)^{-\frac{1}{2} \#T_j} \cdot \frac{1 - (1+8\mu)}{1} \\ &= (1+8\mu)^{-\frac{1}{2} \#T_j} \cdot \frac{1 - (1+8\mu)}{1} \\ &= (1 - 12\mu)^{-\frac{1}{2} \#T_j}; \end{aligned} \quad (14)$$

where the last inequality used the fact that $(1+8\mu) \leq (1 - 12\mu)$ for any $\mu \in (0; \frac{1}{12})$.

For any $\mu \in (0; \frac{1}{4})$, it holds that $\frac{1}{1 - (1+8\mu)} \leq e^{12\mu}$. Using this fact and plugging in $\frac{1}{2} \#T_j = \frac{1}{2} \ln(k) = \frac{1}{2} \ln(k)$ and $T = \frac{1}{2304} \ln(k)$, we can bound the second term of (14) as:

$$(1 - 12\mu)^{-\frac{1}{2} \#T_j} = (1 - 12\mu)^{\frac{1}{48} \ln(k)} = e^{\frac{1}{2} \ln(k) \ln(1 - 12\mu)} = e^{-\frac{1}{2} \ln(k) \ln(1 - 12\mu)} \quad (15)$$

To bound the first term of (14), let us define $g(\mu) = (1 - 12\mu)^{-\frac{1}{2} \#T_j}$. Taking the partial derivative of $\ln g(\mu)$ with respect to μ gives:

$$\begin{aligned} \frac{\partial \ln g(\mu)}{\partial \mu} &= \frac{1}{2} \ln(1 - 12\mu) \\ &= \frac{1}{2} \ln(1 - 12\mu) \cdot \frac{1}{1 - 12\mu} \cdot \frac{1}{(1 - 12\mu)^2} \cdot (-12) \end{aligned} \quad (16)$$

It is not hard to verify that for any $\mu \in (0; \frac{1}{12})$; $\mu \in (0; \frac{1}{4})$, the RHS of (16) is strictly less than 0, which implies

$$g(\mu) \leq g\left(\frac{1}{12}\right) = e^{\ln(1 - \frac{8}{11})^{11}} = e^{\ln(1 - 8)} = 1 - 8:$$

Moreover, observe that

$$\begin{aligned} & (1+8\mu)^{-\frac{1}{2} \#T_j} \\ &= (1+8\mu)^{-\frac{1}{2} \ln(k)} < (1+8\mu)^{-8} \end{aligned}$$

$$1 - (1 - \alpha)^{\frac{1}{\alpha}} \frac{1}{\alpha} \tau_j$$

$$= \frac{1}{(1+8)^2} \frac{1}{64} \frac{(1+8)}{1} \frac{1}{2304} \frac{(1+8)}{2} \ln \frac{k}{18} = \frac{1}{18} \ln \frac{k}{18}$$

Combining the above with (5), we have $\frac{\Pr_{\mathbf{I}}[f_{\mathbf{a}_j} g[W]]}{\Pr_{\mathbf{I}}[W]} = (k)^{5=9} \quad 5=k \text{ for any } 2 \leq j \leq \frac{1}{48}$.

We will now proceed to analyze the probability $\Pr[E_S]$ when A is executed on $I \parallel a_j$. For this purpose, define $\mathbf{1}_S(W)$ to be an indicator function that equals 1 if the history W results in the occurrence of E_S , and 0 otherwise (note that once W is given, the output of A is fully determined, and hence, so is E_S). With this, we can derive:

$$\begin{aligned}
 & \Pr_{I[f]_{a_j g}}[E_B] \\
 &= \Pr_{I[f]_{a_j g}}[E_S] \\
 &= \Pr_{W \text{ s.t. } 1_S(W) = 1} [W] \\
 &= \Pr[W] \frac{5}{k} \\
 &= \Pr[W] \frac{5}{k} \\
 &= (5 = k) \Pr[E_S] \\
 &= 2 = k:
 \end{aligned}$$

Therefore, we have proved that for any $2 \leq k \leq n$ and any ordinary array $a_1, \dots, a_n$, the failure probability of A under instance $(I, f, a_1, \dots, a_n)$ exceeds $\frac{2}{k}$.

Recall that for any $2 \leq k \leq n$, there are at least $\frac{n-k}{2}$ ordinary arms a_j in $U \cap I$. Summing up the failure probabilities of $\mathcal{A}$ on the instances J of all $2 \leq k \leq n$ leads to:

$$\begin{aligned}
& \Pr[V \in J] \\
& \Pr[V = J \mid f_{a_1} g_n f_{a_j} g] \\
& \Pr[V = I \mid f_{a_1} g] \\
& \Pr[f_{a_j} g = J] \\
& \Pr[V = I \mid f_{a_j} g] \\
& \Pr[E_B] \\
& \frac{n}{2} \frac{k}{k} \\
& \frac{n}{k} \frac{1}{k} = \frac{n}{k} = j|_k :
\end{aligned}$$

Algorithm 4: Additive Median Elimination (AME)

```

1 input:  $B, \epsilon, \delta, k$ 
2  $S_1 = B, n_1 = 16, \ell_1 = 8$  and  $\ell = 1$ ;
3 while  $|S_\ell| > 4k$  do
4   sample every arm  $a \in S_\ell$  for  $Q = (12/\epsilon^2) \log(6k/\delta)$  times;
5   for each arm  $a \in S_\ell$  do
6     its empirical value  $\hat{a} =$  the average of the  $Q$  samples from  $a$ ;
7    $a_1, \dots, a_{|S_\ell|}$  = the arms sorted in non-increasing order of their empirical values;
8    $S_{\ell+1} = \{a_1, \dots, a_{|S_\ell|-2}\}$ ;
9    $n_{\ell+1} = 3n_\ell/4$ ;  $\ell_{\ell+1} = \ell_\ell/2$  and  $\ell = \ell + 1$ ;
10 return  $S_\ell$ ;

```

Algorithm 5: Additive Uniform Sampling (AUS)

```

1 input:  $S, \epsilon, \delta, k$ 
2 sample every arm  $a \in S$  for  $Q = (96/\epsilon^2) \log(4|S|/\delta)$  times;
3 for each arm  $a \in S$  do
4   its US-empirical value  $\hat{a}^{US} =$  the average of the  $Q$  samples from  $a$ ;
5  $a_1, \dots, a_{|S|}$  = the arms sorted in non-increasing order of their US-empirical values;
6 return  $f(a_1, \hat{a}_1^{US}); \dots; (a_k, \hat{a}_k^{US})$ 

```

By the pigeon hole principle, there is at least one $i \in [k]$ such that $\Pr_j[V_i \in J] > \epsilon$. As this is not allowed, we conclude that n must take at least $\frac{n}{36864} \frac{1}{\epsilon^2} \ln \frac{k}{\delta}$ samples on the i of at least one $i \in [k-1]$.

Randomized Algorithms. It is standard to extend the above deterministic lower bound to randomized algorithms—such an argument can be found, for example, in [16]. With this, we complete the whole proof of Theorem 2.

B.4 Analysis of Top- k_{add} Arm Selection

The algorithm of k_{add} -AS is very similar to k -AS. As we do not require a multiplicative guarantee here, we simply drop the “guess” stage and adjust the parameters appropriately, as shown in Algorithm 4, 5.

Theorem 3 AME solves the k_{add} -AS problem with expected cost $\frac{n}{2} \log \frac{k}{\delta}$.

Proof: The proof of Theorem 3 directly follows Lemma 1 to Lemma 4, except that we change the multiplicative guarantee into the additive guarantee and utilize the Chernoff bound in the analysis instead.

For the lower bound, as we claimed in Section 3, the k_{add} -AS problem is stronger than $\text{EXPLORE-}k$. Therefore, the lower bound for $\text{EXPLORE-}k$ directly applies to k_{add} -AS.

C Analysis of k -MOST CONNECTED VERTEX

This section is devoted to the proof of Theorem 4.

Theorem 4 For any $\delta \in (0, \frac{1}{12}]$ and $\epsilon \in (0, \frac{1}{48}]$, the following statements are true about any k -MCV algorithm:

when $\deg_k(B) \geq \frac{1}{2} \log \frac{n}{\delta}$, there is an instance on which the algorithm must probe $(\frac{n}{2} \frac{m}{\deg_k(B)} \log \frac{k}{\delta})$ edges in expectation.

when $\deg_k(B) < O(\frac{1}{\epsilon})$, there is an instance on which the algorithm must probe $\Omega(n \epsilon)$ edges in expectation.

Proof of Theorem 4 for Large $\deg_k(B)$. We prove the first branch of the theorem by a reduction from k -AS. Let A be an algorithm solving k -MCV with expected cost C . We will deploy A to

solve the hard k -AS instances created earlier in the proof of Theorem 2 with sample complexity Q , and thereby obtaining a lower bound for Q . Recall that, to prove Theorem 2 in Appendix B.3, we associated each $2 \leq k \leq 1$ with an instance I_k for k -AS that is described by (B, W, E) , (10), and (11). From I_k , we construct a random hidden bipartite graph $G = (B; W; E)$ with $|B| = n$ and $|W| = m$ where $m = \frac{24}{\epsilon} \frac{1}{\delta} \log \frac{n}{\delta}$. For the i -th black vertex, we independently make each edge incident on it solid with probability (a_i) . The graph G thus constructed is fed as an MCV input to A . Finally, we select the arm a_i into the solution of k -AS if and only if A includes the i -th black vertex in its $(k$ -MCV) solution.

We denote by d_i the degree of the i -th black vertex in G ; notice that d_i is a random variable. Set $t = m$. By Chernoff bound, we know:

$$\begin{aligned} & \Pr[d_i \leq (1 - \epsilon/2)(1 + 4\epsilon)t] \\ &= \Pr[d_i \leq (1 - \epsilon/2)(1 + 4\epsilon)t] \\ & \leq \exp\left(-\frac{(\epsilon/2)^2}{2}(1 + 4\epsilon)t\right) \end{aligned}$$

Therefore, if $\frac{24}{\epsilon} \log \frac{n}{\delta}$ and $\epsilon \in (0; \frac{1}{12})$, with probability at least $1 - \frac{1}{n}$, we have $d_i \leq (1 - \epsilon/2)(1 + 4\epsilon)t \leq (1 + 3\epsilon)t$. Similarly, we can show that with probability at least $1 - \frac{1}{n}$, we have $d_i \leq (1 + 7\epsilon)t$ for any $a_i \in I$ and $d_i \leq (1 + \epsilon)t$ for any $a_i \in U \cap I$. Notice that both $(1 - \epsilon/2)(1 + 4\epsilon)t > (1 + \epsilon)t$ and $(1 + 7\epsilon)t > (1 + \epsilon)t$ hold. By the union bound, with probability at least $1 - \frac{1}{n}$, A must select the first vertex and the vertices corresponding to the arms of its $(k$ -MCV) solution set. Thus, we have obtained a $(1 + 2\epsilon)$ -approximate algorithm that solves k -AS on instance I_k .

On the other hand, the proof of Theorem 2 has showed that any $(1 + 2\epsilon)$ -approximate k -AS algorithm must take in expectation $(\frac{n}{2} \frac{1}{\deg_k(B)} \log \frac{k}{2})$ samples on the I_k of at least one $2 \leq k \leq 1$. It thus follows that $Q = (\frac{n}{2} \frac{m}{\deg_k(B)} \log \frac{k}{2})$.

Proof of Theorem 4 for Small $\deg_k(B)$. Let us define the single-vertex problem:

Suppose that $t < \frac{1}{2}$ is a given positive integer. The bipartite graph only has one black vertex b , and $m = \max\{t + 1; 64t\}$ white vertices. The goal is to distinguish whether $\deg(b)$ is t or $t + 1$ with probability at least $\frac{15}{16}$.

We now give a lower bound for the above problem.

Lemma 9 Any algorithm solving the above single-vertex problem must $\Omega(\frac{m}{t})$ probe operations.

Proof: Our proof is based on a sampling lower bound established in [1] for evaluating symmetric functions³. It is showed in [3] that any sampling algorithm for symmetric functions can be simulated by an i.i.d. sampling algorithm under some specific conditions. Therefore, the sampling lower bound can be measured by the distance between two i.i.d. distributions. Formally, we need the following definition:

Definition 1 [Total Variational Distance] Suppose X and Y are two discrete distributions with a common support S . The total variational distance (also known as statistical distance) between X and Y is defined as

$$d_{TV}(X; Y) = \frac{1}{2} \sum_{s \in S} |Pr[X = s] - Pr[Y = s]|$$

Suppose $f(x_1; \dots; x_m) : S^m \rightarrow \mathbb{Z}$ is a symmetric function where S and $\mathbb{Z}$ are arbitrary sets. For any $x \in S^m$, we use U_x to denote the distribution of i.i.d. query outcomes x supported on X , i.e., $U_x(s) = \frac{1}{m} \sum_{i=1}^m 1(x_i = s)$; $8s \in S$ where $1()$ is an indicator function.

Lemma 10 ([3, Theorem 4.21]) For any $\epsilon \in (0; 1)$; $\delta \in (\frac{2}{m}; \frac{1}{8})$, given two inputs $x, y \in S^m$, if a sampling algorithm A can distinguish whether $f = f(x)$ or $f = f(y)$ with probability at least $1 - \delta$

³A function is a symmetric function if permuting the input arguments does not change the function value.

within worst case sampling complexity at most $\frac{1}{8} \frac{P}{d_{TV}(U_x; U_y)} \ln \frac{1}{8}$, it must take at least

$$\frac{1}{8} \frac{P}{d_{TV}(U_x; U_y)} \ln \frac{1}{8}$$

samples in the worst case.

Consider the simple symmetric function $\deg(b) = f(x_1, \dots, x_m) = \sum_{i=1}^m x_i$ where x_i is 1 if the edge connecting b and the i -th white vertex is solid, and 0 otherwise. We use u to denote an input with $f = t$. Then, we have that $U_x(0) = \frac{m-t}{m}$ and $U_x(1) = \frac{t}{m}$. Similarly, we use v to denote an input with $f = t+1$ and we have that $U_y(0) = \frac{m-t-1}{m}$ and $U_y(1) = \frac{t+1}{m}$.

By the definition of d_{TV} , it is easy to see that $d_{TV}(U_x; U_y) = \frac{1}{m}$. Therefore, we conclude that the lower bound of the single vertex problem (sm).

The lemma below shows a reduction from the above problem to 1-MCV.

Lemma 11 For any $\epsilon \in (0, \frac{1}{12})$; $\delta \in (0, \frac{1}{48})$, given an (ϵ, δ) -approximate k -MCV algorithm A with expected sample complexity Q , we can design an algorithm B to solve the single-vertex problem with at most $32Q = n$ probe operations.

Proof: We reduce the single-vertex problem to 1-MCV as follows. First, pick a value t and construct an edge hidden graph G with n black vertices and $n \max\{t+1, 64\}$ white vertices. Denote the set of black vertices $b_1, \dots, b_n$. Randomly choose a black vertex b_i as the pivot vertex. For any b_j with $j \neq i$, we connect it to t white vertices randomly chosen without replacement (i.e., pick a random t -sized subset of white vertices over all the possible choices).

Algorithm B solves the single-vertex problem by simulating algorithm A in solving the 1-MCV problem on G . Whenever A probes a hidden edge b_j with $j \neq i$, B simply does the same. On the other hand, when A probes a hidden edge of the pivot vertex b_i , B probes a hidden edge of the vertex b_i in the single-vertex problem, and passes the result to A . If (i) the pivot vertex is probed less than $\frac{32Q}{n}$ times and at the same time A returns a vertex b_j with $j \neq i$, B decides $\deg(b) = t$. In all other cases B decides $\deg(b) = t+1$.

This finishes the description of B . Its worst case cost is clearly $\frac{32Q}{n}$ (for the single-vertex problem). Next, we show that B decides $\deg(b)$ correctly with probability at least $\frac{15}{16}$.

Consider first the case $\deg(b) = t$. All the black vertices in G have degrees exactly t . Denote by T_b the number of probe operations that A performs on b . As the pivot vertex is indistinguishable from any other vertices, we have $E[T_b] = \frac{Q}{n}$. Markov's inequality then gives:

$$\Pr[T_b \geq \frac{32Q}{n}] \leq \frac{1}{32}.$$

If $T_b < \frac{32Q}{n}$, B errs if and only if A returns b_j . Since b is indistinguishable, the probability that A returns b_j is $\frac{1}{n}$ (recall that every black vertex has the same chance of being the pivot vertex). Therefore, with probability at least $1 - \frac{1}{n} - \frac{1}{32} \geq \frac{15}{16}$, B returns a correct answer.

Next, consider the case $\deg(b) = t+1$. B makes a mistake only if $T_b < \frac{32Q}{n}$ and A returns a vertex b_j with $j \neq i$. Since A is an (ϵ, δ) -approximate 1-MCV algorithm, the probability that B errs is at most $\frac{1}{16}$. Hence, once again, it is correct with probability at least $\frac{15}{16}$.

The second branch of Theorem 4 follows from Lemma 9 and 11.

D Analysis of Top- k_{avg} Arm Selection

In this section we analyze our algorithm for k_{avg} -AS as shown in 6, 7, 8. We first present the analysis of Quartile Elimination in Appendix D.1. Next, we prove Theorem 5 in Appendix D.2. Finally, we prove Theorem 6 in Appendix D.3.

Algorithm 6: QE-AS

```

1 input:  $B, \epsilon, \delta, k$ 
2 for  $\ell = 1; 2; 4; \dots$  do
3    $S = \text{QE}(B; \epsilon; \delta; k)$ ;
4    $f(a_{[1:k]}, \hat{g}^{\text{US}}) = \text{US}(S; \epsilon; (1 - \delta); k)$ ;
5   if  $\hat{g}^{\text{US}} \geq 2$  then
6     return  $f(a_1; \dots; a_k; g)$ ;

```

Algorithm 7: Quartile Elimination (QE)

```

1 input:  $B, \epsilon, \delta, k$ 
2  $S_1 = B, \ell = 1, \ell_{\text{max}} = 8, \ell_{\text{min}} = 1$ , and  $\ell = 1$ ;
3 while  $|S_\ell| > 4k$  do
4   sample every arm  $a \in S_\ell$  for  $Q = (4k - 2)(1 - \delta) \cdot 1 + \frac{\log(2/\delta)}{k}$  times;
5   for each arm  $a \in S_\ell$  do
6     its empirical value  $\hat{a}$  = the average of the  $Q$  samples from  $a$ ;
7    $a_1; \dots; a_{|S_\ell|}$  = the arms sorted in non-increasing order of their empirical values;
8    $S_{\ell+1} = f(a_1; \dots; a_{3|S_\ell|/4}; g)$ ;
9    $\ell_{\text{min}} = 7 \cdot \ell_{\text{min}}/8, \ell_{\text{max}} = \ell_{\text{max}}/2, \ell_{\text{mid}} = (1 - \delta) \cdot \ell_{\text{max}}$ , and  $\ell = \ell_{\text{mid}} + 1$ ;
10 return  $S_\ell$ ;

```

Algorithm 8: Uniform Sampling (US)

```

1 input:  $S; \epsilon; \delta; s; k$ 
2 sample every arm  $a \in S$  for  $Q = (120 - 2)(1 - \delta) \cdot 1 + \frac{\log(4/\delta)}{k}$  times;
3 for each arm  $a \in S$  do
4   its US-empirical value  $\hat{a}^{\text{US}}$  = the average of the  $Q$  samples from  $a$ ;
5  $a_1; \dots; a_{|S|}$  = the arms sorted in non-increasing order of their US-empirical values;
6 return  $f(a_1; \dots; a_k; \hat{g}^{\text{US}} = \frac{1}{k} \sum_{i=1}^k \hat{a}_i^{\text{US}})$ 

```

D.1 Analysis of Quartile Elimination

This subsection serves as a proof for the following lemma:

Lemma 12 In Quartile Elimination, if $\hat{g}^{\text{US}}(B)$, with probability at least $1 - \delta$, QE returns an arm set S satisfying $\hat{g}^{\text{US}}(S) \geq (1 - \delta) \hat{g}^{\text{US}}(B)$.

Focus on the ℓ -th round for a specific ℓ . For any subset S_ℓ , let $\hat{g}^{\text{US}}(X) = \frac{1}{k} \sum_{i=1}^k \hat{a}_i^{\text{US}}(X)$. We use m_ℓ to denote the median of the means of all the arms in S_ℓ and $\hat{a}_\ell$ to denote the largest empirical value of the eliminated arms in this round. For each k , define a Boolean variable $X_{\ell,i}$ which equals 1 if $\hat{a}_i(S_\ell) < m_\ell + \frac{\hat{a}_\ell(S_\ell)}{2}$, and 0 otherwise. Let $X_\ell = \frac{1}{k} \sum_{i=1}^k (\hat{a}_i(S_\ell) - m_\ell) X_{\ell,i}$. Furthermore, define E^A to be the event that $\hat{g}^{\text{US}}(S_\ell) \geq \hat{g}^{\text{US}}(S_\ell)$ holds, and E^B to be the event that $\hat{a}_\ell < m_\ell + \frac{\hat{a}_\ell(S_\ell)}{2}$ holds.

Lemma 13 In Round ℓ , if both E^A and E^B occur, then $\hat{g}^{\text{US}}(S_{\ell+1}) \geq (1 - \delta) \hat{g}^{\text{US}}(S_\ell)$ holds.

Proof: First, we claim that, under E^A and E^B , it holds that

$$\frac{1}{k} \sum_{i=1}^k \hat{a}_i(S_{\ell+1}) \geq \frac{1}{k} \sum_{i=1}^k ((1 - X_{\ell,i}) \hat{a}_i(S_\ell) + X_{\ell,i} m_\ell). \quad (17)$$

To see this, consider an arm $a_i(S_\ell)$ with $i \leq k$. If $X_{\ell,i} = 0$, by definition, we have $\hat{a}_i(S_\ell) \geq m_\ell + \frac{\hat{a}_\ell(S_\ell)}{2}$. Since E^B occurs, we have that

$$\hat{a}_i(S_\ell) \geq m_\ell + \frac{\hat{a}_\ell(S_\ell)}{2} > m_\ell.$$

Therefore, we know that $a_i(S_t)$ has not been eliminated in round t . Since $a_i(S_t)$ is one of the top k arms in S_t , it must be one of the top k -arms in S_{t+1} . Hence, for each arm $a_i(S_t)$ with $i \leq k$ and $X_{t,i} = 0$, $a_i(S_t)$ appears in both LHS and RHS of (17) exactly once. Notice that in Q_t , $\frac{jS_{t,j}}{4}$. After eliminating $\frac{jS_{t,j}}{4}$ arms, there exist at least k arms with means at least $\frac{1}{4}$. Thus, each term in LHS of (17) is at least $\frac{1}{4}$. Therefore, we conclude that our claim (17) is true. Now, we can see that:

$$\begin{aligned} \text{avg}(S_{t+1}) &= \frac{1}{k} \sum_{i=1}^k ((1 - X_{t,i}) a_i(S_t) + X_{t,i} m_i) \\ &= \text{avg}(S_t) - \frac{1}{k} \sum_{i=1}^k X_{t,i} a_i(S_t) \\ (\text{by } E^A) \quad &= (1 - \frac{1}{k}) \text{avg}(S_t): \end{aligned}$$

Lemma 14 Under the condition $\text{avg}(B) \geq \frac{1}{4}$, if for all $j < k$, Quartile Elimination ensures

$$\text{avg}(S_{j+1}) \geq (1 - \frac{1}{k}) \text{avg}(S_j) \quad (18)$$

then, with probability at least $1 - \frac{1}{2^k}$, $\text{avg}(S_{t+1}) \geq (1 - \frac{1}{k}) \text{avg}(S_t)$ holds.

Proof: By Lemma 13, it suffices to prove that with probability no less than $1 - \frac{1}{2^k}$, both E^A and E^B occur.

First we show that $\Pr[E^A] = \Pr[\frac{1}{k} \sum_{i=1}^k X_{t,i} a_i(S_t) \leq \frac{1}{4}] \geq 1 - \frac{1}{2^k}$. From (18), $\frac{1}{k} \sum_{j=1}^k (1 - \frac{1}{k}) \text{avg}(S_1) \geq \text{avg}(S_t)$. For any $i \leq k$, define $Y_{t,i} = \max\{0, a_i(S_t) - \frac{\text{avg}(S_t)}{2}\}$ and $X_{t,i} = \frac{Y_{t,i}}{a_i(S_t)}$. We have:

$$\begin{aligned} \frac{1}{k} \sum_{i=1}^k X_{t,i} a_i(S_t) &= \frac{1}{k} \sum_{i=1}^k Y_{t,i} + \frac{1}{k} \sum_{i=1}^k \frac{\text{avg}(S_t)}{2} X_{t,i} \\ &= \frac{1}{k} \sum_{i=1}^k Y_{t,i} + \frac{\text{avg}(S_t)}{2}. \end{aligned}$$

Hence, it suffices to show $\Pr[E^A] = \Pr[\frac{1}{k} \sum_{i=1}^k Y_{t,i} \leq \frac{\text{avg}(S_t)}{2}] \geq 1 - \frac{1}{2^k}$. Let $Y_t = \frac{1}{k} \sum_{i=1}^k Y_{t,i}$. We have:

$$\begin{aligned} \Pr[\frac{1}{k} \sum_{i=1}^k Y_{t,i} \leq \frac{\text{avg}(S_t)}{2}] &= \Pr[Y_t \leq \frac{\text{avg}(S_t)}{2}] \\ &= \Pr[\sum_{i=1}^k Y_{t,i} \leq \frac{k \text{avg}(S_t)}{2}] \\ &= \Pr[\exp(-\sum_{i=1}^k Y_{t,i}) \geq \exp(-\frac{k \text{avg}(S_t)}{2})] \\ &= \frac{E[\exp(-\sum_{i=1}^k Y_{t,i})]}{\exp(-\frac{k \text{avg}(S_t)}{2})}. \end{aligned} \quad (19)$$

For all $a_i(S_t)$ satisfying $a_i(S_t) - \frac{\text{avg}(S_t)}{2} > 0$ where $i \leq k$, we use U to denote the set of their indexes and $\mathcal{I}$ to denote the collection of all non-empty subsets of U . For any $I \subseteq \mathcal{I}$, denote by E^I the event that both of the following are true: (i) $X_{t,i} = 1$ holds for all $i \in I$, and (ii) $X_{t,i} = 0$

holds for all $i \in [n]$. We have:

$$\begin{aligned}
 & \Pr[E^c] = \Pr\left[\bigwedge_{i \in [n]} (X_i = 1)\right] \\
 & \Pr\left[\bigwedge_{i \in [n]} \hat{\mu}_i(S) < \mu + \frac{\sqrt{\text{avg}(S)}}{2}\right] \\
 & = \Pr\left[\bigwedge_{i \in [n]} \hat{\mu}_i(S) < \mu + \frac{\sqrt{\text{avg}(S)}}{2}\right] \\
 & \leq \exp\left(-\frac{\sum_{i \in [n]} \left(\mu - \hat{\mu}_i(S)\right)^2}{2}\right) \leq \exp\left(-\frac{\sum_{i \in [n]} \left(\mu - \hat{\mu}_i(S)\right)^2}{2}\right) \quad (20)
 \end{aligned}$$

where the last inequality is due to Chernoff bound.

Let $\mu = \frac{\sum_{i \in [n]} \mu_i}{n}$. By (20), we have:

$$\begin{aligned}
 & E[\exp(-\sum_{i \in [n]} \mu_i Y_i)] \\
 & = \sum_{i \in [n]} \exp(-\mu_i Y_i) = \Pr[E^c] + \Pr[Y_i = 0] \\
 & \leq \exp\left(-\frac{\sum_{i \in [n]} \mu_i^2}{2}\right) + \sum_{i \in [n]} \exp(-\mu_i) \\
 & \leq 2^k \exp\left(-\frac{\sum_{i \in [n]} \mu_i^2}{8}\right)
 \end{aligned}$$

Therefore, we get:

$$\begin{aligned}
 & \Pr\left[\frac{1}{k} \sum_{i=1}^k Y_i \leq \mu + \frac{\sqrt{\text{avg}(S)}}{2}\right] \\
 & \leq 2^k \exp\left(-\frac{\sum_{i=1}^k \mu_i^2}{8}\right) \\
 & \leq \frac{2}{e} \exp\left(-\frac{\sum_{i=1}^k \mu_i^2}{8}\right)
 \end{aligned}$$

Next, we prove that $\Pr[E^B] = \Pr\left[\sum_{i \in [n]} \hat{\mu}_i(S) < \mu + \frac{\sqrt{\text{avg}(S)}}{2}\right] \leq \frac{1}{2}$. For any $i \in [n]$, define a Boolean variable Z_i that equals 1 if $\hat{\mu}_i(S) < \mu + \frac{\sqrt{\text{avg}(S)}}{2}$, or 0 otherwise. It suffices to show that at least $\frac{\sum_{i \in [n]} \mu_i}{4}$ arms have empirical values smaller than $\mu + \frac{\sqrt{\text{avg}(S)}}{2}$, i.e., $\sum_{i \in [n]} Z_i \geq \frac{\sum_{i \in [n]} \mu_i}{4}$. We show that with a large probability $\sum_{i \in [n]} Z_i \geq \frac{\sum_{i \in [n]} \mu_i}{4}$ holds. By Chernoff bounds, for any $\epsilon = \frac{1}{4}$, we have:

$$\begin{aligned}
 E[Z_i] &= \Pr\left[\hat{\mu}_i(S) < \mu + \frac{\sqrt{\text{avg}(S)}}{2}\right] \\
 &= \Pr\left[\hat{\mu}_i(S) < \mu + \frac{\sqrt{\text{avg}(S)}}{2}\right] \\
 &\leq \exp\left(-\frac{\sum_{i \in [n]} \mu_i^2}{12}\right)
 \end{aligned}$$

Let $\epsilon = \exp\left(-\frac{\epsilon^2}{12} Q_{\text{avg}}(S)\right)$. It is easy to verify that $\epsilon < \frac{1}{2}$. By Chernoff bound (7), we have

$$\begin{aligned} \Pr & \leq \prod_{i=1}^k \Pr \left[\sum_{j=1}^{|S_i|} Z_{ij} \geq \frac{|S_i|}{4} \right] \\ & \leq \prod_{i=1}^k \exp \left(-\frac{|S_i|}{2} \right) \\ & \leq \exp \left(-\frac{1}{2} \sum_{i=1}^k |S_i| \right) \\ & \leq \exp \left(-\frac{1}{2} Q_{\text{avg}}(S) \right) \\ & \leq \exp \left(-\frac{\epsilon^2}{12} Q_{\text{avg}}(S) \right) \\ & \leq \exp \left(-\frac{\epsilon^2}{12} Q_{\text{avg}}(S) \right) \\ & \leq \exp \left(-\frac{\epsilon^2}{12} Q_{\text{avg}}(S) \right) \end{aligned}$$

where the third to last inequality used the fact that $\frac{1}{2} \leq \frac{1}{4}$ and the second to last inequality used the fact that $\frac{1}{2} \leq \frac{1}{4}$.

Now, suppose QE terminates in L rounds. By summing up the failure probabilities and the errors of all rounds in QE, we have that $\sum_{i=1}^L \epsilon_i \leq 1$ and $\sum_{i=1}^L (1 - \epsilon_i) \geq 1/2$. Thus, we complete the proof of Lemma 12.

D.2 Correctness and Sample Complexity

Analysis of Uniform Sampling. Let us start with:

Lemma 15 In Uniform Sampling, if $s = \frac{1}{\epsilon} \log \frac{1}{\delta}$, then with probability at least $1 - \delta$, US returns an arm set V which satisfies $\text{avg}(V) \geq (1 - \epsilon) \text{avg}(S)$.

Proof: We use $\mathcal{U}$ to denote the collection of all-sized subsets of S . Consider the subset $\mathcal{U}^* \subseteq \mathcal{U}$ with the largest average mean. By definition we have $\text{avg}(\mathcal{U}^*) = \frac{1}{|\mathcal{U}^*|} \sum_{U \in \mathcal{U}^*} \text{avg}(U)$. By Chernoff bound (7),

$$\begin{aligned} \Pr & \leq \prod_{i=1}^k \Pr \left[\sum_{j=1}^{|S_i|} Z_{ij} \geq \frac{|S_i|}{4} \right] \\ & \leq \exp \left(-\frac{|S_i|}{2} \right) \\ & \leq \exp \left(-\frac{1}{2} Q_{\text{avg}}(S) \right) \end{aligned}$$

Consider an arbitrary subset $\mathcal{U} \subseteq \mathcal{U}$. Let $\epsilon = \frac{\text{avg}(\mathcal{U})}{\text{avg}(S)}$. If $\text{avg}(\mathcal{U}) < (1 - \epsilon) \text{avg}(S)$, we obtain:

$$\begin{aligned} \Pr & \leq \prod_{i=1}^k \Pr \left[\sum_{j=1}^{|S_i|} Z_{ij} \geq \frac{|S_i|}{4} \right] \\ & \leq \exp \left(-\frac{|S_i|}{2} \right) \\ & \leq \exp \left(-\frac{1}{2} Q_{\text{avg}}(S) \right) \end{aligned}$$

where the last inequality is due to Chernoff bound (7) (we distinguish $0 < \epsilon < 1$ and $\epsilon = 1$). The calculation details are similar to the proof of Lemma 15.

Notice that in US we have $|S_j| \leq 4k$. Applying the union bound over all the subsets of size k , we assert that the failure probability of US is at most

$$\sum_{k=1}^n \exp \left(-\frac{1}{2} Q_{\text{avg}}(S) \right) \leq \frac{1}{4}$$

Next, we show that with high probability QE-AS terminates at $2 \left[\frac{1}{8} \text{avg}(B); \text{avg}(B) \right]$. If this is not what happens, one of the following two events must have happened:

- (1) Premature QE-AS terminates when $\text{avg}(S) > \text{avg}(B)$.
- (2) Overdue QE-AS does not terminate after invoking US with $\frac{1}{4} \text{avg}(B)$.

We say that QE succeeds if the arm set S it returns satisfies $\text{avg}(S) \leq (1 + \frac{1}{2}) \text{avg}(B)$.

Lemma 16 Both of the following statements are true:

The premature event happens with probability at most

If QE succeeds, the overdue event happens with probability at most

Proof: The second statement directly follows from the proof of Lemma 15. For the first lemma, let us focus on a specific guess $\hat{U} = \frac{\text{avg}(B)}{\text{avg}(U)}$. We use $\mathcal{U}$ to denote the collection of all-sized subsets of S . Consider a specific $U \in \mathcal{U}$. Let $\hat{U} = 2 \frac{\text{avg}(B)}{\text{avg}(U)}$. 1. We claim:

$$\begin{aligned} & \Pr[\hat{U} \text{avg}(U) \geq 2] \\ &= \Pr[\hat{U} \text{avg}(U) \geq (1 + \frac{1}{2}) \text{avg}(U)] \\ &= \frac{4}{2} \exp(-\frac{5}{2}k): \end{aligned} \tag{21}$$

Notice that here we have 1. By Chernoff bound 3),

$$\begin{aligned} & \Pr[\hat{U} \text{avg}(U) \geq (1 + \frac{1}{2}) \text{avg}(U)] \\ & \exp(-\frac{1}{2}k) \leq \frac{4}{4} \exp(-\frac{5}{2}k): \end{aligned}$$

This proves (22) for $\hat{U} < 2$.

When $\hat{U} \geq 2$, by Chernoff bound 4),

$$\begin{aligned} & \Pr[\hat{U} \text{avg}(U) \geq (1 + \frac{1}{2}) \text{avg}(U)] \\ &= \frac{e^{-2Qk}}{2} \\ &= \frac{e^{-2Qk}}{2} \\ &= \frac{1}{2} \exp(-2Qk) \\ &= \frac{1}{2} \exp(-\frac{5}{2}k); \end{aligned}$$

where the second inequality holds since $\frac{1}{2} \leq \frac{1}{\hat{U}}$, and the third inequality holds since $e^{-1} < \frac{1}{2}$. Thus, we complete the proof of 1).

Notice that the size of $\mathcal{U}$ is at most $\frac{1}{k}$. Together, by the union bound, the probability of premature event for a specific $\hat{U}$ is at most $\frac{4}{2} = 2$. We use $\min$ to denote the minimum that is greater than $\text{avg}(B)$. Let $\min = \frac{\min}{k(B)}$. Summing over all $\hat{U} > \text{avg}(B)$, we assert that the probability of premature for QE-AS is at most

$$\frac{4}{2 \min} + \frac{4}{4 \min} + \frac{4}{8 \min} + \dots = \frac{4}{\min} = 4;$$

The rest of the proof is almost the same as that for ME-AS in Appendix A, which leads us to:

Theorem 5 QE-AS solves the k -AS problem with expected cost $\frac{n}{2} \frac{1}{\text{avg}(B)} \left(1 + \frac{\log(1/k)}{k} \right)$.

D.3 Lower Bound

This subsection serves as a proof for:

Theorem 6 For any $\epsilon \in (0, \frac{1}{12})$ and $\delta \in (0, \frac{1}{48})$, given any (ϵ, δ) -approximate algorithm, there is an instance of the k_{avg} -AS problem on which the algorithm must entail

$$\frac{n}{2} \frac{1}{\text{avg}(B)} \left(1 + \frac{\log(1/\delta)}{k} \right) \text{ cost in expectation.}$$

We show that the lower bound of k_{avg} -AS is the maximum of $\frac{n}{2} \frac{1}{\text{avg}(B)} \frac{\log(1/\delta)}{k}$ and

$$\frac{n}{2} \frac{1}{\text{avg}(B)}.$$

First Lower Bound. We reduce 1-AS to this problem. Suppose that for 1-AS, we have a bandit with n^0 arms. Let $\mu = (\mu_1, \dots, \mu_{n^0})$. We associate an instance of k_{avg} -AS where exactly one arm i_1 has mean $(1 + \epsilon)$ whereas all the other arms have mean 1. Without loss of generality, we assume that i_1 has mean $(1 + \epsilon)$. Recall that in Theorem 2, we showed that for any (ϵ, δ) -approximate 1-AS algorithm, distinguishing the single arm with mean $(1 + \epsilon)$ in I takes $((n^0 \log \frac{1}{\delta}) / \epsilon^2)$ samples. We construct a hard instance for k_{avg} -AS based on I .

Let $n = kn^0$. We create n artificial arms as the input to algorithm A and divide them into n^0 groups. Each group contains exactly k arms. The i -th group contains the arms with indexes from $(i-1)k + 1$ to ik . Each time when A attempts to sample an arm in the i -th group, we actually take a sample from a_i and pass the result to A . If there exists an i such that more than $\frac{2k}{3}$ arms returned by A are from group i , we select i as the answer for 1-AS. Otherwise, we select an arbitrary arm. Suppose that A can solve k_{avg} -AS within sample complexity Q . Since $\epsilon < 1/12$, by definition of k_{avg} -AS, any feasible solution must contain at least 3 arms in the first group. Thus, with probability at least $1 - \delta$, we select i_1 as the answer for 1-AS. Hence, A can solve 1-AS on the instance with $n^0 = n/k$ arms and $\mu = \text{avg}(B)$, using Q samples as well. Therefore, we conclude that

$$Q = \frac{n^0}{2} \frac{1}{\log \frac{1}{\delta}} = \frac{n}{2} \frac{1}{\text{avg}(B)} \frac{\log(1/\delta)}{k}.$$

Second Lower Bound. We define the single-arm problem as follows: with probability at least $1 - \delta$, distinguish whether a single-arm has mean $(1 + \epsilon)$.

Lemma 17 For any $\epsilon \in (0, \frac{1}{12})$, any algorithm solving the above single-arm problem must entail $(\frac{1}{2\epsilon^2} \log \frac{1}{\delta})$ cost.

Proof: Our proof is based on a sampling lower bound established in [Theorem 4.7]. We need the following definition:

Definition 2 [Hellinger distance] Suppose X and Y are two discrete distributions with a common support S . The Hellinger distance between X and Y is defined as

$$d_H(X; Y) = \frac{1}{\sqrt{2}} \sqrt{\sum_{s \in S} \frac{X(s)}{Y(s)} (\Pr[X = s])^{1/2} (\Pr[Y = s])^{1/2}}.$$

It is shown in [3] that given $0 < \epsilon < \frac{1}{4}$ and two distributions X, Y on a common support with $d_H^2(X; Y) \geq \frac{1}{2}$, any algorithm distinguishing X and Y with failure error bounded by δ requires at least $\frac{1}{4d_H^2(X; Y)} \ln \frac{1}{\delta}$ samples. We use $\mathcal{U}$ to denote the Bernoulli distribution with mean μ (similarly for $\mathcal{U}_{(1+\epsilon)}$). From Definition 2, we have that

$$\begin{aligned} & d_H^2(\mathcal{U}; \mathcal{U}_{(1+\epsilon)}) \\ &= \frac{1}{2} \frac{p}{(1+\epsilon)} \frac{p}{(1-\epsilon)(1-(1+\epsilon))} \\ &= \frac{1}{16\epsilon^2} : \end{aligned}$$

Therefore, any algorithm distinguishing two Bernoulli distribution with mean $\frac{1}{2}$ and $(1 + \frac{1}{4})$ must take at least $(\frac{1}{2} \frac{1}{1})$ samples.

The second lower bound $(\frac{n}{2} \frac{1}{\deg_{\text{avg}}(B)})$ is established by a reduction from the above single-arm problem. The proof is somewhat similar to the reduction in the proof of Lemma 1. First we show that for any $\epsilon \in (0, \frac{1}{12})$; $\delta \in (0, \frac{1}{48})$, given an (ϵ, δ) -approximate k_{avg} -AS algorithm A with expected sample complexity Q , we can design an algorithm B to solve the single-arm problem using at most $64Q = n$ samples.

Create a set of artificial arms denoted $a_1, \dots, a_n$. We randomly choose a k -sized subset $S \subseteq U$ and an arm $a_i \in S$ as the pivot arm. For any arm $a_j \in S$ with $j \neq i$, we set $(a_j) = (1 + 4\epsilon)$ and for each arm $a_j \in U \setminus S$, we set $(a_j) = \frac{1}{2}$.

Algorithm B solves the single-arm problem by simulating algorithm A . When A samples from an arm a_j with $j \neq i$, B simply do the same. On the other hand, when A samples the pivot arm a_i , B samples from a_i as in the single-arm problem, and passes the result to (i) the pivot arm is sampled less than $\frac{64Q}{n}$ times and at the same time (ii) A returns an arm set not containing a_i , B decides that $(a) = \frac{1}{2}$. In all other cases B decides that $(a) = (1 + 4\epsilon)$.

Consider the case of $(a) = \frac{1}{2}$. We use T_a to denote the number of samples taken from a_i . Let $S^0 = S \cap a_i$. Since a_i is uniformly distributed in $U \setminus S^0$ and $k \leq n/2$, we have:

$$E[T_a] \leq \frac{Q}{n - k + 1} \leq \frac{2Q}{n}.$$

Therefore, $\Pr[T_a \geq \frac{64}{n}] \leq \frac{1}{32}$. If $T_a < \frac{64}{n}$, B errs if and only if $a_i \notin V$ holds. Notice that in this case, any feasible solution must contain at least $2k = 3$ arms with mean $(1 + 4\epsilon)$. Therefore with probability at least $\frac{1}{3}$, A returns at most $\frac{k}{3} + 1$ arms with mean $\frac{1}{2}$. Thus,

$$\Pr[a_i \notin V] \leq \frac{1}{3} + (1 - \frac{1}{3}) \frac{k/3 + 1}{n - k + 1} \leq 0.65.$$

Thus, with probability at least $1 - 0.65 = 0.35$, B outputs the correct answer.

Next, consider the case where $(a) = (1 + 4\epsilon)$. In this case, any feasible solution must contain at least $2k = 3$ arms with mean $(1 + 4\epsilon)$. B errs only if $T_a < \frac{64Q}{n}$ and $a_i \notin V$. Since A is an (ϵ, δ) -approximate k_{avg} -AS algorithm, we have that $\Pr[a_i \notin V] \leq \frac{2}{3}(1 - \epsilon) \leq 0.65$. Therefore, by the lower bound for the single-arm problem, we have that $\frac{n}{2} \frac{1}{\deg_{\text{avg}}(B)} = \frac{n}{2} \frac{1}{\deg_{\text{avg}}(B)}$.

E Analysis of k_{avg} -MOST CONNECTED VERTEX

Our k_{avg} -AS algorithm, combined with the reduction described in Section 1, already settles k_{avg} -MCV with the sample complexity given in Table 1. Next, we prove the following lower bound.

Theorem 7 For any $\epsilon \in (0, \frac{1}{12})$ and $\delta \in (0, \frac{1}{48})$, the following statements are true about any k_{avg} -MCV algorithm:

when $\deg_{\text{avg}}(B) \geq \frac{1}{2\epsilon} \log \frac{n}{k}$, there is an instance on which the algorithm must probe

$$\frac{n}{2\epsilon} \frac{m}{\deg_{\text{avg}}(B)} \left(1 + \frac{\log(1/\delta)}{k} \right)$$

edges in expectation.

when $\deg_{\text{avg}}(B) < O(\frac{1}{\epsilon})$, there is an instance on which the algorithm must probe $\Omega(n)$ edges in expectation.

Case $\deg_{\text{avg}}(B) \geq \frac{1}{2\epsilon} \log \frac{n}{k}$. We first show a lower bound of $(\frac{n}{2\epsilon} \frac{m}{\deg_{\text{avg}}(B)})$, using the hard 1-AS instance constructed in the proof of Theorem 6. Recall that the instance has a set of n^0 arms $a_1, \dots, a_{n^0}$. Construct from $\mathcal{A}$ a hidden bipartite graph $G = (B; W; E)$ where $|B| = kn^0 = n$ and $|W| = m$ ($24 \log \frac{n}{k} = \frac{1}{2\epsilon}$). Divide B into n^0 groups, each having k vertices. For each black vertex in the i th group, we independently make each edge incident on it solid with probability (a_i) .

