
On Sample Complexity Upper and Lower Bounds for Exact Ranking from Noisy Comparisons

Wenbo Ren

Department of Computer Science & Engineering
The Ohio State University
ren.453@osu.edu

Jia Liu

Department of Computer Science
Iowa State University
jialiu@iastate.edu

Ness B. Shroff

Department of Electrical & Computer Engineering and Computer Science & Engineering
The Ohio State University
shroff.11@osu.edu

Abstract

This paper studies the problem of finding the exact ranking from noisy comparisons. A noisy comparison over a set of m items produces a noisy outcome about the most preferred item, and reveals some information about the ranking. By repeatedly and adaptively choosing items to compare, we want to fully rank the items with a certain confidence, and use as few comparisons as possible. Different from most previous works, in this paper, we have three main novelties: (i) compared to prior works, our upper bounds (algorithms) and lower bounds on the sample complexity (aka number of comparisons) require the minimal assumptions on the instances, and are not restricted to specific models; (ii) we give lower bounds and upper bounds on instances with *unequal* noise levels; and (iii) this paper aims at the *exact* ranking without knowledge on the instances, while most of the previous works either focus on approximate rankings or study exact ranking but require prior knowledge. We first derive lower bounds for pairwise ranking (i.e., compare two items each time), and then propose (nearly) *optimal* pairwise ranking algorithms. We further make extensions to listwise ranking (i.e., comparing multiple items each time). Numerical results also show our improvements against the state of the art.

1 Introduction

Background and motivation: Ranking from noisy comparisons has been a canonical problem in the machine learning community, and has found applications in various areas such as social choices [8], web search [10], crowd sourcing [4], and recommendation systems [3]. The main goal of ranking problems is to recover the full or partial rankings of a set of items from noisy comparisons. The items can refer to various things, such as products, movies, pages, and advertisements, and the comparisons refer to tests or queries about the items’ strengths or the users’ preferences. In this paper, we use words “item”, “comparison” and “preference” for simplicity. A comparison involves two (i.e., pairwise) or multiple (i.e., listwise) items, and returns a noisy result about the most preferred one, where “noisy” means that the comparison outcome is random and the returned item may not be the most preferred one. A noisy comparison reveals some information about the ranking of the items. This information can be used to describe users’ preferences, which helps applications such as recommendations, decision making, and advertising, etc. One example is e-commerce: A user’s click or purchase of a product (but not others) is based on a noisy (due to the lack of full information) comparison between several similar products, and one can rank the products based

on the noisy outcomes of the clicks or the purchases to give better recommendations. Due to the wide applications, in this paper, we do not focus on specific applications and regard comparisons as black-box procedures.

This paper studies the *active* (or adaptive) ranking, where the learner adaptively chooses items to compare based on previous comparison results, and returns a ranking when having enough confidence. Previous works [4, 30] have shown that, compared to non-adaptive ranking, active ranking can significantly reduce the number of comparisons needed and achieve a similar confidence or accuracy. In some applications such as news apps, the servers are able to adaptively choose news to present to the users and collect feedbacks, by which they can learn the users' preferences in shorter time compared to non-adaptive methods and may provide better user experience.

We focus on the active *full* ranking problem, that is, to find the *exact* full ranking with a certain confidence level by adaptively choosing the items to compare, and try to use as few comparisons as possible. The comparisons can be either pairwise (i.e., comparing two items each time) or listwise (i.e., comparing more than two items each time). We are interested in the upper and lower bounds on the sample complexity (aka number of comparisons needed). We are also interested in understanding whether using listwise comparisons can reduce the sample complexity.

Models and problem statement: There are n items in total, indexed by $1, 2, 3, \dots, n$. Given a comparison over a set S , each item $i \in S$ has $p_{i,S}$ probability to be returned as the most preferred one (also referred to as i “wins” this comparison), and when a tie happens, we randomly assign one item as the winner, which makes $\sum_{i \in S} p_{i,S} = 1$ for all set $S \subset [n]$. When $|S| = 2$, we say this comparison is pairwise, and when $|S| > 2$, we say listwise. In this paper, a comparison is said to be m -wise if it involves exactly m items (i.e., $|S| = m$). For $m = 2$ and a two-sized set $S = \{i, j\}$, to simplify notation, we define $p_{i,j} := p_{i,S}$ and $p_{j,i} := p_{j,S}$.

Assumptions. In this paper, we make the following assumptions: **A1**) Comparisons are independent across items, sets, and time. We note that the assumption of independence is common in the this area (e.g., [11, 12, 13, 17, 18, 24, 33, 34, 36, 37]). **A2**) There is a unique permutation (r_1, r_2, \dots, r_n) of $[n]$ ¹ such that $r_1 \succ r_2 \succ \dots \succ r_n$, where $i \succ j$ denotes that i ranks higher than j (i.e., i is more preferred than j). We refer to this unique permutation as the *true ranking* or *exact ranking*, and our goal is to recover the true ranking; **A3**) For any set S and item $i \in S$, if i ranks higher than all other items k of S , then $p_{i,S} > p_{k,S}$. For pairwise comparisons, A3 states that $i \succ j$ if and only if $p_{i,j} > 1/2$. We note that for pairwise comparisons, A3 can be viewed as the weak stochastic transitivity [35]. The three assumptions are necessary to make the exact ranking (i.e., finding the unique true ranking) problem meaningful, and thus, we say our assumptions are minimal. Except for the above three assumptions, we do *not* assume any prior knowledge of the $p_{i,S}$ values. We note that any comparison model can be fully described by the comparison probabilities $(p_{i,S} : i \in S, S \subset [n])$.

We further define some notations. Two items i and j are said to be *adjacent* if in the true ranking, there does *not* exist an item k such that $i \succ k \succ j$ or $j \succ k \succ i$. For all items i and j in $[n]$, define $\Delta_{i,j} := |p_{i,j} - 1/2|$, $\Delta_i := \min_{j \neq i} \Delta_{i,j}$, and $\tilde{\Delta}_i := \min\{\Delta_{i,j} : i \text{ and } j \text{ are adjacent}\}$. We adopt the notion of strong stochastic transitivity (SST) [12]: for all items i, j , and k satisfying $i \succ j \succ k$, it holds that $p_{i,k} \geq \max\{p_{i,j}, p_{j,k}\}$. Under the SST condition, we have $\Delta_i = \tilde{\Delta}_i$ for all items i . We note that this paper is not restricted to the SST condition. Pairwise (listwise) ranking refers to ranking from pairwise (listwise) comparisons. In this paper, $f \leq g$ means $f = O(g)$, $f \geq g$ means $f = \Omega(g)$, and $f \simeq g$ means $f = \Theta(g)$. The meanings of $O(\cdot)$, $o(\cdot)$, $\Omega(\cdot)$, $\omega(\cdot)$, and $\Theta(\cdot)$ are standard in the sense of Bachmann-Landau notation with respect to $(n, \delta^{-1}, \epsilon^{-1}, \Delta^{-1}, \eta^{-1}, (\Delta_{i,j}^{-1}, i \neq j))$. For any $a, b \in \mathbb{R}$, define $a \wedge b := \min\{a, b\}$ and $a \vee b := \max\{a, b\}$.

Problem (Exact ranking). *Given $\delta \in (0, 1/2)$ and n items, one wants to determine the true ranking with probability at least $1 - \delta$ by adaptively choosing sets of items to compare.*

Definition 1 (δ -correct algorithms). *An algorithm is said to be δ -correct for a problem if for any input instance of this problem, it, with probability at least $1 - \delta$, returns a correct result in finite time.*

Main results: First, for δ -correct pairwise ranking algorithms with no prior knowledge of the instances, we derive a sample-complexity lower bound of the form $\Omega(\sum_{i \in [n]} \Delta_i^{-2} (\log \log \Delta_i^{-1} + \log(n/\delta)))^2$, which is shown to be *tight* (up to constant factors) under SST and some mild conditions.

¹For any positive integer k , define $[k] := \{1, 2, \dots, k\}$ to simplify notation

²All log in this paper, unless explicitly noted, are natural log.

Second, for pairwise and listwise ranking under the multinomial logit (MNL) model, we derive a model-specific lower bound, which is *tight* (up to constant factors) under some mild conditions, and shows that in the worst case, the listwise lower bound is no lower than the pairwise one.

Third, we propose a pairwise ranking algorithm that requires no prior information and minimal assumptions on the instances, and its sample-complexity upper bound *matches* the lower bounds proved in this paper under the SST condition and some mild conditions, implying that both upper and lower bounds are optimal.

2 Related works

Dating back to 1994, the authors of [16] studied the noisy ranking under the strict constraint that $p_{i,j} \geq 1/2 + \Delta$ for any $i \succ j$, where $\Delta > 0$ is priorly *known*. They showed that any δ -correct algorithm needs $\Theta(n\Delta^{-2} \log(n/\delta))$ comparisons for the worst instances. However, in some cases, it is impossible to either assume the knowledge of Δ or require $p_{i,j} \geq 1/2 + \Delta$ for any $i \succ j$. Also, their bounds only depend on the minimal gap Δ but not $\Delta_{i,j}$'s or Δ_i 's, and hence is not tight in most cases. In contrast, our algorithms require *no knowledge* on the gaps (i.e., $\Delta_{i,j}$'s), and we establish sample-complexity lower bounds and upper bounds that base on *unequal* gaps, which can be much tighter when Δ_i 's vary a lot.

Another line of research is to explore the *probably approximately correct* (PAC) ranking (which aims at finding a permutation (r_1, r_2, \dots, r_n) of $[n]$ such that $p_{r_i, r_j} \geq 1/2 - \epsilon$ for all $i < j$, where $\epsilon > 0$ is a given error tolerance) under various pairwise comparison models [11, 12, 13, 31, 33, 34, 37]. When $\epsilon > 0$, the PAC ranking may not be unique. The authors of [11, 12, 13] proposed algorithms with $O(n\epsilon^{-2} \log(n/\delta))$ upper bound for PAC ranking with tolerance $\epsilon > 0$ under SST and the stochastic triangle inequality³ (STI). When ϵ goes to zero, the PAC ranking reduces to the true ranking. However, when $\epsilon > 0$, we still need some prior knowledge on $(p_{i,j} : i, j \in [n])$ to get the true ranking, as we need to know a lower bound of the values of $\Delta_{i,j}$ to ensure that the PAC ranking equals to the unique true ranking. When $\epsilon = 0$, the algorithms in [11, 12, 13] do not work. Prior to these works, the authors of [37] also studied the PAC ranking. In their work, with $\epsilon = 0$, the unique true ranking can be found by $O(n \log n \cdot \max_{i \in [n]} \{\Delta_i^{-2} \log(n\delta^{-1} \Delta_i^{-1})\})$ comparisons, which is higher than the lower bound and upper bound proved in this paper by at least a log factor.

In contrast, this paper is focused on recovering the *unique* true (exact) ranking, and there are three major motivations. First, in some applications, we prefer to find the exact order, especially in “winner-takes-all” situations. For example, when predicting the winner of an election, we prefer to get the exact result but not the PAC one, as only a few votes can completely change the result. Second, analyzing the exact ranking can help us better understand the instance-wise upper and lower bounds about the ranking problems, while the bounds of PAC ranking (e.g., in [11, 12, 13]) may only work for the worst cases. Third, exact ranking algorithms may better exploit the large gaps (e.g., Δ_i 's) to achieve lower sample complexities. In fact, when finding the PAC ranking, we can perform the exact ranking algorithm and the PAC ranking algorithm parallelly, and return a ranking whenever one of them returns. By this, when ϵ is large, we can benefit from the PAC upper bounds that depend on ϵ^{-2} , and when ϵ is small, we can benefit from the exact ranking bounds that depend on Δ_i^{-2} .

There are also other interesting active ranking works. Authors of [17, 18, 24, 36] studied active ranking under the Borda-Score model, where the Borda-Score of item i is defined as $\frac{1}{n-1} \sum_{j \neq i} p_{i,j}$. We note that the Borda-Score model does not satisfy A2 and A3 and is not comparable with the model in this paper. There are also many works on best item(s) selection, including [1, 5, 7, 21, 28, 29, 34], which are less related to this paper.

3 Lower bound analysis

3.1 Generic lower bound for δ -correct algorithms

In this subsection, we establish a sample-complexity lower bound for pairwise ranking. The lower bound is for δ -correct algorithms, which have performance guarantee for all input instances. There are algorithms that work faster than our lower bound but only return correct results with $1 - \delta$

³Stochastic triangle inequality means that for all items i, j, k with $i \succ j \succ k$, $\Delta_{i,k} \leq \Delta_{i,j} + \Delta_{j,k}$.

confidence for a restricted class of instances, which is discussed in Section A.1 of supplementary material. Theorem 2 states the lower bound, and its full proof is provided in supplementary material. Here we remind that $\tilde{\Delta}_i := \min\{\Delta_{i,j} : i \text{ and } j \text{ are adjacent}\}$.

Theorem 2 (Lower bound for pairwise ranking). *Given $\delta \in (0, 1/12)$ and an instance \mathcal{I} with n items, then the number of comparisons used by a δ -correct algorithm \mathcal{A} with no prior knowledge about the gaps of \mathcal{I} is lower bounded by*

$$\Omega\left(\sum_{i \in [n]} [\tilde{\Delta}_i^{-2}(\log \log \tilde{\Delta}_i^{-1} + \log(1/\delta))] + \min\left\{\sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(1/x_i) : \sum_{i \in [n]} x_i \leq 1\right\}\right). \quad (1)$$

If $\delta \preceq 1/\text{poly}(n)^4$, or $\max_{i,j \in [n]} \{\tilde{\Delta}_i/\tilde{\Delta}_j\} \preceq n^{1/2-p}$ for some constant $p > 0$, then the lower bound becomes

$$\Omega\left(\sum_{i \in [n]} \tilde{\Delta}_i^{-2}(\log \log \tilde{\Delta}_i^{-1} + \log(n/\delta))\right). \quad (2)$$

Remark: (i) When the instance satisfies the SST condition (the algorithm does not need to know this information), the bound in Eq. (2) is tight (up to a constant factor) under the given condition, which will be shown in Theorem 12 later. (ii) The lower bound in Eq. (1) implies an $n \log n$ term in $\min\{\cdot\}$, which can be checked by the convexity of $\log(1/x_i)$ and Jensen's inequality, which yields $\sum_{i \in [n]} \log(1/x_i) \geq n \log(n/\sum_{i \in [n]} x_i) \geq n \log n$. (iii) The lower bound in (2) may not hold if the required conditions do not hold, which will be discussed in Section A.2 of supplementary material.

Proof sketch of Theorem 2. Due to space limitation, we outline the basic idea of the proof here and refer readers to supplementary material for details. Our first step is to use the results in [15, 20, 27] to establish a lower bound for ranking two items. Then, it seems straightforward that the lower bound for ranking n items can be obtained by summing up the lower bounds for ranking $\{q_1, q_2\}$, $\{q_2, q_3\}, \dots, \{q_{n-1}, q_n\}$, where $q_1 \succ q_2 \succ \dots \succ q_n$ is the true ranking. However, Note that to rank q_i and q_j , there may be an algorithm that compares q_i and q_j with other items like q_k , and uses the comparison outcomes over $\{q_i, q_k\}$ and $\{q_j, q_k\}$ to determine the order of q_i and q_j . Since it is unclear to what degree comparing q_i and q_j with other items can help to rank q_i and q_j , the lower bound for ranking n items cannot be simply obtained by summing up the lower bounds for ranking 2 items. To overcome this challenge, our strategy is to construct two problems: \mathcal{P}_1 and \mathcal{P}_2 with decreasing influence of this type of comparisons. Then, we prove that \mathcal{P}_1 reduces to exact ranking and \mathcal{P}_2 reduces to \mathcal{P}_1 . Third, we prove a lower bound on δ -correct algorithms for solving \mathcal{P}_2 , which yields a lower bound for exact ranking. Finally, we use this lower bound to get the desired lower bounds in Eq. (1) and Eq. (2). \square

3.2 Model-specific lower bound

In Section 3.1, we provide a lower bound for δ -correct algorithms that do not require any knowledge of the instances except assumptions A1 to A3. However, in some applications, people may focus on a specific model, and hence, the algorithm may have further knowledge about the instances, such as the model's restrictions. Hence, the lower bound in Theorem 2 may not be applicable any more⁵.

In this paper, we derive a model-specific lower bound for the MNL model. The MNL model can be applied to both pairwise and listwise comparisons. For pairwise comparisons, the MNL model is mathematically equivalent to the Bradley-Terry-Luce (BTL) model [26] and the Plackett-Luce (PL) model [37]. There have been many prior works that focus on active ranking based on this model (e.g., [5, 6, 7, 17, 21, 29, 33, 37]).

Under the MNL model, each item holds a real number representing the users' preference over this item, where the larger the number, the more preferred the item. Specifically, each item i holds a parameter $\gamma_i \in \mathbb{R}$ such that for any set S containing i , $p_{i,S} = \exp(\gamma_i) / \sum_{j \in S} \exp(\gamma_j)$. To simplify notation, we let $\theta_i = \exp(\gamma_i)$, hence, $p_{i,S} = \theta_i / \sum_{j \in S} \theta_j$. We name θ_i as the *preference score* of item i . We define $\Delta_{i,j} := |p_{i,j} - 1/2|$, $\Delta_i := \min_{j \neq i} \Delta_{i,j}$, and we have $\tilde{\Delta}_i = \Delta_i$, i.e., the MNL model satisfies the SST condition.

⁴ $\text{poly}(n)$ means a polynomial function of n , and $\delta \preceq 1/\text{poly}(n)$ means $\delta \preceq n^{-p}$ for some constant $p > 0$.

⁵For example, under a model with $\Delta_{i,j} = \Delta$ for any $i \neq j$ where $\Delta > 0$ is unknown, one may first estimate a lower bound of Δ , and then perform algorithms in [16], yielding a sample complexity lower than Theorem 2.

Theorem 3. [Lower bound for the MNL model] Let $\delta \in (0, 1/12)$ and given a δ -correct algorithm \mathcal{A} with the knowledge that the input instances satisfy the MNL model, let $N_{\mathcal{A}}$ be the number of comparisons conducted by \mathcal{A} , then $\mathbb{E}[N_{\mathcal{A}}]$ is lower bounded by Eq. (1) with a different hidden constant factor. When $\delta \preceq 1/\text{poly}(n)$ or $\max_{i,j \in [n]} \{\Delta_i/\Delta_j\} \preceq n^{1/2-p}$ for some constant $p > 0$, the sample complexity is lower bounded by Eq. (2) with a different hidden constant factor.

Proof sketch. We prove this theorem by Lemmas 4, 5 and 6, which could be of independent interest.

Suppose that there are two coins with unknown *head probabilities* (the probability that a toss produces a head) λ and μ , respectively, and we want to find the more biased one (i.e., the one with the larger head probability). Lemma 4 states a lower bound on the number of heads or tails generated for finding the more biased coin, which works even if λ and μ go to 0. This is in contrast to the lower bounds on the number of tosses given by previous works [20, 23, 27], which go to infinity as λ and μ go to 0.

Lemma 4 (Lower bound on number of heads). Let $\lambda + \mu \leq 1$, $\Delta := |\lambda/(\lambda + \mu) - 1/2|$, and $\delta \in (0, 1/2)$ be given. To find the more biased coin with probability $1 - \delta$, any δ -correct algorithm for this problem produces $\Omega(\Delta^{-2}(\log \log \Delta^{-1} + \log \delta^{-1}))$ heads in expectation.

Now we consider n coins C_1, C_2, \dots, C_n with mean rewards $\mu_1, \mu_2, \dots, \mu_n$, respectively, where for any $i \in [n]$, $\theta_i/\mu_i = c$ for some constant $c > 0$. Define the gaps of coins $\Delta_{i,j}^c := |\mu_i/(\mu_i + \mu_j) - 1/2|$, and $\Delta_i^c := \min_{j \neq i} \Delta_{i,j}^c$. We can check that for all i and j , $\Delta_{i,j}^c = \Delta_{i,j}$, and $\Delta_i = \tilde{\Delta}_i = \Delta_i^c$.

Lemma 5 (Lower bound for arranging coins). For $\delta < 1/12$, to arrange these coins in ascending order of head probabilities, the number of heads generated by any δ -correct algorithm is lower bounded by Eq. (1) with a (possibly) different hidden constant factor.

The next lemma shows that any algorithm that solves a ranking problem under the MNL model can be transformed to solve the pure exploration multi-armed bandit (PEMAB) problem with Bernoulli rewards (e.g., [20, 22, 32]). Previous works [1, 17, 18] have shown that certain types of pairwise ranking problems (e.g., Borda-Score ranking) can also be transformed to PEMAB problems. But in this paper, we make a *reverse connection* that bridges these two classes of problems, which may be of independent interest. We note that in our prior work [31], we proved a similar result.

Lemma 6 (Reducing PEMAB problems to ranking). If there is a δ -correct algorithm that correctly ranks $[n]$ with probability $1 - \delta$ by M expected number of comparisons, then we can construct another δ -correct algorithm that correctly arranges the coins C_1, C_2, \dots, C_n in the order of ascending head probabilities with probability $1 - \delta$ and produces M heads in expectation.

The theorem follows by Lemmas 5 and 6. A full proof can be found in supplementary material. \square

3.3 Discussions on listwise ranking

A listwise comparison compares m ($m > 2$) items and returns a noisy result about the most preferred item. It is an interesting question whether exact ranking from listwise comparisons requires less comparisons. The answer is ‘‘It depends.’’ When every comparison returns the most preferred item with high probability (w.h.p.)⁶, then, by conducting m -wise comparisons, the number of comparisons needed for exact ranking is $\Theta(n \log_m n)$, i.e., there is a $\log m$ reduction, which is stated in Proposition 7. The proof can be found in supplementary material.

Proposition 7 (Listwise ranking with negligible noises). If all comparisons are correct w.h.p., to exactly rank n items w.h.p. by using m -wise comparisons, $\Theta(n \log_m n)$ comparisons are needed.

In general, when the ‘‘w.h.p. condition’’ is violated, listwise ranking does not necessarily require less comparisons than pairwise ranking (in order sense). Here, we give an example. For more general models, it remains an open problem to identify the theoretical limits, which is left for future studies.

Theorem 8. Under the MNL model, given n items with preference scores $\theta_1, \theta_2, \dots, \theta_n$ and $\Delta_{i,j} := |\theta_i/(\theta_i + \theta_j) - 1/2|$, $\tilde{\Delta}_i = \Delta_i := \min_{j \neq i} \Delta_{i,j}$, to correctly rank these n items with probability $1 - \delta$, even with m -wise comparisons for all $m \in \{2, 3, \dots, n\}$, the lower bound is the same as the pairwise ranking (i.e., Theorem 3) with (possibly) different hidden constant factors.

⁶In this paper, ‘‘w.h.p.’’ means with probability at least $1 - n^{-p}$, where $p > 0$ is a sufficiently large constant.

Theorem 8 gives a minimax lower bound for listwise ranking, which is the same as pairwise ranking. The proof is given in supplementary material. In [5], the authors have shown that for top- k item selection under the MNL model, listwise comparisons can reduce the number of comparisons needed compared with pairwise comparisons. However, for exact ranking, listwise comparisons cannot.

4 Algorithm and the upper bound for pairwise ranking

In this section, we establish a (nearly) sample-complexity optimal δ -correct algorithm for exact ranking, where whether the word “nearly” can be deleted depends on the structures of the instances. The algorithm is based on Binary Search proposed in [16] with upper bound $O(n\Delta_{\min}^{-2} \log(n/\delta))$, where $\Delta_{\min} := \min_{i \neq j} \Delta_{i,j}$. Binary Search has two limitations: (i) it requires the knowledge of Δ_{\min} *a priori* to run, and (ii) it does not utilize the unequal noise levels.

In this paper, we propose a technique named *Attempting with error prevention* and establish a corresponding insertion subroutine that attempts to insert an item i into a sorted list with a guessing Δ_i -value, while preventing errors from happening if the guess is not well chosen. If the guess is small enough, this subroutine correctly inserts the item with a large probability, and if not, this subroutine will, with a large probability, not insert the item into a wrong position. By attempting to insert item i with diminishing guesses of Δ_i , this subroutine finally correctly inserts item i with a large confidence.

To implement the technique “Attempting with error prevention”, we first need to construct a useful subroutine called Attempting-Comparison (ATC), which attempts to rank two items with ϵ , a guess of $\Delta_{i,j}$. Then, by ATC, we establish Attempting-Insertion (ATI), which also adopts this technique.

Subroutine 1 Attempting-Comparison(i, j, ϵ, δ) (ATC)

Initialize: $\forall t$, let $b^t = \sqrt{\frac{1}{2t} \log \frac{\pi^2 t^2}{3\delta}}$; $b^{max} \leftarrow \lceil \frac{1}{2\epsilon^2} \log \frac{2}{\delta} \rceil$; $w_i \leftarrow 0$;

- 1: **for** $t \leftarrow 1$ to b^{max} **do**
 - 2: Compare i and j once; Update $w_i \leftarrow w_i + 1$ if i wins; Update $\hat{p}_i^t \leftarrow w_i/t$;
 - 3: **if** $\hat{p}_i^t > 1/2 + b^t$ **then return** i ;
 - 4: **if** $\hat{p}_i^t < 1/2 - b^t$ **then return** j ;
 - 5: **end for**
 - 6: **return** i **if** $\hat{p}_i^t > 1/2$; **return** j **if** $\hat{p}_i^t < 1/2$; and **return** a random item **if** $\hat{p}_i^t = 1/2$;
-

Lemma 9 (Theoretical Performance of ATC). *ATC terminates after at most $b^{max} = O(\epsilon^{-2} \log(1/\delta))$ comparisons and returns the more preferred item with probability at least $1/2$. Further, if $\epsilon \leq \Delta_{i,j}$, then ATC returns the more preferred item with probability at least $1 - \delta$.*

Next, to establish insertion subroutine ATI, we introduce preference interval trees [16] (PIT). A PIT is constructed from a sorted list of items. For a sorted list of items S with size l , without loss of generality, we assume that $r_1 \succ r_2 \succ \dots \succ r_l$. We introduce two artificial items $-\infty$ and $+\infty$, where $-\infty$ is such that $p_{i,-\infty} = 1$ for any item i , and $+\infty$ is such that $p_{i,+\infty} = 0$ for any item i .

Preference Interval Tree [16]. A preference interval tree constructed from the sorted list S satisfies the following conditions: (i) It is a binary tree with depth $\lceil 1 + \log_2(|S| + 1) \rceil$. (ii) Each node u holds an interval $(u.\text{left}, u.\text{right})$ where $u.\text{left}, u.\text{right} \in S \cup \{-\infty, +\infty\}$, and if u is non-leaf, it holds an item $u.\text{mid}$ satisfying $u.\text{right} \succ u.\text{mid} \succ u.\text{left}$. (iii) A node i is in the interval (j, k) if and only if $k \succ i \succ j$. (iv) The root node is with interval $(-\infty, +\infty)$. From left to right, the leaf nodes are with intervals $(-\infty, r_l), (r_l, r_{l-1}), (r_{l-1}, r_{l-2}), \dots, (r_2, r_1), (r_1, +\infty)$. (v) Each non-leaf node u has two children $u.\text{lchild}$ and $u.\text{rchild}$ such that $u.\text{left} = u.\text{lchild}.\text{left}$, $u.\text{right} = u.\text{rchild}.\text{right}$ and $u.\text{mid} = u.\text{lchild}.\text{right} = u.\text{rchild}.\text{left}$.

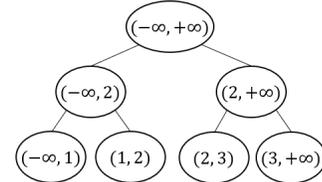


Figure 1: An example of PIT, constructed from a sorted list with three items $3 \succ 2 \succ 1$.

Based on the notion of PIT, we present insertion subroutine ATI in Subroutine 2. ATI runs a random walk on the PIT to insert i into S . Let X be the point that moves on the tree. We say a leaf u_0 *correct*

if the item i belongs to $(u_0.\text{left}, u_0.\text{right})$. Define $d(X) :=$ the distance (i.e., the number of edges) between X and u_0 . At each round of the subroutine, if all comparisons give correct results, we say this round is *correct*, otherwise we say *incorrect*. For each correct round, either $d(X)$ is decreased by 1 or the counter of u_0 is increased by 1. The subroutine inserts i into u_0 if u_0 is counted for $1 + \frac{5}{16}t^{\max}$ times. Thus, after t^{\max} rounds, the subroutine correctly inserts i into S if the number of correct rounds is no less than $\frac{21}{32}t^{\max} + \frac{h}{2}$, where $h = \lceil 1 + \log_2(|S| + 1) \rceil$ is the depth of the tree. If guessing $\epsilon \leq \Delta_i$, then each round is correct with probability at least q , making the subroutine correctly insert item i with probability at least $1 - \delta$.

For all $\epsilon > 0$, each round is incorrect with probability at most $1/2$, and thus, by concentration inequalities, we can also show that with probability at least $1 - \delta$, i will not be placed into any leaf node other than u_0 . That is, if $\epsilon > \Delta_i$, the subroutine either correctly inserts i or returns *unsure* with probability at least $1 - \delta$. The choice of parameters guarantees the sample complexity. Lemma 10 states its theoretical performance, and the proof is relegated to the supplementary material.

Subroutine 2 Attempting-Insertion(i, S, ϵ, δ) (ATI).

Initialize: Let T be a PIT constructed from S ;
 $h \leftarrow \lceil 1 + \log_2(1 + |S|) \rceil$, the depth of T ;
For all leaf nodes u of T , initialize $c_u \leftarrow 0$;
Set $t^{\max} \leftarrow \lceil \max\{4h, \frac{512}{25} \log \frac{2}{\delta}\} \rceil$ and $q \leftarrow \frac{15}{16}$;

- 1: $X \leftarrow$ the root node of T ;
- 2: **for** $t \leftarrow 1$ to t^{\max} **do**
- 3: **if** X is the root node **then**
- 4: **if** $\text{ATC}(i, X.\text{mid}, \epsilon, 1 - q) = i$ **then** $X \leftarrow X.\text{right}$; #i.e., ATC returns $i \succ X.\text{mid}$
- 5: **else** $X \leftarrow X.\text{left}$;
- 6: **else if** X is a leaf node **then**
- 7: **if** $\text{ATC}(i, X.\text{left}, \epsilon, 1 - \sqrt{q}) = i \wedge \text{ATC}(i, X.\text{right}, \epsilon, 1 - \sqrt{q}) = X.\text{right}$ **then**
- 8: $c_X \leftarrow c_X + 1$;
- 9: **if** $c_X > b^t := \frac{1}{2}t + \sqrt{\frac{t}{2} \log \frac{\pi^2 t^2}{3\delta}} + 1$ **then**
- 10: Insert i into the corresponding interval of X and **return** *inserted*;
- 11: **else if** $c_X > 0$ **then** $c_X \leftarrow c_X - 1$
- 12: **else** $X \leftarrow X.\text{parent}$
- 13: **else**
- 14: **if** $\text{ATC}(i, X.\text{left}, \epsilon, 1 - \sqrt[3]{q}) = X.\text{left} \vee \text{ATC}(i, X.\text{right}, \epsilon, 1 - \sqrt[3]{q}) = i$ **then**
- 15: $X \leftarrow X.\text{parent}$;
- 16: **else if** $\text{ATC}(i, X.\text{mid}, \epsilon, 1 - \sqrt[3]{q}) = i$ **then** $X \leftarrow X.\text{rchild}$;
- 17: **else** $X \leftarrow X.\text{lchild}$;
- 18: **end for**
- 19: **if** there is a leaf node u with $c_u \geq 1 + \frac{5}{16}t^{\max}$ **then**
- 20: Insert i into the corresponding interval of u and **return** *inserted*;
- 21: **else return** *unsure*;

Lemma 10 (Theoretical performance of ATI). *Let $\delta \in (0, 1)$. ATI returns after $O(\epsilon^{-2} \log(|S|/\delta))$ comparisons and, with probability at least $1 - \delta$, correctly inserts i or returns *unsure*. Further, if $\epsilon \leq \Delta_i$, it correctly inserts i with probability at least $1 - \delta$.*

By Lemma 10, we can see that the idea ‘‘Attempting with error prevention’’ is successfully implemented. Thus, by repeatedly attempting to insert an item with diminishing guess ϵ with proper confidences for the attempts, one can finally correctly insert i with probability $1 - \delta$. We use this idea to establish the insertion subroutine Iterative-Attempting-Insertion (IAI, Subroutine 3), and then use it to establish the ranking algorithm Iterative-Insertion-Ranking (IIR, Algorithm 4). Their theoretical performances are stated in Lemma 11 and Theorem 12, respectively, and their proofs are given in supplementary material.

Lemma 11 (Theoretical Performance of IAI). *With probability at least $1 - \delta$, IAI correctly inserts i into S , and conducts at most $O(\Delta_i^{-2}(\log \log \Delta_i^{-1} + \log(|S|/\delta)))$ comparisons.*

Theorem 12 (Theoretical Performance of IIR). *With probability at least $1 - \delta$, IIR returns the exact ranking of $[n]$ and conducts at most $O(\sum_{i \in [n]} \Delta_i^{-2}(\log \log \Delta_i^{-1} + \log(n/\delta)))$ comparisons.*

Subroutine 3 Iterative-Attempting-Insertion (IAI). **Algorithm 4** Iterative-Insertion-Ranking (IIR).

Input parameters: (i, S, δ) ;

Initialize: For all $\tau \in \mathbb{Z}^+$, set $\epsilon_\tau = 2^{-(\tau+1)}$ and $\delta_\tau = \frac{6\delta}{\pi^2\tau^2}$; $t \leftarrow 0$; $Flag \leftarrow \textit{unsure}$;

- 1: **repeat** $t \leftarrow t + 1$;
 - 2: $Flag \leftarrow \text{ATI}(i, S, \epsilon_t, \delta_t)$;
 - 3: **until** $Flag = \textit{inserted}$
-

Input: $S = [n]$, and confidence $\delta > 0$;

- 1: $Ans \leftarrow$ the list containing only $S[1]$;
 - 2: **for** $t \leftarrow 2$ to $|S|$ **do**
 - 3: $\text{IAI}(S[t], Ans, \delta/(n-1))$;
 - 4: **end for**
 - 5: **return** Ans ;
-

Remark: We can see that the upper bounds of IIR depend on the values of $(\Delta_i, i \in [n])$ while the lower bounds given in Theorem 2 depend on the values of $(\tilde{\Delta}_i, i \in [n])$. Without SST, it is possible $\tilde{\Delta}_i < \Delta_i$, but if SST holds, then our algorithm is optimal up to a constant factor given $\delta \leq 1/\text{poly}(n)$, or $\max_{i,j \in [n]} \tilde{\Delta}_i / \tilde{\Delta}_j \leq O(n^{1/2-p})$ for some constant $p > 0$. According to [11, 12, 13], ranking without the SST condition can be much harder than that with SST, and it remains an open problem whether our upper bound is tight or not when the SST condition does not hold.

5 Numerical results

In this section, we provide numerical results to demonstrate the efficacy of our proposed IIR algorithm. We compare IIR with: (i) Active-Ranking (AR) [17], which focuses on the Borda-Score model and is not directly comparable to our algorithm. We use it as an example to show that although Borda-Ranking may be the same as exact ranking, for finding the exact ranking, the performance of Borda-Score algorithms is not always as good as that for finding the Borda-Ranking⁷; (ii) PLPAC-AMPR [37], an algorithm for PAC ranking under the MNL model. By setting the parameter $\epsilon = 0$, it can find the exact ranking with $O((n \log n) \max_{i \in [n]} \Delta_i^{-2} \log(n \Delta_i^{-1} \delta^{-1}))$ comparisons, higher than our algorithm by at least a log factor; (iii) UCB + Binary Search of [16]. In the Binary Search algorithm of [16], a subroutine that ranks two items with a constant confidence is required. In [16], it assumes the value of $\Delta_{\min} = \min_{i \in [n]} \Delta_i$ is priorly known, and the subroutine is simply comparing two items for $\Theta(\Delta_{\min}^{-2})$ times and returns the item that wins more. In this paper, the value of Δ_{\min} is not priorly known, and here, we use UCB algorithms such as LUCB [25] to play the role of the required subroutine. The UCB algorithms that we use include Hoeffding-LUCB [19, 25], KL-LUCB [2, 25], and lil'UCB [20]. For Hoeffding-LUCB and KL-LUCB, we choose $\gamma = 2$. For lil'UCB, we choose $\epsilon = 0.01$, $\beta = 1$, and $\lambda = (\frac{2+\beta}{\beta})^2$ ⁸.

Instances. The experiments are conducted on three different types of instances. To simplify notation, we use $r_1 \succ r_2 \succ \dots \succ r_n$ to denote the true ranking, and let $\Delta = 0.1$. (i) Type-Homo: For any $r_i \succ r_j$, $p_{r_i, r_j} = 1/2 + \Delta$. (ii) Type-MNL: The preference score of r_i (i.e., θ_{r_i}) is generated by taking an independent instance of Uniform($[0.9 * 1.5^{n-i}, 1.1 * 1.5^{n-i}]$). By this, for any i , Δ_i is around 0.1. (iii) Type-Random: For any $r_i \succ r_j$, p_{r_i, r_j} is generated by taking an independent instance of Uniform($[0.5 + 0.8\Delta, 0.5 + 1.5\Delta]$). By this, for any i , Δ_i is around 0.1. We let Δ_i 's be close to 0.1 in order to decrease the influence of Δ_i 's on sample complexities and show how the sample complexities of the algorithms grow with n .

The numerical results for these three types are presented in Figure 2 (a)-(c), respectively. For all simulations, we input $\delta = 0.01$. Every point of every figure is averaged over 100 independent trials. In every figure, for the same n -value, the algorithms are tested on an identical input instance.

From Figure 2, we can see that our algorithm significantly outperforms the existing algorithms. We can also see that the sample complexity of IIR scales with $n \log n$, which is consistent with our theoretical results. There are some insights about the practical performance of IIR. First, in Lines 3 and 4 of ATC and Lines 9 and 10 of ATI, we use LUCB-like [25] designs to allow the algorithms return before completing all required iterations, which does not improve the theoretical upper bound but can improve the practical performance. Second, in the theoretical analysis, we only show that

⁷For instance, when $p_{r_i, r_j} = 1/2 + \Delta$ for all $i < j$, the Borda-Score of item r_i is $\frac{1}{n-1} \sum_{j \neq i} p_{r_i, r_j} = 1/2 + \frac{n+1-2i}{n-1} \Delta$, and $\Delta_{r_i} = \Theta(1/n)$. Thus, by [17], the sample complexity of AR is at least $O(n^3 \log n)$.

⁸We do not choose the combination ($\epsilon = 0$, $\beta = 1$, and $\lambda = 1 + 10/n$) that has a better practical performance because this combination does not have theoretical guarantee, making the comparison in some sense unfair.

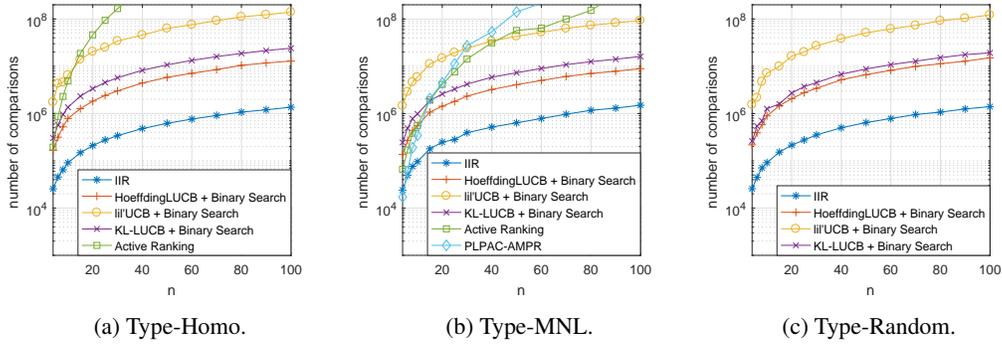


Figure 2: Comparisons between IIR and existing methods.

ATI correctly inserts an item i with high probability when inputting $\epsilon \leq \Delta_i$, but the algorithm may return before ϵ being that small, making the practical performance better than what the theoretical upper bound suggests.

6 Conclusion

In this paper, we investigated the theoretical limits of exact ranking with minimal assumptions. We do not assume any prior knowledge of the comparison probabilities and gaps, and derived the lower bounds and upper bound for instances with unequal noise levels. We also derived the model-specific pairwise and listwise lower bound for the MNL model, which further shows that in the worst case, listwise ranking is no more efficient than pairwise ranking in terms of sample complexity. The iterative-insertion-ranking (IIR) algorithm proposed in this paper indicates that our lower bounds are optimal under strong stochastic transitivity (SST) and some mild conditions. Numerical results also suggest that our ranking algorithm outperforms existing works in the literature.

Acknowledgments

This work has been supported in part by NSF grants ECCS-1818791, CCF-1758736, CNS-1758757, CNS-1446582, CNS-1901057; ONR grant N00014-17-1-2417; AFRL grant FA8750-18-1-0107, and by Institute for Information & communications Technology Promotion (IITP) grant funded by the Korea government (MSIT), (2017-0-00692, Transport-aware Streaming Technique Enabling Ultra Low-Latency AR/VR Services).

References

- [1] Agarwal, A., Agarwal, S., Assadi, S., and Khanna, S. (2017). Learning with limited rounds of adaptivity: Coin tossing, multi-armed bandits, and ranking from pairwise comparisons. In *In Conference on Learning Theory*.
- [2] Arratia, R. and Gordon, L. (1989). Tutorial on large deviations for the Binomial distribution. *Bulletin of Mathematical Biology*.
- [3] Baltrunas, L., Makcinskas, T., and Ricci, F. (2010). Group recommendations with rank aggregation and collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems*. ACM.
- [4] Chen, X., Bennett, P. N., Collins-Thompson, K., and Horvitz, E. (2013). Pairwise ranking aggregation in a crowdsourced setting. In *In ACM Conference on Web Search and Data Mining*. ACM.
- [5] Chen, X., Li, Y., and Mao, J. (2018). A nearly instance optimal algorithm for top-k ranking under the multinomial logit model. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Algorithms*. SIAM.
- [6] Chen, Y., Fan, J., Ma, C., and Wang, K. (2017). Spectral method and regularized MLE are both optimal for top-k ranking. *stat*.
- [7] Chen, Y. and Suh, C. (2015). Spectral MLE: Top-k rank aggregation from pairwise comparisons. In *International Conference on Machine Learning*.
- [8] Conitzer, V. and Sandholm, T. (2005). Communication complexity of common voting rules. In *Proceedings of the 6th ACM conference on Electronic commerce*. ACM.
- [9] Cover, T. and Thomas, J. (1991). *Elements of Information Theory*. John Wiley & Sons.
- [10] Dwork, C., Kumar, R., Naor, M., and Sivakumar, D. (2001). Rank aggregation methods for the web. In *In Proceedings of the 10th international conference on World Wide Web*. ACM.
- [11] Falahatgar, M., Hao, Y., Orlitsky, A., Pichapati, V., and Ravindrakumar, V. (2017a). Maxing and ranking with few assumptions. In *In Advances in Neural Information Processing Systems*.
- [12] Falahatgar, M., Jain, A., Orlitsky, A., Pichapati, V., and Ravindrakumar, V. (2018). The limits of maxing, ranking, and preference learning. In *Proceedings of the 35th International Conference on Machine Learning*. PMLR.
- [13] Falahatgar, M., Orlitsky, A., Pichapati, V., and Suresh, A. T. (2017b). Maximum selection and ranking under noisy comparisons. In *International Conference on Machine Learning*.
- [14] Fano, R. M. (1961). Transmission of information. *Bell System technical journal*.
- [15] Farrell, R. H. (1964). Asymptotic behavior of expected sample size in certain one sided tests. *The Annals of Mathematical Statistics*.
- [16] Feige, U., Raghavan, P., Peleg, D., and Upfal, E. (1994). Computing with noisy information. *SIAM Journal on Computing*.
- [17] Heckel, R., Shah, N. B., Ramchandran, K., and Wainwright, M. J. (2016). Active ranking from pairwise comparisons and when parametric assumptions don't help. *arXiv preprint arXiv:1606.08842*.
- [18] Heckel, R., Simchowitz, M., Ramchandran, K., and Wainwright, M. J. (2018). Approximate ranking from pairwise comparisons. In *International Conference on Artificial Intelligence and Statistics*.
- [19] Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*.
- [20] Jamieson, K., Malloy, M., Nowak, R., and Bubeck, S. (2014). lil'UCB: An optimal exploration algorithm for multi-armed bandits. In *In Conference on Learning Theory*.

- [21] Jang, M., Kim, S., Suh, C., and Oh, S. (2017). Optimal sample complexity of m-wise data for top-k ranking. In *In Advances in Neural Information Processing Systems*.
- [22] Kalyanakrishnan, S. and Stone, P. (2010). Efficient selection of multiple bandit arms: Theory and practice. In *International Conference on Machine Learning*.
- [23] Kalyanakrishnan, S., Tewari, A., Auer, P., and Stone, P. (2012). PAC subset selection in stochastic multi-armed bandits. In *ICML*.
- [24] Katariya, S., Jain, L., Sengupta, N., Evans, J., and Nowak, R. (2018). Adaptive sampling for coarse ranking. In *International Conference on Artificial Intelligence and Statistics*.
- [25] Kaufmann, E. and Kalyanakrishnan, S. (2013). Information complexity in bandit subset selection. In *Conference on Learning Theory*.
- [26] Luce, R. D. (2012). *Individual choice behavior: A theoretical analysis*. Courier Corporation.
- [27] Mannor, S. and Tsitsiklis, J. N. (2004). The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research*.
- [28] Mohajer, S. and Suh, C. (2016). Active top-k ranking from noisy comparisons. In *Communication, Control, and Computing (Allerton), 2016 54th Annual Allerton Conference on*. IEEE.
- [29] Negahban, S., Oh, S., and Shah, D. (2016). Rank centrality: Ranking from pairwise comparisons. In *Operations Research*.
- [30] Pfeiffer, T., Xi, A., Gao, A., Mao, Y., Chen, and Rand, D. G. (2012). Adaptive polling for information aggregation. In *Twenty-sixth AAAI Conference on Artificial Intelligence*.
- [31] Ren, W., Liu, J., and Shroff, N. B. (2018). PAC ranking from pairwise and listwise queries: Lower bounds and upper bounds. *arXiv preprint arXiv:1806.02970*.
- [32] Ren, W., Liu, J., and Shroff, N. B. (2019). Exploring k out of top ρ fraction of arms in stochastic bandits. In *The 22nd International Conference on Artificial Intelligence and Statistics*.
- [33] Saha, A. and Gopalan, A. (2019a). Active ranking with subset-wise preferences. In *Proceedings of Machine Learning Research*, Proceedings of Machine Learning Research. PMLR.
- [34] Saha, A. and Gopalan, A. (2019b). From PAC to instance-optimal sample complexity in the Plackett-Luce model. *arXiv preprint arXiv:1903.00558*.
- [35] Shah, N., Balakrishnan, S., Guntuboyina, A., and Wainwright, M. (2016). Stochastically transitive models for pairwise comparisons: Statistical and computational issues. In *International Conference on Machine Learning*.
- [36] Shah, N. B. and Wainwright, M. J. (2017). Simple, robust and optimal ranking from pairwise comparisons. *Journal of machine learning research*.
- [37] Szörényi, B., Busa-Fekete, R., Paul, A., and Hüllermeier, E. (2015). Online rank elicitation for Plackett-Luce: A dueling bandits approach. In *In Advances in Neural Information Processing Systems*.

Supplementary material

A Further discussions

A.1 Non- δ -correct algorithms

In Section 1, we define the notion of δ -correct algorithms, which return correct results with probability at least $1 - \delta$ for any input instances satisfying assumptions A1 to A3 (defined in Section 1). It is reasonable to consider δ -correct algorithms since we may not want an algorithm that performs pretty well on some instances but badly on others. However, to give better insights about δ -correct algorithms and the lower bounds in Theorem 2, we give an algorithm that is not δ -correct and has sample complexity lower than Theorem 2 for a specific class of instances.

Example 13 (A non- δ -correct algorithm). *\mathcal{A} is an algorithm for ranking 3 items. It views each pair of items as a coin, and calls KL-LUCB [25] to find the pair (i, j) with the largest $p_{i,j}$ -value. Then, it claims that i is the most preferred item and j is the worst. Obviously, \mathcal{A} is not δ -correct for ranking 3 items. However, for an instance with $p_{r_1, r_2} = 1/2 + \Delta$, $p_{r_1, r_3} = 1 - \Delta$, and $p_{r_2, r_3} = 1 - 2\Delta$, where $r_1 \succ r_2 \succ r_3$ is the unknown true ranking and $\Delta \in (0, 1/6)$ is unknown, with probability at least $1 - \delta$, algorithm \mathcal{A} finds its true ranking by using $O(\Delta^{-1} \log(\Delta^{-1} \delta^{-1}))$ comparisons.*

To see this upper bound, we first define some notations. For $p, q \in [0, 1]$, the KL-Divergence [9] between them is defined as $d(q, p) := D_{KL}(q||p) = q \log \frac{q}{p} + (1 - q) \log \frac{1-q}{1-p}$. The Chernoff-Information [25] between them is defined as $d^*(q, p) := d(z^*, p) = d(z^*, q)$, where z^* is the unique z such that $d(z, p) = d(z, q)$. According to [25, Theorem 3], the algorithm KL-LUCB distinguishes two coins (Bernoulli arms) with mean rewards λ and μ by taking $O(\frac{1}{d^*(\lambda, \mu)} \log \frac{1}{\delta d^*(\lambda, \mu)})$ samples. In this instance, we observe that for a constant $c > 1$, $d(c\Delta, \Delta) = \Theta(\Delta)$. Thus, we have $d^*(1 - 2\Delta, 1 - \Delta) = d^*(2\Delta, \Delta) = \Theta(\Delta)$. Hence, KL-LUCB distinguishes p_{r_1, r_3} and p_{r_2, r_3} by $O(\Delta^{-1} \log(\delta^{-1} \Delta^{-1}))$ comparisons. Since the gap between p_{r_1, r_2} and p_{r_1, r_3} is even larger, they can also be distinguished by the above number of comparisons. This shows the upper bound, which suggests that the Δ_i^{-2} term is not necessary for non- δ -correct algorithms.

We note that \mathcal{A} does not need any information of this instance a priori to run. Although it is not δ -correct, it can solve this class of instances with sample complexity lower than Theorem 2. However, in general, this algorithm may be of no sense as it only works for a restricted class of instances. This is the reason why we want to bound the sample complexity of δ -correct algorithms but not that of arbitrary ones, as there may always exist non- δ -correct algorithms that have extremely good performance on some restricted class of instances.

A.2 An instance where Eq. (2) does not hold as a lower bound

When δ is a positive constant and $\max_{i,j \in [n]} \tilde{\Delta}_i / \tilde{\Delta}_j \succ \sqrt{n}$, the lower bound given in Eq. (2) may not hold. In this subsection, we give an example such that Eq. (2) does not hold as a lower bound.

Example 14 (An example that Eq. (2) does not hold as a lower bound). *Assume that $r_1 \succ r_2 \succ \dots \succ r_n$ is the unknown true ranking. Suppose $\delta = 1/4$, $\Delta_{r_1, r_2} = n^{-10}$ and $\Delta_{r_i, r_j} = 0.01$ for all $\{r_i, r_j\} \neq \{r_1, r_2\}$. For this instance, there is a $(1/4)$ -correct algorithm that finds its true ranking with confidence $3/4$ by $O(n^{20} \log \log n + n^2 \log n)$ comparisons, which is lower than Eq. (2): $\Omega(n^{20} \log n + n \log n)$. This implies that Eq. (2) does not hold as a lower bound in this case.*

To see the upper bound, we can view each pair as a coin (aka Bernoulli arms), and then use lil'UCB [20] to find the pair with the least gap (i.e., $\Delta_{i,j}$) with confidence $11/12$. According to [20], this step takes $O(n)$ comparisons. Then, we rank the pair with the smallest gap with $11/12$ confidence. This step takes $O(\Delta_{r_1, r_2}^{-2} \log \log \Delta_{r_1, r_2}^{-1}) = O(n^{20} \log \log n)$ comparisons. Finally, we rank all other pairs with $1 - \frac{1}{12n^2}$ confidence for each, and this step takes $O(n^2 \log n)$ comparisons. After ranking all pairs of items, the true ranking is found, and thus, the total sample complexity is $O(n^{20} \log \log n + n^2 \log n)$.

For this instance, the lower bound in Eq. (2) is $\Omega(n^{20} \log n + n \log n)$, higher than the upper bound. Thus, when the given condition does not hold, the lower bound in Eq. (2) may not hold. However, there is at most a log gap, and the lower bound in Eq. (1) does not need this condition.

B Proofs

B.1 Proof of Theorem 2

Theorem 2 (Lower bound for pairwise ranking). *Given $\delta \in (0, 1/12)$ and an instance \mathcal{I} with n items, then the number of comparisons used by a δ -correct algorithm \mathcal{A} with no prior knowledge about the gaps of \mathcal{I} is lower bounded by*

$$\Omega\left(\sum_{i \in [n]} [\tilde{\Delta}_i^{-2}(\log \log \tilde{\Delta}_i^{-1} + \log(1/\delta))] + \min\left\{\sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(1/x_i) : \sum_{i \in [n]} x_i \leq 1\right\}\right). \quad (1)$$

If $\delta \preceq 1/\text{poly}(n)^9$, or $\max_{i,j \in [n]} \{\tilde{\Delta}_i/\tilde{\Delta}_j\} \preceq n^{1/2-p}$ for some constant $p > 0$, then the lower bound becomes

$$\Omega\left(\sum_{i \in [n]} \tilde{\Delta}_i^{-2}(\log \log \tilde{\Delta}_i^{-1} + \log(n/\delta))\right). \quad (2)$$

Proof. **Step 1** is to prove the lower bound for ranking two items, which is stated in Lemma 15. In the proof of Lemma 15, we will make use of the results in [15, 20, 27]. The proof can be found in Section B.12

Lemma 15 (Lower bound for ranking two items). *Let $\delta \in (0, 1/4)$ and δ -correct algorithm \mathcal{A}_2 be given. Let $T_{\mathcal{A}_2}(\Delta_{i,j})$ be the number of comparisons conducted by \mathcal{A}_2 under the $\Delta_{i,j}$ -values. To rank i and j with error probability no more than δ , there is a universal constant $c_{lb2} > 0$ such that*

$$\limsup_{\Delta_{i,j} \rightarrow 0} \frac{\mathbb{E}[T_{\mathcal{A}_2}(\Delta_{i,j})]}{\Delta_{i,j}^{-2}(\log \log \Delta_{i,j}^{-2} + \log \delta^{-1})} \geq c_{lb2}. \quad (3)$$

Step 2 is to define problems \mathcal{P}_1 and \mathcal{P}_2 . Let (r_1, r_2, \dots, r_n) be a given permutations of $[n]$ and assume that $q_1 \succ q_2 \succ \dots \succ q_n$ is the unknown true ranking. Assume that n is odd (when n is even, we can prove the same results similarly), and say $n = 2m + 1$. A pair (r_i, r_j) is said to be *significant* if there exists an k in $[m]$ such that $\{r_i, r_j\} = \{r_{2k-1}, r_{2k}\}$, and *insignificant* otherwise.

Define a set $\Pi := \{0, 1\}^m$. For any $\vec{\pi} = (\pi_1, \pi_2, \dots, \pi_m) \in \Pi$, define a corresponding hypothesis $\mathcal{H}_{\vec{\pi}}$ that claims: (i) the true ranking of $[n]$ is $s_1 \succ s_2 \succ \dots \succ s_n$; (ii) $s_n = r_n$; (iii) for any $k \in [m]$, $(s_{2k-1}, s_{2k}) = (r_{2k-1}, r_{2k})$ if $\pi_k = 1$, and $(s_{2k-1}, s_{2k}) = (r_{2k}, r_{2k-1})$ otherwise; (iv) for any insignificant pair (r_i, r_j) , the probability that r_i wins a comparison over the pair (r_i, r_j) is $p_{r_i, r_j}^{\vec{\pi}} = p_{r_i, r_j}$; (v) For any $k \in [m]$ and the corresponding significant pair (r_{2k-1}, r_{2k}) , the probability that r_{2k-1} wins a comparison over the pair (r_{2k-1}, r_{2k}) is $p_{r_{2k-1}, r_{2k}}^{\vec{\pi}} = 1/2 + \Delta_{r_{2k-1}, r_{2k}}$ if $\pi_k = 1$, and is $(1/2 - \Delta_{r_{2k-1}, r_{2k}})$ otherwise. In other words, $\mathcal{H}_{\vec{\pi}}$ claims a true ranking that is almost the same as $r_1 \succ r_2 \succ \dots \succ r_n$ but the positions of (r_{2k-1}, r_{2k}) are interchanged for all $k \in [m]$ such that $\pi_k = 0$. E.g., for $n = 3$ and $\vec{\pi} = (0)$, $\mathcal{H}_{\vec{\pi}}$ claims that the true ranking is $r_2 \succ r_1 \succ r_3$, $p_{r_1, r_2}^{\vec{\pi}} = 1/2 - \Delta_{r_1, r_2}$, $p_{r_1, r_3}^{\vec{\pi}} = p_{r_1, r_3}$, and $p_{r_2, r_3}^{\vec{\pi}} = p_{r_2, r_3}$.

We further assume that there is a $\vec{\pi}_0 \in \Pi$ such that $\mathcal{H}_{\vec{\pi}_0}$ is true, and each $\vec{\pi} \in \Pi$ has the same prior probability to be $\vec{\pi}_0$.

Problem \mathcal{P}_1 . Knowing the fact that there exists a $\vec{\pi}_0 \in \Pi$ such that $\mathcal{H}_{\vec{\pi}_0}$ is true, we want to find $\vec{\pi}_0$ with confidence $1 - \delta$, and use as few comparisons as possible.

Next, we start defining problem \mathcal{P}_2 . An instance of \mathcal{P}_2 involves $\binom{n}{2}$ coins, and each is indexed by an element of $\{(i, j) : i, j \in [n] \wedge i < j\}$. We use $C_{i,j}$ to denote the coin indexed by (i, j) . For each coin $C_{i,j}$, each toss of it gives a head with probability $\mu_{i,j}$, and gives a tail with probability $1 - \mu_{i,j}$. We name $\mu_{i,j}$ as the *head probability* of coin $C_{i,j}$. We assume that the outcomes of tosses are independent across coins and time. Similar to the items, coin $C_{i,j}$ is said to be *significant* if there is a k such that $(i, j) = (2k - 1, 2k)$, and is *insignificant* otherwise. We assume that for all insignificant coins $C_{i,j}$, $\mu_{i,j} = p_{r_i, r_j}$, and for all significant coins $C_{2k-1, 2k}$, $\mu_{2k-1, 2k} = 1/2 + \Delta_{r_{2k-1}, r_{2k}}$ or $1/2 - \Delta_{r_{2k-1}, r_{2k}}$, either has a prior probability $1/2$ to be true.

Problem \mathcal{P}_2 . With probability $\geq 1 - \delta$, we want to find whether $\mu_{2k-1, 2k} > 1/2$ for all $k \in [m]$.

⁹*poly*(n) means a polynomial function of n , and $\delta \preceq 1/\text{poly}(n)$ means $\delta \preceq n^{-p}$ for some constant $p > 0$.

Step 3 is to show the following lemma, which states that \mathcal{P}_2 can be reduced to \mathcal{P}_1 , and \mathcal{P}_1 can be reduced to exact ranking. Its proof can be found in Section B.13.

Lemma 16 (Reductions). *With the above definitions, (i) if the true ranking of $[n]$ is found, with no more comparisons, one can get the solution of \mathcal{P}_1 , and (ii) if an algorithm solves \mathcal{P}_1 with N expected number of comparison, there is another algorithm that solves \mathcal{P}_2 with N expected number of tosses.*

Step 4 is to prove the following lemma regarding the lower bound of problem \mathcal{P}_2 . Its proof can be found in Section B.14

Lemma 17. *For $\delta \in (0, 1/12)$, the expected number of tosses needed for solving \mathcal{P}_2 is at least*

$$\Omega\left(\sum_{k \in [m]} \Delta_{q_{2k-1}, q_{2k}}^{-2} \cdot \log \log \Delta_{q_{2k-1}, q_{2k}}^{-1} + \min\left\{\sum_{k \in [m]} \Delta_{q_{2k-1}, q_{2k}}^{-2} \cdot \log(\delta_k^{-1}) : \sum_{k \in [m]} \delta_k \leq 2\delta\right\}\right). \quad (4)$$

Step 5 is to prove the lower bound given in Eq. (1). Lemma 16 proves that we can reduce \mathcal{P}_2 to \mathcal{P}_1 and reduce \mathcal{P}_1 to exact ranking. Lemma 17 states a lower bound on \mathcal{P}_2 . Thus, by Lemmas 16 and 17, we have that the sample complexity of exact ranking is lower bounded by (4).

We can construct a similar problem to \mathcal{P}_2 , and by the similar steps as in the proof of Lemma 17, we have that the sample complexity of exact ranking is also lower bounded by

$$\Omega\left(\sum_{k \in [m]} \Delta_{q_{2k}, q_{2k+1}}^{-2} \log \log \Delta_{q_{2k}, q_{2k+1}}^{-1} + \min\left\{\sum_{k \in [m]} \Delta_{q_{2k}, q_{2k+1}}^{-2} \log(1/\delta_k) : \sum_{k \in [m]} \delta_k \leq 2\delta\right\}\right). \quad (5)$$

We recall that $q_1 \succ q_2 \succ \dots \succ q_n$ is the true ranking. Since for any $i \in [n]$, $\tilde{\Delta}_{q_i} = \Delta_{q_i, q_{i-1}} \wedge \Delta_{q_i, q_{i+1}}$, we have

$$\begin{aligned} \mathbb{E}N_{\mathcal{A}}(\mathcal{I}) &\succeq \sum_{k \in [m]} \Delta_{q_{2k-1}, q_{2k}}^{-2} \log \log \Delta_{q_{2k-1}, q_{2k}}^{-1} + \sum_{k \in [m]} \Delta_{q_{2k}, q_{2k+1}}^{-2} \log \log \Delta_{q_{2k}, q_{2k+1}}^{-1} \\ &\geq \sum_{k \in [m]} \max\{\Delta_{q_{2k-1}, q_{2k}}^{-2} \log \log \Delta_{q_{2k-1}, q_{2k}}^{-1}, \Delta_{q_{2k}, q_{2k+1}}^{-2} \log \log \Delta_{q_{2k}, q_{2k+1}}^{-1}\} \\ &= \sum_{k \in [m]} \tilde{\Delta}_{q_{2k}}^{-2} \log \log \tilde{\Delta}_{q_{2k}}^{-1} \\ &\stackrel{(a)}{\geq} \frac{1}{3} \sum_{i=1}^n \tilde{\Delta}_i^{-2} \log \log \tilde{\Delta}_i^{-1}, \end{aligned} \quad (6)$$

where (a) holds because for any $k \in [m]$, $\tilde{\Delta}_{q_{2k+1}} = \Delta_{q_{2k}, q_{2k+1}} \wedge \Delta_{q_{2k+1}, q_{2k+2}} \geq \tilde{\Delta}_{q_{2k}} \wedge \tilde{\Delta}_{q_{2k+2}}$.

We also have

$$\begin{aligned} &\min\left\{\sum_{k \in [m]} \Delta_{q_{2k-1}, q_{2k}}^{-2} \log \delta_k^{-1} : \sum_{k \in [m]} \delta_k \leq 2\delta\right\} + \min\left\{\sum_{k \in [m]} \Delta_{q_{2k}, q_{2k+1}}^{-2} \log \delta_k^{-1} : \sum_{k \in [m]} \delta_k \leq 2\delta\right\} \\ &= \min\left\{\sum_{k \in [m]} [\Delta_{q_{2k-1}, q_{2k}}^{-2} \log(1/\delta_k) + \Delta_{q_{2k}, q_{2k+1}}^{-2} \log(1/\delta'_k)] : \sum_{k \in [m]} \delta_k \leq 2\delta, \sum_{k \in [m]} \delta'_k \leq 2\delta\right\} \\ &\geq \min\left\{\sum_{k \in [m]} [\Delta_{q_{2k-1}, q_{2k}}^{-2} + \Delta_{q_{2k}, q_{2k+1}}^{-2}] \log \frac{1}{\delta \sqrt{\delta'}} : \sum_{k \in [m]} \delta_k \leq 2\delta, \sum_{k \in [m]} \delta'_k \leq 2\delta\right\} \\ &\geq \min\left\{\sum_{k \in [m]} \tilde{\Delta}_{q_{2k}}^{-2} \log \frac{1}{\delta_k \sqrt{\delta'_k}} : \sum_{k \in [m]} \delta_k \leq 2\delta, \sum_{k \in [m]} \delta'_k \leq 2\delta\right\} \\ &\geq \min\left\{\sum_{k \in [m]} \tilde{\Delta}_{q_{2k}}^{-2} \log \frac{1}{\delta_k \sqrt{\delta'_k}} : \sum_{k \in [m]} \delta_k \vee \delta'_k \leq 4\delta\right\} \\ &\geq \min\left\{\sum_{k \in [m]} \tilde{\Delta}_{q_{2k}}^{-2} \log(1/\delta_k) : \sum_{k \in [m]} \delta_k \leq 4\delta\right\} \\ &\geq \min\left\{\frac{1}{3} \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(1/x_i) : \sum_{i \in [n]} x_i \leq 12\delta\right\}. \end{aligned} \quad (7)$$

By (7), first, we obtain that, for all $\delta \in (1, 1/12)$,

$$\mathbb{E}N_{\mathcal{A}}(\mathcal{I}) \succeq \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(1/\delta). \quad (8)$$

Also, since $\delta < 1/12$, we obtain the lower bound

$$\begin{aligned} \mathbb{E}N_{\mathcal{A}}(\mathcal{I}) &\succeq \min\left\{ \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(1/x_i) : \sum_{i \in [n]} x_i \leq 12\delta \right\} \\ &\geq \min\left\{ \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(1/x_i) : \sum_{i \in [n]} x_i \leq 1 \right\}. \end{aligned} \quad (9)$$

The lower bound in Eq. (1) follows from summing up Equations (6), (8), and (9). This prove the lower bound in Eq. (2).

Step 6 is to deduce the lower bound in Eq. (2) from Eq. (1).

Case 1. We consider the cases where $\delta \preceq 1/\text{poly}(n)$. We observe that, when $\delta \preceq 1/\text{poly}(n)$, $\log(1/\delta) \succeq \log n$. Thus, in Eq. (9), setting all $x_i = 1/n$, we have

$$\min\left\{ \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(1/x_i) : \sum_{i \in [n]} x_i \leq 1 \right\} \leq \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log n \preceq \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(1/\delta).$$

This means that the term $\min\{\dots\}$ is dominated by the term $\sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(1/\delta)$. We also have $\sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(1/\delta) \simeq \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(n/\delta)$ since $\log \delta^{-1} \succeq \log n$. Thus,

$$\sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(1/\delta) + \min\left\{ \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(1/x_i) : \sum_{i \in [n]} x_i \leq 1 \right\} \simeq \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(n/\delta),$$

which implies that when $\delta = 1/\text{poly}(n)$, the lower bound in (2) holds.

Case 2. We consider the case where $\max_{i,j \in [n]} \{\tilde{\Delta}_i/\tilde{\Delta}_j\} \leq c \cdot n^{1/2-p}$ for some constants $c, p > 0$. When this condition holds, for any x_1, x_2, \dots, x_n with $\sum_{i \in [n]} x_i \leq 1$, we have

$$\begin{aligned} \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(1/x_i) &= \sum_{j \in [n]} \tilde{\Delta}_j^{-2} \sum_{i \in [n]} \frac{\tilde{\Delta}_i^{-2}}{\sum_{j \in [n]} \tilde{\Delta}_j^{-2}} \cdot \log(1/x_i) \\ &\stackrel{(a)}{\geq} \sum_{j \in [n]} \tilde{\Delta}_j^{-2} \cdot \log \frac{1}{\sum_{i \in [n]} x_i \cdot \frac{\tilde{\Delta}_i^{-2}}{\sum_{j \in [n]} \tilde{\Delta}_j^{-2}}} \\ &\geq \sum_{j \in [n]} \tilde{\Delta}_j^{-2} \cdot \log \frac{1}{\sum_{i \in [n]} x_i \frac{1}{\sum_{j \in [n]} (c \cdot n^{-1/2+p})^2}} \\ &\stackrel{(b)}{\geq} \sum_{j \in [n]} \tilde{\Delta}_j^{-2} \log \left[\sum_{i \in [n]} (c \cdot n^{-1/2+p})^2 \right] \\ &\geq \sum_{j \in [n]} \tilde{\Delta}_j^{-2} \log(c^2 n^{2p}) \\ &\succeq \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log n, \end{aligned}$$

where (a) is due to the convexity of the functions $(\log(1/x_i), i \in [n])$, and (b) is due to $\sum_{k \in [n]} \delta_k \leq 1$. Thus, in this case,

$$\begin{aligned} \mathbb{E}N_{\mathcal{A}}(\mathcal{I}) &\succeq \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log \log \Delta_i^{-1} + \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log n + \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log \delta^{-1} \\ &= \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log \log \Delta_i^{-1} + \sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(n/\delta), \end{aligned}$$

which is the lower bound in (2). This completes the proof of (2) and Theorem 2. \square

B.2 Proof of Theorem 3

Theorem 3. [Lower bound for the MNL model] Let $\delta \in (0, 1/12)$ and given a δ -correct algorithm \mathcal{A} with the knowledge that the input instances satisfy the MNL model, let $N_{\mathcal{A}}$ be the number of comparisons conducted by \mathcal{A} , then $\mathbb{E}[N_{\mathcal{A}}]$ is lower bounded by Eq. (1) with a different hidden constant factor. When $\delta \preceq 1/\text{poly}(n)$ or $\max_{i,j \in [n]} \{\Delta_i/\Delta_j\} \preceq n^{1/2-p}$ for some constant $p > 0$, the sample complexity is lower bounded by Eq. (2) with a different hidden constant factor.

Proof. We prove this theorem by Lemmas 4, 5 and 6, which could be of independent interest. The proofs of these three lemmas can be found in Sections B.3, B.4, and B.5

Suppose that there are two coins with unknown *head probabilities* (the probability that a toss produces a head) λ and μ , respectively, and we want to find the more biased one (i.e., the one with the larger head probability). Lemma 4 states a lower bound on the number of heads or tails generated for finding the more biased coin, which works even if λ and μ go to 0. This is in contrast to the lower bounds on the number of tosses given by previous works [20, 23, 27], which go to infinity as λ and μ go to 0.

Lemma 4 (Lower bound on number of heads). Let $\lambda + \mu \leq 1$, $\Delta := |\lambda/(\lambda + \mu) - 1/2|$, and $\delta \in (0, 1/2)$ be given. To find the more biased coin with probability $1 - \delta$, any δ -correct algorithm for this problem produces $\Omega(\Delta^{-2}(\log \log \Delta^{-1} + \log \delta^{-1}))$ heads in expectation.

Now we consider n coins C_1, C_2, \dots, C_n with mean rewards $\mu_1, \mu_2, \dots, \mu_n$, respectively, where for any $i \in [n]$, $\theta_i/\mu_i = c$ for some constant $c > 0$. Define the gaps of coins $\Delta_{i,j}^c := |\mu_i/(\mu_i + \mu_j) - 1/2|$, and $\Delta_i^c := \min_{j \neq i} \Delta_{i,j}^c$. We can check that for all i and j , $\Delta_{i,j}^c = \Delta_{j,i}^c$, and $\Delta_i = \tilde{\Delta}_i = \Delta_i^c$.

Lemma 5 (Lower bound for arranging coins). For $\delta < 1/12$, to arrange these coins in ascending order of head probabilities, the number of heads generated by any δ -correct algorithm is lower bounded by Eq. (1) with a (possibly) different hidden constant factor.

The next lemma shows that any algorithm solves a ranking problem under the MNL model can be transformed to solve the pure exploration multi-armed bandit (PEMAB) problem with Bernoulli rewards. Previous works [1, 17, 18] have shown that certain types of pairwise ranking problems (e.g., Borda-Score ranking) can also be transformed to PEMAB problems. But in this paper, we make a *reverse connection* that bridges these two classes of problems, which may be of independent interest.

Lemma 6 (Reducing PEMAB problems to ranking). If there is a δ -correct algorithm that correctly ranks $[n]$ with probability $1 - \delta$ by M expected number of comparisons, then we can construct another δ -correct algorithm that correctly arranges the coins C_1, C_2, \dots, C_n in the order of ascending head probabilities with probability $1 - \delta$ and produces M heads in expectation.

Combining Lemmas 5 and 6, we have that $\mathbb{E}[N_{\mathcal{A}}]$ is lower bounded by Eq. (1) with a different hidden constant factor. Then, by the same steps as the Step 6 of the proof of Theorem 2, we have that when $\delta \preceq 1/\text{poly}(n)$ or $\max_{i,j \in [n]} \{\Delta_i/\Delta_j\} \preceq n^{1/2-p}$ for some constant $p > 0$, $\mathbb{E}[N_{\mathcal{A}}]$ is lower bounded by Eq. (2) with a different hidden constant factor. This completes the proof. We omit the repetition for brevity and note that under the pairwise MNL model, $\Delta_i = \tilde{\Delta}_i$ for any item i , as the pairwise MNL model satisfies the SST condition. \square

B.3 Proof of Lemma 4

Lemma 4 (Lower bound on number of heads). Let $\lambda + \mu \leq 1$, $\Delta := |\lambda/(\lambda + \mu) - 1/2|$, and $\delta \in (0, 1/2)$ be given. To find the more biased coin with probability $1 - \delta$, any δ -correct algorithm for this problem produces $\Omega(\Delta^{-2}(\log \log \Delta^{-1} + \log \delta^{-1}))$ heads in expectation.

Proof. By contradiction, suppose that there is an algorithm \mathcal{A} that does not satisfy the stated lower bound. We will show a contradiction to Lemma 15.

Given a coin with head probability $p = 1/2 + \eta$, where $\eta \in (-1/4, 0) \cup (0, 1/4)$ is unknown, we will use \mathcal{A} to construct an algorithm to recover the value of $\text{sign}(\eta)$, i.e. the sign of η . Choose an $\alpha \in (0, 1)$. We recall that a p -coin denotes a coin such that each toss of it produces a head with probability p , and a tail otherwise.

Now, we construct two i.i.d. sequences of random variables: $\{X^t\}_{t=1}^\infty$ and $\{Y^t\}_{t=1}^\infty$.

Sequence $\{X^t\}_{t=1}^\infty$ is generated as follows: For any $t \in \mathbb{Z}^+$, with probability α , we toss the p -coin, and assign $X^t = 1$ if the toss gives a head, and assign $X^t = 0$ otherwise. With probability $1 - \alpha$, we assign $X^t = 0$.

Sequence $\{Y^t\}_{t=1}^\infty$ is generated as follows: For any $t \in \mathbb{Z}^+$, with probability α , we toss the p -coin, and assign $Y^t = 1$ if the toss gives a tail, and assign $Y^t = 0$ otherwise. With probability $1 - \alpha$, we assign $Y^t = 0$.

As a result, $(X^t, t \in \mathbb{Z}^+)$ are i.i.d. Bernoulli(λ), and $(Y^t, t \in \mathbb{Z}^+)$ are i.i.d. Bernoulli(μ), respectively. Thus, we can view that X^t 's are generated by a λ -coin and Y^t 's are generated by a μ -coin, where $\lambda = \alpha(1/2 + \eta)$ and $\mu = \alpha(1/2 - \eta)$. We check that $|\lambda/(\lambda + \mu) - 1/2| = \eta$.

Next, we use algorithm \mathcal{A} to find the more biased one of $(X^t, t \in \mathbb{Z}^+)$ and $(Y^t, t \in \mathbb{Z}^+)$. If the result is X^t 's, then we decide $\eta > 0$, and if the result is Y^t 's, then we decide $\eta < 0$. According to the assumption, \mathcal{A} finds the results with probability at least $1 - \eta$ and the number of times t such that $X^t = 1$ or $Y^t = 1$ is at most $o(\eta^{-2}(\log \log \eta^{-1} + \log \delta^{-1}))$ in expectation. For each t with $X^t = 1$ or $Y^t = 1$, the p -coin is tossed for at most 4 times in expectation (since $1/4 < p < 3/4$).

Thus, we can determine whether $\eta < 0$ or $\eta > 0$ (equivalent to ranking two items i and j with $p_{i,j} = 1/2 + \eta$) by $o(\eta^{-2}(\log \log \eta^{-1} + \log \delta^{-1}))$ tosses in expectation, contradicting Lemma 15. Thus, such an algorithm \mathcal{A} does not exist. This completes the proof of Lemma 4. \square

B.4 Proof of Lemma 5

Lemma 5 (Lower bound for arranging coins). *For $\delta < 1/12$, to arrange these coins in ascending order of head probabilities, the number of heads generated by any δ -correct algorithm is lower bounded by Eq. (1) with a (possibly) different hidden constant factor.*

Proof. To prove this lemma, we need to show the following lower bound:

$$\Omega\left(\sum_{i \in [n]} [\tilde{\Delta}_i^{-2}(\log \log \tilde{\Delta}_i^{-1} + \log(1/\delta))] + \min\left\{\sum_{i \in [n]} \tilde{\Delta}_i^{-2} \log(1/x_i) : \sum_{i \in [n]} x_i \leq 1\right\}\right).$$

The proof is similar to that of Lemma 17. We assume that the true order of these coins is (q_1, q_2, \dots, q_n) , and $n = 2m + 1$ is odd. When n is even, we can prove the results in similar steps.

To arrange the coins in the ascending order of head probabilities, one at least needs to distinguish the orders of the pairs $(q_1, q_2), (q_3, q_4), \dots, (q_{2m-1}, q_{2m})$. For any k in $[m]$, to order q_{2k-1} and q_{2k} with probability $1 - \delta_k$, by Lemma 4, any δ -correct algorithm generates $\Omega(\Delta_{q_{2k-1}, q_{2k}}^{-2}(\log \log \Delta_{q_{2k-1}, q_{2k}}^{-1} + \log \delta_k^{-1}))$ heads in expectation. Thus, by the same steps as in the proof of Lemma 17, we obtain a lower bound as follows:

$$\Omega\left(\sum_{k \in [m]} \Delta_{q_{2k-1}, q_{2k}}^{-2} \log \log \Delta_{q_{2k-1}, q_{2k}}^{-1} + \min\left\{\sum_{k \in [m]} \Delta_{q_{2k-1}, q_{2k}}^{-2} \log \delta_k^{-1} : \sum_{k \in [m]} \delta_k \leq 2\delta\right\}\right).$$

Also, to find to orders of the pairs $(q_2, q_3), (q_4, q_5), \dots, (q_{2m}, q_{2m+1})$, there is another lower bound shown below:

$$\Omega\left(\sum_{k \in [m]} \Delta_{q_{2k}, q_{2k+1}}^{-2} \log \log \Delta_{q_{2k}, q_{2k+1}}^{-1} + \min\left\{\sum_{k \in [m]} \Delta_{q_{2k}, q_{2k+1}}^{-2} \log \delta_k^{-1} : \sum_{k \in [m]} \delta_k \leq 2\delta\right\}\right).$$

By the same steps as the Step 5 of the proof of Theorem 2, we can get the desired lower bound. We omit the repetition for brevity. This completes the proof. \square

B.5 Proof of Lemma 6

Lemma 6 (Reducing PEMAB problems to ranking). *If there is a δ -correct algorithm that correctly ranks $[n]$ with probability $1 - \delta$ by M expected number of comparisons, then we can construct another δ -correct algorithm that correctly arranges the coins C_1, C_2, \dots, C_n in the order of ascending head probabilities with probability $1 - \delta$ and produces M heads in expectation.*

Proof. To prove this lemma, consider the following procedure \mathcal{A}_c .

Algorithm Procedure \mathcal{A}_c

Input: Two coins C_i and C_j with unknown head probabilities μ_i and μ_j , respectively;

- 1: **repeat**
 - 2: Randomly choose a coin C_w and toss it;
 - 3: Let $s \leftarrow 1$ if the the toss gives a head, and $s \leftarrow 0$ otherwise;
 - 4: **until** $s = 1$
 - 5: **return** C_w ;
-

Claim 18. Procedure \mathcal{A}_c returns coin C_i with probability $\mu_i/(\mu_i + \mu_j)$ and returns C_j otherwise.

Proof of Claim 18. Let T be the number of tosses conducted before \mathcal{A}_c returns, and X be the coin it returns. By using conditional probability, we have that for all $t \geq 1$ and i in $[m]$,

$$\begin{aligned} \mathbb{P}\{T = t, X = C_i\} &= \prod_{\tau=1}^{t-1} \mathbb{P}\{T > \tau \mid T > \tau - 1\} \cdot \mathbb{P}\{T = t, X = C_i \mid T > t - 1\} \\ &= (\mathbb{P}\{T > 1\})^{t-1} \cdot \mathbb{P}\{T = 1, X = C_i\} \\ &= \left(1 - \frac{1}{2}(\mu_i + \mu_j)\right)^{t-1} \cdot \frac{1}{2}\mu_i, \end{aligned}$$

$$\begin{aligned} \mathbb{P}\{X = C_i\} &= \sum_{t=1}^{\infty} \mathbb{P}\{T = t, X = C_i\} \\ &= \sum_{t=1}^{\infty} \left(1 - \frac{1}{2}(\mu_i + \mu_j)\right)^{t-1} \cdot \frac{1}{2}\mu_i = \frac{\mu_i}{\mu_i + \mu_j}, \end{aligned}$$

and the proof of Claim 18 is complete. □

By Claim 18, we see that the probabilities that \mathcal{A}_c return arms are with the same form as the MNL model. For a ranking algorithm \mathcal{A} , we substitute the input with these n arms and use the procedure \mathcal{A}_c to imitate the comparisons. Whenever the algorithm wants a comparison over C_i and C_j , we call procedure \mathcal{A}_c with input C_i and C_j . If \mathcal{A}_c returns C_i , then we tell \mathcal{A} that C_i wins the comparison, and otherwise, tell \mathcal{A}_c that C_j wins the comparison. Since \mathcal{A}_c returns the arms with probabilities with the same form as the MNL model, \mathcal{A} does not notice any abnormal and work as usual.

For each call of \mathcal{A}_c , there is exactly one head generated. Thus, by this modification, \mathcal{A} arranges these $[n]$ coins in the order of ascending head probabilities with confidence $1 - \delta$, and generates M heads in expectation.

This completes the proof of Lemma 6. □

B.6 Proof of Proposition 7

Proposition 7 (Listwise ranking with negligible noises). *If all comparisons are correct w.h.p., to exactly rank n items w.h.p. by using m -wise comparisons, $\Theta(n \log_m n)$ comparisons are needed.*

Proof. Lower Bound. The proof of the lower bound leverages techniques from information theory. Let X, Y be two discrete random variables (i.e., with at most countably infinite choices of values), and Ω_X, Ω_Y be their sample spaces, respectively. We first briefly introduce some terms of information theory. More information about the information theory can be found in standard texts (e.g., [9]).

Define

$$\begin{aligned} p_x &:= \mathbb{P}\{X = x\}, & p_y &:= \mathbb{P}\{Y = y\}, \\ p_{x,y} &:= \mathbb{P}\{X = x, Y = y\}, & p_{x|y} &:= \mathbb{P}\{X = x \mid Y = y\}. \end{aligned}$$

The information entropy of X is defined as

$$H(X) := \sum_{x \in \Omega_X} p_x \log(1/p_x),$$

and the information entropy of Y is defined as

$$H(Y) := \sum_{y \in \Omega_Y} p_y \log(1/p_y).$$

The joint entropy of X and Y is

$$H(X, Y) := \sum_{x \in \Omega_X, y \in \Omega_Y} p_{x,y} \log(1/p_{x,y}).$$

The conditional entropy of X given $Y = y$ is

$$H(X | Y = y) := \sum_{x \in \Omega_X} p_{x|y} \log(1/p_{x|y}),$$

and the conditional entropy of X given Y is

$$H(X | Y) := \sum_{y \in \Omega_Y} p_y H(X | Y = y).$$

The mutual information of X and Y is

$$I(X; Y) := \sum_{x \in \Omega_X, y \in \Omega} p_{x,y} \log \frac{p_{x,y}}{p_x p_y}.$$

Given another discrete random variable Z , the conditional mutual information of X and Y given Z is

$$I(X; Y | Z) := I(X; Y, Z) - I(X; Z).$$

We further have the following facts [9]

$$\begin{aligned} H(X) &\leq \log |\Omega_X|, \\ H(X | Y) &\leq H(X) \leq H(X, Y), \\ H(X, Y) &= H(Y) + H(X | Y) = H(X) + H(Y | X), \\ I(X; Y) &= H(X) - H(X | Y), \\ I(X; Y | Z) &\leq I(X; Y). \end{aligned}$$

Also, if X is determined by Y , then

$$H(X | Y) = 0.$$

With the above introduction of information, we show the following fact that is used in the proof.

Fact 19 (Fano's Inequality [14]). *To recover the value of X from Y with error probability no more than δ , it must hold that*

$$H(X|Y) \leq H(\delta) + \delta \log(|\Omega_X| - 1).$$

The key idea to prove the lower bound is to show that if the expected number of samples conducted is lower than the lower bound, then Fano's Inequality will not be satisfied.

From now on, we assume that all the comparisons are correct and choose $\delta = 1/4$. We reuse some notation and let X be the ranking of the n items. Before any comparison, we have no information about it, and thus, each ranking has the same probability to be the correct one. Since there are $n!$ possible permutations in total, we have that $H(X) = \log(n!) \simeq n \log n$.

Let \mathcal{A} be an algorithm that adaptively selects the sets to compare and determine whether to stop by past comparison outcomes, let N be the number of comparisons conducted till termination (i.e., stopping time). Let $\vec{S} = (S_1, S_2, \dots, S_N)$ be the sequence of sets that the algorithm compares. Let $\vec{Y} = (Y_1, Y_2, \dots, Y_N)$ be the sequence of comparison outcomes generated by the algorithm. For

any t , S_t is of the form $(S_t[1], S_t[2], \dots, S_t[m])$, which consists of the items compared in the t -th comparison. The value of Y_t is in $\{1, 2, \dots, m\}$, where $Y_t = i$ means the winner of the t -th comparison is $S_t[i]$. We assume that \mathcal{A} is deterministic, i.e., the value of S_t is determined by $(Y_1, Y_2, \dots, Y_{t-1})$ and $(S_1, S_2, \dots, S_{t-1})$, and N is determined by \vec{Y} and \vec{S} . We have

$$\begin{aligned} I(X; \vec{S} | \vec{Y}, N) &\leq I(X; \vec{S} | \vec{Y}) \leq H(\vec{S} | \vec{Y}) \\ &= H(S_1 | \vec{Y}) + H(S_2 | \vec{Y}) + \dots + H(S_N | \vec{Y}) \\ &\leq H(S_1) + H(S_2 | Y_1) + \dots + H(S_N | Y_1, Y_2, \dots, Y_{N-1}) \\ &= 0. \end{aligned} \tag{10}$$

Also, for any t -th comparison, there are at most m different choices of values for Y_t , and thus, $H(Y_t) \leq \log m$. For any $n \in \mathbb{Z}^+$, when $N = n$, the number of choices of values of \vec{Y} is at most m^n , so $H(\vec{Y} | N = n) \leq n \log m$, which implies that

$$H(\vec{Y} | N) = \sum_{n=1}^{\infty} \mathbb{P}\{N = n\} H(\vec{Y} | N = n) \leq \mathbb{E}N \log m. \tag{11}$$

Now, we bound $H(N)$ by $\mathbb{E}N$. Define a random variable R such that $R = 0$ if $N < 2\mathbb{E}N$ and $R = k$ if $2^k \mathbb{E}N \leq N < 2^{k+1} \mathbb{E}N$ for any $k \in \mathbb{Z}^+$. By Markov's Inequality, we have that for $k \in \mathbb{Z}^+$,

$$\mathbb{P}\{R = k\} = \mathbb{P}\{2^k \mathbb{E}N \leq N < 2^{k+1} \mathbb{E}N\} \leq \mathbb{P}\{N \geq 2^k \mathbb{E}N\} \leq 2^{-k}, \tag{12}$$

Use p_k to denote $\mathbb{P}\{R = k\}$. By analyzing the function $p \log(1/p)$, $p \in [0, 1]$, it holds that

$$H(R) = p_0 \log(1/p_0) + \sum_{k=1}^{\infty} p_k \log(1/p_k) \leq 2/e + \sum_{k=2}^{\infty} 2^{-k} \log(2^k) \leq 2/e + (3/2) \log 2. \tag{13}$$

Noting that $H(N | N \in S) \leq \log |S|$ for all sets S , we have

$$\begin{aligned} H(N) &= H(R) + H(N | R) \\ &= H(R) + \mathbb{P}\{N < 2\mathbb{E}N\} H(N | N < 2\mathbb{E}N) \\ &\quad + \sum_{k=1}^{\infty} \mathbb{P}\{2^k \mathbb{E}N \leq N < 2^{k+1} \mathbb{E}N\} H(N | 2^k \mathbb{E}N \leq N < 2^{k+1} \mathbb{E}N) \\ &\stackrel{(a)}{\leq} 2/e + (3/2) \log 2 + \log(2\mathbb{E}N) + \sum_{i=1}^{\infty} 2^{-i} \log(2^i \mathbb{E}N) \\ &\leq 2/e + \log(24\mathbb{E}^2 N), \end{aligned} \tag{14}$$

where (a) is due to (12) and (13).

By (10) (11) (14), we have

$$\begin{aligned} H(X | N, \vec{Y}, \vec{A}) &= H(X) - I(X; N, \vec{Y}, \vec{A}) \\ &\geq H(X) - H(N, \vec{Y}, \vec{A}) \\ &= H(X) - (H(N) + H(\vec{Y} | N) + H(\vec{A} | N, \vec{Y})) \\ &\geq \log(n!) - (2/e + \log(24\mathbb{E}^2 N) + \mathbb{E}N \log m + 0). \end{aligned} \tag{15}$$

By Fano's Inequality, to recover X with probability at least $1/4$, it must hold that

$$H(X | N, \vec{Y}, \vec{A}) \leq H(1/4) + (1/4) \log(n! - 1),$$

which, along with (15) and $\log(n!) = \Theta(n \log n)$, implies that

$$\mathbb{E}N = \Omega(n \log_m n).$$

Algorithm 6 ListwiseMerge(A_1, A_2, \dots, A_m, m)

```
1:  $Ans \leftarrow$  an empty list to store the result;
2: For all  $i$  in  $[m]$ , Let  $I_i \leftarrow 1$  be the index of  $A_i$ ;
3: while  $\exists i \in [m], I_i \leq |A_i|$  do
4:    $B \leftarrow \{A_i[I_i] : I_i \leq |A_i|\}$ ;
5:   Conduct a listwise comparison over  $B$ , and let  $A_j[I_j]$  be the winner;
6:   Push  $A_j[I_j]$  to the end of  $Ans$ ;  $I_j \leftarrow I_j + 1$ ;
7: end while
8: return  $Ans$ 
```

Algorithm 7 ListwiseMergeSort(S, m) (LWMS)

```
1: if  $|S| = 1$  then
2:   return  $S$ ; # No need to do anything
3: end if
4: Divide  $S$  into  $m$  sets  $A_1, A_2, \dots, A_m$  such that  $|A_i| \leq \lceil |S|/m \rceil$  for all  $i \in [m]$ ;
5: for  $i \in [m]$  do
6:    $A_i \leftarrow$  ListwiseMergeSort( $A_i, m$ )
7: end for
8: return ListwiseMerge( $A_1, A_2, \dots, A_m, m$ );
```

For randomized algorithms, its sample complexity is no less than that of the fastest deterministic algorithm, and thus, satisfies the same lower bound. This proves the lower bound.

Upper Bound. To see the upper bound, consider the following ListwiseMergeSort (LWMS) algorithm, which is presented in Algorithm 7. LWMS is similar to the binary merge-sort. Algorithm 6 ListwiseMerge is the subroutine of LWMS, which merges m sorted lists of items.

Lemma 20 (Theoretical upper bound of LWMS). *Algorithm LWMS correctly ranks n items with high probability using $O(n \log_m n)$ comparisons.*

Proof. We use $T_s(x)$ to denote the number of comparisons needed to rank (sort) x items, and use $T_m(x)$ to denote the number of comparisons needed to merge m sorted lists with x items in total. In the algorithm ListwiseMerge, since after each comparison, a new item is added to the result Ans , we have that $T_m(x) \leq x$. Also, we have that $T_s(1) = 0$, and for all $t \geq 1$, $T_s(m^t) = mT_s(m^{t-1}) + T_m(m^t)$. It then follows that $T_s(m^t) \leq tm^t$, which implies $T_s(n) = O(n \log_m n)$. This completes the proof. \square

This completes the proof of Proposition 7. \square

B.7 Proof of Theorem 8

Theorem 8. *Under the MNL model, given n items with preference scores $\theta_1, \theta_2, \dots, \theta_n$ and $\Delta_{i,j} := |\theta_i/(\theta_i + \theta_j) - 1/2|$, $\tilde{\Delta}_i = \Delta_i := \min_{j \neq i} \Delta_{i,j}$, to correctly rank these n items with probability $1 - \delta$, even with m -wise comparisons for all $m \in \{2, 3, \dots, n\}$, the lower bound is the same as the pairwise ranking (i.e., Theorem 3) with (possibly) different hidden constant factors.*

Proof. Let n coins C_1, C_2, \dots, C_n with unknown head probabilities $\mu_1, \mu_2, \dots, \mu_n$ be given, where μ_i/θ_i is a fixed constant for all $i \in [n]$. We only need to show that the reduction from PEMAB problems to exact ranking stated in Lemma 6 still holds for listwise comparisons under the MNL model.

Consider the the following procedure: \square

Claim 21. *Procedure \mathcal{A}'_c returns a coin C_{r_i} with probability $\mu_{r_i}/\sum_{j=1}^m \mu_{r_j}$.*

Algorithm Procedure \mathcal{A}'_c

Input: Totally m coins $C_{r_1}, C_{r_2}, \dots, C_{r_m}$ with unknown head probabilities $\mu_{r_1}, \mu_{r_2}, \dots, \mu_{r_m}$;

- 1: **repeat**
 - 2: Randomly choose a coin C_w and toss it;
 - 3: Let $s \leftarrow 1$ if the toss gives a head, and let $s \leftarrow 0$ otherwise;
 - 4: **until** $s = 1$
 - 5: **return** C_w ;
-

Proof of Claim 21. Let T be the number of tosses conducted before \mathcal{A}'_c returns, and X be the coin \mathcal{A}'_c returns. By using conditional probability, we have that for all $t \geq 1$ and i in $[m]$,

$$\begin{aligned} \mathbb{P}\{T = t, X = C_{r_i}\} &= \prod_{\tau=1}^{t-1} \mathbb{P}\{T > \tau \mid T > \tau - 1\} \cdot \mathbb{P}\{T = t, X = C_{r_i} \mid T > t - 1\} \\ &= (\mathbb{P}\{T > 1\})^{t-1} \mathbb{P}\{T = 1, X = C_{r_i}\} \\ &= \left(1 - \frac{1}{m} \sum_{j=1}^m \mu_{r_j}\right)^{t-1} \cdot \frac{1}{m} \mu_{r_i}, \\ \mathbb{P}\{X = C_{r_i}\} &= \sum_{t=1}^{\infty} \mathbb{P}\{T = t, X = C_{r_i}\} \\ &= \sum_{t=1}^{\infty} \left(1 - \frac{1}{m} \sum_{j=1}^m \mu_{r_j}\right)^{t-1} \cdot \frac{1}{m} \mu_{r_i} = \frac{\mu_{r_i}}{\sum_{j=1}^m \mu_{r_j}}, \end{aligned}$$

and the proof of the claim is complete. \square

The proof of Theorem 8 is complete by Lemma 5 and the same steps as in the proof of Theorem 3, the pairwise lower bound for the MNL model.

B.8 Proof of Lemma 9

Lemma 9 (Theoretical Performance of ATC). *ATC terminates after at most $b^{max} = O(\epsilon^{-2} \log(1/\delta))$ comparisons and returns the more preferred item with probability at least $1/2$. Further, if $\epsilon \leq \Delta_{i,j}$, then ATC returns the more preferred item with probability at least $1 - \delta$.*

Proof. Without loss of generality, we assume $i \succ j$. Since the for loop runs at most $b^{max} = \lceil \frac{1}{2} \epsilon^{-2} \log(2\delta^{-1}) \rceil$ iterations and each iteration performs one comparison, the subroutine returns after at most $O(\epsilon^{-2} \log \delta^{-1})$ comparisons. Since the return condition of items i and j are symmetric and $i \succ j$, by this symmetry, ATC returns j with probability no more than $1/2$.

Now we consider the case where $p_{i,j} \geq 1/2 + \epsilon$, and it remains to prove that ATC returns i with probability at least $1 - \delta$. Define $b^t := \sqrt{\frac{1}{2t} \log \frac{\pi^2 t^2}{3\delta}}$. Let \mathcal{E}_t^{out} be the event that $\hat{p}_i^t \leq p_{i,j} - b^t$, and define $\mathcal{E}^{out} := \bigcup_{t=1}^{\infty} \mathcal{E}_t^{out}$. We have

$$\mathbb{P}\{\mathcal{E}^{out}\} \stackrel{(a)}{\leq} \sum_{t=1}^{\infty} \mathbb{P}\{\mathcal{E}_t^{out}\} \stackrel{(b)}{\leq} \sum_{t=1}^{\infty} \left[\exp\left(-2t(b^t)^2\right) \right] \leq \sum_{t=1}^{\infty} \frac{3\delta}{\pi^2 t^2} \leq \frac{\delta}{2}, \quad (16)$$

where (a) is due to the union bound and (b) is due to the Chernoff-Hoeffding Inequality [19].

Assume that \mathcal{E}^{out} does not happen, and we have that for all t , $\hat{p}_i^t > 1/2 + \epsilon - b^t \geq 1/2 - b^t$. Thus, ATC does not return j during the for loop with probability at least $1 - \delta/2$.

After the for loop, by Chernoff-Hoeffding Inequality and $b^{max} = \lceil \frac{1}{2\epsilon^2} \log \frac{2}{\delta} \rceil$, we have

$$\mathbb{P}\left\{\hat{p}_i^{b^{max}} \leq 1/2\right\} \leq \exp\left\{-2b^{max}(p_{i,j} - 1/2)^2\right\} \leq \exp\left\{-2b^{max}\epsilon^2\right\} \leq \delta/2, \quad (17)$$

which implies that the last line of ATC returns i with probability at least $1 - \delta/2$. This completes the proof of Lemma 9. \square

B.9 Proof of Lemma 10

Lemma 10 (Theoretical performance of ATI). *Let $\delta \in (0, 1)$. ATI returns after $O(\epsilon^{-2} \log(|S|/\delta))$ comparisons and, with probability at least $1 - \delta$, correctly inserts i or returns unsure. Further, if $\epsilon \leq \Delta_i$, it correctly inserts i with probability at least $1 - \delta$.*

Proof. (I) We first prove the sample complexity. We observe that for a constant $\delta_0 \in (0, 1/2)$, a call of $\text{ATC}(i, j, \epsilon, \delta_0)$ returns after at most $O(\epsilon^{-2})$ comparisons by Lemma 9. In ATI, for each iteration, there are at most three calls of ATC and all the calls are with constant confidence. Also, ATI returns after at most $t^{\max} = O(h + \log \delta^{-1})$ iterations, where $h = 1 + \lceil \log_2(1 + |S|) \rceil = O(\log |S|)$. Thus, the number of comparisons is at most $3t^{\max} \cdot O(\epsilon^{-2}) = O(\epsilon^{-2} \log(|S|/\delta))$. This completes the proof sample complexity.

(II) We prove that ATI does not insert i into a wrong place with probability at least $1/2$. A round (or iteration) is said to be *correct* if during this round, all calls of ATC return the more preferred item, and is said to be *incorrect* otherwise. A leaf node u is said to be *correct* if $i \in (u.\text{left}, u.\text{right})$, i.e., i belongs to the corresponding interval of u . A leaf node u is said to be *incorrect* if it is not correct.

For any round t , we define an event \mathcal{E}_{il}^t such that

$$\mathcal{E}_{il}^t := \{X = \text{some incorrect leaf node at the beginning of round } t \text{ and } c_X \geq 1\}. \quad (18)$$

We assume that for some round t , \mathcal{E}_{il}^t happens, which implies that $i \succ u.\text{right}$ or $u.\text{left} \succ i$, i.e., i does not belong to the interval of u . By Lemma 9 the property of ATC, it holds that

$$\begin{aligned} \mathbb{P}\{\text{ATC}(i, u.\text{right}, \epsilon, \delta) = i \mid i \succ u.\text{right}\} &\geq 1/2, \\ \mathbb{P}\{\text{ATC}(i, u.\text{left}, \epsilon, \delta) = u.\text{left} \mid u.\text{left} \succ i\} &\geq 1/2. \end{aligned}$$

which implies that for any round t ,

$$\mathbb{P}\{\text{round } t \text{ is correct} \mid \mathcal{E}_{il}^t\} \geq 1/2. \quad (19)$$

For any t , define

$$\begin{aligned} R_1^t &:= |\{\tau \leq t : \text{round } \tau \text{ is correct, and } \mathcal{E}_{il}^\tau \text{ happens}\}|, \\ W_1^t &:= |\{\tau \leq t : \text{round } \tau \text{ is incorrect, and } \mathcal{E}_{il}^\tau \text{ happens}\}|. \end{aligned}$$

For any incorrect leaf node u and any round t , the counter c_u is increased by one during this round if and only if \mathcal{E}_{il}^t happens and this round is incorrect. Also, for any round t , given \mathcal{E}_{il}^t , if this round is correct, then the counter c_u is decreased by one. Thus, for any incorrect leaf node u , at the end of any round t , the value of c_u is at most

$$c_u(t) \leq 1 + W_1^t - R_1^t.$$

After the for loop, ATI incorrectly inserts i if and only if some incorrect leaf node u is counted for $\frac{5}{16}t^{\max} + 1$ times, i.e., $c_u \geq \frac{5}{16}t^{\max} + 1$, which implies $W_1^{t^{\max}} - R_1^{t^{\max}} \geq \frac{5}{16}t^{\max}$. Thus, by the fact that $W_1^{t^{\max}} + R_1^{t^{\max}} \leq t^{\max}$, and Eq. 19, we obtain

$$\begin{aligned} &\mathbb{P}\left\{W_1^{t^{\max}} - R_1^{t^{\max}} \geq \frac{5}{16}t^{\max}\right\} \\ &= \mathbb{P}\left\{W_1^{t^{\max}} \geq \frac{1}{2}\left(R_1^{t^{\max}} + W_1^{t^{\max}} + \frac{5}{16}t^{\max}\right)\right\} \\ &\stackrel{(a)}{\leq} \sup_{K \leq t^{\max}} \mathbb{P}\left\{\frac{W_1^{t^{\max}}}{W_1^{t^{\max}} + R_1^{t^{\max}}} \geq \frac{1}{2} + \frac{5}{32} \cdot \frac{t^{\max}}{W_1^{t^{\max}} + R_1^{t^{\max}}} \mid W_1^{t^{\max}} + R_1^{t^{\max}} = K\right\} \\ &\stackrel{(b)}{\leq} \sup_{K \leq t^{\max}} \exp\left\{-2K \left(\frac{5t^{\max}}{32K}\right)^2\right\} \\ &= \exp\left\{-2t^{\max} \left(\frac{5t^{\max}}{32t^{\max}}\right)^2\right\} \leq \delta/2, \end{aligned} \quad (20)$$

where (a) is due to $R_1^{t^{max}} + W_1^{t^{max}} \leq t^{max}$, and (b) follows from Chernoff-Hoeffding Inequality. This proves that with probability at least $1 - \delta/2$, i is not inserted into a wrong place by the second last line.

Then, during the for loop, for any $t \leq t^{max}$, by (19) and Chernoff-Hoeffding Inequality, we have that at the end of the t -th round, the probability that X equals to an incorrect leaf node and $c_X > \frac{1}{2}t + \sqrt{\frac{t}{2} \log \frac{\pi^2 t^2}{3\delta}} + 1$ is at most

$$\begin{aligned} \mathbb{P} \left\{ W_1^t - R_1^t \geq \frac{1}{2}t + \sqrt{\frac{t}{2} \log \frac{\pi^2 t^2}{3\delta}} \right\} &\leq \mathbb{P} \left\{ W_1^t \geq \frac{1}{2}t + \sqrt{\frac{t}{2} \log \frac{\pi^2 t^2}{3\delta}} \right\} \\ &\leq \exp \left\{ -\frac{2}{t} \left(\frac{1}{2}t + \sqrt{\frac{t}{2} \log \frac{\pi^2 t^2}{3\delta}} - \frac{t}{2} \right)^2 \right\} \leq \frac{3\delta}{\pi^2 t^2}. \end{aligned}$$

Since

$$\sum_{t=1}^{\infty} \frac{3\delta}{\pi^2 t^2} \leq \delta/2,$$

during the for loop, with probability at least $1 - \delta/2$, ATI does not insert i into a wrong place. This, along with Eq. (20), proves that with probability at least $1 - \delta$, ATI does not insert i into a wrong place. This completes the proof of the first part of Lemma 10.

(III) In this part, we assume $\epsilon \leq \Delta_i$ and we prove the second part of Lemma 10. For any round t , by Lemma 9 and the choice of input parameters of the calls of ATC, this round is correct with probability at least q . Here, we define R as the number of correct rounds before termination, and let W be the number of incorrect rounds before termination.

Let u_0 be the correct node. Define the distance between two nodes u and v as $d(u, v) :=$ the length of the shortest path from u to the v , i.e., the number of edges between u and v . During each correct round, either $d(X, u_0)$ is decreased by one or the value of c_{u_0} is increased by one, i.e., $c_{u_0} - d(X, u_0)$ is increased by one. During each incorrect round, either $d(X, u_0)$ is increased by one or the value of c_{u_0} is decreased by one, i.e., $c_{u_0} - d(X, u_0)$ is decreased by one. Since the distance between the start node (i.e., the root node) and u_0 is at most $h - 1$, we always have

$$R - W \leq h - 1 + (c_{u_0} - d(X, u_0)).$$

After the for loop, if $c_{u_0} \geq \frac{5}{16}t^{max} + 1$, then ATI correctly inserts i . Thus, if $R - W \geq h + \frac{5}{16}t^{max}$, then ATI correctly inserts i .

Assume that ATI does not return during the for loop, and then, we have $R + W = t^{max}$. For all t , round t is correct with probability at least q by Lemma 9 and the choices of input parameters of the calls of ATC, hence, by $t^{max} \geq \max\{4h, \frac{512}{25} \log \frac{2}{\delta}\}$ and $q = 15/16$, we have

$$\begin{aligned} \mathbb{P} \left\{ R - W < h + \frac{5}{16}t^{max} \right\} &\stackrel{(a)}{\leq} \mathbb{P} \left\{ R - W < \left(\frac{1}{4} + \frac{5}{16} \right) t^{max} \right\} \\ &= \mathbb{P} \left\{ R - (t^{max} - R) < \left(\frac{1}{4} + \frac{5}{16} \right) t^{max} \right\} \\ &= \mathbb{P} \left\{ R < \frac{25}{32}t^{max} \right\} \\ &\stackrel{(b)}{\leq} \exp \left\{ -2t^{max} \left(q - \frac{25}{32} \right)^2 \right\} \leq \frac{\delta}{2}, \end{aligned}$$

where (a) is due to $t^{max} \geq 4h$ and (b) follows from Chernoff-Hoeffding Inequality.

In conclusion, when $\epsilon \leq \Delta_i$, if ATI does not return during the for loop, then it will, with probability at least $1 - \delta/2$, insert i into a correct position by the second last line (after the for loop). Also, by part (II), with probability at least $1 - \delta/2$, ATI does not insert i into a wrong position during the for loop. Thus, when $\epsilon \leq \Delta_i$, ATI correctly inserts the input item i with probability at least $1 - \delta$. This proves the second part of Lemma 10, and along with parts (I) and (II), completes the proof. \square

B.10 Proof of Lemma 11

Lemma 11 (Theoretical Performance of IAI). *With probability at least $1 - \delta$, IAI correctly inserts i into S , and conducts at most $O(\Delta_i^{-2}(\log \log \Delta_i^{-1} + \log(|S|/\delta)))$ comparisons.*

Proof. Define events

$$\begin{aligned}\mathcal{E}_1^t &:= \{\epsilon_t > \Delta_i \text{ and IAI does not insert } i \text{ into a wrong position}\}, \\ \mathcal{E}_2^t &:= \{\epsilon_t \leq \Delta_i \text{ and IAI correctly inserts } i\},\end{aligned}$$

and the bad event

$$\mathcal{E}^{bad} := \bigcup_{t=1}^{\infty} (\mathcal{E}_1^t \cup \mathcal{E}_2^t)^c.$$

By the union bound and Lemma 10, we have

$$\mathbb{P}\{\mathcal{E}^{bad}\} \leq \sum_{t=1}^{\infty} \mathbb{P}\left\{(\mathcal{E}_1^t \cup \mathcal{E}_2^t)^c\right\} \leq \sum_{t=1}^{\infty} \delta_t = \sum_{t=1}^{\infty} \frac{6\delta}{\pi^2 t^2} = \delta.$$

In this proof, we assume that \mathcal{E}^{bad} does not happen.

Correctness. We first prove the correctness. By the definition of \mathcal{E}^{bad} , for all t such that $\epsilon_t > \Delta_i$, IAI does not insert i into a wrong position, and when $\epsilon_t \leq \Delta_i$, IAI correctly inserts i . Since $\lim_{t \rightarrow \infty} \epsilon_t = 0$, there is a t^* such that $\epsilon_{t^*} \leq \Delta_i$. Thus, when \mathcal{E}^{bad} does not happen, IAI correctly inserts i . Since \mathcal{E}^{bad} happens with probability at most δ , the correctness follows.

Sample complexity. Second, we prove the sample complexity. Let τ be the integer such that $\epsilon_{\tau} \leq \Delta_i < \epsilon_{\tau-1}$. By the definition of \mathcal{E}^{bad} , when \mathcal{E}^{bad} does not happen, IAI correctly inserts i and returns before the end of the τ -th round.

By $\epsilon_{\tau-1} = 2^{-\tau}$ and $\epsilon_{\tau-1} > \Delta_i$, we have $\tau < \log_2 \Delta_i^{-1}$. For $1 \leq t \leq \tau$, by Lemma 10, the t -th round of IAI conducts at most $O(\epsilon_t^{-2} \log(|S| \cdot \delta_t^{-1}))$ comparisons. Thus, given \mathcal{E}^{bad} does not happen, the number of comparisons conducted by IAI is at most

$$\begin{aligned}O\left(\sum_{t=1}^{\tau} \epsilon_t^{-2} \log(|S|/\delta_t)\right) &\stackrel{(a)}{=} O\left(\sum_{t=1}^{\tau} (2^{t+1})^2 \log(\pi^2 t^2 |S|/(6\delta))\right) \\ &= O\left(\sum_{t=1}^{\tau} 4^t \cdot \log(|S|/\delta)\right) \\ &= O(4^{\tau} \cdot \log(|S|/\delta)) \\ &\stackrel{(b)}{=} O\left(4^{\log_2(1/\Delta_i)} \cdot \log(|S| \cdot \log(1/\Delta_i)/\delta)\right) \\ &= O\left(\Delta_i^{-2} (\log \log \Delta_i^{-1} + \log(|S|/\delta))\right),\end{aligned}$$

where (a) follows from $\epsilon_t = 2^{t+1}$ and $\delta_t = \frac{6\delta}{\pi^2 t^2}$, and (b) is due to $\tau < \log_2(1/\Delta_i)$. This proves the sample complexity.

The proof of Lemma 11 is complete. \square

B.11 Proof of Theorem 12

Theorem 12 (Theoretical Performance of IIR). *With probability at least $1 - \delta$, IIR returns the exact ranking of $[n]$ and conducts at most $O(\sum_{i \in [n]} \Delta_i^{-2}(\log \log \Delta_i^{-1} + \log(n/\delta)))$ comparisons.*

Proof. At iteration t for each $t \in \{2, 3, \dots, n\}$, by Lemma 11, with probability at least $1 - \delta/(n-1)$, the call of IAI correctly inserts $S[t]$ into Ans , and uses at most $O(\Delta_{S[t]}^{-2}(\log \log \Delta_{S[t]}^{-1} + \log(n/\delta)))$ comparisons. The desired sample complexity follows by summing up the upper bounds for $t \in \{2, 3, \dots, n\}$. For correctness, if all calls of IAI are correct (which happens with probability at least $1 - \delta$ by the union bound), then IIR correctly returns the true ranking. This completes the proof. \square

B.12 Proof of Lemma 15

Lemma 15 (Lower bound for ranking two items). *Let $\delta \in (0, 1/4)$ and δ -correct algorithm \mathcal{A}_2 be given. Let $T_{\mathcal{A}_2}(\Delta_{i,j})$ be the number of comparisons conducted by \mathcal{A}_2 under the $\Delta_{i,j}$ -values. To rank i and j with error probability no more than δ , there is a universal constant $c_{lb2} > 0$ such that*

$$\limsup_{\Delta_{i,j} \rightarrow 0} \frac{\mathbb{E}[T_{\mathcal{A}_2}(\Delta_{i,j})]}{\Delta_{i,j}^{-2} (\log \log \Delta_{i,j}^{-2} + \log \delta^{-1})} \geq c_{lb2}. \quad (3)$$

Proof. We will invoke the results for pure exploration multi-armed bandit (PEMAB) problems, and we refer to [27] as a reference for details about PEMAB. Assume that there is an arm a , and whenever it is pulled for the t -th time, it gives an i.i.d. reward Y^t . Further assume that for $t \in \mathbb{Z}^+$, Y^t is a Gaussian random variable with mean η and variance 1. We assume that $|\eta| \leq 1/2$ and $\eta \neq 0$. Let \mathcal{B} be a δ -correct algorithm that has no knowledge of η and is able to tell whether $\eta > 0$ with probability $1 - \delta$ for any non-zero η -value. Let $T_{\mathcal{B}}(\eta)$ be the number of pulls \mathcal{B} uses before termination under the given η -value. The authors of [20, 15] have shown that

$$\limsup_{|\eta| \rightarrow 0} \frac{\mathbb{E}[T_{\mathcal{B}}(\eta)]}{\eta^{-2} \log \log \eta^{-2}} \geq 2 - 4\delta. \quad (21)$$

In this proof, we reduce the problem of distinguishing whether $\eta > 0$ to the problem of ranking two items. For any $t \in \mathbb{Z}^+$, if $0 < \eta < 1/2$, we have

$$\mathbb{P}\{Y^t \geq 0\} = \int_{-\eta}^{\infty} e^{-\frac{x^2}{2}} dx \geq \frac{1}{2} + \eta \cdot e^{-\frac{\eta^2}{2}} \geq \frac{1}{2} + \eta \cdot e^{-1/8},$$

and if $-1/2 < \eta < 0$, we have

$$\mathbb{P}\{Y^t < 0\} = \int_{-\infty}^{-\eta} e^{-\frac{x^2}{2}} dx \geq \frac{1}{2} + |\eta| \cdot e^{-\frac{\eta^2}{2}} \geq \frac{1}{2} + |\eta| \cdot e^{-1/8}.$$

For each t , we let $Z^t = 2 \cdot \mathbb{1}\{Y^t \geq 0\} - 1$. When $\eta > 0$, Z^t is with probability at least $1/2 + \eta \cdot e^{-1/8}$ to be 1, and when $\eta < 0$, it is with probability at least $1/2 + |\eta| \cdot e^{-1/8}$ to be -1 . Thus, we can view that $(Z^t, t \in \mathbb{Z}^+)$ are generated by tossing a coin with $\mathbb{P}\{Y^t \geq 0\}$ head probability, and we have $|\mathbb{P}\{Y^t \geq 0\} - 1/2| \geq e^{1/8}|\eta|$. Assume \mathcal{A}_2 can ranking two items i and j with probability $1 - \delta$ by $T_{\mathcal{A}_2}(\Delta_{i,j})$ expected number of comparisons, then it can find whether $\eta > 0$ by at most $T_{\mathcal{A}_2}(\eta \cdot e^{-1/8})$ expected number of pulls of the arm a . Thus, by (21), we have

$$\limsup_{\Delta_{i,j} \rightarrow 0} \frac{\mathbb{E}[T_{\mathcal{A}_2}(\Delta_{i,j})]}{\Delta_{i,j}^{-2} \log \log \Delta_{i,j}^{-2}} \geq e^{-1/4}(2 - 4\delta). \quad (22)$$

Then, by the previous work [27], we obtain another lower bound on ranking two items, i.e., $\Omega(\Delta_{i,j}^{-2} \log \delta^{-1})$. Summing up this lower bound and (22), we obtain the desired lower bound. This completes the proof. \square

B.13 Proof of Lemma 16

Lemma 16 (Reductions). *With the above definitions, (i) if the true ranking of $[n]$ is found, with no more comparisons, one can get the solution of \mathcal{P}_1 , and (ii) if an algorithm solves \mathcal{P}_1 with N expected number of comparison, there is another algorithm that solves \mathcal{P}_2 with N expected number of tosses.*

Proof. We first prove the reduction from \mathcal{P}_1 to exact ranking. Given an instance of \mathcal{P}_1 , we simply use an exact ranking algorithm to find its true ranking. By the assumptions made in the construction of \mathcal{P}_1 , the comparison probabilities under the correct hypothesis \mathcal{H}_{π^0} is exactly the same as the corresponding ranking instance. Thus, by the found true ranking, we can find the true hypothesis with no more comparisons. This completes the first part of Lemma 16.

Secondly, we prove the reduction from \mathcal{P}_2 to \mathcal{P}_1 . Assume that n is odd and $n = 2m + 1$, and when n is even, we can prove the same results by similar steps. Let \mathcal{B} be an arbitrary δ -correct algorithm for

\mathcal{P}_1 . Let the $\binom{n}{2}$ coins satisfying the restrictions of \mathcal{P}_2 be given. We construct n virtual items indexed by r_1, r_2, \dots, r_n , where (r_1, r_2, \dots, r_n) is a permutation of $[n]$. With these n items, we construct 2^m hypotheses as defined in the construction of Problem \mathcal{P}_1 (i.e., $\mathcal{H}_{\vec{\pi}}, \vec{\pi} \in \{0, 1\}^m$). Then, we send these n items and the hypotheses as the input to algorithm \mathcal{B} . Whenever \mathcal{B} wants a comparison over the pair (r_i, r_j) , we toss the coin $C_{i,j}$. If the toss gives a head, we tell \mathcal{B} that the winner of the comparison is r_i , and if the toss gives a tail, we tell \mathcal{B} that the winner is r_j . Since the values of the head probabilities $\mu_{i,j}$ are lawful for the comparison probabilities of Problem \mathcal{P}_1 , \mathcal{B} does not notice any abnormal and works as usual. Finally, \mathcal{B} terminates and returns a $\vec{\pi} \in \Pi$.

For any $k \in [m]$, if $\pi(k) = 1$, then we return $\mu_{2k-1, 2k} > 1/2$, and otherwise, we return $\mu_{2k-1, 2k} < 1/2$. If \mathcal{B} returns a correct hypothesis for these n virtual items, one can determine whether $\mu_{2k-1, 2k} > 1/2$ for any $k \in [n]$ by no more tosses of coins. Moreover, for any (i, j) , the head probability $\mu_{i,j}$ of problem \mathcal{P}_1 equals to p_{r_i, r_j} , the comparison probability of problem \mathcal{P}_2 . This completes the second part of Lemma 16. The proof is complete. \square

B.14 Proof of Lemma 17

Lemma 17. For $\delta \in (0, 1/12)$, the expected number of tosses needed for solving \mathcal{P}_2 is at least

$$\Omega\left(\sum_{k \in [m]} \Delta_{q_{2k-1}, q_{2k}}^{-2} \cdot \log \log \Delta_{q_{2k-1}, q_{2k}}^{-1} + \min\left\{\sum_{k \in [m]} \Delta_{q_{2k-1}, q_{2k}}^{-2} \cdot \log(\delta_k^{-1}) : \sum_{k \in [m]} \delta_k \leq 2\delta\right\}\right). \quad (4)$$

Proof. In \mathcal{P}_2 , the tosses of the coins are independent across time and coins. Also, whether one coin has head probability larger than $1/2$ is independent of other coins. Thus, \mathcal{P}_2 is simply a problem such that given m coins with head probability not equal to $1/2$, to identify all the coins with head probabilities larger than $1/2$, and the total error probability is no more than δ .

Given a coin with non- $1/2$ head probability, deciding whether the head probability is larger than $1/2$ is equivalent to the problem of ranking two items, as a toss of a coin with head probability η can be viewed as a comparison of items i and j with $p_{i,j} = \eta$. Thus, for coin $C_{2k-1, 2k}$, to find whether $\mu_{2k-1, 2k} > 1/2$ with at most δ_k error probability, the expected number of tosses is at least

$$c \Delta_{q_{2k-1}, q_{2k}}^{-2} (\log \log \Delta_{q_{2k-1}, q_{2k}}^{-1} + \log(1/\delta_k)),$$

where $c > 0$ is a universal constant, and here, we note that $|\mu_{2k-1, 2k} - 1/2| = \Delta_{q_{2k-1}, q_{2k}}$ for any $k \in [m]$ due to the constructions of \mathcal{P}_1 and \mathcal{P}_2 .

Let δ_k be the error probability incurred by determining whether $\mu_{2k-1, 2k} > 1/2$. To solve \mathcal{P}_2 with confidence $1 - \delta$, it is necessary that

$$\prod_{k \in [m]} (1 - \delta_k) \geq 1 - \delta.$$

We also have that for $\delta \in (0, 1/2)$,

$$\begin{aligned} \sum_{k \in [m]} \delta_k &\leq -\sum_{k \in [m]} \log(1 - \delta_k) = -\log \prod_{k \in [m]} (1 - \delta_k) \\ &\leq -\log(1 - \delta) = \log(1 + \delta/(1 - \delta)) \\ &\leq \delta/(1 - \delta) \leq 2\delta. \end{aligned}$$

Thus, the lower bound of \mathcal{P}_2 is at least

$$\Omega\left(\sum_{k \in [m]} \Delta_{r_{2k-1}, r_{2k}}^{-2} \log \log \Delta_{r_{2k-1}, r_{2k}}^{-1} + \min\left\{\sum_{k \in [m]} \Delta_{r_{2k-1}, r_{2k}}^{-2} \log \delta_k^{-1} : \sum_{k \in [m]} \delta_k \leq 2\delta\right\}\right).$$

This completes the proof of Lemma 17. \square