Maxing and Ranking with Few Assumptions

Moein Falahatgar Yi Hao Alon Orlitsky Venkatadheeraj Pichapati Vaishakh Ravindrakumar
University of California, San Deigo
{moein,yih179,alon,dheerajpv7,vaishakhr}@ucsd.edu

Abstract

PAC maximum selection (maxing) and ranking of n elements via random pairwise comparisons have diverse applications and have been studied under many models and assumptions. With just one simple natural assumption: strong stochastic transitivity, we show that maxing can be performed with linearly many comparisons yet ranking requires quadratically many. With no assumptions at all, we show that for the Borda-score metric, maximum selection can be performed with linearly many comparisons and ranking can be performed with $\mathcal{O}(n \log n)$ comparisons.

1 Introduction

1.1 Motivation

Maximum selection (maxing) and sorting using pairwise comparisons are among the most practical and fundamental algorithmic problems in computer science. As is well-known, maxing requires $n-1$ comparisons, while sorting takes $\Theta(n \log n)$ comparisons.

The probabilistic version of this problem, where comparison outcomes are random, is of significant theoretical interest as well, and it too arises in many applications and diverse disciplines. In sports, pairwise games with random outcomes are used to determine the best, or the order, of teams or players. Similarly *Trueskill* [1] matches video gamers to create their ranking. It is also used for a variety of online applications such as to learn consumer preferences with the popular *A/B tests*, in recommender systems [2], for ranking documents from user clickthrough data [3, 4], and more. The popular crowd sourcing website GIFGIF [5] shows how pairwise comparisons can help associate emotions with many animated GIF images. Visitors are presented with two images and asked to select the one that better corresponds to a given emotion. For these reasons, and because of its intrinsic theoretical interest, the problem received a fair amount of attention.

1.2 Terminology and previous results

One of the first studies in the area, [6] assumed n totally-ordered elements, where each comparison errs with the same, known, probability $\alpha < \frac{1}{2}$. It presented a maxing algorithm that uses $\mathcal{O}(\frac{n}{\alpha^2} \log \frac{1}{\delta})$ comparisons to output the maximum with probability $\ge 1-\delta$, and a ranking algorithm that uses $\mathcal{O}(\frac{n}{\alpha^2} \log \frac{n}{\delta})$ comparisons to output the ranking with probability $\ge 1-\delta$.

These results have been and continue to be of great interest. Yet this model has two shortcomings. It assumes that there is only one random comparison probability, α, and that its value is known. In practice, comparisons have different, and arbitrary, probabilities, and they are not known in advance. To address more realistic scenarios, researchers considered more general probabilistic models.

Consider a set of n elements, without loss of generality $[n] \stackrel{\text{def}}{=} \{1, 2, \dots, n\}$. A *probabilistic model*, or *model* for short, is an assignment of a *preference probability* $p_{i,j} \in [0,1]$ for every $i \neq j \in [n]$, reflecting the probability that i is *preferred* when compared with j. We assume that repeated comparisons are independent and that there are no "draws", hence $p_{j,i} = 1 - p_{i,j}$.

31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA.

If $p_{i,j} \ge \frac{1}{2}$, we say that i is *preferable* to j and write $i \ge j$. Element i is *maximal* in a model if $i \ge j$ for all $j \ne i$. And a permutation $\ell_1, \ldots, \ell_n$ is a *ranking* if $\ell_i \ge \ell_j$ for all $i \le j$. Observe that the first element of any ranking is always maximal. For example, for $n = 3$, $p_{1,2} = 1/2$, $p_{1,3} = 1/3$, and $p_{2,3} = 2/3$, we have $1 \ge 2$, $2 \ge 1$, $3 \ge 1$, and $2 \ge 3$. Hence 2 is the unique maximum, and 2,3,1 is the unique ranking. We seek algorithms that without knowing the underlying model, use pairwise comparisons to find a maximal element and a ranking.

Two concerns spring to mind. First, there may be two elements i, j with $p_{i,j}$ arbitrarily close to half, requiring arbitrarily many comparisons just to determine which is preferable to the other. This concern has a common remedy, that we also adopt. The *PAC* paradigm, e.g. [7, 8], that requires the algorithm's output to be only *Probably Approximately Correct*.

Let $\tilde{p}_{i,j} \stackrel{\text{def}}{=} p_{i,j} - \frac{1}{2}$ be the *centered* preference probability. Note that $\tilde{p}_{i,j} \ge 0$ iff i is preferable to j. If $\tilde{p}_{i,j} \ge -\epsilon$ we say that i is ϵ-preferable to j. For $0 < \epsilon < 1/2$, an element $i \in [n]$ is *ϵ-maximum* if it is ϵ-preferable to all other elements, namely, $\tilde{p}_{i,j} \ge -\epsilon \;\forall j \ne i$. Given $\epsilon > 0, \frac{1}{2} \ge \delta > 0$, a PAC maxing algorithm must output an ϵ-maxima with probability $\ge 1 - \delta$, henceforth abbreviated *with high probability (WHP)*. Similarly, a permutation $\ell_1, \ldots, \ell_n$ of $\{1, \ldots, n\}$ is an *ϵ-ranking* if ℓ_i is ϵ-preferable to ℓ_j for all $i \le j$, and a PAC ranking algorithm must output an ϵ-ranking WHP. Note that in this paper, we consider $\delta \le \frac{1}{2}$, the more practical regime. For larger values of δ, one can use our algorithms with $\delta = \frac{1}{2}$.

The second concern is that not all models have a ranking, or even a maximal element. For example, for $p_{1,2} = p_{2,3} = p_{3,1} = 1$, or the more opaque yet interesting non-transitive coins [9], each element is preferable to the cyclically next, hence there is no maximal element and no ranking.

A standard approach, that again we too will adopt, to address this concern is to consider structured models. The simplest may be parametric models, of which one of the more common is *Placket Luce* (PL) [10, 11], where each element i is associated with an unknown positive number a_i and $p_{i,j} = \frac{a_i}{a_i + a_j}$. [12] derived a PAC maxing algorithm that uses $\mathcal{O}(\frac{n}{\epsilon^2} \log \frac{n}{\epsilon\delta})$ comparisons and a PAC ranking algorithm that uses $\mathcal{O}(\frac{n}{\epsilon^2} \log n \log \frac{n}{\epsilon\delta})$ comparisons for any PL model. Related results for the *Mallows model* under a non-PAC paradigm were derived by [13].

But significantly more general, and more realistic, non-parametric, models may also have maxima and rankings. A model is *strongly stochastically transitive (SST)*, if $i \ge j$ and $j \ge k$ imply $p_{i,k} \ge \max(p_{i,j}, p_{j,k})$. By simple induction, every SST model has a maximum element and a ranking. And one additional property, that is perhaps more difficult to justify, has proved helpful in constructing maxing and sorting PAC algorithms. A tournament satisfies the *stochastic triangle inequality* if $i \ge j$ and $j \ge k$ imply that $\tilde{p}_{i,k} \le \tilde{p}_{i,j} + \tilde{p}_{j,k}$.

In Section 4 we show that if a model has a ranking, then an ϵ-ranking can be found WHP via $\mathcal{O}(\frac{n^2}{\epsilon^2} \log \frac{n}{\delta})$ comparisons. For all models that satisfy both SST and triangle inequality, [7] derived a PAC maxing algorithm that uses $\mathcal{O}(\frac{n}{\epsilon^2} \log \frac{n}{\epsilon\delta})$ comparisons. [14] eliminated the $\log \frac{n}{\epsilon}$ factor and showed that $\mathcal{O}\big(\frac{n}{\epsilon^2} \log \frac{1}{\delta}\big)$ comparisons suffice and are optimal, and constructed a nearly-optimal PAC ranking algorithm that uses $\mathcal{O}(\frac{n \log n (\log \log n)^3}{\epsilon^2})$ comparisons for all $\delta \ge \frac{1}{n}$, off by a factor of $\mathcal{O}((\log \log n)^3)$ from optimum. Lower-bounds follow from an analogy to [15, 6]. Observe that since the PL model satisfies both SST and triangle inequality, these results also improve the corresponding PL results.

Finally, we consider models that are not SST, or perhaps don't have maximal elements, rankings, or even their ϵ-equivalents. In all these cases, one can apply a weaker order relation. The *Borda score* $s(i) \stackrel{\text{def}}{=} \frac{1}{n} \sum_j p_{i,j}$ is the probability that i is preferable to another, randomly selected, element. Element i is *Borda maximal* if $s(i) = \max_j s(j)$, and *ϵ-Borda maximal* if $s(i) \ge \max_j s(j) - \epsilon$. A PAC Borda-maxing algorithm outputs an ϵ-Borda maximal element WHP (with probability $\ge 1 - \delta$). Similarly, a *Borda ranking* is a permutation $i_1, \ldots, i_n$ such that for all $1 \le j \le n - 1$, $s(i_j) \ge s(i_{j+1})$. An *ϵ-Borda ranking* is a permutation where for all $1 \le j \le k \le n$, $s(i_j) \ge s(i_k) - \epsilon$. A PAC Borda-ranking algorithm outputs an ϵ-Borda ranking WHP.

Recall that Borda scores apply to all models. As noted in [16, 17, 8, 18] considering elements with nearly identical Borda scores shows that exact Borda-maxing and ranking requires arbitrarily many comparisons. [8] derived a PAC Borda ranking, and therefore also maxing, algorithms that use

$\mathcal{O}(\frac{n^2}{\epsilon^2})$ comparisons. [19] derived a $\mathcal{O}(\frac{n\log n}{\epsilon^2}\log(\frac{n}{\delta}))$ PAC Borda ranking algorithm for restricted setting. However note that several simple models, including $p_{1,2} = p_{2,3} = p_{3,1} = 1$ do not belong to this model.

[20, 21, 22] considered deterministic adversarial versions of this problem that has applications in [23]. Finally, we note that all our algorithms are adaptive, where each comparison is chosen based on the outcome of previous comparisons. Non-adaptive algorithms were discussed in [24, 25, 26, 27].

2 Results and Outline

Our goal is to find the minimal assumptions that enable efficient algorithms for these problems. In particular, we would like to see if we can eliminate the somewhat less-natural triangle inequality. With two algorithmic problems: maxing and ranking, and one property–SST and one special metric–Borda scores, the puzzle consists of four main questions.

1) With just SST (and no triangle inequality) are there: a) PAC maxing algorithms with $\mathcal{O}(n)$ comparisons? b) PAC ranking algorithms with near $\mathcal{O}(n\log n)$ comparisons? 2) With no assumptions at all, but for the Borda-score metric, are there: a) PAC Borda-maxing algorithms with $\mathcal{O}(n)$ comparisons? b) PAC Borda-ranking algorithms with near $\mathcal{O}(n\log n)$ comparisons?

We essentially resolve all four questions. 1a) Yes. In Section 3, Theorem 6, we use SST alone to derive a $\mathcal{O}\big(\frac{n}{\epsilon^2}\log\frac{1}{\delta}\big)$ comparisons PAC maxing algorithm. Note that this is the same complexity as with triangle inequality, and it matches the lower bound. 1b) No. In Section 4, Theorem 7, we show that there are SST models where any PAC ranking algorithm with $\epsilon \le 1/4$ requires $\Omega(n^2)$ comparisons. This is significantly higher than the roughly $\mathcal{O}(n\log n)$ comparisons needed with triangle inequality, and is close to the $\mathcal{O}(n^2\log n)$ comparisons required without any assumptions. 2a) Yes. In Section 5, Theorem 8, we derive a PAC Borda maxing algorithm that without any model assumptions requires $\mathcal{O}\big(\frac{n}{\epsilon^2}\log\frac{1}{\delta}\big)$ comparisons which is order optimal. 2b) Yes. In Section 5, Theorem 9, we derive a PAC Borda ranking algorithm that without any model assumptions requires $\mathcal{O}\big(\frac{n}{\epsilon^2}\log\frac{n}{\delta}\big)$ comparisons.

Beyond the theoretical results sections, in Section 6, we provide experiments on simulated data. In Section 7, we discuss the results.

3 Maxing

3.1 SEQ-ELIMINATE

Our main building block is a simple, though sub-optimal, algorithm SEQ-ELIMINATE that sequentially eliminates one element from input set to find an ϵ-maximum under SST.

SEQ-ELIMINATE uses $\mathcal{O}\big(\frac{n}{\epsilon^2}\log\frac{n}{\delta}\big)$ comparisons and w.p.$\ge 1-\delta$, finds an ϵ-maximum. Even for simpler models [15] we know that an algorithm needs $\Omega\big(\frac{n}{\epsilon^2}\log\frac{1}{\delta}\big)$ comparisons to find an ϵ-maximum w.p.$\ge 1-\delta$. Hence the number of comparisons used by SEQ-ELIMINATE is optimal up to a constant factor when $\delta \le \frac{1}{n}$ but can be $\log n$ times the lower bound for $\delta = \frac{1}{2}$.

By SST, any element that is ϵ-preferable to absolute maximum element of S is an ϵ-maximum of S. Therefore if we can reduce S to a subset S' of size $\mathcal{O}(\frac{n}{\log n})$ that contains an absolute maximum of S using $\mathcal{O}\big(\frac{n}{\epsilon^2}\log\frac{1}{\delta}\big)$ comparisons, we can then use SEQ-ELIMINATE to find an ϵ-maximum of S' and the number of comparisons is optimal up to constants. We provide one such reduction in subsection 3.2.

Sequential elimination techniques have been used before [13] to find an absolute maximum. In such approaches, a running element is maintained, and is compared and replaced with a competing element in S if the latter is found to be better with confidence $\ge 1-\delta/n$. Note that if the running and competing elements are close to each other, this technique can take an arbitrarily long time to declare the winner. But since we are interested in finding only an ϵ-maximum, SEQ-ELIMINATE circumvents this issue. We later show that SEQ-ELIMINATE needs to update the running element r with the competing element c if $\tilde{p}_{c,r} \ge \epsilon$ and retain r if $\tilde{p}_{c,r} \le 0$. If $0 < \tilde{p}_{c,r} < \epsilon$, replacing or retaining r doesn't affect the performance of SEQ-ELIMINATE significantly. Thus, in other words we've reduced the problem to testing whether $\tilde{p}_{c,r} \le 0$ or $\tilde{p}_{c,r} \ge \epsilon$.

Assuming that testing problem always returns the right answer, since SEQ-ELIMINATE never replaces the running element with a worse element, either the output is the absolute maximum b^* or b^* is never the running element. If b^* is eliminated against running element r then $\tilde{p}_{b^*,r} \le \epsilon$ and hence r is an ϵ-maximum and since the running element only gets better, the output is an ϵ-maximum.

We first present a testing procedure COMPARE that we use to update the running element in SEQ-ELIMINATE.

3.1.1 COMPARE

COMPARE$(i, j, \epsilon_l, \epsilon_u, \delta)$ takes two elements i and j, and two biases $\epsilon_u > \epsilon_l$, and with confidence $\ge 1 - \delta$, determines whether $\tilde{p}_{i,j}$ is $\le \epsilon_l$ or $\ge \epsilon_u$.

For this, COMPARE compares the two elements $2/(\epsilon_u - \epsilon_l)^2 \log(2/\delta)$ times. Let $\hat{p}_{i,j}$ be the fraction of times i beats j, and let $\hat{\tilde{p}}_{i,j} \stackrel{\text{def}}{=} \hat{p}_{i,j} - \frac{1}{2}$. If $\hat{\tilde{p}}_{i,j} < (\epsilon_l + \epsilon_u)/2$, COMPARE declares $\tilde{p}_{i,j} \le \epsilon_l$ (returns 1), and otherwise it declares $\tilde{p}_{i,j} \ge \epsilon_u$ (returns 2).

Due to lack of space, we present the algorithm COMPARE in Appendix A.1 along with certain improvements for better performance in practice .

In the Lemma below, we bound the number of comparisons used by COMPARE and prove its correctness. Proof is in A.2.

Lemma 1. *For $\epsilon_u > \epsilon_l$, COMPARE$(i, j, \epsilon_l, \epsilon_u, \delta)$ uses $\le \frac{2}{(\epsilon_u - \epsilon_l)^2} \log \frac{2}{\delta}$ comparisons and if $\tilde{p}_{i,j} \le \epsilon_l$, then w.p.$\ge 1 - \delta$, it returns 1, else if $\tilde{p}_{i,j} \ge \epsilon_u$, w.p.$\ge 1 - \delta$, it returns 2.*

Now we present SEQ-ELIMINATE that uses the testing subroutine COMPARE and finds an ϵ-maximum.

3.1.2 SEQ-ELIMINATE Algorithm

SEQ-ELIMINATE takes a variable set S, selects a random *running element* $r \in S$ and repeatedly uses COMPARE$(c, r, 0, \epsilon, \delta/n)$ to compare r to a random *competing* element $c \in S \setminus r$. If COMPARE returns 1 i.e., deems $\tilde{p}_{c,r} \le 0$, it retains r as the running element and eliminates c from S, but if COMPARE returns 2 i.e., deems $\tilde{p}_{c,r} \ge \epsilon$, it eliminates r from S and updates c as the new running element.

Algorithm 1 SEQ-ELIMINATE

inputs
  Set S, bias ϵ, confidence δ
$n \leftarrow |S|$
$r \leftarrow$ a random $c \in S$, $S = S \setminus \{r\}$
while $S \ne \emptyset$ **do**
  Pick a random $c \in S$, $S = S \setminus \{c\}$.
  if COMPARE$(c, r, 0, \epsilon, \frac{\delta}{n}) = 2$ **then**
    $r \leftarrow c$
  end if
end while
return r

We now bound the number of comparisons used by SEQ-ELIMINATE(S, ϵ, δ) and prove its correctness. Proof is in A.3.

Theorem 2. SEQ-ELIMINATE(S, ϵ, δ) *uses $\mathcal{O}\left(\frac{|S|}{\epsilon^2} \log \frac{|S|}{\delta}\right)$ comparisons, and w.p.$\ge 1 - \delta$ outputs an ϵ-maximum.*

3.2 Reduction

Recall that, for $\delta \le \frac{1}{n}$, SEQ-ELIMINATE is order-wise optimal. For $\delta \ge \frac{1}{n}$, here we present a reduction procedure that uses $\mathcal{O}\left(\frac{n}{\epsilon^2}\log\frac{1}{\delta}\right)$ comparisons and w.p.$\ge 1-\delta$, outputs a subset S' of size $\mathcal{O}(\sqrt{n\log n})$ and an element a such that either a is a $2\epsilon/3$-maximum or S' contains an absolute maximum of S. Combining the reduction with SEQ-ELIMINATE results in an order-wise optimal algorithm.

We form the reduced subset S' by pruning S. We compare each element $e \in S$ with an anchor element a, test whether $\tilde{p}_{e,a} \le 0$ or $\tilde{p}_{e,a} \ge 2\epsilon/3$ using COMPARE, and retain all elements e for which COMPARE returns the second hypothesis. For S' to be of size $\mathcal{O}(\sqrt{n\log n})$ we'd like to pick an anchor element that is among the top $\mathcal{O}(\sqrt{n\log n})$ elements. But this can be computationally hard and we show that it suffices to pick an anchor that is not $\epsilon/3$-preferable to at most $\mathcal{O}(\sqrt{n\log n})$ elements in S.

An element a is called an (ϵ, n')-*good anchor* if a is not ϵ-preferable to at most n' elements, i.e., $|\{e : e \in S \text{ and } \tilde{p}_{e,a} > \epsilon\}| \le n'$.

We now present the subroutine PICK-ANCHOR that finds a good anchor element.

3.2.1 Picking Anchor Element

PICK-ANCHOR$(S, n', \epsilon, \delta)$ uses $\mathcal{O}\left(\frac{n}{n'\epsilon^2}\log\frac{1}{\delta}\log\frac{n}{n'\delta}\right)$ comparisons and w.p.$\ge 1-\delta$, outputs an (ϵ, n')-good anchor element. PICK-ANCHOR first picks randomly a set Q of $\frac{n}{n'}\log\frac{2}{\delta}$ elements from S without replacement. This ensures that w.p.$\ge 1-\delta$, Q contains at least one of the top n' elements. We then use SEQ-ELIMINATE to find an ϵ-maximum of Q.

Let the absolute maximum element of Q be denoted as q^*. Now an ϵ-maximum of Q is ϵ-preferable to q^*. Further, if Q contains an element in the top n' elements, there exists $n-n'$ elements worse than q^* in S. Thus by SST, the ϵ-maximum of Q is also ϵ-preferable to these $n-n'$ elements and hence the output of PICK-ANCHOR is an (ϵ, n')-good anchor element. PICK-ANCHOR is shown in appendix A.4

We now bound the number of comparisons used by PICK-ANCHOR and prove its correctness. Proof is in A.5.

Lemma 3. *PICK-ANCHOR$(S, n', \epsilon, \delta)$ uses $\mathcal{O}\left(\frac{n}{n'\epsilon^2}\log\frac{1}{\delta}\log\frac{n}{n'\delta}\right)$ comparisons and w.p.$\ge 1-\delta$, outputs an (ϵ, n')-good anchor element.*

Remark 4. *Note that PICK-ANCHOR$(S, cn, \epsilon, \delta)$ uses $\mathcal{O}_c\left(\frac{1}{\epsilon^2}\left(\log\frac{1}{\delta}\right)^2\right)$ comparisons where the constant depends only on c but not on n. Hence it is advantageous to use this method to pick near-maximum element when n is large.*

We now present PRUNE that takes an anchor element as input and prunes the set S using the anchor.

3.2.2 Pruning

Given an (ϵ_l, n')-good anchor element a, w.p.$\ge 1-\delta/2$, PRUNE$(S, a, n', \epsilon_l, \epsilon_u, \delta)$ outputs a subset S' of size $\le 2n'$. Further, any element e that is at least ϵ_u-better than a i.e., $\tilde{p}_{e,a} \ge \epsilon_u$ is in S' w.p.$\ge 1-\delta/2$.

PRUNE prunes S in multiple rounds. In each round t, for every element e in S, PRUNE tests whether $\tilde{p}_{e,a} \le \epsilon_l$ or $\tilde{p}_{e,a} \ge \epsilon_u$ using COMPARE$(e, a, \epsilon_l, \epsilon_u, \delta/2^{t+1})$ and eliminates e if the first hypothesis i.e., $\tilde{p}_{e,a} \le \epsilon_l$ is returned. By Lemma 1, an element e that is ϵ_u better than a i.e., $\tilde{p}_{e,a} \ge \epsilon_u$ passes the t^{th} round of pruning w.p.$\ge 1-\delta/2^{t+1}$. Thus by union bound, the probability that such an element is not present in the pruned set is $\le \sum_{t=1}^{\infty}\delta/2^{t+1} \le \delta/2$.

Now for element e that is not ϵ_l-better than a i.e., $\tilde{p}_{e,a} \le \epsilon_l$, by Lemma 1, the first hypothesis is returned w.p.$\ge 1-\delta/4$. Hence w.h.p., the number of such elements (not ϵ_l-better elements) is reduced by a factor of δ after each round. Since a is an (ϵ_l, n')-good anchor element, there are at most n' elements atleast ϵ_l-better than a. Thus the number of elements left in the pruned set after round t is at most $n' + n\delta^t$. Thus PRUNE succeeds eventually in reducing the size to $\le 2n'$ (in $\le \log_{1/\delta}\frac{n}{n'}$ rounds).

Algorithm 2 PRUNE

inputs
  Set S, element a, size n', lower bias ϵ_l, upper bias ϵ_u, confidence δ.
$t \leftarrow 1$
$S_1 \leftarrow S$
while $|S_t| > 2n'$ and $t < \log^2 n$ **do**
  Initialize: $Q_t \leftarrow \emptyset$
  for e in S_t **do**
    if COMPARE$(e, a, \epsilon_l, \epsilon_u, \delta/2^{t+1}) = 1$ **then**
      $Q_t \leftarrow Q_t \cup \{e\}$
    end if
  end for
  $S_{t+1} \leftarrow S_t \setminus Q_t$
  $t \leftarrow t + 1$
end while
return S_t.

We now bound the number of comparisons used by PRUNE and prove its correctness. Proof is in A.6.

Lemma 5. *If $n' \geq \sqrt{6n \log n}$, $\delta \geq \frac{1}{n}$ and a is an (ϵ_l, n')-good anchor element, then w.p.$\geq 1 - \frac{\delta}{2}$, PRUNE$(S, a, n', \epsilon_l, \epsilon_u, \delta)$ uses $\mathcal{O}\Big(\frac{n}{(\epsilon_u - \epsilon_l)^2} \log \frac{1}{\delta}\Big)$ comparisons and outputs a set of size less than $2n'$. Further if a is not an ϵ_u-maximum of S then w.p.$\geq 1 - \frac{\delta}{2}$, the output set contains an absolute maximum element of S.*

3.3 Full Algorithm

We now present the main algorithm, OPT-MAXIMIZE that w.p.$\geq 1-\delta$, uses $\mathcal{O}\big(\frac{n}{\epsilon^2} \log \frac{1}{\delta}\big)$ comparisons and outputs an ϵ-maximum. For $\delta \leq \frac{1}{n}$, SEQ-ELIMINATE uses $\mathcal{O}(\frac{n}{\epsilon^2} \log \frac{1}{\delta})$ comparisons and hence we directly use SEQ-ELIMINATE. Below we assume $\delta > \frac{1}{n}$.

Here OPT-MAXIMIZE first finds an $(\epsilon/3, \sqrt{6n \log n})$-good anchor element a using PICK-ANCHOR$(S, \sqrt{6n \log n}, \epsilon/3, \frac{\delta}{4})$. Then using PRUNE$(S, a, \sqrt{6n \log n}, \epsilon/3, 2\epsilon/3, \frac{\delta}{4})$ with a, OPT-MAXIMIZE prunes S to a subset S' of size $\leq 2\sqrt{6n \log n}$ such that if a is not a $2\epsilon/3$ maximum i.e. $\tilde{p}_{b^*,a} > 2\epsilon/3$, S' contains the absolute maximum b^* w.p.$\geq 1 - \delta/2$. OPT-MAXIMIZE then checks if a is a $2\epsilon/3$ maximum by using COMPARE$(e, a, 2\epsilon/3, \epsilon, \delta/(4n))$ for every element $e \in S'$. If COMPARE returns first hypothesis for every $e \in S'$ then OPT-MAXIMIZE outputs a or else OPT-MAXIMIZE outputs SEQ-ELIMINATE$(S', \epsilon, \frac{\delta}{4})$.

Note that only one of these three cases is possible: (1) a is a $2\epsilon/3$-maximum, (2) a is not an ϵ-maximum and (3) a is an ϵ-maximum but not a $2\epsilon/3$-maximum. In case (1), since a is a $2\epsilon/3$-maximum, by Lemma 1, w.p.$\geq 1 - \delta/4$, COMPARE returns the first hypothesis for every $e \in S'$ and OPT-MAXIMIZE outputs a. In both cases (2) and (3), as stated above, w.p.$\geq 1 - \delta/2$, S' contains the absolute maximum b^*. Now in case (2) since a is not an ϵ-maximum, by Lemma 1, w.p.$\geq 1 - \delta/(4n)$, COMPARE$(b^*, a, 2\epsilon/3, \epsilon, \delta/(4n))$ returns the second hypothesis. Thus OPT-MAXIMIZE outputs SEQ-ELIMINATE$(S', \epsilon, \delta/4)$, which w.p.$\geq 1 - \delta/4$, returns an ϵ-maximum of S' (recall that an ϵ-maximum of S' is an ϵ-maximum of S if S' contains b^*). Finally in case (3), OPT-MAXIMIZE either outputs a or SEQ-ELIMINATE$(S', \epsilon, \delta/4)$ and either output is an ϵ-maximum w.p.$\geq 1 - \delta$. In the below Theorem, we bound comparisons used by OPT-MAXIMIZE and prove its correctness. Proof is in A.7.

Theorem 6. *W.p.$\geq 1 - \delta$, OPT-MAXIMIZE(S, ϵ, δ) uses $\mathcal{O}(\frac{n}{\epsilon^2} \log \frac{1}{\delta})$ comparisons and outputs an ϵ-maximum.*

4 Ranking

Recall that [14] considered a model with both SST and stochastic triangle inequality and derived an ϵ-ranking with $\mathcal{O}\Big(\frac{n \log n (\log \log n)^3}{\epsilon^2}\Big)$ comparisons for $\delta = \frac{1}{n}$. By constrast, we consider a more

Algorithm 3 OPT-MAXIMIZE

inputs
    Set S, bias ϵ, confidence δ.
if $\delta \le \frac{1}{n}$ **then**
    return SEQ-ELIMINATE(S, ϵ, δ)
end if
$a \leftarrow$ PICK-ANCHOR$(S, \sqrt{6n \log n}, \epsilon/3, \frac{\delta}{4})$
$S' \leftarrow$ PRUNE$(S, a, \sqrt{6n \log n}, \epsilon/3, 2\epsilon/3, \frac{\delta}{4})$
for element e in S' **do**
    if COMPARE$(e, a, \frac{2\epsilon}{3}, \epsilon, \frac{\delta}{4n}) = 2$ **then**
        return SEQ-ELIMINATE$(S', \epsilon, \frac{\delta}{4})$
    end if
end for
return a

general model without stochastic triangle inequality and show that even a $1/4$-ranking with just SST takes $\Omega(n^2)$ comparisons for $\delta \le \frac{1}{8}$.

To establish the lower bound, we reduce the problem of finding $1/4$-ranking to finding a coin with bias 1 among $\frac{n(n-1)}{2} - 1$ other fair coins. For this, we consider the following model with n elements $\{a_1, a_2, ..., a_n\}$: $\tilde{p}_{a_1,a_n} = \frac{1}{2}$, $\tilde{p}_{a_i,a_j} = \mu (0 < \mu < 1/n^{10})$, when $i < j$ and $(i, j) \ne (1, n)$. Note that this model satisfies SST but not stochastic triangle inequality. Also note that any ranking where a_1 precedes a_n is an $1/4$-ranking and thus the algorithm only needs to order a_1 and a_n correctly. Now the output of a comparison between any two elements other than a_1 and a_n is essentially a fair coin toss (since μ is very small). Thus if we output a ranking without querying comparison between a_1 and a_n, then the ranking is correct w.p.$\approx \frac{1}{2}$ since a_1 and a_n must necessarily be ordered correctly. Now if an algorithm uses only $n^2/20$ comparisons then the probability that the algorithm queried at least one comparison between a_1 and a_n is less than $\frac{1}{2}$ and hence cannot achieve a confidence of $\frac{7}{8}$. Proof sketch in B.1.

Theorem 7. *There exists a model that satisfies SST for which any algorithm requires $\Omega(n^2)$ comparisons to find a $1/4$-ranking with probability $\ge 7/8$.*

We also present a trivial ϵ-ranking algorithm in Appendix B.2 that for any stochastic model with ranking (Weak Stochastic Transitivity), uses $\mathcal{O}(\frac{n^2}{\epsilon^2} \log \frac{n}{\delta})$ comparisons and outputs an ϵ-ranking w.p.$\ge 1 - \delta$.

5 Borda Scores

We show that for general models, using $\mathcal{O}(\frac{n}{\epsilon^2} \log \frac{1}{\delta})$ comparisons w.p.$\ge 1 - \delta$, we can find an ϵ-Borda maximum and using $\mathcal{O}(\frac{n}{\epsilon^2} \log \frac{n}{\delta})$ comparisons w.p.$\ge 1 - \delta$, we can find an ϵ-Borda ranking.

Recall that Borda score $s(e)$ of an element e is the probability that e is preferable to an element picked randomly from S i.e., $s(e) = \frac{1}{n} \sum_{f \in S} \tilde{p}_{e,f}$. We first make a connection between Borda scores of elements and the traditional multi armed bandit setting. In the Bernoulli multi armed setting, every arm a is associated with a parameter $q(a)$ and pulling that arm results in a reward $\mathrm{B}(q(a))$, a Bernoulli random variable with parameter $q(a)$. Observe that we can simulate our pairwise comparisons setting as a traditional bandit arms setting by comparing an element with a random element where in our setting, for every element e, the associated parameter is $s(e)$. Thus PAC optimal algorithms derived under traditional bandit setting work for PAC Borda score setting too. [28] and several others derived a PAC maximum arm selection algorithms that use $\mathcal{O}(\frac{n}{\epsilon^2} \log \frac{1}{\delta})$ comparisons and find an arm with parameter at most ϵ less than the highest. This implies an ϵ-Borda maxing algorithm with the same complexity. Proof follows from reduction to Bernoulli multi-armed bandit setting.

Theorem 8. *There exists an algorithm that uses $\mathcal{O}(\frac{n}{\epsilon^2} \log \frac{1}{\delta})$ comparisons and w.p.$\ge 1 - \delta$, outputs an ϵ-Borda maximum.*

For ϵ-Borda ranking, we note that if we compare an element e with $\frac{2}{\epsilon^2}\log\frac{2n}{\delta}$ random elements, w.p. $\geq 1-\delta/n$, the fraction of times e wins approximates the Borda score of e to an additive error of $\frac{\epsilon}{2}$. Ranking based on these approximate scores results in an ϵ-Borda ranking. We present BORDA-RANKING in C.1 that uses $\frac{2n}{\epsilon^2}\log\frac{2n}{\delta}$ comparisons and w.p.$\geq 1-\delta$ outputs an ϵ-Borda ranking. Proof in C.1.

Theorem 9. BORDA-RANKING(S,ϵ,δ) *uses* $\frac{2n}{\epsilon^2}\log\frac{2n}{\delta}$ *comparisons and w.p.*$\geq 1-\delta$ *outputs an* ϵ*-Borda ranking.*

6 Experiments

In this section we validate the performance of our algorithms using simulated data. Since we essentially derived a negative result for ϵ-ranking, we consider only our ϵ-maxing algorithms - SEQ-ELIMINATE and OPT-MAXIMIZE for experiments. All results are averaged over 100 runs.

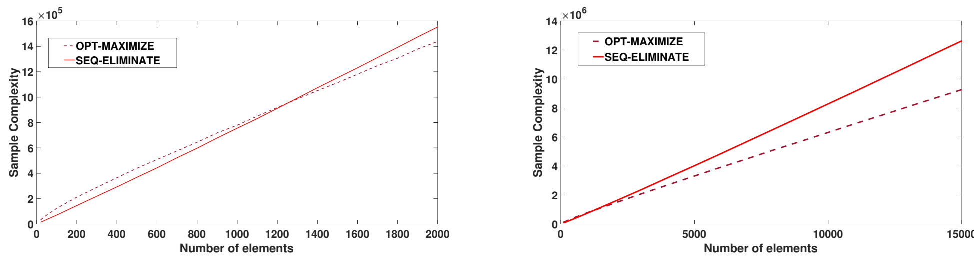

(a) small values of n (b) large values of n

Figure 1: Comparison of SEQ-ELIMINATE and OPT-MAXIMIZE

Similar to [14, 7], we consider the stochastic model $p_{i,j} = 0.6\ \forall i < j$. We use maxing algorithms to find 0.05-maximum with error probability $\delta = 0.1$. Note that $i = 1$ is the unique 0.05-maximum under this model. In Figure 1, we compare the performance of SEQ-ELIMINATE and OPT-MAXIMIZE over different ranges of n. Figures 1(a), 1(b) show that for small n i.e., $n \leq 1300$ SEQ-ELIMINATE performs well and for large n i.e., $n \geq 1300$, OPT-MAXIMIZE performs well. Since we are using $\delta = 0.1$, the experiment suggests that for $\delta \gtrapprox \frac{1}{n^{1/3}}$, OPT-MAXIMIZE uses fewer comparisons as compared to SEQ-ELIMINATE. Hence it would be beneficial to use SEQ-ELIMINATE for $\delta \leq \frac{1}{n^{1/3}}$ and OPT-MAXIMIZE for higher values of δ. In further experiments, we use $\delta = 0.1$ and $n < 1000$ so we use SEQ-ELIMINATE for better performance.

We compare SEQ-ELIMINATE with **BTM-PAC** [7], **KNOCKOUT** [14], **MallowsMPI** [13], and **AR** [16] . **KNOCKOUT** and **BTM-PAC** are PAC maxing algorithms for models with SST and stochastic triangle inequality requirements. **AR** finds an element with maximum Borda score. **Mallows** finds the absolute best element under Weak Stochastic Transitivity.

We again consider the model: $p_{i,j} = 0.6\ \forall i < j$ and try to find a 0.05-maximum with error probability $\delta = 0.1$. Note that this model satisfies both SST and stochastic triangle inequality and under this model all these algorithms can find an ϵ-maximum. From Figure 2(a), we can see that **BTM-PAC** performs worse for even small values of n and from Figure 2(b), we can see that **AR** performs worse for higher values of n. One possible reason is that **BTM-PAC** is tailored for reducing regret in the bandit setting and in the case of **AR**, Borda scores of elements become approximately the same with increasing number of elements, leading to more comparisons. For this reason, we drop **BTM-PAC** and **AR** for further experiments.

We also tried **PLPAC** [12] but it fails to achieve required accuracy of $1-\delta$ since it is designed primarily for Plackett-Luce. For example, we considered the previous setting $p_{i,j} = 0.6\ \forall i < j$ with $n = 100$ and tried to find a 0.09-maximum with $\delta = 0.1$. Even though **PLPAC** used almost same number of comparisons (57237) as SEQ-ELIMINATE (56683), PLPAC failed to find 0.09-maxima 20 out of 100 runs whereas SEQ-ELIMINATE found the maximum in all 100 runs.

In figure 3, we compare algorithms SEQ-ELIMINATE, **KNOCKOUT** [14] and **MallowsMPI** [13] for models that do not satisfy stochastic triangle inequality. In Figure 3(a), we consider the stochastic model $p_{1,j} = \frac{1}{2} + \tilde{q}\ \forall j \leq n/2$, $p_{1,j} = 1\ \forall j > n/2$ and $p_{i,j} = \frac{1}{2} + \tilde{q}\ \forall 1 < i < j$ where $\tilde{q} \leq 0.05$ and we pick $n = 10$. Observe that this model satisfies SST but not stochastic triangle inequality. Here

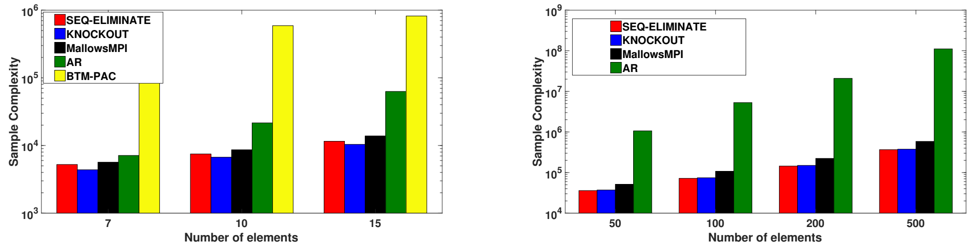

(a) small values of n (b) large valuesof n

Figure 2: Comparison of Maxing Algorithms with Stochastic Triangle Inequality

again, we try to find a 0.05-maximum with $\delta = 0.1$. Note that any $i \le n/2$ is a 0.05 maximum. From Figure 3(a), we can see that **MallowsMPI** uses more comparisons as $\tilde{q}$ decreases since **MallowsMPI** is not a PAC algorithm and tries to find the absolute maximum. Even though **KNOCKOUT** performs better than **MallowsMPI**, it fails to output a 0.05 maximum with probability 0.12 for $\tilde{q} = 0.001$ and 0.26 for $\tilde{q} = 0.0001$. Thus **KNOCKOUT** can fail when the model doesn't satisfy stochastic triangle inequality. We give an explanation for this behavior in Appendix D. By constrast, even for $\tilde{q} = 0.0001$, SEQ-ELIMINATE outputted a 0.05 maximum in all runs and outputted the abosulte maximum in 76% of trials. We can also see that SEQ-ELIMINATE uses much fewer comparisons compared to the other two algorithms.

In Figure 3(b), we compare SEQ-ELIMINATE and **MallowsMPI** on the Mallows model, a model which doesn't satisfy stochastic triangle inequality. Mallows model can be specified with one parameter ϕ. We consider $n = 10$ elements and find a 0.05-maximum with error probablility $\delta = 0.05$. From Figure 3(b) we can see that the performance of **MallowsMPI** gets worse as ϕ approaches 1, since comparison probabilities get close to $\frac{1}{2}$ whereas SEQ-ELIMINATE is not affected.

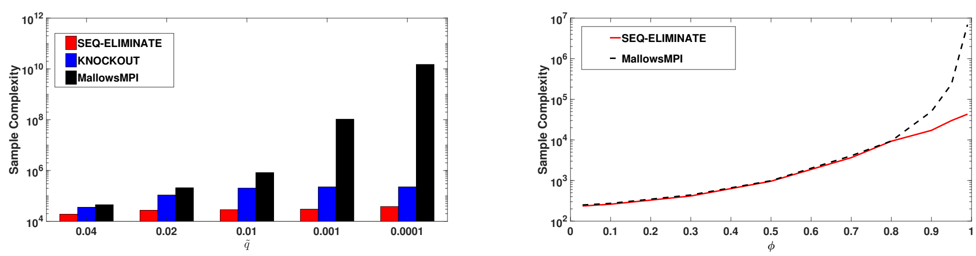

(a) No Triangle Inequality (b) Mallows Model

Figure 3: Comparison of SEQ-ELIMINATE and MALLOWSMPI over Mallows Model

One more experiment is presented in Appendix E.

7 Conclusion

We extended the study of PAC maxing and ranking to general models which satisfy SST but not stochastic triangle inequality. For PAC maxing, we derived an algorithm with linear complexity. For PAC ranking, we showed a negative result that any algorithm needs $\Omega(n^2)$ comparisons. We thus showed that removal of stochastic triangle inequality constraint does not affect PAC maxing but affects PAC ranking. We also ran experiments over simulated data and showed that our PAC maximum selection algorithms are better than other maximum selection algorithms.

For unconstrained models, we derived algorithms for PAC Borda maxing and PAC Borda ranking by making connections with traditional multi-armed bandit setting.

Acknowledgments

We thank NSF for supporting this work through grants CIF-1564355 and CIF-1619448.

References

[1] Ralf Herbrich, Tom Minka, and Thore Graepel. Trueskill: a bayesian skill rating system. In *Proceedings of the 19th International Conference on Neural Information Processing Systems*, pages 569–576. MIT Press, 2006. 1.1

[2] Jialei Wang, Nathan Srebro, and James Evans. Active collaborative permutation learning. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 502–511. ACM, 2014. 1.1

[3] Filip Radlinski and Thorsten Joachims. Active exploration for learning rankings from clickthrough data. In *Proceedings of the 13th ACM SIGKDD*, pages 570–579. ACM, 2007. 1.1

[4] Filip Radlinski, Madhu Kurup, and Thorsten Joachims. How does clickthrough data reflect retrieval quality? In *Proceedings of the 17th ACM conference on Information and knowledge management*, pages 43–52. ACM, 2008. 1.1

[5] `http://www.gif.gf/`. 1.1

[6] Uriel Feige, Prabhakar Raghavan, David Peleg, and Eli Upfal. Computing with noisy information. *SIAM Journal on Computing*, 23(5):1001–1018, 1994. 1.2

[7] Yisong Yue and Thorsten Joachims. Beat the mean bandit. In *Proc. of the ICML*, pages 241–248, 2011. 1.2, 6

[8] Róbert Busa-Fekete, Balázs Szörényi, and Eyke Hüllermeier. Pac rank elicitation through adaptive sampling of stochastic pairwise preferences. In *AAAI*, 2014. 1.2

[9] Wikipedia. Nontransitive dice — Wikipedia, the free encyclopedia. `http://en.wikipedia.org/w/index.php?title=Nontransitive%20dice&oldid=779713322`, 2017. [Online; accessed 19-May-2017]. 1.2

[10] Robin L Plackett. The analysis of permutations. *Applied Statistics*, pages 193–202, 1975. 1.2

[11] R Duncan Luce. *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2005. 1.2

[12] Balázs Szörényi, Róbert Busa-Fekete, Adil Paul, and Eyke Hüllermeier. Online rank elicitation for plackett-luce: A dueling bandits approach. In *NIPS*, pages 604–612, 2015. 1.2, 6

[13] Róbert Busa-Fekete, Eyke Hüllermeier, and Balázs Szörényi. Preference-based rank elicitation using statistical models: The case of mallows. In *Proc. of the ICML*, pages 1071–1079, 2014. 1.2, 3.1, 6, 6

[14] Moein Falahatgar, Alon Orlitsky, Venkatadheeraj Pichapati, and Ananda Theertha Suresh. Maximum selection and ranking under noisy comparisons. *arXiv preprint arXiv:1705.05366*, 2017. 1.2, 4, 6, 6, A.1, D

[15] Yuan Zhou, Xi Chen, and Jian Li. Optimal pac multiple arm identification with applications to crowdsourcing. In *International Conference on Machine Learning*, pages 217–225, 2014. 1.2, 3.1

[16] Reinhard Heckel, Nihar B Shah, Kannan Ramchandran, and Martin J Wainwright. Active ranking from pairwise comparisons and when parametric assumptions don't help. *arXiv preprint arXiv:1606.08842*, 2016. 1.2, 6

[17] Tanguy Urvoy, Fabrice Clerot, Raphael Féraud, and Sami Naamane. Generic exploration and k-armed voting bandits. In *Proc. of the ICML*, pages 91–99, 2013. 1.2

[18] Kevin Jamieson, Sumeet Katariya, Atul Deshpande, and Robert Nowak. Sparse dueling bandits. In *Artificial Intelligence and Statistics*, pages 416–424, 2015. 1.2

[19] David Timothy Lee, Ashish Goel, Tanja Aitamurto, and Helene Landemore. Crowdsourcing for participatory democracies: Efficient elicitation of social choice functions. In *Second AAAI Conference on Human Computation and Crowdsourcing*, 2014. 1.2

[20] Jayadev Acharya, Ashkan Jafarpour, Alon Orlitsky, and Ananda Theertha Suresh. Sorting with adversarial comparators and application to density estimation. In *Information Theory (ISIT), 2014 IEEE International Symposium on*, pages 1682–1686. IEEE, 2014. 1.2

[21] Jayadev Acharya, Moein Falahatgar, Ashkan Jafarpour, Alon Orlitsky, and Ananda Theertha Suresh. Maximum selection and sorting with adversarial comparators and an application to density estimation. *arXiv preprint arXiv:1606.02786*, 2016. 1.2

[22] Miklós Ajtai, Vitaly Feldman, Avinatan Hassidim, and Jelani Nelson. Sorting and selection with imprecise comparisons. *ACM Transactions on Algorithms (TALG)*, 12(2):19, 2016. 1.2

[23] Jayadev Acharya, Ashkan Jafarpour, Alon Orlitsky, and Ananda Theertha Suresh. Near-optimal-sample estimators for spherical gaussian mixtures. *NIPS*, 2014. 1.2

[24] Arun Rajkumar and Shivani Agarwal. A statistical convergence perspective of algorithms for rank aggregation from pairwise data. In *Proc. of the ICML*, pages 118–126, 2014. 1.2

[25] Sahand Negahban, Sewoong Oh, and Devavrat Shah. Iterative ranking from pair-wise comparisons. In *NIPS*, pages 2474–2482, 2012. 1.2

[26] Sahand Negahban, Sewoong Oh, and Devavrat Shah. Rank centrality: Ranking from pairwise comparisons. *Operations Research*, 2016. 1.2

[27] Minje Jang, Sunghyun Kim, Changho Suh, and Sewoong Oh. Top-k ranking from pairwise comparisons: When spectral ranking is optimal. *arXiv preprint arXiv:1603.04153*, 2016. 1.2

[28] Yuan Zhou and Xi Chen. Optimal pac multiple arm identification with applications to crowdsourcing. 2014. 5

A Maxing

A.1 COMPARE Algorithm

Motivated by a related algorithm in [14], we describe an adaptive version of COMPARE that stops when it is confident about the result, even if the number of comparisons is less than that specified in subsection 3.1.1. If $\tilde{p}_{i,j}$ is far outside (ϵ_l, ϵ_u), this adaptive algorithm will terminate much sooner.

To do so, COMPARE maintains a varying confidence interval $\hat{c}$ such that w.p. $\geq 1-\delta$, $|\hat{\tilde{p}}_{i,j} - \tilde{p}_{i,j}| < \hat{c}$ after any number of comparisons. If at any time before the above number of comparisons, $|\hat{\tilde{p}}_{i,j} - \frac{\epsilon_l+\epsilon_u}{2}| > \hat{c}$, COMPARE simply returns the result based on the current $\hat{\tilde{p}}_{i,j}$.

Algorithm 4 COMPARE

inputs
  element i, element j, bias lower limit $\epsilon_l \geq 0$, bias upper limit $\epsilon_u > \epsilon_l$, confidence δ
initialize
  $\epsilon_m = (\epsilon_l + \epsilon_u)/2$, $\hat{\tilde{p}}_{i,j} \leftarrow 0$, $\hat{c} \leftarrow \frac{1}{2}$, $t \leftarrow 0$, $w \leftarrow 0$
while $|\hat{\tilde{p}}_{i,j} - \epsilon_m| \leq \hat{c}$ and $t \leq \frac{2}{(\epsilon_u-\epsilon_l)^2} \log \frac{2}{\delta}$ **do**
  Compare i and j
  if i wins **then**
    $w \leftarrow w + 1$
  end if
  $t \leftarrow t + 1$
  $\hat{\tilde{p}}_{i,j} \leftarrow \frac{w}{t} - \frac{1}{2}$, $\hat{c} \leftarrow \sqrt{\frac{1}{2t} \log \frac{4t^2}{\delta}}$
end while
if $\hat{\tilde{p}}_{i,j} \leq \epsilon_m$ **then**
  return 1
end if
return 2

A.2 Proof of Lemma 1

We prove Lemma by dividing it into smaller parts. We first bound the comparisons used by COMPARE.

Lemma 10. *For $\epsilon_u > \epsilon_l \geq 0$, COMPARE$(i, j, \epsilon_l, \epsilon_u, \delta)$ uses $\leq \frac{2}{(\epsilon_u-\epsilon_l)^2} \log \frac{2}{\delta}$ comparisons.*

Proof. Notice that COMPARE$(i, j, \epsilon_l, \epsilon_u, \delta)$ compares elements i and j for at most $m = \frac{2}{(\epsilon_u-\epsilon_l)^2} \log \frac{2}{\delta}$ times and hence the Lemma follows. □

We show that under the first hypothesis namely $\tilde{p}_{i,j} \leq \epsilon_l$, w.p.$\geq 1-\delta$, COMPARE$(i, j, \epsilon_l, \epsilon_u, \delta)$ returns 1.

Lemma 11. *For $\epsilon_u > \epsilon_l \geq 0$, if $\tilde{p}_{i,j} \leq \epsilon_l$, then w.p.$\geq 1-\delta$, COMPARE$(i, j, \epsilon_l, \epsilon_u, \delta)$ outputs 1.*

Proof. Let $\hat{p}^t_{i,j}$ and $\hat{c}^t$ denote $\hat{p}_{i,j}$ and $\hat{c}$ respectively after t comparisons between i and j during COMPARE$(i, j, \epsilon_l, \epsilon_u, \delta)$. COMPARE$(i, j, \epsilon_l, \epsilon_u, \delta)$ outputs 2 only if $\hat{p}^t_{i,j} > \frac{1}{2} + \frac{\epsilon_l+\epsilon_u}{2} + \hat{c}^t$ for any $t < m = \frac{2}{(\epsilon_l-\epsilon_u)^2} \log \frac{2}{\delta}$ or if $\hat{p}^m_{i,j} > \frac{1}{2} + \frac{\epsilon_l+\epsilon_u}{2}$. We bound the probability of either of these events by $\frac{\delta}{2}$ and the result follows from the union bound.

By Hoeffding's inequality,

$$Pr\left(\hat{p}^t_{i,j} > \frac{1}{2} + \frac{\epsilon_l + \epsilon_u}{2} + \hat{c}^t\right) \leq Pr\left(\hat{p}^t_{i,j} > \frac{1}{2} + \epsilon_l + \hat{c}^t\right) \leq e^{-2t(\hat{c}^t)^2} = e^{-\log \frac{4t^2}{\delta}} = \frac{\delta}{4t^2}.$$

By the union bound, $Pr\left(\exists t \text{ s.t. } \hat{p}^t_{i,j} > \frac{1}{2} + \frac{\epsilon_l+\epsilon_u}{2} + \hat{c}^t\right) \leq \sum_t \frac{\delta}{4t^2} \leq \frac{\delta}{2}$.

Similarly, by Hoeffding's inequality,

$$Pr\left(\hat{p}^m_{i,j} > \frac{1}{2} + \frac{\epsilon_l + \epsilon_u}{2}\right) \leq e^{-2m((\epsilon_u-\epsilon_l)/2)^2} = e^{-\log \frac{2}{\delta}} = \frac{\delta}{2}.$$ □

We now show that under the second hypothesis namely $\tilde{p}_{i,j} \geq \epsilon_u$, w.p.$\geq 1-\delta$, COMPARE$(i, j, \epsilon_l, \epsilon_u, \delta)$ returns 2.

Lemma 12. *For $\epsilon_u > \epsilon_l \geq 0$, if $\tilde{p}_{i,j} \geq \epsilon_u$, then w.p.$\geq 1-\delta$,* COMPARE$(i,j,\epsilon_l,\epsilon_u,\delta)$ *outputs* 2.

Proof. Let $\hat{p}^t_{i,j}$ and $\hat{c}^t$ denote $\hat{p}_{i,j}$ and $\hat{c}$ respectively after t comparisons between i and j during COMPARE$(i,j,\epsilon_l,\epsilon_u,\delta)$. COMPARE$(i,j,\epsilon_l,\epsilon_u,\delta)$ outputs 1 only if $\hat{p}^t_{i,j} < \frac{1}{2} + \frac{\epsilon_l+\epsilon_u}{2} - \hat{c}^t$ for any $t < m = \frac{2}{(\epsilon_u-\epsilon_l)^2}\log\frac{2}{\delta}$ or if $\hat{p}^m_{i,j} \leq \frac{1}{2} + \frac{\epsilon_l+\epsilon_u}{2}$. We bound the probability of either of these events by $\frac{\delta}{2}$ and the result follows from the union bound.

By Hoeffding's inequality,

$$Pr\Big(\hat{p}^t_{i,j} < \frac{1}{2} + \frac{\epsilon_l+\epsilon_u}{2} - \hat{c}^t\Big) \leq Pr\Big(\hat{p}^t_{i,j} < \frac{1}{2} + \epsilon_u - \hat{c}^t\Big) \leq e^{-2t(\hat{c}^t)^2} = e^{-\log\frac{4t^2}{\delta}} = \frac{\delta}{4t^2}.$$

By the union bound, $Pr\big(\exists t \text{ s.t. } \hat{p}^t_{i,j} < \frac{1}{2} + \frac{\epsilon_l+\epsilon_u}{2} - \hat{c}^t\big) \leq \sum_t \frac{\delta}{4t^2} \leq \frac{\delta}{2}$.

Similarly, by Hoeffding's inequality,

$$Pr\Big(\hat{p}^m_{i,j} \leq \frac{1}{2} + \frac{\epsilon_l+\epsilon_u}{2}\Big) \leq e^{-2m((\epsilon_u-\epsilon_l)/2)^2} = e^{-\log\frac{2}{\delta}} = \frac{\delta}{2}. \qquad \square$$

Thus proof of Lemma 1 follows from Lemmas 10, 11 and 12.

A.3 Proof of Theorem 2

Proof. We first bound the total number of comparisons. Before each call of COMPARE (step 7), SEQ-ELIMINATE eliminates an element in step 6, hence COMPARE is called for exactly $n-1$ times. Further observe that COMPARE$(i,j,0,\epsilon,\delta/n)$ always uses less than $\frac{2}{\epsilon^2}\log\frac{2n}{\delta}$ comparisons. Hence the total comparisons used by SEQ-ELIMINATE(S,ϵ,δ) is

$$\leq \sum_{k=1}^{n-1} \frac{2}{\epsilon^2}\log\frac{2n}{\delta} = \mathcal{O}\Big(\frac{n}{\epsilon^2}\log\frac{n}{\delta}\Big).$$

We now show that w.p.$\geq 1-\delta$, SEQ-ELIMINATE(S,ϵ,δ) outputs an ϵ-maximum. Let r^t, c^t denote the running and competing elements respectively before t^{th} run of COMPARE. Then by Lemmas 11 and 12, for any t, w.p. $\geq 1-\frac{\delta}{n}$,

$$\tilde{p}_{r^{t+1},r^t} \geq 0, \tag{1}$$

$$\tilde{p}_{r^{t+1},c^t} > -\epsilon. \tag{2}$$

Further, by the union bound the probability that Equations 1 and 2 do not hold for some $1 \leq t \leq n$ is $\leq \delta$. Now let b^* be the absolute maximum element i.e., $\tilde{p}_{b^*,e} \geq 0 \ \forall e \in S$. Then, either b^* is set as the running element before the first run of COMPARE i.e., $r^1 = b^*$ or b^* is the competing element at the t^{th} run of COMPARE for some $1 \leq t \leq n$ i.e., $c^t = b^*$. We show that in both cases, the output is an ϵ-maximum.

If $r^1 = b^*$, then by equation 1, future running elements are either b^* or better than b^*. Since b^* is the absolute maximum, future running elements must be b^* and hence b^* is the output.

If for some t, $c^t = b^*$, then by equation 2, $\tilde{p}_{r^{t+1},b^*} > -\epsilon$. Further, by equation 1, $\tilde{p}_{r^l,r^{t+1}} \geq 0 \quad \forall l \geq t+1$. Hence by strong stochastic transitivity, $\tilde{p}_{r^n,b^*} > -\epsilon$. Again, by strong stochastic transitivity, $\tilde{p}_{r^n,e} > -\epsilon \ \forall e \in S$. Hence, the output is ϵ-maximum. $\square$

A.4 PICK-ANCHOR algorithm

Algorithm 5 PICK-ANCHOR

1: **inputs**
2: Set S of size n, size n', bias ϵ, confidence δ.
3: Form a set Q by selecting $\min\big(\frac{n}{n'}\log\frac{2}{\delta}, n\big)$ random elements from S without replacement.
4: **return** SEQ-ELIMINATE$\big(Q,\epsilon,\frac{\delta}{2}\big)$

A.5 Proof of Lemma 3

Proof. We first bound the number of comparisons used by PICK-ANCHOR$(S, n', \epsilon, \delta)$. Since $|Q| \le \frac{n}{n'} \log \frac{2}{\delta} = \mathcal{O}(\frac{n}{n'} \log \frac{1}{\delta})$, Theorem 2 implies that the number of comparisons used by PICK-ANCHOR is

$$
\begin{aligned}
= \frac{2|Q|}{\epsilon^2} \log \frac{2|Q|}{\delta} &= \mathcal{O}\left(\frac{\frac{n}{n'} \log \frac{1}{\delta}}{\epsilon^2} \log \frac{\frac{n}{n'} \log \frac{1}{\delta}}{\delta} \right) \\
&= \mathcal{O}\left(\frac{n}{n'\epsilon^2} \log \frac{1}{\delta} \left(\log \frac{n}{n'\delta} + \log \log \frac{1}{\delta} \right) \right) \\
&= \mathcal{O}\left(\frac{n}{n'\epsilon^2} \log \frac{1}{\delta} \log \frac{n}{n'\delta} \right).
\end{aligned}
$$

We show that w.p.$\ge 1 - \delta$, the output element is an (ϵ, n')-good anchor element. We first show that Q contains atleast one of the top n' elements. We then show that the output element defeats one of the top n' elements with probability $\ge \frac{1}{2} - \epsilon$. Hence by strong stochastic transitivity, the output element defeats every element outside the top n' elements with probability $\ge \frac{1}{2} - \epsilon$.

The probability that Q does not contain an element in top n' elements is $\le \left(1 - \frac{n'}{n}\right)^{\frac{n}{n'} \log \frac{2}{\delta}} \le \frac{\delta}{2}$. Note that the above statement is true even when size of Q is n. Let the best element in Q be denoted as q^*. By Theorem 2, w.p.$\ge 1 - \delta/2$, the output element o of SEQ-ELIMINATE$(Q, \epsilon, \frac{\delta}{2})$ is an ϵ-maximum of Q. Hence $w.p. \ge 1 - \delta/2$, $\tilde{p}_{q^*,o} \le \epsilon$ and therefore by strong stochastic transitivity, for element e worse than q^*, $\tilde{p}_{e,o} \le \tilde{p}_{q^*,o} \le \epsilon$. Since, w.p.$\ge 1 - \delta/2$, the number of elements that are better than q^* is less than n', by the union bound, w.p.$\ge 1 - \delta$, o is an (ϵ, n')-good anchor element. □

A.6 Proof of Lemma 5

We prove the Lemma by dividing it into three parts. We first show that an element e that is ϵ_u better than anchor a i.e., $\tilde{p}_{e,a} \ge \epsilon_u$, is part of PRUNE$(S, a, n', \epsilon_l, \epsilon_u)$ w.p.$\ge 1 - \delta/2$.

Lemma 13. *If $\tilde{p}_{e,a} \ge \epsilon_u$, then w.p.$\ge 1 - \delta/2$, the output set of* PRUNE$(S, a, n', \epsilon_l, \epsilon_u, \delta)$ *contains e.*

Proof. e is not part of output set only if $e \in Q_t$ for some t. Q_t will contain e only if S_t contains e and COMPARE$(e, a, \epsilon_l, \epsilon_u, \frac{\delta}{2^{t+1}})$ returns 1. By Lemma 12, since $\tilde{p}_{e,a} \ge \epsilon_u$, probability that COMPARE$(e, a, \epsilon_l, \epsilon_u, \frac{\delta}{2^{t+1}})$ returns 1 is $\le \frac{\delta}{2^{t+1}}$. Hence the probability that Q_t contains e is $\le \frac{\delta}{2^{t+1}}$ and therefore by the union bound the probability that output set does not contain e is $\le \sum_{t=1}^{\infty} \frac{\delta}{2^{t+1}} \le \frac{\delta}{2}$. □

For $\frac{1}{n} \le \delta \le \frac{n'}{n}$, and if a is a good anchor element, we show that first round of pruning itself will reduce the set size to $2n'$ and hence bound the number of comparisons used by PRUNE.

Lemma 14. *If $n' \ge 8 \log^2 n$, $\frac{1}{n} \le \delta \le \frac{n'}{n}$ and a is an (ϵ_l, n')-good anchor element then w.p.$\ge 1 - \delta/2$,* PRUNE$(S, a, n', \epsilon_l, \epsilon_u, \delta)$ *uses $\mathcal{O}\left(\frac{n}{(\epsilon_u - \epsilon_l)^2} \log \frac{1}{\delta}\right)$ comparisons and outputs a set of size at most $2n'$.*

Proof. If $n' \ge \frac{n}{2}$, the lemma is trivial. So let $n' < \frac{n}{2}$. Let the elements that defeat a with probability $\ge \frac{1}{2} + \epsilon_l$ i.e., elements in set $\{e | \tilde{p}_{e,a} \ge \epsilon_l\}$ be called good elements and the remaining elements be bad elements. Note that the number of bad elements in S_1 is $\ge n - n'$. We show that the number of bad elements in S_2 is $\le n'$. An element e is part of S_2 only if COMPARE$(e, a, \epsilon_l, \epsilon_u, \delta/4)$ returns 2. By Lemma 11, each bad element in S_1 appears in S_2 w.p.$\le \delta/4$. Therefore by the Chernoff bound, the probability that there are more than n' bad elements in S_2 is

$$
\le e^{-(n-n')D(\frac{n'}{n-n'} || \delta/4)} \le e^{-\frac{n}{2} D(\frac{n'}{n} || \delta/4)} \le e^{-\frac{n}{2} \frac{n'}{2n}} = e^{-n'/4} \le \frac{1}{n^2} \le \frac{\delta}{2}.
$$

Since the number of good elements in S_1 is $\le n'$, their size in S_2 is also $\le n'$. Hence w.p.$\ge 1 - \delta/2$, $|S_2| \le 2n'$ and therefore PRUNE stops after first iteration. Noting that PRUNE ran only for one iteration $t = 1$, COMPARE$(e, a, \epsilon_l, \epsilon_u, \delta/4)$ uses $\mathcal{O}(\frac{1}{(\epsilon_u - \epsilon_l)^2} \log \frac{1}{\delta})$ comparisons. □

We now bound the number of comparisons used by PRUNE for higher values of δ by showing that after each round, the number of elements reduces roughly by a factor of δ.

Lemma 15. *If $n' > \sqrt{6n \log n}$, $\delta \ge \frac{n'}{n}$ and a is an (ϵ_l, n')-good anchor element, then w.p.$\ge 1 - \frac{\delta}{2}$,* PRUNE$(S, a, n', \epsilon_l, \epsilon_u, \delta)$ *uses $\mathcal{O}\left(\frac{n}{(\epsilon_u - \epsilon_l)^2} \log \frac{1}{\delta}\right)$ comparisons and outputs a set of size less than $2n'$.*

Proof. As before let the elements that defeat a with probability $\geq \frac{1}{2} + \epsilon_l$ i.e., elements in set $\{e|\tilde{p}_{e,a} \geq \epsilon_l\}$ be called good elements and the remaining elements be bad elements. The number of good elements in S_1 is $\leq n'$ and number of bad elements in S_1 is $\geq n - n'$. We first show that in each iteration the number of bad elements decreases by atleast a factor of δ until it falls below n'. We then bound the number of rounds it takes for number of bad elements to fall below n'. Using this bound on number of rounds, we separately bound the number of comparisons used over bad and good elements.

Note that for every bad element e, COMPARE$(e, a, \epsilon_l, \epsilon_u, \delta')$ outputs 2 with probability $\leq \delta' \leq \delta/4$. Hence, if at the beginning of the round, the number of bad elements is more than n', the probability that number of bad elements does not reduce by at least a factor of δ is

$$\leq e^{-n' D(\delta \| \delta/4)} \leq e^{-n'\delta/2} \leq e^{-\frac{(n')^2}{2n}} \leq \frac{1}{n^3}$$

where the last inequality follows from $n' \geq \sqrt{6n \log n}$.

Now if the number of bad elements reduces by δ after each round, then the number of bad elements falls below n' in $t = 2\log_{\frac{1}{\delta}} \frac{n}{n'} \leq n$ rounds. Thus by the union bound, w.p.$\geq 1 - \frac{1}{n^2}$, the number of bad elements reduces by δ until the size becomes less than n'. Henceforth we assume this and bound the number of comparisons used.

We first bound the number of comparisons taken by PRUNE over bad elements. Number of bad elements in S_t is $\leq n\delta^{t-1}$. Since COMPARE$(e, a, \epsilon_l, \epsilon_u, \delta')$ uses $\frac{2}{(\epsilon_u - \epsilon_l)^2} \log \frac{1}{\delta'}$, the number of comparisons used by PRUNE over bad elements is

$$\begin{aligned}
&\leq \sum_{t=1}^{2\log_{1/\delta} \frac{n}{n'}} \frac{2n\delta^{t-1}}{(\epsilon_u - \epsilon_l)^2} \log \frac{2^{t+1}}{\delta} \\
&\leq \frac{2n}{(\epsilon_u - \epsilon_l)^2} \sum_{t=1}^{2\log_{1/\delta} \frac{n}{n'}} \left(\delta^{t-1} \log \frac{1}{\delta} + (t+1)(\delta)^{t-1} \log 2 \right) \\
&= \mathcal{O}\left(\frac{n}{(\epsilon_u - \epsilon_l)^2} \log \frac{1}{\delta} \right).
\end{aligned}$$

The last equality follows from the fact that if $\delta \leq 1/2$ (if $\delta > 1/2$, we can choose $\delta = 1/2$) then $\sum_t \delta^{t-1}$ and $\sum_t (t+1)\delta^{t-1}$ are bounded.

Now we bound the number of comparisons used by PRUNE over good elements. The number of comparisons used by PRUNE over good elements is

$$\begin{aligned}
&\leq \sum_{t=1}^{2\log_{1/\delta} \frac{n}{n'}} \frac{n'}{(\epsilon_u - \epsilon_l)^2} \log \frac{2^{t+1}}{\delta} \\
&\leq \frac{n'}{(\epsilon_u - \epsilon_l)^2} \sum_{t=1}^{2\log_{1/\delta} \frac{n}{n'}} \left(\log \frac{1}{\delta} + (t+1) \log 2 \right) \\
&\leq \frac{n'}{(\epsilon_u - \epsilon_l)^2} \left(\left(2\log_{1/\delta} \frac{n}{n'} \right) \log \frac{1}{\delta} + \left(2\log_{1/\delta} \frac{n}{n'} \right)^2 \right) \\
&= \mathcal{O}\left(\frac{n}{(\epsilon_u - \epsilon_l)^2} \log \frac{1}{\delta} \right).
\end{aligned}$$ □

Proof of Lemma 5 follows from Lemmas 13, 14 and 15.

A.7 Proof of Theorem 6

We prove the theorem by breaking it into parts. We first show that if anchor element a, the output of PICK-ANCHOR is a $2\epsilon/3$-maximum then w.p.$\geq 1 - \delta/4$, OPT-MAXIMIZE outputs a.

Lemma 16. *If a, the output of step 6 in* OPT-MAXIMIZE(S, ϵ, δ) *is a $\frac{2\epsilon}{3}$-maximum of S, then w.p.$\geq 1 - \frac{\delta}{4}$,* OPT-MAXIMIZE$(S, \epsilon, \delta)$ *outputs a.*

Proof. a is not returned only if COMPARE in step 9 of OPT-MAXIMIZE returns 2. Since a is $\frac{2\epsilon}{3}$-maximum of S, $\tilde{p}_{e,a} \leq \frac{2\epsilon}{3}$, $\forall e \in S$. Then by Lemma 11, the probability that a single call of COMPARE$(e, a, 2\epsilon/3, \epsilon, \frac{\delta}{4n})$ returns 2 is $\leq \frac{\delta}{4n}$. Hence by the union bound, the probability that COMPARE returns 1 for all calls in step 9 of OPT-MAXIMIZE is $\geq 1 - \delta/4$. Therefore the Lemma follows. □

We now bound the number of comparisons used by OPT-MAXIMIZE in steps 1-6 and also prove some properties of PRUNE 's output set and anchor element.

Lemma 17. *For $\delta \geq \frac{1}{n}$, w.p.$\geq 1-\delta/2$, steps 1-6 in* OPT-MAXIMIZE(S,ϵ,δ) *uses $\mathcal{O}\left(\frac{n}{\epsilon^2}\log\frac{1}{\delta}\right)$ comparisons, outputs a set S' of size at most $\sqrt{24n\log n}$ and either a is a $2\epsilon/3$-maximum element or S' contains the absolute maximum element.*

Proof. By Lemma 3, w.p.$\geq 1-\delta/4$, PICK-ANCHOR$(S,\sqrt{6n\log n},\frac{\epsilon}{3},\frac{\delta}{4})$ uses $\mathcal{O}(\frac{\sqrt{n\log n}}{\epsilon^2}\log\frac{1}{\delta})$ comparisons and outputs an $(\epsilon/3,\sqrt{6n\log n})$-good anchor element. From now we assume that a, the output of PICK-ANCHOR$(S,\sqrt{6n\log n},\frac{\epsilon}{3},\frac{\delta}{4})$ is an $(\epsilon/3,\sqrt{6n\log n})$-good achor element.

By Lemma 5, w.p.$\geq 1-\delta/4$, PRUNE$(S,a,\sqrt{6n\log n},\epsilon/3,2\epsilon/3,\delta/4)$ uses $\mathcal{O}(\frac{n}{\epsilon^2}\log\frac{1}{\delta})$ comparisons, outputs a set of size at most $\sqrt{24n\log n}$ and if a is not an $2\epsilon/3$-maximum, then S' contains the absolute maximum.

And the Lemma follows by using the union bound. □

We now bound the number of comparisons used by OPT-MAXIMIZE during steps 8-13 assuming that either anchor element a is $2\epsilon/3$-maximum or S' contains the absolute maximum of S.

Lemma 18. *For $\delta \geq \frac{1}{n}$, if a, the output of step 6 and S', the output of step 7 are such that either a is $2\epsilon/3$-maximum of S or S' contains the absolute maximum element of S, then steps 8-13 of* OPT-MAXIMIZE(S,ϵ,δ) *uses $\mathcal{O}(\frac{|S'|}{\epsilon^2}\log\frac{n}{\delta})$ comparisons and w.p.$\geq 1-\delta/2$, outputs an ϵ-maximum.*

Proof. We first bound the number of comparisons. Each COMPARE$(e,a,2\epsilon/3,\epsilon,\frac{\delta}{4n})$ uses $\mathcal{O}(\frac{1}{\epsilon^2}\log\frac{n}{\delta})$ comparisons and hence over all elements of S', COMPARE uses at most $\mathcal{O}(\frac{|S'|}{\epsilon^2}\log\frac{n}{\delta})$ comparisons. Further SEQ-ELIMINATE$(S',\epsilon,\delta/4)$ uses $\mathcal{O}(\frac{|S'|}{\epsilon^2}\log\frac{n}{\delta})$ comparisons by Theorem 2.

If a is a $\frac{2\epsilon}{3}$-maximum, then the result follows by Lemma 16.

Let a not be an $\frac{2\epsilon}{3}$-maximum. Then S' contains the absolute maximum denoted here by b^*. Notice that by strong stochastic transitivity, an ϵ-maximum of S' is an ϵ-maximum of S since $b^* \in S'$. By Theorem 2, w.p.$\geq 1-\frac{\delta}{4}$, SEQ-ELIMINATE$(S',\epsilon,\delta/4)$ outputs an ϵ-maximum. Now if $\tilde{p}_{b^*,a} > \epsilon$, then w.p.$\geq 1-\frac{\delta}{4n}$, COMPARE$(b^*,a,2\epsilon/3,\epsilon,\frac{\delta}{4n})$ returns 2 and hence a is not returned but SEQ-ELIMINATE$(S',\epsilon,\delta/4)$ is returned. If $\tilde{p}_{b^*,a} \leq \epsilon$, then a is an ϵ-maximum and hence returning a also results in an ϵ-maximum output. Lemma then follows by the union bound. □

Theorem 6 then follows from Theorem 2 and Lemmas 17 and 18.

B Ranking

B.1 Proof sketch for Theorem 7

Proof sketch. Consider the model where $\tilde{p}_{a_1,a_n} = 1/2$, $\tilde{p}_{a_i,a_j} = (0 <)\mu(\ll 1/n^{10})$, when $i < j$ and $(i,j) \neq (n,1)$. This model has an order: $a_1 > a_2 > \cdots > a_{n-1} > a_n$ i.e., $\tilde{p}_{a_i,a_j} > 0\ \forall i < j$. Further this model satisfies strong stochastic transitivity since $\tilde{p}_{a_i,a_k} \geq \max(\tilde{p}_{a_i,a_j},\tilde{p}_{a_j,a_k})\ \forall i<j<k$.

We prove the Lemma by reducing the above model to the model where μ is replaced by 0. Note that new model does not satisfy strong stochastic transitivity but helps us in proving the Lemma.

Note that μ is so small that if we consider a model where we replace μ with 0, the comparisons behave essentially similarly. More formally, let model M_μ be the model we consider and M_0 be the model when μ is replaced with 0. Let S denote a sequence of comparisons where each element of the sequence includes the elements compared and its outcome. Further, for each sequence S, let $P_\mu(S)$ and $P_0(S)$ denote the probability of sequence S under models M_μ and M_0 respectively. Now consider a sequence S of comparisons of length $\leq n^2/20$. Then

$$\frac{P_0(S)}{P_\mu(S)} \geq \left(\frac{1/2}{1/2+\mu}\right)^{n^2/20} \geq e^{-n^2/(10n^{10})} \geq \frac{6}{7}$$

Thus the probability of any sequence of length $\leq \frac{n^2}{20}$ is approximately same under both models. Hence if there is an algorithm that uses $\frac{n^2}{20}$ comparisons and w.p.$\geq 7/8$ produces an $1/4$-ranking under M_μ model then applying same algorithm over M_0 model produces an $1/4$-ranking w.p.$\geq \frac{7}{8}\cdot\frac{6}{7} = \frac{3}{4}$.

We now show that there exists no algorithm that uses $\frac{n^2}{20}$ comparisons and w.p.$\geq \frac{3}{4}$ generates a $1/4$-ranking under M_0, thus proving the Lemma. It is easy to see that any ordering outputted without querying the comparison between a_1 and a_n is a $1/4$-ranking w.p. exactly 1/2 since no order between a_1 and a_n can be deduced. Since the pair (a_1, a_n) is one random pair among $\binom{n}{2}$ pairs, the probability that the algorithm asks a comparison between this pair with $n^2/20$ comparisons is $< \frac{1}{2}$. So the probability that the output order contains a_1 and a_n in the right order is $< \frac{1}{2} + \frac{1}{2} \cdot \frac{1}{2} = \frac{3}{4}$. □

B.2 Ranking Algorithm

We present STRONG-TRANSITIVITY-RANKING that uses $\mathcal{O}(\frac{n^2}{\epsilon^2} \log \frac{n}{\delta})$ comparisons and w.p.$\geq 1-\delta$ outputs an ϵ-ranking. STRONG-TRANSITIVITY-RANKING achieves this by approximating each $\tilde{p}_{i,j}$ with $\hat{p}_{i,j}$ to an additive error of $\frac{\epsilon}{2}$. We first argue that there is an element e such that $\hat{p}_{e,j} \geq \frac{1}{2} - \epsilon/2 \ \forall j \in S$ and such an element is an ϵ-maximum. Observe that if there is any element e such that $\hat{p}_{e,j} \geq \frac{1}{2} - \epsilon/2 \ \forall j \in S$ then $p_{e,j} \geq \frac{1}{2} - \epsilon \ \forall j \in S$ and hence e is an ϵ-maximum of S. Further recall that for the absolute maximum a^*, $\tilde{p}_{a^*,j} \geq \frac{1}{2} \ \forall j \in S$ and hence $\hat{p}_{a^*,j} \geq \frac{1}{2} - \epsilon/2 \ \forall j \in S$. Therefore there will be at least one element e s.t. $\hat{p}_{e,j} \geq \frac{1}{2} - \epsilon/2$ and such an element will be an ϵ-maximum of S. We find one such element, delete it from S and add it to the end of the ordered output set. We continue this process until we run out of elements in S. Since at every step we are adding an ϵ-maximum of the remaining set, the ordered output set will be an ϵ-ranking. We first present a subroutine ESTIMATE-PROBABILITY that compares two elements a and b for $\frac{1}{2\epsilon^2} \log \frac{2}{\delta}$ and w.p.$\geq 1-\delta$ approximates $p(i,j)$ to an additive error of ϵ.

Algorithm 6 ESTIMATE-PROBABILITY

inputs
  element i, element j, bias ϵ, confidence δ.
Compare i and j for $\frac{1}{2\epsilon^2} \log \frac{2}{\delta}$ times.
return Fraction of times i won

Lemma 19. ESTIMATE-PROBABILITY(i,j,ϵ,δ) *uses* $\frac{1}{2\epsilon^2} \log \frac{2}{\delta}$ *comparisons and w.p.*$\geq 1-\delta$ *approximates* $p_{i,j}$ *to an additive error of* ϵ.

Proof. Proof follows from Hoeffding's inequality. □

Algorithm 7 STRONG-TRANSITIVITY-RANKING

inputs
  Set S, bias ϵ, confidence δ
for every pair $\{i,j\}$ such that $i,j \in S$ **do**
  $\hat{p}_{i,j} \leftarrow$ ESTIMATE-PROBABILITY$(i,j,\epsilon/2,\delta/n^2)$
  $\hat{p}_{j,i} \leftarrow 1 - p(i,j)$
end for
ordered set $T \leftarrow \varnothing$
while $|S| > 0$ **do**
  if $\exists e$ s.t. $\hat{p}_{e,f} \geq \frac{1}{2} - \epsilon \ \forall f \in S$ **then**
    Add e at the end of T
    $S = S \setminus \{e\}$
  else
    Add S at the end of T
    return T
  end if
end while
return T

Lemma 20. STRONG-TRANSITIVITY-RANKING(S,ϵ,δ) *uses* $\mathcal{O}(\frac{n^2}{\epsilon^2} \log \frac{n}{\delta})$ *comparisons and w.p.*$\geq 1-\delta$ *returns an* ϵ*-ranking.*

Proof. STRONG-TRANSITIVITY-RANKING calls ESTIMATE-PROBABILITY for $\mathcal{O}(n^2)$ times, once for each pair and each $EP(i,j,\epsilon/2,\delta/n^2)$ uses $\mathcal{O}(\frac{1}{\epsilon^2} \log \frac{n}{\delta})$ comparisons and hence bound on comparisons follow.

W.p.$\geq 1-\delta/n^2$, ESTIMATE-PROBABILITY$(i,j,\epsilon/2,\delta/n^2)$ approximates $p_{i,j}$ with $\hat{p}_{i,j}$ such that $|p_{i,j}-\hat{p}_{i,j}| \leq \frac{\epsilon}{2}$. By the union bound, w.p.$\geq 1-\delta$, $|p_{i,j}-\hat{p}_{i,j}| \leq \frac{\epsilon}{2}$ $\forall i,j,\in S$. From here we assume that $|p_{i,j}-\hat{p}_{i,j}| \leq \frac{\epsilon}{2}$ $\forall i,j,\in S$ and show that the output is an ϵ-ranking. Let S^t denote the set of remaining elements in S after t elements are removed from S. We first show that for $0 \leq t \leq n-1$, there is one element e such that $\hat{p}_{e,j} \geq \frac{1}{2}-\epsilon$ $\forall j \in S^t$ and such an element is an ϵ-maximum of S^t. Observe that if there is an element e such that $\hat{p}_{e,j} \geq \frac{1}{2}-\epsilon/2$ $\forall j \in S^t$ then $p_{e,j} \geq \frac{1}{2}-\epsilon$ $\forall j \in S^t$ and hence e is an ϵ-maximum of S^t. Further recall that for the absolute maximum a^{t*} of S^t, $p_{a^{t*},j} \geq \frac{1}{2}$ $\forall j \in S^t$ and hence $\hat{p}_{a^{t*},j} \geq \frac{1}{2}-\epsilon/2$ $\forall j \in S^t$. Therefore there will be at least one element e s.t. $\hat{p}_{e,j} \geq \frac{1}{2}-\epsilon/2$ and such an element will be an ϵ-maximum of S^t. STRONG-TRANSITIVITY-RANKING deletes one such element from S^t and adds it to the end of the ordered output set. Since for every t, STRONG-TRANSITIVITY-RANKING adds an ϵ-maximum of S^t to the output set, the Lemma follows. ☐

C Borda Scores

C.1 Ranking Algorithm for Borda Scores

Algorithm 8 ESTIMATE-BORDA-SCORE

inputs
  set S, element e, bias ϵ, confidence δ.
Initialize: $w \leftarrow 0$, $\hat{s} \leftarrow \frac{1}{2}$, $m \leftarrow \frac{1}{2\epsilon^2}\log\frac{2}{\delta}$.
for $k=1$ **to** $k=m$ **do**
  Compare e with random element $\in S$
  if e wins **then**
    $w \leftarrow w+1$
  end if
  $\hat{s} = \frac{w}{k}$
end for
return $\hat{s}$

Lemma 21. ESTIMATE-BORDA-SCORE(S,a,ϵ,δ) *uses* $\frac{1}{2\epsilon^2}\log\frac{2}{\delta}$ *comparisons and w.p.*$\geq 1-\delta$ *approximates* $s(a)$ *to an additive error of* ϵ.

Proof. Proof follows from properties of ESTIMATE-BORDA-SCORE and Hoeffding's inequlity. ☐

Algorithm 9 BORDA-RANKING

inputs
  set S, bias ϵ, confidence δ.
Initialize: $b_e \leftarrow \frac{1}{2}$ for all $e \in S$
for element e in S **do**
  $b_e \leftarrow$ ESTIMATE-BORDA-SCORE$(S,e,\frac{\epsilon}{2},\frac{\delta}{n})$
end for
Rank S according to b_e.
return S.

Theorem 22. BORDA-RANKING(S,ϵ,δ) *uses* $\frac{2n}{\epsilon^2}\log\frac{2n}{\delta}$ *comparisons and w.p.*$\geq 1-\delta$ *outputs an* ϵ*-Borda ranking.*

Proof. BORDA-RANKING calls ESTIMATE-BORDA-SCORE for exactly n times and each call of ESTIMATE-BORDA-SCORE$(S,e,\epsilon/2,\delta/n)$ uses $\frac{2}{\epsilon^2}\log\frac{2n}{\delta}$ comparisons and hence the bound on comparisons follows.

Note that w.p.$\geq 1-\delta/n$, ESTIMATE-BORDA-SCORE$(S,e,\epsilon/2,\delta/n)$ approximates the Borda score of e to an additive error of $\epsilon/2$. Let the approximate Borda score of element e be b_e. By the union bound, w.p.$\geq 1-\delta$, BORDA-RANKING(S,ϵ,δ) approximates all Borda scores to an additive error of $\epsilon/2$. From here, we assume that $|b_e - s(e)| \leq \frac{\epsilon}{2}$ and show that ranking based on approximate Borda scores results in an ϵ-Borda ranking.

If an element e appears before element f in the output ranking then $b_e \geq b_f$. Since $|b_e - s(e)| \leq \frac{\epsilon}{2}$ and $|b_f - s(f)| \leq \frac{\epsilon}{2}$, $s(e)-s(f) = (b_e - b_f) + (s(e)-b_e) + (b_f - s(f)) \leq \epsilon$. Hence the Lemma follows. ☐

D Why Knockout Fails

We will show that KNOCKOUT proposed in [14] fails under SST model without stochastic triangle inequality constraint.

Consider the model where $\tilde{p}_{a_1,a_j} = \mu \ \forall j < n/2$, $\tilde{p}_{a_1,a_j} = \frac{1}{2} \ \forall j \geq n/2$ and $\tilde{p}_{a_i,a_j} = \mu \ \forall 1 < i < j$ for some $0 < \mu < \frac{1}{n^{10}}$. Observe that this model satisfies SST but not stochastic triangle inequality. Under this model, a_1 is the absolute maximum and any element in the set $\{a_i | i < n/2\}$ is a $1/4$-maximum. We show that under this model, w.p.$\geq 1/16$, KNOCKOUT$(S, 1/4, 1/16)$ fails to find a $1/4$-maximum.

KNOCKOUT pairs elements randomly in each round and compares each pair for a certain number of times and the winners proceed to the next round until there is only one element left. Observe that in the first round a_1 can get paired with an element from set $\{a_i | 1 < i < n/2\}$ w.p.$\approx 1/2$ and if that happens a_1 can lose the tie w.p.$\approx 1/2$. Hence a_1 can get eliminated in the first round w.p.$\approx 1/4$. Once a_1 is eliminated, in the second round, the elements will be approximately half from the first half of the original set and half from the second half. Since these elements are almost incomparable (comparisons between any two elements is now approximately a Bernoulli random variable with parameter $1/2$), each element is almost equally likely to be the final output. Therefore w.p.$\approx 1/8$, the output can be an element from second half of the set and hence not a $1/4$-maximum.

E Additional Experiment

To show why PAC maximum algorithms could be preferred to absolute maximum algorithms, once again we compare SEQ-ELIMINATE, **KNOCKOUT** and **MallowsMPI** for comparison probability values close to $1/2$. In Figure 4, we consider the stochastic model, $p_{1,j} = 0.6 \ \forall j > 1$ and $p_{i,j} = 0.5 + \tilde{q} \ \forall 1 < i < j$ where $\tilde{q} \ll 0.05$ with $n = 15$. Again, we find a 0.05-maximum with error probability $\delta = 0.1$. From Figure 4, we can observe that performance of **MallowsMPI** gets much worse as $\tilde{q}$ decreases whereas SEQ-ELIMINATE and **KNOCKOUT** do not get affected since they are PAC maxing algorithms. Also observe that SEQ-ELIMINATE performs much better than **KNOCKOUT**.

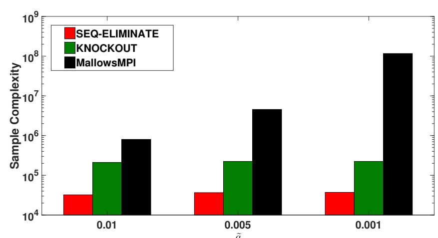

Figure 4: Comparison of Maximum Selection Algorithms for probability values close to $1/2$