Ranking Wily People Who Rank Each Other

Anson Kahng
Computer Science Department
Carnegie Mellon University

Yasmine Kotturi
Human-Computer Interaction Institute
Carnegie Mellon University

Chinmay Kulkarni
Human-Computer Interaction Institute
Carnegie Mellon University

David Kurokawa
Computer Science Department
Carnegie Mellon University

Ariel D. Procaccia
Computer Science Department
Carnegie Mellon University

Abstract

We study rank aggregation algorithms that take as input the opinions of players over their peers, represented as rankings, and output a social ordering of the players (which reflects, e.g., relative contribution to a project or fit for a job). To prevent strategic behavior, these algorithms must be *impartial*, i.e., players should not be able to influence their own position in the output ranking. We design several randomized algorithms that are impartial and closely emulate given (non-impartial) rank aggregation rules in a rigorous sense. Experimental results further support the efficacy and practicability of our algorithms.

1 Introduction

Our work is primarily motivated by two applications. The first is ordering authors on a scientific paper based on contribution. In medical research, for example, author lists tend to be long, and the difference between being listed second or third on an important paper can make or break a career. We would like to employ an algorithm that receives from each author a ranking, which reflects his own opinion regarding the relative contribution of different authors, and outputs an aggregate ranking.

The second application is online labor markets, such as Upwork or Freelancer. In the bigger markets, employers typically receive dozens of applications for a job, but there is an embarrassment of riches, because employers do not have the knowledge required to accurately evaluate applicants. For the past year we have been building a prototype of a new online labor market, where applicants for a job — who are well-suited to evaluate applications for that same job — rank each other.[1] We would like to implement a mechanism that aggregates these rankings into a single ranking that is then shown to the employer.

As the reader has no doubt realized by now, the foregoing applications share a common problem, which gets in the way of applying standard rank aggregation rules: strategic behavior. Specifically, in these relatively high-stakes scenarios, it is likely that a player would try to improve his own

[1]To be more accurate, they submit pairwise comparisons of other applicants, but, as we discuss in Section 7, that is essentially the same setting.

position in the output ranking by manipulating his reported ranking. For example, he might weaken a strong contender for the top position by ranking him last. Our goal, therefore, is to design rank aggregation rules that are *impartial*, in the sense that the position of a player in the output ranking is completely independent of the report of that player.

1.1 Our Approach and Results

On a high level, our approach is to design *randomized* rank aggregation rules that are impartial and closely emulate standard rank aggregation rules that are not impartial. Specifically, we focus on providing impartial approximations to the important class of *pairwise* rules, which, as input, only require information about the fraction of players ranking any one player above another. Our theoretical results crucially depend on the notion of approximation — or measure of error — in question.

In Section 4, we introduce the k-PARTITE algorithm, which, in a nutshell, randomly partitions the players into subsets, builds a probability distribution over the positions of members of one subset based on the aggregate opinion of members of other subsets, and then generates a distribution over rankings that is consistent with these distributions over positions. We prove that k-PARTITE is impartial, and, when used in conjunction with any pairwise rule, it provides small *backward error* with respect to that rule: With high probability, k-PARTITE places each player in the same position that the given pairwise rule would have placed him had the input rankings been slightly perturbed.

In Section 5, we present the COMMITTEE algorithm. It randomly chooses a subset of players, who serve as the eponymous committee. Each committee member is positioned based on the opinions of other committee members, and then all other players are ordered by the committee. The key idea is that, to avoid conflicts and achieve impartiality, each committee member has slots that are reserved for him, and he is inserted into the reserved slot that most closely matches the aggregate opinion of other committee members. We prove that COMMITTEE provides *mixed error* guarantees with respect to any given pairwise rule, that is, with high probability, COMMITTEE places each player in a position that is *close* to where the given pairwise rule would have placed him had the input rankings been slightly perturbed. Taking on some *forward error* — a mismatch between the

positions — allows for improved backward error compared to k-PARTITE.

In Section 6, we empirically measure the performance of our impartial algorithms with respect to the popular *Kemeny rule*, which is defined via a natural optimization objective. The experimental results demonstrate that our impartial algorithms, when coupled with the Kemeny rule, output near-optimal rankings with respect to the Kemeny objective, despite the impartiality constraint.

1.2 Related Work

At this point there is a significant body of work on the design of impartial mechanisms (de Clippel, Moulin, and Tideman 2008; Alon et al. 2011; Holzman and Moulin 2013; Bousquet, Norin, and Vetta 2014; Tamura and Ohseto 2014; Berga and Gjorgjiev 2014; Fischer and Klimm 2014; Mackenzie 2015), including several papers in major AI conferences (Kurokawa et al. 2015; Aziz et al. 2016). We only elaborate on the papers that are most closely related to ours.

The paper of de Clippel et al. (2008) introduced the notion of impartiality, in the context of dividing *credit* for a joint project. Specifically, the output of their mechanism is the fraction of the total credit each player receives, and impartiality means that a player cannot affect his own share of the credit. This mechanism is deployed on the fair division website Spliddit.org, where one of the suggested applications is ordering authors on scientific papers. However, an impartial credit division mechanism does *not* induce an impartial ranking mechanism, because, when players are sorted by credit, a player can improve his own position by decreasing another player's share.

Berga and Gjorgjiev (2014) study the impartial rank aggregation problem from an axiomatic viewpoint, but focus on deterministic rules and a stronger notion of impartiality (which we discuss in Section 7). Their results suggest that deterministic impartial rank aggregation methods are severely limited, and support our focus on randomized algorithms.[2]

On a technical level, our k-PARTITE algorithm is reminiscent of an algorithm of Alon et al. (2011), in that it randomly partitions the players into subsets, and the outcome of players in one subset is only determined by players in other subsets. But the details of the algorithm, and its analysis, are completely different.

2 Preliminaries

In this section we introduce terminology and notations that are standard in computational social choice (Brandt et al. 2016), as well as the formal instantiation of the concept of impartiality in our setting.

2.1 Rankings and Aggregation

For any $k \in \mathbb{N}$, let $[k] = \{1, \ldots, k\}$. Our setting involves a set of players $[n] = \{1, \ldots, n\}$. The opinions of players are represented as rankings over $[n]$, which we think of as *permutations*. Let Π represent the set of all permutations of $[n]$, and let Π^n represent the set of all *input profiles*. For any $\sigma \in \Pi$, let $\sigma(j)$ be the player at position j in σ and let $\sigma^{-1}(i)$ be the position of player i in the ranking σ (where position 1 is the highest and position n is the lowest).

A *deterministic rank aggregation rule* (also known as a *social welfare function*) is a function $f : \Pi^n \to \Pi$, which takes in an input profile and returns a ranking. A *randomized* rank aggregation rule returns a probability distribution over rankings. We sometimes find it convenient to slightly abuse notation and think of the domain of a rank aggregation rule as $\Pi^n \times 2^{[n]}$ — for $\vec{\sigma} = (\sigma_1, \ldots, \sigma_n)$ and $X \subseteq [n]$, $f(\vec{\sigma}, X)$ is the application of the rule to the input profile $(\sigma_i)_{i \in X}$.

2.2 Pairwise Rank Aggregation Rules

An input profile $\vec{\sigma} = (\sigma_1, \ldots, \sigma_n)$ induces a *pairwise comparison matrix* $A(\vec{\sigma})$, where

$$A(\vec{\sigma})_{ij} = \frac{|\{k \in [n] : \sigma_k^{-1}(i) < \sigma_k^{-1}(j)\}|}{n}.$$

In words, the (i, j) entry is the fraction of players who rank i above j. Let Ω be the set of pairwise comparison matrices. Therefore, we can think of $A : \Pi^n \to \Omega$ as a function that takes in an input profile and returns its associated pairwise comparison matrix. As before, we will also use the notation $A(\vec{\sigma}, X)$, for a subset of players $X \subseteq [n]$, to denote the pairwise comparison matrix associated with the rankings of the players in X.

Some rank aggregation rules only require the information encoded in the pairwise comparison matrix to compute their output. Formally, a deterministic *pairwise* rank aggregation rule is a function $f : \Omega \to \Pi$. We denote the class of all deterministic pairwise rules by $\mathcal{P}$.[3]

We pay special attention to two popular pairwise rules:

- *The Borda Rule:* Given $\vec{\sigma} \in \Pi^n$, the score of each player i is $\sum_{j=1}^{n}(n - \sigma_j^{-1}(i))$ (that is, each player awards $n - k$ points to the player in position k), and the players are ranked by non-increasing score. It may not be immediately apparent that Borda is a pairwise rule — proving this well-known fact is left to the curious reader as an easy exercise.

- *The Kemeny Rule:* The *Kendall tau distance* d_{KT} between two rankings $\sigma, \tau \in \Pi$ is the number of pairs of players on which the two rankings disagree. Given $\vec{\sigma} \in \Pi^n$, the Kemeny rule returns a ranking in $\operatorname{argmin}_{\tau \in \Pi} \sum_{i=1}^{n} d_{KT}(\tau, \sigma_i)$. Computing the output of the Kemeny rule is hard, but can be done in practice using integer programming or heuristic algorithms (Conitzer, Davenport, and Kalagnanam 2006).

Other well-known rules, such as Copeland and Maximin, are also pairwise.

2.3 Impartiality

Recall that we are interested in designing rank aggregation rules that are impartial, that is, no player i can affect his probability of being ranked in position j, for all $i, j \in [n]$. Formally:

Definition 2.1. A (possibly randomized) rank aggregation rule f is *impartial* if for all $i \in [n]$, all input profiles $(\sigma_1, \ldots, \sigma_n) \in \Pi^n$, and all $\tilde{\sigma}_i \in \Pi$, it holds that $\vec{x} = \vec{y}$, where x_j is the probability i is ranked in position j in $f(\sigma_1, \ldots, \sigma_{i-1}, \sigma_i, \sigma_{i+1}, \ldots, \sigma_n)$, and y_j is the probability i is ranked in position j in $f(\sigma_1, \ldots, \sigma_{i-1}, \tilde{\sigma}_i, \sigma_{i+1}, \ldots, \sigma_n)$.

In an alternative model, we may assume that each player i has a value v_{ij} for being ranked in position j, and then impartiality would mean no player can affect his *expected* value for the outcome, regardless of his value function. Although this definition may seem weaker than Definition 2.1 at first glance, it is easy to verify that the two definitions are, in fact, equivalent.

3 Measures of Error

Given that our goal is to approximate rank aggregation rules, the measure of error is critical to the statement of the formal problem. To define appropriate notions, we adapt concepts that are standard in scientific computing (e.g., in numerical stability analysis): forward error, backward error, and mixed error. We view these imported definitions as part of our conceptual contribution.

Definition 3.1. Let f be a rank aggregation rule. A rank aggregation rule g is said to have (Δ_P, Δ_F) *forward error* with respect to f if for every input profile $\vec{\sigma} \in \Pi^n$, the probability that for all $i \in [n]$ it holds that

$$\frac{\left| f(\vec{\sigma})^{-1}(i) - g(\vec{\sigma})^{-1}(i) \right|}{n} < \Delta_F$$

is at least $1 - \Delta_P$.

Intuitively, a low amount of forward error implies that every player i is placed near his correct rank (as determined by f) with high probability. Unfortunately, as the next theorem states, impartial rank aggregation rules cannot approximate the Borda rule. Since Borda is a pairwise rule, the theorem rules out the possibility of approximating all pairwise rules.

Theorem 3.2. *For all $n \geq 2$ and $\varepsilon > 0$, there exists no impartial rank aggregation rule g that gives a $(1/2 - \varepsilon, 1/3)$ forward error with respect to the Borda rule f.*

Proof. For $n = 2$, a direct analysis (which we omit) gives the result. Let us therefore consider only the case $n \geq 3$. Let g be an impartial rank aggregation rule.

Suppose we have the input profile $\vec{\sigma}$ where $i \neq 2$ gives the ranking $(i-1, \ldots, n, 1, \ldots, i-2)$. Note that if player 2 continued this trend and gave the ranking $1, \ldots, n$ then all players would have the same Borda score.

Now let us consider player 2 in more depth, and define the probability vector $\vec{x} \in [0, 1]^n$, where x_i denotes the probability player 2 will be in position i when g determines the ranking. By impartiality we know that $\vec{x}$ does not depend on

the ranking of player 2. As $\vec{x}$ is a probability vector, we must have one of the following.

Case 1: The first $\lfloor n/2 \rfloor$ entries of $\vec{x}$ sum to at most $1/2$. In this case, if player 2 has the ranking $(2, 1, 3, 4, 5, \ldots, n)$ in $\vec{\sigma}$, then $f(\vec{\sigma})^{-1}(2) = 1$.

Case 2: The last $\lfloor n/2 \rfloor$ entries of $\vec{x}$ sum to at most $1/2$. In this case, if player 2 has the ranking $(1, 3, 2, 4, 5, \ldots, n)$ in $\vec{\sigma}$, then $f(\vec{\sigma})^{-1}(2) = n$.

In either case, we find that with probability at least $1/2$, g will place 2 in a position at distance at least $\lfloor n/2 \rfloor$ from f's placement. That is, with probability at least $1/2$ we have $\left| f(\vec{\sigma})^{-1}(2) - g(\vec{\sigma})^{-1}(2) \right| \geq \lfloor n/2 \rfloor \geq n/3$, giving at best a forward error of $(1/2, 1/3)$. $\qquad \square$

With this impossibility in hand, we set our sights on an alternate error measure, which is well defined only with respect to *pairwise* rank aggregation rules. For this definition and throughout the paper, we use the Frobenius norm and denote $\|A\|_\infty = \max_{i,j} |A_{i,j}|$.

Definition 3.3. Let $f \in \mathcal{P}$. A rank aggregation rule g is said to have (Δ_P, Δ_B) *backward error* with respect to f if for every input profile $\vec{\sigma} \in \Pi^n$ the probability that for all $i \in [n]$ there exists a matrix $\tilde{A} \in \Omega$ such that

1. $\left\| A(\vec{\sigma}) - \tilde{A} \right\|_\infty < \Delta_B$, and

2. $f(\tilde{A})^{-1}(i) = g(\vec{\sigma})^{-1}(i)$,

is at least $1 - \Delta_P$.

Intuitively, a low amount of backward error implies that every player i is placed in a rank that had the players altered their opinions slightly, i would be in the correct rank (according to f) with high probability. See Section 7 for further discussion of this concept.

Finally, we define a third measure of error, which, in a sense, is a union of the two previous notions.

Definition 3.4. Let $f \in \mathcal{P}$. A rank aggregation rule g is said to have $(\Delta_P, \Delta_B, \Delta_F)$ *mixed error* with respect to f if for every input profile $\vec{\sigma} \in \Pi^n$, the probability that for all $i \in [n]$ there exists a matrix $\tilde{A} \in \Omega$ such that

1. $\left\| A(\vec{\sigma}) - \tilde{A} \right\|_\infty < \Delta_B$, and

2. $\frac{\left| f(\tilde{A})^{-1}(i) - g(\vec{\sigma})^{-1}(i) \right|}{n} < \Delta_F$,

is at least $1 - \Delta_P$.

4 The k-PARTITE Algorithm

We now introduce and analyze our first impartial rule, k-PARTITE, which is formally given as Algorithm 1. As it appears somewhat opaque, it is best to understand its ideas when we assume that all the X_i are the same size, i.e., k divides n, $|X_i| = n/k$, and $\gamma_i = k$ for all $i \in [k]$. Slight adjustments are made when this is not the case, which for purposes of intuition can be safely ignored.

First, players are randomly split into k groups of equal size $X_1, \ldots, X_k$, and then each such group separately ranks all n players producing rankings τ_i. The crux of the algorithm is the construction of the matrix Z, which, in turn, is

the sum of $Z^{(i)}$ matrices. Intuitively, the $Z^{(i)}$ matrix represents X_i's contribution to Z, and its (a, b) entry indicates the probability that a should be placed in position b overall. Specifically, each player not in X_i is placed in his exact position dictated by τ_i with probability $1/k$, and in all positions that the players in X_i themselves were assigned to in τ_i with probability $1/(n(k-1))$. This information is encoded as the only non-zero entries in $Z^{(i)}$ — each column then sums to $1/k$, each row representing a player in X_i is zero, and all other rows sum to $1/(k-1)$. As we show in Appendix A, Z is doubly stochastic (its rows and columns sum to 1); hence we can apply the Birkhoff-von Neumann Theorem (Birkhoff 1946; von Neumann 1953) to sample from this distribution and remain faithful to the probabilities.

---

input: $f \in \mathcal{P}$ and $\vec{\sigma} \in \Pi^n$

1: Randomly split all n players into k groups $X_1, \ldots, X_k$ where $|X_i| \in \{\lfloor n/k \rfloor, \lceil n/k \rceil\}$
2: **for** $i = 1, \ldots, k$ **do**
3:     $\tau_i \leftarrow f(\vec{\sigma}, X_i)$
4:     $\gamma_i \leftarrow n/|X_i|$
5:     Let $Z^{(i)} \in \mathbb{R}^{n \times n}$ where

$$Z^{(i)}_{a,b} \leftarrow \begin{cases} \frac{1}{\gamma_i} & \text{if } a \notin X_i \text{ and } \tau_i(b) = a \\ \frac{1}{\gamma_i(\gamma_i-1)|X_i|} & \text{if } a \notin X_i \text{ and } \tau_i(b) \in X_i \\ 0 & \text{otherwise} \end{cases}$$

6: **end for**
7: $Z \leftarrow \sum_{i \in [k]} \frac{\frac{n}{|X_i|} - 1}{k - 1} Z^{(i)}$
8: Sample a ranking σ such that a is ranked in position b with probability $Z_{a,b}$
9: **return** σ

Algorithm 1: k-PARTITE

Our goal is to prove the following theorem, which states the guarantees of k-PARTITE.

Theorem 4.1. k-PARTITE *is impartial, and, for every $f \in \mathcal{P}$ and $\vec{\sigma} \in \Pi^n$, if $k = \lfloor (n/\ln n)^{1/3} \rfloor$, it gives at most*

$$(4/k, 4/k) \in \left(O\left(\left(\frac{\ln n}{n} \right)^{1/3} \right), O\left(\left(\frac{\ln n}{n} \right)^{1/3} \right) \right)$$

backward error with respect to f.

Note that, in particular, the error goes to 0 as n grows. Turning to the proof, it is clear that the algorithm is impartial because of the inability of any player i to affect the i^{th} row of the Z matrix. We therefore focus on establishing the error bound. To this end, we first prove several lemmas.

Lemma 4.2. *If t players $X = \{x_1, \ldots, x_t\}$ are sampled without replacement from $[n]$ with input profile $\vec{\sigma} \in \Pi^n$, then*

$$\mathbb{P}\left[\|A(\vec{\sigma}) - A(\vec{\sigma}, X)\|_\infty \geq \varepsilon \right] < n^2 \exp\left(-\frac{t\varepsilon^2}{2} \right).$$

Proof.

$$\mathbb{P}\left[\|A(\vec{\sigma}) - A(\vec{\sigma}, X)\|_\infty \geq \varepsilon \right]$$

$$\leq \sum_{i<j} \mathbb{P}\left[|A(\vec{\sigma})_{i,j} - A(\vec{\sigma}, X)_{i,j}| \geq \varepsilon \right]$$

$$\leq \sum_{i<j} 2 \exp\left(-\frac{t\varepsilon^2}{2} \right) = 2\binom{n}{2} \exp\left(-\frac{t\varepsilon^2}{2} \right)$$

$$< n^2 \exp\left(-\frac{t\varepsilon^2}{2} \right),$$

where the first transition follows from the union bound, and the second transition follows from Hoeffding's Inequality. $\square$

Lemma 4.3. *For every $f \in \mathcal{P}$, $\vec{\sigma} \in \Pi^n$, and $\varepsilon > 0$, k-PARTITE gives at most*

$$\left(1 - \left(\frac{k-2}{k-1} \right) \left(1 - n^2 k \exp\left(-\frac{\lfloor n/k \rfloor \varepsilon^2}{2} \right) \right), \varepsilon \right)$$

backward error to f.

Proof. Observe that

$$\mathbb{P}\left[\exists i \in [k] \text{ s.t. } \|A(\vec{\sigma}) - A(\vec{\sigma}, X_i)\|_\infty \geq \varepsilon \right]$$

$$\leq \sum_{i=1}^{k} \mathbb{P}\left[\|A(\vec{\sigma}) - A(\vec{\sigma}, X_i)\|_\infty \geq \varepsilon \right]$$

$$\leq \sum_{i=1}^{k} n^2 \exp\left(-\frac{|X_i| \varepsilon^2}{2} \right)$$

$$\leq \sum_{i=1}^{k} n^2 \exp\left(-\frac{\lfloor n/k \rfloor \varepsilon^2}{2} \right)$$

$$= n^2 k \exp\left(-\frac{\lfloor n/k \rfloor \varepsilon^2}{2} \right),$$

where the second inequality follows from Lemma 4.2.

Further observe that, by Lines 5 and 7 of k-PARTITE, for any player a he is placed directly where one of the X_i places him with probability

$$\sum_{i \in [k]:\ a \notin X_i} \frac{\frac{n}{|X_i|} - 1}{k - 1} \frac{1}{\gamma_i} = \sum_{i \in [k]:\ a \notin X_i} \frac{\frac{n}{|X_i|} - 1}{k - 1} \frac{1}{\frac{n}{|X_i|}}$$

$$= 1 - \frac{1}{n(k-1)} \sum_{i \in [k]:\ a \notin X_i} |X_i|$$

$$\geq 1 - \frac{1}{n(k-1)} \sum_{i=1}^{k} |X_i|$$

$$= 1 - \frac{1}{n(k-1)} n$$

$$= \frac{k-2}{k-1}.$$

Now, if for all $i \in [n]$, $\|A(\vec{\sigma}) - A(\vec{\sigma}, X_i)\|_\infty < \varepsilon$, and each player a is placed in the position that some X_{i_a} places him, then we can set $\tilde{A} = A(\vec{\sigma}, X_{i_a})$ for all $a \in [n]$ to satisfy the conditions of backward error. Moreover, these events are independent. We conclude that k-PARTITE gives at most

$$\left(1 - \left(\frac{k-2}{k-1} \right) \left(1 - n^2 k \exp\left(-\frac{\lfloor n/k \rfloor \varepsilon^2}{2} \right) \right), \varepsilon \right)$$

backward error, as stated. $\qquad\square$

Proof of Theorem 4.1. From Lemma 4.3 it suffices to show that if we have $\varepsilon = 4/k$,

$$1 - \left(\frac{k-2}{k-1}\right)\left(1 - n^2 k \exp\left(-\frac{\lfloor n/k \rfloor \varepsilon^2}{2}\right)\right) \leq \frac{4}{k}.$$

Observe that

$$
\begin{aligned}
n^2 k \exp&\left(-\frac{\lfloor n/k \rfloor \varepsilon^2}{2}\right) \\
&\leq n^2 k \exp\left(-\frac{(n/k-1)\varepsilon^2}{2}\right) \\
&= n^2 k \exp\left(-\frac{(n/k-1)(4/k)^2}{2}\right) \\
&= n^2 k \exp\left(\frac{8}{k^2}\right)\exp\left(-\frac{8n}{k^3}\right) \\
&\leq n^2 \left(n^{1/3}\right)\exp\left(\frac{8}{2^2}\right)\exp\left(-\frac{8n}{\frac{n}{\ln n}}\right) \\
&= e^2 n^{-17/3} \\
&\leq n^{-2}.
\end{aligned}
$$

Thus we see that

$$
\begin{aligned}
1 &- \left(\frac{k-2}{k-1}\right)\left(1 - n^2 k \exp\left(-\frac{\lfloor n/k \rfloor \varepsilon^2}{2}\right)\right) \\
&= \frac{1}{k-1} + \left(\frac{k-2}{k-1}\right)\left(n^2 k \exp\left(-\frac{\lfloor n/k \rfloor \varepsilon^2}{2}\right)\right) \\
&\leq \frac{1}{k-1} + n^{-2} \\
&\leq 2/k + 2/k \\
&= 4/k.
\end{aligned}
$$

$\qquad\square$

A natural question is why we insist on what appears to be a somewhat convoluted algorithm, instead of a more natural approach such as the impartial NAIVE-BIPARTITE, formally given as Algorithm 2. The reason is that this algorithm does not even guarantee tolerable *mixed* error in general. Indeed, consider $f \in \mathcal{P}$ that is defined as follows. Let $X \subseteq \{2, \ldots, n\}$ be the set of players such that at least one player ranks i above 1; return the ranking starting with the players of X ordered lexicographically, followed by the players of $[n] \setminus (\{1\} \cup X)$ ordered lexicographically, and player 1 inserted into position $\lfloor n/3 \rfloor$ overall (shifting appropriately). Now consider the input profile where i reports the ranking $(i, 1, 2, \ldots, i-1, i+1, \ldots, n)$. Then NAIVE-BIPARTITE will always return a ranking where player 1 is placed first or second — as he will always top his set. This means that the algorithm cannot even provide a mixed error of $(1/2, 1, 1/4)$.

5 The COMMITTEE Algorithm

k-PARTITE demonstrates that there exist impartial rules that accurately imitate any $f \in \mathcal{P}$. Observe, however, that the

---

input: $f \in \mathcal{P}$ and $\vec{\sigma} \in \Pi^n$

1: Randomly split the n players into two sets X and Y where $|X| = \lceil \frac{n}{2} \rceil$ and $|Y| = \lfloor \frac{n}{2} \rfloor$
2: $\tau_1 \leftarrow f(\vec{\sigma}, X)$ restricted to the players only in Y
3: $\tau_2 \leftarrow f(\vec{\sigma}, Y)$ restricted to the players only in X
4: σ interlaces τ_1 and τ_2, that is,

$$\sigma(i) \leftarrow \begin{cases} \tau_1\left((i+1)/2\right) & \text{if } i \text{ is odd} \\ \tau_2\left(i/2\right) & \text{if } i \text{ is even} \end{cases}$$

5: **return** σ

Algorithm 2: NAIVE-BIPARTITE

---

algorithm is somewhat hamstrung by the fact that a player must be (with high probability) ranked in *exactly* the location that a small perturbation of the input rankings would give.

To allow more flexibility, we focus on mixed error, and consider COMMITTEE, given as Algorithm 3. Intuitively, this algorithm selects a random committee $X = \{x_1, \ldots, x_k\}$, which then determines the entire ranking. First, for each committee member x_i, we determine their rank using only the rankings given by the remaining $k - 1$ members. However, as directly placing each committee member in this fashion may cause collisions (i.e., multiple members may be assigned the same rank) we restrict placement of x_i to only the positions $i, i + k, i + 2k, \ldots$ Specifically, we assign x_i to the closest such position to the rank given to x_i by the other committee members. There are then k of the n positions assigned. Second, the committee ranks all of the n players, and the non-committee members are placed in the order ranked by the committee in the remaining $n - k$ slots.

---

input: $f \in \mathcal{P}$ and $\vec{\sigma} \in \Pi^n$

1: Randomly select a subset $X = \{x_1, \ldots, x_k\} \subseteq [n]$
2: **for** $i = 1, \ldots, k$ **do**
3:    $c \leftarrow \arg\min_{j \in \{i, i+k, \ldots\}} \left| j - f\left(\vec{\sigma}, X \setminus \{x_i\}\right)^{-1}(x_i) \right|$
4:    $\sigma(c) \leftarrow x_i$
5: **end for**
6: $\tau \leftarrow f(\vec{\sigma}, X)$
7: $j \leftarrow 1$
8: **for** $i = 1, \ldots, n$ **do**
9:    **if** $\tau(i) \notin X$ **then**
10:      **while** $\sigma(j)$ is occupied **do**
11:        $j \leftarrow j + 1$
12:      **end while**
13:      $\sigma(j) \leftarrow \tau(i)$
14:    **end if**
15: **end for**
16: **return** σ

Algorithm 3: COMMITTEE

The algorithm yields the following guarantees

Theorem 5.1. COMMITTEE *is impartial, and, for every $f \in \mathcal{P}$, $\vec{\sigma} \in \Pi^n$, and $\varepsilon > 0$, if $k = 1 + \frac{2}{\varepsilon^2} \ln\left(\frac{n^3}{\varepsilon}\right)$, it gives at most $(\varepsilon, \varepsilon, (k+1)/n)$ mixed error with respect to f.*

Importantly, this theorem allows for an incomparable error to Theorem 4.1. That is, we can reduce the backward error so long as we are willing to take on some forward error. For example, setting ε appropriately gives at most $\left(n^{-2/5}, n^{-2/5}, 2/n + (34/5)n^{-1/5}\ln n\right)$ mixed error.

As in the case of k-PARTITE, the impartiality of COMMITTEE is obvious, because the position of each committee member is only determined by other committee members, and the position of non-committee members is determined by committee members. The proof of Theorem 5.1 therefore focuses on establishing the stated mixed error guarantee; it is relegated to the full version of the paper.[4]

6 Experiments

Our theoretical results indicate that impartial rules like k-PARTITE and COMMITTEE are likely to yield rankings that are close, in a sense (backward error or mixed error), to the output of a given rank aggregation rule. In this section we investigate a more natural metric, which is beyond the reach of our theory, and empirically demonstrate that our rules perform well with respect to this metric, too.

In our experiments, we focus on the Kemeny rule (see Section 2), as it is defined via an optimization problem, so we can use its objective function as our measure. Specifically, we interpret the Kemeny rule as maximizing the number of *agreements* with the input rankings, that is, given $\vec{\sigma} \in \Pi^n$, it chooses a ranking $\tau \in \Pi$ that maximizes

$$\text{Kem}(\tau, \vec{\sigma}) = \sum_{i=1}^{n}\left(\binom{m}{2} - d_{KT}(\tau, \sigma_i)\right). \quad (1)$$

We quantify the error of an impartial rule by comparing how well it does with respect to measure (1) with the performance of the optimal ranking returned by the Kemeny rule. In more detail, let f be the Kemeny rule, and let g be an impartial rule; given an input profile $\vec{\sigma}$, we are interested in the *Kemeny approximation ratio* $\text{Kem}(g(\vec{\sigma}), \vec{\sigma})/\text{Kem}(f(\vec{\sigma}), \vec{\sigma})$; this ratio is upper-bounded by 1 due to the definition of the Kemeny rule. We use the number of agreements, instead of the number of disagreements, as our measure because in cases where the number of disagreements is very small, the ratio would be misleadingly large.

The input profiles are generated according to the popular Mallows (1957) model. In this model, there is a base ranking of the alternatives τ^*, and rankings are drawn i.i.d. from a probability distribution over Π, defined by

$$\Pr[\sigma \mid \tau^*] = \frac{\phi^{d_{KT}(\sigma, \tau^*)}}{\sum_{\sigma' \in \Pi} \phi^{d_{KT}(\sigma', \tau^*)}},$$

for a *dispersion* parameter $\phi \in [0, 1]$. Note that $\phi = 1$ corresponds to uniformly random rankings (and therefore input rankings disagree on pairs of alternatives with probability $1/2$), whereas $\phi = 0$ means, by convention, that all rankings coincide with the base ranking, that is, there is unanimous agreement. We empirically study the Kemeny approximation ratio of the impartial rules NAIVE-BIPARTITE, COMMITTEE, and k-PARTITE, for multiple values of ϕ, each of which represents a different level of agreement.

---
[4]Available from http://procaccia.info/research.

Throughout our experiments, we let $k = n/4, n/8$ for COMMITTEE and let $k = 4, 8$ for k-PARTITE. The intuition behind these choices is that the size of the initial committee and each subset in the partition should grow with n, and choosing these values of k works reasonably in practice. We ran experiments with $n \in \{8, 16, 24, 32, 40\}$ players and $\phi \in \{0, 0.1, 0.3, 0.5, 0.7, 0.9, 1\}$.

Our results for the three impartial rules — NAIVE-BIPARTITE, COMMITTEE, and k-PARTITE — and $\phi = 0.3$ are shown in Figures 1(a), 1(b), and 1(c), respectively; the results for other values of ϕ, which can be found in the full version of the paper, are qualitatively similar. In each figure, the x axis shows the values of n, and the y axis shows the Kemeny approximation ratio obtained by the impartial rule. All three impartial rules perform well as n increases. As a baseline, a straightforward calculation shows that with $\phi = 0.3$, a ranking drawn from the Mallows model would agree with the base ranking (which is typically the output of Kemeny when the input profile is drawn from Mallows) with probability 0.77 on any given pair of alternatives, and that probability is 0.5 if the latter ranking is replaced with a uniformly random ranking, so the (impartial) rule that chooses a ranking uniformly at random gives a Kemeny approximation ratio of $0.5/0.77 = 0.65$.

On a high level, the three impartial rules achieve excellent Kemeny approximations despite their very different theoretical guarantees. But k-PARTITE has by far the highest variance. This phenomenon is due to the fact that if the position of a player is chosen from a setting in which he was not placed in his exact place as prescribed by players in some partition, he is essentially placed in a random location in the final ranking. In NAIVE-BIPARTITE and COMMITTEE, this is not an issue, as players are always placed in some sense close to a position in which a subset of players believes they belong. As can be seen in the full version of the paper, this phenomenon is more pronounced for lower values of ϕ because placing players in a random location is penalized more heavily when the population is generally sharply clustered around a certain ranking.

Furthermore, the performance of our impartial algorithms depends on the specific choice of k (except for NAIVE-BIPARTITE, which does not depend on any parameter k). The following observations can be seen in the full version of the paper. COMMITTEE performs better with the smaller value of k, but this effect lessens as ϕ increases (i.e., it helps more when players have generally similar opinions). With small ϕ, because most committees will agree on a consistent ranking, the additional error from inserting players into larger buckets leads to a noticeable difference. However, as players start to disagree more, the benefit of getting better estimates from larger committees counteracts the insertion error. k-PARTITE acts similarly: the larger value of k leads to better performance for low ϕ, but this effect again lessens as ϕ increases. This is because each player is placed exactly where one of the groups places him with probability $(k-2)/(k-1)$, which increases with k, but being placed in one of these positions is most beneficial when players generally agree. As opinions become increasingly random and diffuse, groups disagree more strongly about where to place
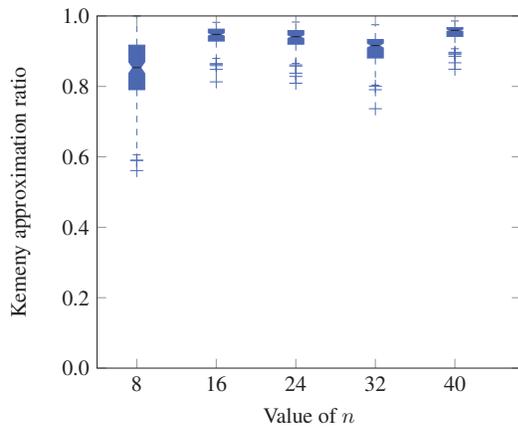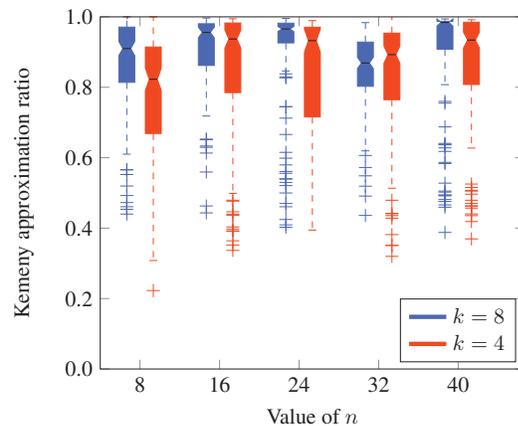
(a) NAIVE-BIPARTITE
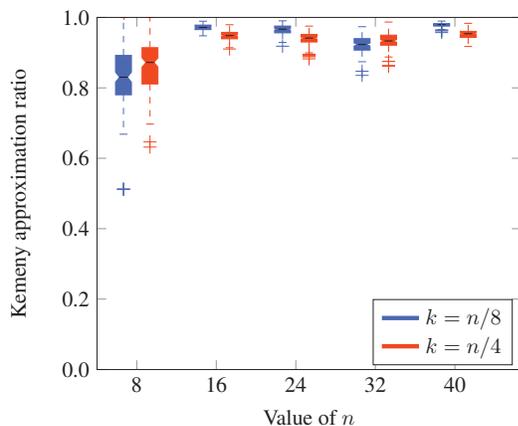
(b) COMMITTEE

(c) k-PARTITE

Figure 1: Kemeny approximation ratio of three impartial mechanisms for $\phi = 0.3$. The median of each boxplot is marked with a black line, the edges of each box denote the quartile values, and the whiskers extend to data within 1.5 times the interquartile range from the edges of each box.

a specific player.

7 Discussion

We conclude with a discussion of several important points that were referenced earlier.

Pairwise comparisons. In practice, players would often be asked to compare pairs of players, rather than giving a complete ranking. Crucially, our results seamlessly extend to that setting. Indeed, standard measure concentration inequalities show that a small number of random caparisons are sufficient to accurately estimate the pairwise comparison matrix (see, e.g., Procaccia and Shah 2016). Because we focus on pairwise rank aggregation rules, this means that the input to the rule is qualitatively unaffected by the transition from complete rankings to pairwise comparisons.

Strong impartiality. A natural notion of impartiality that is stronger than ours (call it *strong impartiality*) requires that a player would not be able to affect the *subset* of players that are ranked above himself. In the context of online labor markets, for example, the rationale is that an employer is more likely to select the applicant in position k if the applicants in positions $1, \ldots, k-1$ are relatively weak, so it is not just the applicant's position that determines his chances of getting the job. Unfortunately, strong impartiality seems too stringent to admit reasonable rules. In fact, we can prove that no strongly impartial rule can give a $(1/3, 1/7, 1/10)$ *mixed* error with respect to Borda, but this statement requires the additional assumption that strongly impartial randomized rules are distributions over strongly impartial deterministic rules.[5] We believe that a similar statement holds without the additional assumption.

Flipping the quantifiers. Our definitions of backward error and mixed error include the words "for all $i \in [n]$ there exists a matrix $\tilde{A} \in \Omega$." That is, for each player we can find a pairwise comparison matrix close to the original one, such that our rule puts i in a position that is identical or close to that in which the given rule would put i on the input $\tilde{A}$. The definitions would be even more compelling if the quantifiers were flipped, i.e., "there exists a matrix $\tilde{A} \in \Omega$ such that for all $i \in [n]$." Under this alternative formulation, we are not allowed to tailor $\tilde{A}$ to i, but rather there is one pairwise comparison matrix that achieves the desired property for every i. It remains an open problem whether our theorems (or variants thereof) still hold under these more demanding notions of error, and, if not, whether these notions are feasible at all.

Acknowledgments

---

[5]Note that COMMITTEE and NAIVE-BIPARTITE can be represented as distributions over impartial (not strongly impartial) deterministic rules, but k-PARTITE cannot.

References

Alon, N.; Fischer, F.; Procaccia, A. D.; and Tennenholtz, M. 2011. Sum of us: Strategyproof selection from the selectors. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 101–110.

Aziz, H.; Lev, O.; Mattei, N.; Rosenschein, J. S.; and Walsh, T. 2016. Strategyproof peer selection: Mechanisms, analyses, and experiments. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, 397–403.

Berga, D., and Gjorgjiev, R. 2014. Impartial social rankings. Manuscript.

Birkhoff, G. 1946. Three observations on linear algebra. *Universidad Nacional de Tucumán, Revista A* 5:147–151.

Bousquet, N.; Norin, S.; and Vetta, A. 2014. A near-optimal mechanism for impartial selection. In *Proceedings of the 10th Conference on Web and Internet Economics (WINE)*, 133–146.

Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.

Conitzer, V.; Davenport, A.; and Kalagnanam, H. 2006. Improved bounds for computing Kemeny rankings. In *Proceedings of the 21st AAAI Conference on Artificial Intelligence (AAAI)*, 620–626.

de Clippel, G.; Moulin, H.; and Tideman, N. 2008. Impartial division of a dollar. *Journal of Economic Theory* 139:176–191.

Fischer, F., and Klimm, M. 2014. Optimal impartial selection. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, 803–820.

Holzman, R., and Moulin, H. 2013. Impartial nominations for a prize. *Econometrica* 81(1):173–196.

Kurokawa, D.; Lev, O.; Morgenstern, J.; and Procaccia, A. D. 2015. Impartial peer review. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 582–588.

Mackenzie, A. 2015. Symmetry and impartial lotteries. *Games and Economic Behavior* 94:15–28.

Mallows, C. L. 1957. Non-null ranking models. *Biometrika* 44:114–130.

Moulin, H. 2014. Personal communication.

Procaccia, A. D., and Shah, N. 2016. Optimal aggregation of uncertain preferences. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, 608–614.

Tamura, S., and Ohseto, S. 2014. Impartial nomination correspondences. *Social Choice and Welfare* 43:47–54.

von Neumann, J. 1953. A certain zero-sum two-person game equivalent to the optimal assignment problem. In Kuhn, W., and Tucker, A. W., eds., *Contributions to the Theory of Games*, volume 2. Princeton University Press. 5–12.

A Z is Doubly Stochastic

In this appendix, we verify that the matrix Z constructed in k-PARTITE (Algorithm 1) is indeed doubly stochastic.

Let us first consider the row sums of the $Z^{(i)}$. Denote the sum of the a^{th} row by

$$S_a^{(i)} = \sum_{b=1}^{n} Z_{ab}^{(i)}.$$

If $a \in X_i$, clearly $S_a^{(i)} = 0$. Otherwise for $a \notin X_i$, we find that

$$S_a^{(i)} = \frac{1}{\gamma_i} + |X_i| \frac{1}{\gamma_i(\gamma_i - 1)|X_i|} = \frac{1}{\gamma_i - 1}.$$

We therefore find that the a^{th} row of Z sums to

$$\begin{aligned}
\sum_{i \in [k]} \left(\frac{\frac{n}{|X_i|} - 1}{k - 1} \right) S_a^{(i)} &= \sum_{i \in [k]:\, a \notin X_i} \frac{\frac{n}{|X_i|} - 1}{k - 1} \frac{1}{\gamma_i - 1} \\
&= \sum_{i \in [k]:\, a \notin X_i} \frac{\gamma_i - 1}{k - 1} \frac{1}{\gamma_i - 1} \\
&= \sum_{i \in [k]:\, a \notin X_i} \frac{1}{k - 1} \\
&= 1.
\end{aligned}$$

Next we consider the column sums of the $Z^{(i)}$, denoted $T^{(i)}_b$. If $\tau_i(b) \notin X_i$ we have that the column has only one non-zero entry with a value of $1/\gamma_i = |X_i|/n$. Otherwise, if $\tau_i(b) \in X_i$, we find that

$$\begin{aligned}
T_b^{(i)} &= (n - |X_i|) \frac{1}{\gamma_i(\gamma_i - 1)|X_i|} \\
&= (\gamma_i |X_i| - |X_i|) \frac{1}{\gamma_i(\gamma_i - 1)|X_i|} \\
&= \frac{1}{\gamma_i} = \frac{|X_i|}{n}.
\end{aligned}$$

Therefore, all columns of $Z^{(i)}$ sum to $|X_i|/n$, and we find that the b^{th} column of Z sums to

$$\begin{aligned}
\sum_{i \in [k]} \left(\frac{\frac{n}{|X_i|} - 1}{k - 1} \right) T_b^{(i)} &= \sum_{i \in [k]} \frac{\frac{n}{|X_i|} - 1}{k - 1} \frac{|X_i|}{n} \\
&= \frac{1}{k - 1} \sum_{i \in [k]} \left(1 - \frac{|X_i|}{n} \right) \\
&= \frac{1}{k - 1} \left(k - \sum_{i \in [k]} \frac{|X_i|}{n} \right) \\
&= \frac{1}{k - 1} (k - 1) \\
&= 1.
\end{aligned}$$