
Ranking from Stochastic Pairwise Preferences: Recovering Condorcet Winners and Tournament Solution Sets at the Top

Arun Rajkumar
Suprovat Ghoshal
Indian Institute of Science

ARUN_R@CSA.IISC.ERNET.IN
SUPROVAT.GHOSHAL@CSA.IISC.ERNET.IN

Lek-Heng Lim
University of Chicago

LEKHENG@GALTON.UCHICAGO.EDU

Shivani Agarwal
Indian Institute of Science

SHIVANI@CSA.IISC.ERNET.IN

Abstract

We consider the problem of ranking n items from stochastically sampled pairwise preferences. It was shown recently that when the underlying pairwise preferences are acyclic, several algorithms including the Rank Centrality algorithm, the Matrix Borda algorithm, and the SVM-RankAggregation algorithm succeed in recovering a ranking that minimizes a global pairwise disagreement error (Rajkumar and Agarwal, 2014). In this paper, we consider settings where pairwise preferences can contain cycles. In such settings, one may still like to be able to recover ‘good’ items at the top of the ranking. For example, if a Condorcet winner exists that beats every other item, it is natural to ask that this be ranked at the top. More generally, several tournament solution concepts such as the top cycle, Copeland set, Markov set and others have been proposed in the social choice literature for choosing a set of winners in the presence of cycles. We show that existing algorithms can fail to perform well in terms of ranking Condorcet winners and various natural tournament solution sets at the top. We then give alternative ranking algorithms that provably rank Condorcet winners, top cycles, and other tournament solution sets of interest at the top. In all cases, we give finite sample complexity bounds for our algorithms to recover such winners. As a by-product of our analysis, we also obtain an improved sample complexity bound for the Rank Centrality algorithm to recover an optimal ranking under a Bradley-Terry-Luce (BTL) condition, which answers an open question of Rajkumar and Agarwal (2014).

1. Introduction

There has been much interest in recent years in designing algorithms for aggregating pairwise preferences to rank a set of items (Fürnkranz & Hüllermeier, 2010; Ailon, 2011; Lu & Boutilier, 2011; Jiang et al., 2011; Jamieson & Nowak, 2011; Negahban et al., 2012; Osting et al., 2013; Wauthier et al., 2013; Busa-Fekete et al., 2014a;b; Rajkumar & Agarwal, 2014). Indeed, the need for aggregating pairwise preferences arises in many domains where fully ordered preferences may be hard to obtain, e.g. in customer surveys, recommender systems, sports team rankings etc.

In many such settings, particularly those related to applications such as customer surveys and recommender systems, it is common to assume that the pairwise comparisons follow some (unknown) underlying statistical model. An important question that arises then is, as we observe more and more data from the underlying statistical model, do the algorithms used converge to a ‘good’ ranking solution?

Recently, Rajkumar & Agarwal (2014) showed that when the underlying pairwise preferences are acyclic, several algorithms including the Rank Centrality (RC), Matrix Borda (MB) and SVM-RankAggregation (SVM-RA) algorithms can converge to an optimal ranking solution in terms of minimizing a pairwise disagreement error. However, in practice, pairwise preferences often contain cycles. In such settings, minimizing the pairwise disagreement error is typically NP-hard, but one may still want to converge to a ranking that places ‘good’ items at the top.

In this paper, we take ‘good’ items to refer to Condorcet winners (which beat all other items) when they exist, and more generally, to tournament solution sets such as top cycles and Copeland and Markov sets, which have been used to define ‘winners’ in the computational social choice literature (Moulin, 1986; De Donder et al., 2000; Laslier, 1997; Brandt et al., 2015). We show that when preferences contain cycles, the RC, MB and SVM-RA algorithms can fail to rank such ‘good’ items at the top; we then propose three

Table 1. Ranking behavior of various algorithms for ranking from stochastic pairwise preferences under various conditions on the underlying statistical model. Here \mathcal{P}^{BTL} is the class of preferences following a Bradley–Terry–Luce model, \mathcal{P}^{LN} the class of preferences following a certain ‘low-noise’ model, \mathcal{P}^{DAG} the class of acyclic preferences, and \mathcal{P}^{CW} the class of preferences that admit a Condorcet winner; it is known that $\mathcal{P}^{\text{BTL}} \subset \mathcal{P}^{\text{LN}} \subset \mathcal{P}^{\text{DAG}} \subset \mathcal{P}^{\text{CW}}$ (see Section 2.3 and Figures 1 and 2). The first three algorithms were analyzed with respect to pairwise disagreement in (Rajkumar & Agarwal, 2014). Results established in this paper are highlighted in gray.

Algorithm	Minimizes pairwise disagreement?			Ranks Condorcet winner at top (in \mathcal{P}^{CW})?	Ranks top cycle at top (in \mathcal{P})?	Ranks Copeland set at top (in \mathcal{P})?	Ranks Markov set at top (in \mathcal{P})?
	\mathcal{P}^{BTL}	\mathcal{P}^{LN}	\mathcal{P}^{DAG}				
Rank Centrality	✓	×	×	×	×	×	×
Matrix Borda	✓	✓	×	×	×	×	×
SVM-RankAggregation	✓	✓	✓	×	×	×	×
Matrix Copeland	✓	✓	✓	✓	✓	✓	×
Unweighted Markov	×	×	×	✓	✓	×	✓
Parametrized Markov	✓	×	×	✓	✓	×	✓

$\mathcal{P}_n^{\text{TC}(k)} = \{\mathbf{P} \in \mathcal{P}_n : \text{TC}(\mathbf{P}) = k\}$
$\mathcal{P}_n^{\text{CW}} = \{\mathbf{P} \in \mathcal{P}_n : \mathbf{P} \text{ has a Condorcet winner}\} = \mathcal{P}_n^{\text{TC}(1)}$
$\mathcal{P}_n^{\text{DAG}} = \{\mathbf{P} \in \mathcal{P}_n : G_{\mathbf{P}} = ([n], E_{\mathbf{P}}) \text{ is a DAG}\}$
$\mathcal{P}_n^{\text{LN}} = \{\mathbf{P} \in \mathcal{P}_n : i \succ_{\mathbf{P}} j \implies \sum_{k=1}^n p_{ki} > \sum_{k=1}^n p_{kj}\}$
$\mathcal{P}_n^{\text{BTL}} = \{\mathbf{P} \in \mathcal{P}_n : \exists \mathbf{w} \in \mathbb{R}_+^n \text{ s.t. } p_{ij} = \frac{w_j}{w_i + w_j} \forall i \neq j\}$

Figure 1. Definitions of various conditions on \mathbf{P} .

new algorithms, namely the Matrix Copeland, Unweighted Markov, and Parametrized Markov algorithms, that provably rank such good items at the top (see Table 1 for a summary). In all cases, we provide explicit sample complexity bounds for these algorithms to recover (with high probability) a ranking that places the desired elements at the top. As a by-product of our analysis of the Parametrized Markov algorithm, we also obtain a tighter sample complexity bound for the RC algorithm, thereby answering an open question of Rajkumar & Agarwal (2014).

2. Preliminaries and Background

2.1. Problem Setup

Let $[n] = \{1, \dots, n\}$ denote the set of n items to be ranked, and $\binom{[n]}{2} = \{(i, j) \in [n] \times [n] : i < j\}$ the set of $\binom{n}{2}$ pairs that can be compared. We assume that there is an underlying (unknown) probability distribution $\mu \in \Delta_{\binom{[n]}{2}}$ such that each time a comparison is to be made, a pair (i, j) is selected with probability μ_{ij} , and that for each $i < j$, there is an (unknown) parameter $p_{ij} \in [0, 1]$ such that when items i and j are compared, item j ‘beats’ (or is preferred to) item i with probability p_{ij} . For each $i < j$, define $p_{ji} = 1 - p_{ij}$; also define $p_{ii} = 0$ for every i . Denote by $\mathbf{P} = [p_{ij}]$ the resulting (unknown) pairwise preference matrix. Given a finite sample of pairwise comparisons $S = ((i_1, j_1, y_1), \dots, (i_m, j_m, y_m)) \in \left(\binom{[n]}{2} \times \{0, 1\}\right)^m$ drawn according to (μ, \mathbf{P}) as above (where pairs (i_k, y_k)




Figure 2. Relationships between various conditions on \mathbf{P} .

are drawn iid according to μ , and given (i_k, j_k) , $y_k \sim \text{Bernoulli}(p_{i_k, j_k})$, the goal is to construct a ranking or permutation $\sigma \in \mathcal{S}_n$ that ranks ‘good’ elements at the top. Most algorithms we consider operate on an empirical pairwise comparison matrix $\hat{\mathbf{P}}$ constructed from S as follows:

$$\hat{p}_{ij} = \begin{cases} m_{ij}^{(1)} / m_{ij} & \text{if } i < j \text{ and } m_{ij} > 0 \\ 1 - (m_{ji}^{(1)} / m_{ji}) & \text{if } i > j \text{ and } m_{ji} > 0 \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

where $m_{ij} = \sum_{k=1}^m \mathbf{1}(i_k = i, j_k = j)$;

$$m_{ij}^{(1)} = \sum_{k=1}^m \mathbf{1}(i_k = i, j_k = j, y_k = 1).$$

Binary Relation and Tournament Induced by \mathbf{P} . We will assume throughout that the matrix \mathbf{P} does not involve ties, i.e., $p_{ij} \neq \frac{1}{2}$ for all i, j , and will find it convenient to define the binary relation $\succ_{\mathbf{P}}$ on $[n]$ as follows:

$$i \succ_{\mathbf{P}} j \iff p_{ij} < \frac{1}{2} \quad \text{for all } i, j \in [n].$$

We will also define the edge set associated with \mathbf{P} as

$$E_{\mathbf{P}} = \{(i, j) \in [n] \times [n] : i \succ_{\mathbf{P}} j\},$$

and the induced graph $G_{\mathbf{P}} = ([n], E_{\mathbf{P}})$. Note that under the assumption $p_{ij} \neq \frac{1}{2}$ for all i, j , the graph $G_{\mathbf{P}}$ is always a complete directed graph, i.e., a tournament.

2.2. Measuring Ranking Performance

We will consider several notions of ‘goodness’ of a ranking with respect to the underlying pairwise preferences \mathbf{P} .

Pairwise Disagreement. Define the *pairwise disagreement* (PD) error of a permutation $\sigma \in \mathcal{S}_n$ w.r.t. \mathbf{P} as

$$\text{er}_{\mathbf{P}}^{\text{PD}}[\sigma] = \sum_{i \neq j} \mathbf{1}(p_{ij} < \frac{1}{2}) \cdot \mathbf{1}(\sigma(i) > \sigma(j)).$$

The pairwise disagreement error measures the number of pairs on which σ disagrees with \mathbf{P} , and is a natural measure of overall ranking performance; a variant of this was studied in (Rajkumar & Agarwal, 2014). In general, however, even if \mathbf{P} is known directly, identifying an optimal ranking w.r.t. PD error is computationally hard (it corresponds to the NP-hard minimum feedback arc set problem). Thus one can hope to recover an optimal ranking w.r.t. PD error only under certain restrictive conditions on \mathbf{P} ; indeed, all conditions studied in (Rajkumar & Agarwal, 2014) assumed $G_{\mathbf{P}}$ is acyclic.

In this paper, we are interested in algorithms that can ensure ‘good’ elements are ranked at the top even when the graph $G_{\mathbf{P}}$ contains cycles; we formalize the notion of ‘good’ below in terms of Condorcet winners, top cycles, and other tournament solution sets (see also Figure 3).

Condorcet Winners. An element $i \in [n]$ is said to be a *Condorcet winner* w.r.t. \mathbf{P} if it beats all other elements under \mathbf{P} , i.e., if $i \succ_{\mathbf{P}} j \forall j \neq i$. A Condorcet winner does not always exist, but when it does, it is unique, and it is then natural to want this element to be ranked at the top. When a Condorcet winner exists, we will denote it by $\text{CW}(\mathbf{P})$.

Top Cycles. More generally, one can ask that elements of the top cycle of \mathbf{P} be ranked at the top. The *top cycle* of \mathbf{P} , denoted $\text{TC}(\mathbf{P})$, is defined as the smallest set of elements $W \subseteq [n]$ such that every element in W beats every element not in W under \mathbf{P} , i.e. such that $i \succ_{\mathbf{P}} j \forall i \in W, j \notin W$. A Condorcet winner corresponds to a top cycle of size 1, i.e. when $|\text{TC}(\mathbf{P})| = 1$, we have $\text{TC}(\mathbf{P}) = \{\text{CW}(\mathbf{P})\}$. The top cycle is also referred to as the Smith set in voting theory literature (Smith, 1973).

Copeland and Markov Sets. The top cycle is one form of *tournament solution set*, which selects a set of elements considered to be ‘winners’ in a tournament (Laslier, 1997; Brandt et al., 2015). One can also consider other notions of tournament solution sets associated with \mathbf{P} (or more specifically, associated with the tournament $G_{\mathbf{P}}$ induced by \mathbf{P}), and ask that elements of the desired tournament solution set be ranked at the top. Two such tournament solution sets that we will be interested in are the Copeland set and the Markov set, both of which are ‘score-based’ solution sets that select all elements maximizing a certain ‘score’. The *Copeland set* of \mathbf{P} , denoted $\text{CO}(\mathbf{P})$, selects elements with maximal out-degree in the (unweighted) tournament $G_{\mathbf{P}}$: $\text{CO}(\mathbf{P}) = \text{argmax}_{i \in [n]} d(i)$, where $d(i) = \sum_j \mathbf{1}(i \succ_{\mathbf{P}} j)$. The *Markov set* of \mathbf{P} , denoted $\text{MA}(\mathbf{P})$, selects elements with highest stationary probability under a certain Markov chain on the tournament $G_{\mathbf{P}}$: $\text{MA}(\mathbf{P}) = \text{argmax}_{i \in [n]} \pi_i$,




Figure 3. **Left:** A tournament on $n = 4$ nodes in which there are cycles but node 1 is a Condorcet winner. **Middle:** A tournament on $n = 4$ nodes in which there is no Condorcet winner. Here this set of nodes $\{1, 2, 3\}$ forms a top cycle of size 3; in this case this set also corresponds to the Copeland and Markov sets. **Right:** A tournament on $n = 5$ nodes in which the top cycle $\{1, 2, 3, 4\}$ is a strict superset of the Copeland set $\{1, 3\}$ and Markov set $\{1\}$.

where π is the stationary probability vector of the Markov chain $\tilde{\mathbf{P}}$ defined as $\tilde{p}_{ij} = \frac{1}{n}$ for all $(i, j) : j \succ_{\mathbf{P}} i$ and $\tilde{p}_{ii} = 1 - \sum_{j \neq i} \tilde{p}_{ij}$ for all i . It is known that the Copeland and Markov sets are refinements of the top cycle, i.e. that $\text{CO}(\mathbf{P}) \subseteq \text{TC}(\mathbf{P})$ and $\text{MA}(\mathbf{P}) \subseteq \text{TC}(\mathbf{P})$, and that these containments can be strict (see Figure 3 for illustrations).

2.3. Conditions on the Pairwise Preference Matrix \mathbf{P}

As noted above, we are interested in understanding the ranking behavior of various algorithms under various conditions on the pairwise preference matrix \mathbf{P} . Let \mathcal{P}_n denote the set of all such matrices on n items not involving ties:

$$\mathcal{P}_n = \left\{ \mathbf{P} \in [0, 1]^{n \times n} : p_{ij} = 1 - p_{ji} \forall i \neq j ; p_{ij} \neq \frac{1}{2} \forall i, j ; p_{ii} = 0 \forall i \right\}.$$

For each $k \in [n]$, define $\mathcal{P}_n^{\text{TC}(k)}$ to be the set of preference matrices \mathbf{P} in \mathcal{P}_n that have a top cycle of size k .¹ Clearly, $\mathcal{P}_n^{\text{TC}(1)}$ is the set of preference matrices in \mathcal{P}_n that have a Condorcet winner; we will also denote this by $\mathcal{P}_n^{\text{CW}}$. Next, define $\mathcal{P}_n^{\text{DAG}}$ to be the set of preference matrices \mathbf{P} in \mathcal{P}_n for which the directed graph $G_{\mathbf{P}}$ induced by the binary relation $\succ_{\mathbf{P}}$ is acyclic,² and define $\mathcal{P}_n^{\text{LN}}$ to be the set of matrices in \mathcal{P}_n satisfying the ‘low-noise’ (LN) condition (Rajkumar & Agarwal, 2014). Finally, define $\mathcal{P}_n^{\text{BTL}}$ to be the set of preference matrices in \mathcal{P}_n that follow a Bradley-Terry-Luce (BTL) model. Formal definitions of all these sets are given in Figure 1. It can be verified that $\mathcal{P}_n^{\text{TC}(1)}, \dots, \mathcal{P}_n^{\text{TC}(n)}$ form a partition of \mathcal{P}_n , and that

$$\mathcal{P}_n^{\text{BTL}} \subset \mathcal{P}_n^{\text{LN}} \subset \mathcal{P}_n^{\text{DAG}} \subset \mathcal{P}_n^{\text{CW}} = \mathcal{P}_n^{\text{TC}(1)},$$

with each containment above being strict (see Figure 2). When n is clear from context, we will write $\mathcal{P}, \mathcal{P}^{\text{BTL}}$, etc.

3. Related Work and Existing Results

As noted above, there has been much work in recent years on developing algorithms for ranking a set of items from

¹Note that under a strict tournament (where there are no ties), there cannot be a top cycle of size 2, i.e. $\mathcal{P}_n^{\text{TC}(2)} = \emptyset$.

²The DAG condition is equivalent to the ‘generalized low-noise’ (GLN) condition studied in (Rajkumar & Agarwal, 2014).

pairwise preferences; below we discuss three specific algorithms that form the backdrop to our work. Before doing so, we point out that our work differs from that of (Brandt & Fischer, 2007), where any ranking algorithm that given $\mathbf{P} \in \mathcal{P}_n$ produces a score vector $\mathbf{f} \in \mathbb{R}^n$ is considered to induce a corresponding tournament solution $\text{TS}(\mathbf{P}) = \text{argmax}_{i \in [n]} f_i$; in particular, Brandt & Fischer study properties of the tournament solution induced in this manner by PageRank scores. Our work also differs from (Braverman & Mossel, 2009; Wauthier et al., 2013), where pairs are sampled only once and the conditions on \mathbf{P} are significantly stronger than those assumed here.

Three representative ranking algorithms studied in recent years that are most related to our work are the Rank Centrality (RC) algorithm (Negahban et al., 2012), the Matrix Borda (MB) algorithm (Rajkumar & Agarwal, 2014; Borda, 1781), and the SVM-RankAggregation (SVM-RA) algorithm (Rajkumar & Agarwal, 2014). RC produces a ranking by sorting the stationary probabilities of a Markov chain constructed from pairwise data. MB ranks items based on the average probability of an item being preferred over the rest of the items. SVM-RA constructs a binary classification dataset from the pairwise data and ranks items based on the maximum margin hyperplane for this dataset. Detailed descriptions of these algorithms are given in the appendix; below we summarize results of (Rajkumar & Agarwal, 2014) which establish convergence of these algorithms to an optimal permutation w.r.t. PD error within \mathcal{P}^{BTL} , \mathcal{P}^{LN} , and \mathcal{P}^{DAG} , respectively (see Table 2 for notation).³

Theorem 1 (Rank Centrality minimizes PD error in \mathcal{P}^{BTL}). Let $\mu_{\min} > 0$. Let $\mathbf{P} \in \mathcal{P}_n^{\text{BTL}}$. Let $\boldsymbol{\pi}$ be the stationary probability vector of the Markov chain $\bar{\mathbf{P}}$ defined as $\bar{p}_{ij} = \frac{p_{ij}}{n} \forall i \neq j$ and $\bar{p}_{ii} = 1 - \sum_{j \neq i} \bar{p}_{ij} \forall i$, and let $r_{\min} = \min_{i,j:\pi_i \neq \pi_j} |\pi_i - \pi_j|$. Let $\delta \in (0, 1]$. If

$$m \geq \max \left(\frac{9216 n}{r_{\min}^2 \mu_{\min}^2 p_{\min}^2} \left(\frac{1}{p_{\min}} - 1 \right)^3 \ln \left(\frac{16n^2}{\delta} \right), B_\mu \right),$$

then with probability at least $1 - \delta$ (over S), the permutation $\hat{\sigma}_{\text{RC}}$ produced by running the RC algorithm on $\hat{\mathbf{P}}$ satisfies

$$\hat{\sigma}_{\text{RC}} \in \text{argmin}_{\sigma \in \mathcal{S}_n} \text{er}_{\mathbf{P}}^{\text{PD}}[\sigma].$$

Theorem 2 (Matrix Borda minimizes PD error in \mathcal{P}^{LN}). Let $\mu_{\min} > 0$. Let $\mathbf{P} \in \mathcal{P}_n^{\text{LN}}$. Let $f_i^* = \frac{1}{n} \sum_{k=1}^n p_{ki} \forall i$, and let $r_{\min} = \min_{i,j:f_i^* \neq f_j^*} |f_i^* - f_j^*|$. Let $\delta \in (0, 1]$. If

$$m \geq \max \left(\frac{1152}{r_{\min}^2 \mu_{\min}^2} \ln \left(\frac{4n^2}{\delta} \right), B_\mu \right),$$

then with probability at least $1 - \delta$ (over S), the permutation $\hat{\sigma}_{\text{MB}}$ produced by running the MB algorithm on $\hat{\mathbf{P}}$ satisfies

$$\hat{\sigma}_{\text{MB}} \in \text{argmin}_{\sigma \in \mathcal{S}_n} \text{er}_{\mathbf{P}}^{\text{PD}}[\sigma].$$

³We note that the form of the PD error considered in (Rajkumar & Agarwal, 2014) is slightly different from that considered here, but within \mathcal{P}^{DAG} both have the same sets of minimizers.

Table 2. Useful quantities associated with μ and \mathbf{P} .

μ_{\min}	$\min_{i < j} \mu_{ij}$
B_μ	$3 \left(\frac{12}{\mu_{\min}^2} + 3 \right) \ln \left(\frac{12}{\mu_{\min}^2} + 3 \right)$
p_{\min}	$\min_{i \neq j} p_{ij}$
γ_{\min}	$\min_{i \neq j} \left p_{ij} - \frac{1}{2} \right $
γ_{TC}	$\min_{i \in \text{TC}(\mathbf{P}), j \notin \text{TC}(\mathbf{P})} \left p_{ij} - \frac{1}{2} \right $

Theorem 3 (SVM-RA minimizes PD error in \mathcal{P}^{DAG}). Let $\mu_{\min} > 0$. Let $\mathbf{P} \in \mathcal{P}_n^{\text{DAG}}$. Let $\boldsymbol{\alpha} \in \mathbb{R}^n$ be such that $i \succ_{\mathbf{P}} j \implies \sum_{k=1}^n \alpha_k p_{ki} > \sum_{k=1}^n \alpha_k p_{jk}$,⁴ and let $r_{\min}^\alpha = \min_{i \neq j} \frac{|\sum_{k=1}^n \alpha_k (P_{ki} - P_{kj})|}{\|\boldsymbol{\alpha}\|_2}$. Let $\delta \in (0, 1]$. If

$$m \geq \max \left(\frac{2048 n}{(r_{\min}^\alpha)^2 \mu_{\min}^2} \ln \left(\frac{16n^3}{\delta} \right), \frac{128}{\gamma_{\min}^2 \mu_{\min}^2} \ln \left(\frac{8n^2}{\delta} \right), B_\mu \right),$$

then with probability at least $1 - \delta$ (over S), the permutation $\hat{\sigma}_{\text{SVM-RA}}$ produced by running SVM-RA on $\hat{\mathbf{P}}$ satisfies

$$\hat{\sigma}_{\text{SVM-RA}} \in \text{argmin}_{\sigma \in \mathcal{S}_n} \text{er}_{\mathbf{P}}^{\text{PD}}[\sigma].$$

We will make use of the following result of (Rajkumar & Agarwal, 2014) on concentration of entries of $\hat{\mathbf{P}}$ around \mathbf{P} :

Lemma 4. Let $\mu_{\min} > 0$ and $\mathbf{P} \in \mathcal{P}_n$. Fix any $i \neq j$. Let $0 < \epsilon < 4\sqrt{2}$. If $m \geq \max \left(\frac{128}{\epsilon^2 \mu_{\min}^2} \ln \left(\frac{4}{\delta} \right), B_\mu \right)$, then with probability at least $1 - \delta$ (over S), $|p_{ij} - \hat{p}_{ij}| < \epsilon$.

4. Performance of RC, MB and SVM-RA Algorithms Under Cyclic Preferences

It is easy to see that for $\mathbf{P} \in \mathcal{P}_n^{\text{CW}}$,

$$\hat{\sigma} \in \text{argmin}_{\sigma \in \mathcal{S}_n} \text{er}_{\mathbf{P}}^{\text{PD}}[\sigma] \implies \hat{\sigma}^{-1}(1) = \text{CW}(\mathbf{P}).$$

Therefore it follows from Theorems 1–3 that (when given a sufficiently large sample, with high probability) the RC algorithm ranks the Condorcet winner at the top for $\mathbf{P} \in \mathcal{P}^{\text{BTL}}$; the MB algorithm does so for $\mathbf{P} \in \mathcal{P}^{\text{LN}}$; and the SVM-RA algorithm does so for $\mathbf{P} \in \mathcal{P}^{\text{DAG}}$. However, as the following examples show, outside \mathcal{P}^{DAG} , these algorithms can fail to rank the Condorcet winner, top cycle, or Copeland or Markov sets at the top (in fact RC and MB can fail to do so even in $\mathcal{P}^{\text{LN}} \setminus \mathcal{P}^{\text{BTL}}$ and $\mathcal{P}^{\text{DAG}} \setminus \mathcal{P}^{\text{LN}}$, respectively; see Examples 6–7 in the supplementary material).

Example 1 (For $\mathbf{P} \in \mathcal{P}^{\text{CW}} \setminus \mathcal{P}^{\text{DAG}}$, RC, MB and SVM-RA algorithms can fail to rank the Condorcet winner/Copeland set/Markov set at the top). Let $n = 7$, and consider

$$\mathbf{P} = \begin{bmatrix} 0 & 0.49 & 0.49 & 0.49 & 0.49 & 0.49 & 0.49 \\ 0.51 & 0 & 0.1 & 0.1 & 0.1 & 0.1 & 0.1 \\ 0.51 & 0.9 & 0 & 0.4 & 0.4 & 0.6 & 0.6 \\ 0.51 & 0.9 & 0.6 & 0 & 0.4 & 0.6 & 0.4 \\ 0.51 & 0.9 & 0.6 & 0.6 & 0 & 0.4 & 0.6 \\ 0.51 & 0.9 & 0.4 & 0.4 & 0.6 & 0 & 0.4 \\ 0.51 & 0.9 & 0.4 & 0.6 & 0.4 & 0.6 & 0 \end{bmatrix}.$$

⁴Such an $\boldsymbol{\alpha}$ always exists by equivalence of the DAG and GLN conditions (Rajkumar & Agarwal, 2014).

The tournament induced by \mathbf{P} is shown in Figure 4 (left). It can be verified that $\mathbf{P} \in \mathcal{P}^{\text{CW}} \setminus \mathcal{P}^{\text{DAG}}$, with $\text{CW}(\mathbf{P}) = 1$ and correspondingly $\text{CO}(\mathbf{P}) = \text{MA}(\mathbf{P}) = \{1\}$. However the permutations σ_{RC} , σ_{MB} and $\sigma_{\text{SVM-RA}}$ produced by running RC, MB and SVM-RA on the above preference matrix \mathbf{P} do not rank item 1 at the top; instead, it can be verified that $\sigma_{\text{RC}}^{-1}(1) = \sigma_{\text{MB}}^{-1}(1) = 2$ and $\sigma_{\text{SVM-RA}}^{-1}(1) = 7$.

Example 2 (For $\mathbf{P} \notin \mathcal{P}^{\text{CW}}$, RC, MB and SVM-RA algorithms can fail to rank the top cycle/Copeland set/Markov set at the top). Let $n = 6$, and consider

$$\mathbf{P} = \begin{bmatrix} 0 & 0.3 & 0.7 & 0.4 & 0.45 & 0.45 \\ 0.7 & 0 & 0.3 & 0.4 & 0.45 & 0.45 \\ 0.3 & 0.7 & 0 & 0.4 & 0.45 & 0.45 \\ 0.6 & 0.6 & 0.6 & 0 & 0.1 & 0.1 \\ 0.55 & 0.55 & 0.55 & 0.9 & 0 & 0.6 \\ 0.55 & 0.55 & 0.55 & 0.9 & 0.4 & 0 \end{bmatrix}.$$

The tournament induced by \mathbf{P} is shown in Figure 4 (middle). It can be verified that $\mathbf{P} \in \mathcal{P}^{\text{TC}(3)}$, with $\text{TC}(\mathbf{P}) = \{1, 2, 3\} = \text{CO}(\mathbf{P}) = \text{MA}(\mathbf{P})$. However the permutations σ_{RC} , σ_{MB} and $\sigma_{\text{SVM-RA}}$ produced by running RC, MB and SVM-RA on the above matrix \mathbf{P} do not rank these tournament solution sets at the top; indeed, it can be verified that $\sigma_{\text{RC}} = \sigma_{\text{MB}} = (4\ 1\ 2\ 3\ 6\ 5)$ and $\sigma_{\text{SVM-RA}} = (4\ 2\ 6\ 1\ 3\ 5)$.

The above examples lead us to consider alternative ranking algorithms that rank Condorcet winners, top cycles, and Copeland/Markov sets at the top even when the underlying preference matrix \mathbf{P} induces a cyclic tournament. Proofs of all results can be found in the supplementary material.

5. Recovering Condorcet Winners, Top Cycles and Copeland Sets at the Top

Let us start by considering the Copeland set. Recall that this is defined as the set of items with maximal out-degree in the underlying preference-induced tournament $G_{\mathbf{P}}$. A natural approach to ranking the Copeland set at the top would therefore be to order items by their observed out-degrees in the empirical comparison matrix $\hat{\mathbf{P}}$; as this empirical matrix approaches the true preference matrix \mathbf{P} , one would expect items in the Copeland set to appear at the top. The corresponding algorithm is shown in Algorithm 1; we term this the Matrix Copeland (MC) algorithm due to its similarity to the Copeland voting rule (Copeland, 1951). The following result shows this algorithm indeed ranks the Copeland set at the top for any preference matrix and also minimizes the pairwise disagreement error when the preference matrix is acyclic:

Theorem 5 (Matrix Copeland ranks Copeland set at top for a general \mathbf{P} and minimizes PD error if $\mathbf{P} \in \mathcal{P}^{\text{DAG}}$). Let $\mu_{\min} > 0$. Let $\mathbf{P} \in \mathcal{P}_n$ and let $\delta \in (0, 1]$. If

$$m \geq \max \left(\frac{512}{\gamma_{\min}^2 \mu_{\min}^2} \ln \left(\frac{4n^2}{\delta} \right), B_{\mu} \right),$$

then with probability at least $1 - \delta$ (over S), the permutation $\hat{\sigma}_{\text{MC}}$ produced by running the MC algorithm on $\hat{\mathbf{P}}$ satisfies

Algorithm 1 Matrix Copeland (MC) Algorithm

Input: Pairwise comparison matrix $\bar{\mathbf{P}} \in [0, 1]^{n \times n}$ satisfying the following conditions:

- (i) for all $i \neq j$: $\bar{p}_{ij} + \bar{p}_{ji} = 1$ or $\bar{p}_{ij} = \bar{p}_{ji} = 0$
- (ii) for every i : $\bar{p}_{ii} = 0$

- For $i = 1$ to n : $\bar{f}_i = \frac{1}{n} \sum_{j=1}^n \mathbf{1}(\bar{p}_{ji} > \frac{1}{2})$

Output: Permutation $\bar{\sigma}_{\text{MC}} \in \text{argsort}(\bar{\mathbf{f}})$

$$\hat{\sigma}_{\text{MC}}(i) < \hat{\sigma}_{\text{MC}}(j) \quad \text{for all } i \in \text{CO}(\mathbf{P}), j \notin \text{CO}(\mathbf{P}).$$

Moreover, if $\mathbf{P} \in \mathcal{P}^{\text{DAG}}$, then

$$\hat{\sigma}_{\text{MC}} \in \text{argmin}_{\sigma \in S_n} \text{er}_{\mathbf{P}}^{\text{PD}}[\sigma]$$

In fact, as the following result shows, the MC algorithm also recovers the top cycle (and therefore the Condorcet winner whenever it exists) at the top:

Theorem 6 (Matrix Copeland ranks top cycle at top). Let $\mu_{\min} > 0$. Let $\mathbf{P} \in \mathcal{P}_n^{\text{TC}(k)}$ for some $k \in [n - 1]$. Let $\delta \in (0, 1]$. If

$$m \geq \max \left(\frac{512}{\gamma_{\text{TC}}^2 \mu_{\min}^2} \ln \left(\frac{4n^2}{\delta} \right), B_{\mu} \right),$$

then with probability at least $1 - \delta$ (over S), the permutation $\hat{\sigma}_{\text{MC}}$ produced by running the MC algorithm on $\hat{\mathbf{P}}$ satisfies

$$\hat{\sigma}_{\text{MC}}(i) < \hat{\sigma}_{\text{MC}}(j) \quad \text{for all } i \in \text{TC}(\mathbf{P}), j \notin \text{TC}(\mathbf{P}).$$

While the MC algorithm is good for recovering the Copeland set and the top cycle or the Condorcet winner when it exists, as the following example illustrates, it does not necessarily rank the Markov set at the top:

Example 3 (Matrix Copeland can fail to rank Markov set at top). Let $n = 8$, and consider

$$\mathbf{P} = \begin{bmatrix} 0 & 0.51 & 0.4 & 0.6 & 0.6 & 0.6 & 0.4 & 0.4 \\ 0.49 & 0 & 0.49 & 0.49 & 0.6 & 0.6 & 0.49 & 0.49 \\ 0.6 & 0.51 & 0 & 0.6 & 0.4 & 0.6 & 0.6 & 0.4 \\ 0.4 & 0.51 & 0.4 & 0 & 0.6 & 0.4 & 0.4 & 0.4 \\ 0.4 & 0.4 & 0.6 & 0.4 & 0 & 0.4 & 0.6 & 0.6 \\ 0.4 & 0.4 & 0.4 & 0.6 & 0.6 & 0 & 0.4 & 0.4 \\ 0.6 & 0.51 & 0.4 & 0.6 & 0.4 & 0.6 & 0 & 0.4 \\ 0.6 & 0.51 & 0.6 & 0.6 & 0.4 & 0.6 & 0.6 & 0 \end{bmatrix}.$$

The tournament induced by \mathbf{P} is shown in Figure 4 (right). It can be verified that $\text{MA}(\mathbf{P}) = \{5\}$. However the permutation σ_{MC} produced by running the MC algorithm on the above matrix \mathbf{P} does not rank item 5 at the top; indeed, it can be verified that $\sigma_{\text{MC}} = (2\ 4\ 6\ 5\ 1\ 7\ 3\ 8)$.

6. Recovering Condorcet Winners, Top Cycles and Markov Sets at the Top

The Matrix Copeland algorithm considered above can be viewed as running the Matrix Borda algorithm on a transformed input matrix $\bar{\mathbf{H}}$ where $\bar{h}_{ij} = \mathbf{1}(\bar{p}_{ij} > \frac{1}{2}) \forall i \neq j$, which thresholds the preferences \bar{p}_{ij} at $\frac{1}{2}$ to ‘amplify’ them to 0 or 1, and has the effect of emphasizing more preferred




Figure 4. Tournaments induced by the preferences matrices in Examples 1–3. An edge from one group of nodes to another indicates that the graph contains edges from all nodes in the first group to all those in the second group. **Left:** Here $n = 7$ and $\text{TC}(\mathbf{P}) = \text{CO}(\mathbf{P}) = \text{MA}(\mathbf{P}) = \{1\}$. **Middle:** Here $n = 6$ and $\text{TC}(\mathbf{P}) = \text{CO}(\mathbf{P}) = \text{MA}(\mathbf{P}) = \{1, 2, 3\}$. **Right:** Here $n = 8$ and $\text{TC}(\mathbf{P}) = \{1, 2, 3, 4, 5, 6, 7, 8\}$, $\text{CO}(\mathbf{P}) = \{2, 4, 6\}$, and $\text{MA}(\mathbf{P}) = \{5\}$.

elements more strongly. As we saw above, this yielded desirable properties in terms of ranking Condorcet winners, top cycles and Copeland sets (though not Markov sets) at the top, while also maintaining the PD optimality properties of the MB algorithm in \mathcal{P}^{LN} (in fact, MC recovers a PD-optimal ranking for all \mathbf{P} in the larger set \mathcal{P}^{DAG}).

In order to recover Markov sets at the top, one might consider running the Rank Centrality algorithm, which ranks items according to the stationary vector associated with a certain Markov chain associated with the input matrix, on a similarly transformed matrix. Below we start by considering in Section 6.1 the Unweighted Markov (UM) ranking algorithm, which does exactly this: it runs the RC algorithm on the same transformed matrix $\bar{\mathbf{H}}$ described above. As we will see below, this algorithm indeed ranks Condorcet winners, top cycles and Markov sets (though not Copeland sets) at the top. Unfortunately, however, as we will see, the UM algorithm fails to maintain the PD optimality properties of the RC algorithm in \mathcal{P}^{BTL} .

In Section 6.2, we propose instead the Parametrized Markov (PM) algorithm, which runs the RC algorithm on a matrix $\bar{\mathbf{H}}^c$ obtained by applying a soft transformation parametrized by a real number $c \geq 1$ to the input matrix \mathbf{P} . This also amplifies the probabilities in the input matrix, but not to the extremes of 0 and 1 (the hard transform $\bar{\mathbf{H}}$ above corresponds to $c = \infty$). We will see that for suitable choices of c , the resulting PM algorithm can be made to both rank Condorcet winners, top cycles and Markov sets at the top, and maintain the PD optimality properties of the RC algorithm in \mathcal{P}^{BTL} . As a by-product of our analysis, we also obtain an improved sample complexity bound for the RC algorithm to recover a PD-optimal ranking, thus answering an open question of Rajkumar & Agarwal (2014).

6.1. Unweighted Markov Algorithm

The UM algorithm is shown in Algorithm 2. We first show this algorithm indeed recovers the Markov set at the top:

Theorem 7 (Unweighted Markov ranks Markov set at top). *Let $\mu_{\min} > 0$. Let $\mathbf{P} \in \mathcal{P}_n$. Let $\delta \in (0, 1]$. If*

$$m \geq \max \left(\frac{512}{\gamma_{\min}^2 \mu_{\min}^2} \ln \left(\frac{4n^2}{\delta} \right), B_\mu \right),$$

then with probability at least $1 - \delta$ (over S), the permutation

Algorithm 2 Unweighted Markov (UM) Algorithm

Input: Pairwise comparison matrix $\bar{\mathbf{P}} \in [0, 1]^{n \times n}$ satisfying the following conditions:

- (i) for all $i \neq j$: $\bar{p}_{ij} + \bar{p}_{ji} = 1$ or $\bar{p}_{ij} = \bar{p}_{ji} = 0$
- (ii) for every i : $\bar{p}_{ii} = 0$

- Define $\bar{\mathbf{H}} \in [0, 1]^{n \times n}$ as

$$\bar{h}_{ij} = \begin{cases} \mathbf{1}(\bar{p}_{ij} > \frac{1}{2}) & \text{if } i \neq j \\ 0 & \text{otherwise.} \end{cases}$$

- Run Rank Centrality algorithm on input matrix $\bar{\mathbf{H}}$; obtain stationary probability vector $\bar{\pi}$

Output: Permutation $\bar{\sigma}_{\text{UM}} \in \text{argsort}(\bar{\pi})$

$\hat{\sigma}_{\text{UM}}$ produced by running the UM algorithm on $\hat{\mathbf{P}}$ satisfies

$$\hat{\sigma}_{\text{UM}}(i) < \hat{\sigma}_{\text{UM}}(j) \quad \text{for all } i \in \text{MA}(\mathbf{P}), j \notin \text{MA}(\mathbf{P}).$$

In fact, as with the MC algorithm, the UM algorithm also recovers the top cycle (and therefore the Condorcet winner whenever it exists) at the top:

Theorem 8 (Unweighted Markov ranks top cycle at top).

Let $\mu_{\min} > 0$. Let $\mathbf{P} \in \mathcal{P}_n^{\text{TC}(k)}$ for some $k \in [n - 1]$. Let $\delta \in (0, 1]$. If

$$m \geq \max \left(\frac{512}{\gamma_{\min}^2 \mu_{\min}^2} \ln \left(\frac{4n^2}{\delta} \right), B_\mu \right),$$

then with probability at least $1 - \delta$ (over S), the permutation $\hat{\sigma}_{\text{UM}}$ produced by running the UM algorithm on $\hat{\mathbf{P}}$ satisfies

$$\hat{\sigma}_{\text{UM}}(i) < \hat{\sigma}_{\text{UM}}(j) \quad \text{for all } i \in \text{TC}(\mathbf{P}), j \notin \text{TC}(\mathbf{P}).$$

Again, as with the MC algorithm, the following example shows that while UM is good for recovering the Markov set and the top cycle or the Condorcet winner when it exists, it does not necessarily rank the Copeland set at the top:

Example 4 (Unweighted Markov can fail to rank Copeland set at top). *Let $n = 8$, and consider again the matrix \mathbf{P} considered in Example 3. It can be verified that $\text{CO}(\mathbf{P}) = \{2, 4, 6\}$. However the permutation σ_{UM} produced by running the UM algorithm on \mathbf{P} does not rank this set at the top; indeed, it can be verified that $\sigma_{\text{UM}} = (5 \ 2 \ 4 \ 6 \ 7 \ 3 \ 8 \ 1)$.*

Unfortunately, however, unlike the MC algorithm, the UM algorithm fails to maintain the PD optimality properties of the RC algorithm in \mathcal{P}^{BTL} :

Example 5 (Unweighted Markov can fail to minimize PD error in \mathcal{P}^{BTL}). Let $n = 3$, and consider

$$\mathbf{P} = \begin{bmatrix} 0 & 0.4 & 0.25 \\ 0.6 & 0 & 0.33 \\ 0.75 & 0.67 & 0 \end{bmatrix}.$$

It can be verified that $\mathbf{P} \in \mathcal{P}^{\text{BTL}}$ (consider the score vector $\mathbf{w} = (3, 2, 1)^\top$), and that the PD error is uniquely minimized by $\sigma = (1\ 2\ 3)$. However the permutation σ_{UM} produced by running the UM algorithm on \mathbf{P} is $\sigma_{\text{UM}} = (1\ 3\ 2)$, and therefore this does not minimize the PD error.

Next we consider an alternative algorithm that will achieve the best properties of both the UM and RC algorithms.

6.2. Parametrized Markov Algorithm

We now consider the Parametrized Markov (PM) ranking algorithm shown in Algorithm 3, which effectively applies the RC algorithm to a matrix $\bar{\mathbf{H}}^c$ obtained by applying a ‘soft’ transform $g_c : [0, 1] \rightarrow [0, 1]$, parametrized by $c \geq 1$, to the entries of the input matrix \mathbf{P} :

$$g_c(p) = \frac{p^c}{p^c + (1-p)^c}.$$

See Figure 5 for an illustration. With $c = 1$, one recovers the RC algorithm; with $c = \infty$, one recovers the UM algorithm. The PM algorithm can therefore be viewed as interpolating between these two extremes. As we show below, with a suitable choice of c , the PM algorithm gives the best of both worlds: as with RC, it minimizes the PD error in \mathcal{P}^{BTL} ; and as with UM, it ranks Condorcet winners, top cycles, and Markov sets at the top.




Figure 5. The function g_c .

6.2.1. PM ALGORITHM MINIMIZES PD ERROR IN \mathcal{P}^{BTL}

We first show that for any choice of $1 \leq c < \infty$, the PM algorithm minimizes the PD error in \mathcal{P}^{BTL} :

Theorem 9 (Parametrized Markov minimizes PD error in \mathcal{P}^{BTL}). Let $\mu_{\min} > 0$. Let $\mathbf{P} \in \mathcal{P}_n^{\text{BTL}}$. Let $r_{\min} = \min_{i \neq j, k \neq i, j} |p_{ik} - p_{jk}|$. Let $\delta \in (0, 1]$. Let $1 \leq c < \infty$. If

$$m \geq \max \left(\frac{512}{\min(r_{\min}^2, \gamma_{\min}^2) \mu_{\min}^2} \ln \left(\frac{4n^2}{\delta} \right), B_\mu \right),$$

then with probability at least $1 - \delta$ (over S), the permutation $\hat{\sigma}_{\text{PM}}$ produced by running the PM algorithm on $\hat{\mathbf{P}}$ satisfies

$$\hat{\sigma}_{\text{PM}} \in \operatorname{argmin}_{\sigma \in \mathcal{S}_n} \operatorname{er}_{\mathbf{P}}^{\text{PD}}[\sigma].$$

The proof of Theorem 9 in fact establishes PD-optimality of the PM algorithm (and therefore the RC algorithm, which is a special case with $c = 1$) for a slightly larger set of preference matrices than \mathcal{P}^{BTL} , namely, for \mathbf{P} satisfying the restricted low-noise (RLN) property, defined as

Algorithm 3 Parametrized Markov (PM) Algorithm

Input: Pairwise comparison matrix $\bar{\mathbf{P}} \in [0, 1]^{n \times n}$ satisfying the following conditions:
 (i) for all $i \neq j$: $\bar{p}_{ij} + \bar{p}_{ji} = 1$ or $\bar{p}_{ij} = \bar{p}_{ji} = 0$
 (ii) for every i : $\bar{p}_{ii} = 0$

Parameter: $c \geq 1$

• Define $\bar{\mathbf{H}}^c \in [0, 1]^{n \times n}$ as

$$\bar{h}_{ij}^c = \begin{cases} g_c(\bar{p}_{ij}) & \text{if } i \neq j \\ 0 & \text{otherwise.} \end{cases}$$

• Run Rank Centrality algorithm on input matrix $\bar{\mathbf{H}}^c$; obtain stationary probability vector $\bar{\pi}^c$

Output: Permutation $\bar{\sigma}_{\text{PM}} \in \operatorname{argsort}(\bar{\pi}^c)$

$$\mathcal{P}_n^{\text{RLN}} = \left\{ \mathbf{P} \in \mathcal{P}_n : \forall i \neq j \neq k : i \succ_{\mathbf{P}} j \implies p_{kj} < p_{ki} \right\}.$$

It is not hard to show that $\mathcal{P}_n^{\text{BTL}} \subset \mathcal{P}_n^{\text{RLN}} \subset \mathcal{P}_n^{\text{LN}}$. The proof of Theorem 9 then follows from the following three observations: (a) for $\mathbf{P} \in \mathcal{P}_n^{\text{RLN}}$, running the RC algorithm on \mathbf{P} yields a PD-optimal permutation; (b) for large enough sample size, if \mathbf{P} is in $\mathcal{P}_n^{\text{RLN}}$, then with high probability so is $\hat{\mathbf{P}}$; and (c) the RLN property is preserved by the g_c transform (see supplementary material for details).

Remark (Improved sample complexity bound for RC algorithm). As noted above, the sample complexity bound in Theorem 9 holds for any $1 \leq c < \infty$, and so in particular it holds for the RC algorithm. Comparing this with the previous bound of (Rajkumar & Agarwal, 2014) in Theorem 1, we see that the bound in Theorem 9 is in general considerably tighter; in particular, it removes an extraneous factor of n present in the earlier bound. This answers in the affirmative an open question of (Rajkumar & Agarwal, 2014).

6.2.2. FOR SUITABLE c , PM ALGORITHM RECOVERS MARKOV SETS AND TOP CYCLES AT THE TOP

We now show that for sufficiently large (but finite) c , which can be chosen based on the observed empirical comparison matrix $\hat{\mathbf{P}}$, the PM algorithm also recovers Markov sets and top cycles (and therefore Condorcet winners) at the top:

Theorem 10 (Parametrized Markov ranks Markov set at top). Let $\mu_{\min} > 0$. Let $\mathbf{P} \in \mathcal{P}_n$. Let $\beta_2(\hat{\mathbf{P}}) = \hat{\pi}_{\max} - \max_{k \notin \text{MA}(\hat{\mathbf{P}})} \hat{\pi}_k$ where $\hat{\pi}$ is obtained by running UM on $\hat{\mathbf{P}}$. Let $\delta \in (0, 1]$. Let $\hat{\alpha}_{\min} = \min_{i \succ_{\hat{\mathbf{P}}} j} \frac{\hat{p}_{ji}}{\hat{p}_{ij}}$. If $c \geq \frac{\ln(4n/\beta_2(\hat{\mathbf{P}}))}{\ln(\hat{\alpha}_{\min})}$ and

$$m \geq \max \left(\frac{512}{\gamma_{\min}^2 \mu_{\min}^2} \ln \left(\frac{4n^2}{\delta} \right), B_\mu \right),$$

then with probability at least $1 - \delta$ (over S), the permutation $\hat{\sigma}_{\text{PM}}$ produced by running the PM algorithm on $\hat{\mathbf{P}}$ satisfies

$$\hat{\sigma}_{\text{PM}}(i) < \hat{\sigma}_{\text{PM}}(j) \quad \text{for all } i \in \text{MA}(\mathbf{P}), j \notin \text{MA}(\mathbf{P}).$$

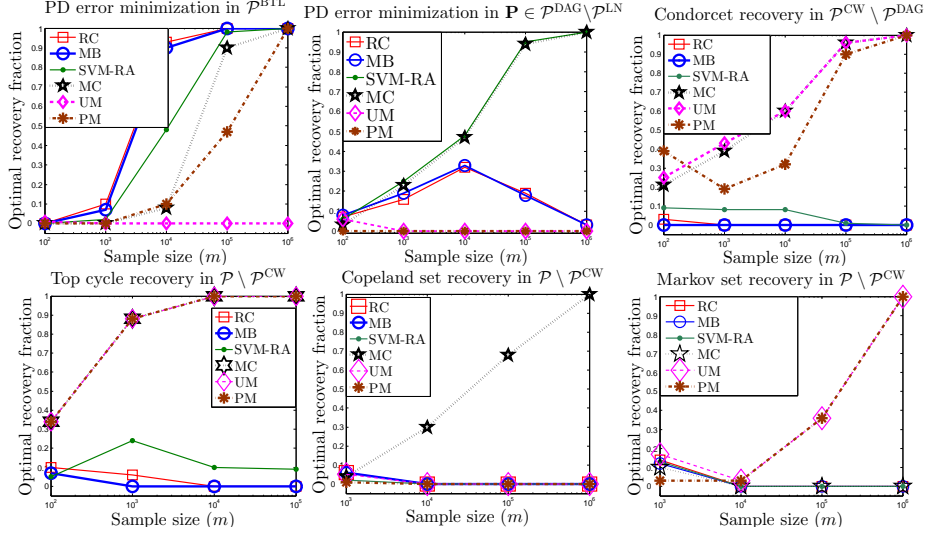


Figure 6. Experimental results. See text for details.

Theorem 11 (Parametrized Markov ranks top cycle at top).

Let $\mu_{\min} > 0$. Let $\mathbf{P} \in \mathcal{P}_n^{\text{TC}(k)}$ for some $k \in [n-1]$. Let $\hat{\alpha}_{\min} = \min_{i>j} \frac{\hat{p}_{ji}}{\hat{p}_{ij}}$. Let $\delta \in (0, 1]$. If $c \geq \frac{(n+1)\ln(4n)}{\ln(\hat{\alpha}_{\min})}$ and

$$m \geq \max\left(\frac{512}{\gamma_{\min}^2 \mu_{\min}^2} \ln\left(\frac{4n^2}{\delta}\right), B_{\mu}\right),$$

then with probability at least $1-\delta$ (over S), the permutation $\hat{\sigma}_{\text{PM}}$ produced by running the PM algorithm on $\hat{\mathbf{P}}$ satisfies

$$\hat{\sigma}_{\text{PM}}(i) < \hat{\sigma}_{\text{PM}}(j) \quad \text{for all } i \in \text{TC}(\mathbf{P}), j \notin \text{TC}(\mathbf{P}).$$

The proofs of Theorems 10–11 make use of the Cho-Meyer perturbation bound for Markov chains (Cho & Meyer, 2001) to bound the difference between the stationary vector π^c and π . (see supplementary material for details).

7. Experiments

We conducted experiments to compare the performance of different algorithms, including existing algorithms (RC, MB, SVM-RA) and the proposed algorithms (MC, UM, PM), both in terms of minimizing PD error and recovering ‘good’ items at the top (the parameter c for PM algorithm was chosen as the maximum of the values prescribed by Theorems 10 and 11). In all our experiments, we generated 100 training samples S (as described in Section 2.1) from an underlying preference matrix \mathbf{P} of interest (using a uniform distribution μ on pairs) each of a number of sizes m , ran the various algorithms on the samples, and for each algorithm, counted the fraction of times (out of 100) that the returned permutation satisfied a desired property, such as PD error minimization or top cycle recovery at the top.

PD error minimization. For this we used two acyclic preference matrices: a matrix $\mathbf{P} \in \mathcal{P}_{10}^{\text{BTL}}$ generated from a score vector \mathbf{w} drawn uniformly from $[0, 1]^{10}$, and the matrix $\mathbf{P} \in (\mathcal{P}_5^{\text{DAG}} \setminus \mathcal{P}_5^{\text{LN}})$ described in (Rajkumar & Agarwal, 2014). The results are shown in the first two plots in Figure 6. As expected, for the first \mathbf{P} , all algorithms except

UM recover an optimal ranking w.r.t. PD error, while for the second \mathbf{P} , only SVM-RA and MC do so.

Condorcet winner and top cycle recovery. For this we used two preference matrices with cycles: the matrix $\mathbf{P} \in \mathcal{P}_7^{\text{CW}} \setminus \mathcal{P}_7^{\text{DAG}}$ described in Example 1 (see also Figure 4, left), and the matrix $\mathbf{P} \in \mathcal{P}_6^{\text{TC}(3)} \subset (\mathcal{P}_6 \setminus \mathcal{P}_6^{\text{CW}})$ described in Example 2 (see also Figure 4, middle). The results are shown in the 3rd and 4th plots in Figure 6. In this case, only the new algorithms considered in this paper (MC, UM, PM) recover the Condorcet winner and the top cycle at the top.

Copeland set recovery. For this we used the cyclic preference matrix $\mathbf{P} \in (\mathcal{P}_8 \setminus \mathcal{P}_8^{\text{CW}})$ described in Example 3 (see also Figure 4, right). The results are shown in the 5th plot in Figure 6. As expected, only the MC algorithm successfully recovers the Copeland set at the top.

Markov set recovery. For this we again used the cyclic matrix $\mathbf{P} \in (\mathcal{P}_8 \setminus \mathcal{P}_8^{\text{CW}})$ described in Example 3, for which $\text{MA}(\mathbf{P}) \neq \text{CO}(\mathbf{P})$. The results are shown in the 6th plot in Figure 6. As expected, in this case only the UM and PM algorithms correctly recover the Markov set at the top.

8. Conclusion

In this paper, we investigated convergence properties of algorithms for ranking from pairwise preferences when preferences can contain cycles. When a globally optimal ranking is hard to compute, it is natural to ask that ‘good’ items, such as Condorcet winners or suitable tournament solution sets, be ranked at the top. We showed that several ranking algorithms developed in recent years fail in this respect, and designed new algorithms that rank Condorcet winners, top cycles, Copeland and Markov sets at the top, while retaining or improving guarantees of previous algorithms for acyclic preferences. Future work includes exploring alternative complexity parameters for our bounds and considering similar objectives in an active ranking setting.

Acknowledgements. Thanks to the anonymous reviewers for their helpful comments. AR is supported by a Microsoft Research India PhD Fellowship. The work of LHL is supported by AFOSR FA9550-13-1-0133, NSF DMS-1209136, and NSF DMS-1057064. SA thanks DST and Indo-US Science & Technology Forum for their support.

References

- Ailon, Nir. Active learning ranking from pairwise preferences with almost optimal query complexity. In *Advances in Neural Information Processing Systems*, 2011.
- Borda, Jean-Charles de. Mémoire sur les élections au scrutin. *Histoire de l'Académie Royale des Sciences*, 1781.
- Brandt, Felix and Fischer, Felix. PageRank as a weak tournament solution. In *Proceedings of the 3rd International Workshop on Internet and Network Economics*, 2007.
- Brandt, Felix, Brill, Markus, and Harrenstein, Paul. Tournament solutions. In *Handbook of Computational Social Choice*. Cambridge University Press, To appear, 2015.
- Braverman, Mark and Mossel, Elchanan. Sorting from noisy information. *arXiv preprint arXiv:0910.1191*, 2009.
- Busa-Fekete, Róbert, Hüllermeier, Eyke, and Szörényi, Balázs. Preference-based rank elicitation using statistical models: The case of Mallows. In *Proceedings of the 31st International Conference on Machine Learning*, 2014a.
- Busa-Fekete, Róbert, Szörényi, Balázs, and Hüllermeier, Eyke. PAC rank elicitation through adaptive sampling of stochastic pairwise preferences. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, 2014b.
- Cho, Grace E. and Meyer, Carl D. Comparison of perturbation bounds for the stationary distribution of a Markov chain. *Linear Algebra and its Applications*, 335(1–3): 137–150, 2001.
- Copeland, A. H. A ‘reasonable’ social welfare function. In *Seminar on Mathematics in Social Sciences, University of Michigan*, 1951.
- De Donder, Philippe, Le Breton, Michel, and Truchon, Michel. Choosing from a weighted tournament. *Mathematical Social Sciences*, 40(1):85–109, 2000.
- Fürnkranz, Johannes and Hüllermeier, Eyke. *Preference Learning*. Springer, 2010.
- Jamieson, Kevin G. and Nowak, Rob. Active ranking using pairwise comparisons. In *Advances in Neural Information Processing Systems*, 2011.
- Jiang, Xiaoye, Lim, Lek-Heng, Yao, Yuan, and Ye, Yinyu. Statistical ranking and combinatorial Hodge theory. *Mathematical Programming*, 127(1):203–244, 2011.
- Laslier, J.-F. *Tournament Solutions and Majority Voting*. Springer-Verlag, 1997.
- Lu, Tyler and Boutilier, Craig. Learning Mallows models with pairwise preferences. In *Proceedings of the 28th International Conference on Machine Learning*, 2011.
- Moulin, Hervé. Choosing from a tournament. *Social Choice and Welfare*, 3(4):271–291, 1986.
- Negahban, Sahand, Oh, Sewoong, and Shah, Devavrat. Iterative ranking from pair-wise comparisons. In *Advances in Neural Information Processing Systems*, 2012.
- Osting, Braxton, Brune, Christoph, and Osher, Stanley. Enhanced statistical rankings via targeted data collection. In *Proceedings of the 30th International Conference on Machine Learning*, 2013.
- Rajkumar, Arun and Agarwal, Shivani. A statistical convergence perspective of algorithms for rank aggregation from pairwise data. In *Proceedings of the 31st International Conference on Machine Learning*, 2014.
- Seneta, E. Perturbation of the stationary distribution measured by ergodicity coefficients. *Advances in Applied Probability*, pp. 228–230, 1988.
- Smith, John H. Aggregation of preferences with variable electorate. *Econometrica: Journal of the Econometric Society*, pp. 1027–1041, 1973.
- Wauthier, Fabian L., Jordan, Michael I., and Jojic, Nebojsa. Efficient ranking from pairwise comparisons. In *Proceedings of the 30th International Conference on Machine Learning*, 2013.

Ranking from Stochastic Pairwise Preferences: Recovering Condorcet Winners and Tournament Solution Sets at the Top

Supplementary Material

Before we give the appendix material, we note the following: In general, all the algorithms described (both previous algorithms and those proposed in this paper) can be run either on a true preference matrix \mathbf{P} as input (where $P_{ij} + P_{ji} = 1$ for all $i \neq j$) or on an empirical preference matrix $\hat{\mathbf{P}}$ as input (where there might be some (i, j) pairs for which $\hat{P}_{ij} = \hat{P}_{ji} = 0$). The notation $\bar{\mathbf{P}}$ used in describing the algorithms allows for both possibilities. All the theorems we give apply to the empirical preference matrix and therefore use the notation $\bar{\mathbf{P}}$; on the other hand, the counter-examples we give in Examples show that even if the corresponding algorithms are run on the true matrix \mathbf{P} , the algorithms can fail to satisfy the desired property, and therefore use the notation \mathbf{P} .

A. Supplement to Section 3 (Related Work and Existing Results)

For completeness, here we include descriptions of the Rank Centrality (RC), Matrix Borda (MB) and SVM-RankAggregation (SVM-RA) algorithms that form the backdrop to our work. For the SMV-RA algorithm, we will need the following definition:

Definition 12 ($\bar{\mathbf{P}}$ -Induced Dataset (Rajkumar & Agarwal, 2014)). Let $\bar{\mathbf{P}} \in [0, 1]^{n \times n}$ satisfy the following conditions: (i) for all $i \neq j$: $\bar{p}_{ij} + \bar{p}_{ji} = 1$ or $\bar{p}_{ij} = \bar{p}_{ji} = 0$ and (ii) for every i : $\bar{p}_{ii} = 0$. Define the $\bar{\mathbf{P}}$ -induced dataset $S_{\bar{\mathbf{P}}} = \{\mathbf{v}_{ij}, z_{ij}\}_{i < j: \bar{p}_{ij} + \bar{p}_{ji} = 1}$ as consisting of the vectors $\mathbf{v}_{ij} = (\bar{\mathbf{p}}_i - \bar{\mathbf{p}}_j) \in \mathbb{R}^n$ for every pair (i, j) such that $i < j$ and $\bar{p}_{ij} + \bar{p}_{ji} = 1$, where $\bar{\mathbf{p}}_i$ denotes the i -th column of $\bar{\mathbf{P}}$, together with binary labels $z_{ij} = \text{sign}(\bar{p}_{ji} - \bar{p}_{ij}) \in \{\pm 1\}$.

Algorithm 4 Rank Centrality (RC) Algorithm ((Negahban et al., 2012))

Input: Pairwise comparison matrix $\bar{\mathbf{P}} \in [0, 1]^{n \times n}$ satisfying the following conditions:

- (i) for all $i \neq j$: $\bar{p}_{ij} + \bar{p}_{ji} = 1$ or $\bar{p}_{ij} = \bar{p}_{ji} = 0$
- (ii) for every i : $\bar{p}_{ii} = 0$

Construct an empirical Markov chain with transition probability matrix $\check{\bar{\mathbf{P}}}$ as follows:

$$\check{\bar{p}}_{ij} = \begin{cases} \frac{1}{n} \bar{p}_{ij} & \text{if } i \neq j \\ 1 - \frac{1}{n} \sum_{k \neq i} \bar{p}_{ik} & \text{if } i = j. \end{cases}$$

Compute $\bar{\boldsymbol{\pi}}$, the stationary probability vector of $\check{\bar{\mathbf{P}}}$

Output: Permutation $\bar{\sigma}_{\text{RC}} \in \text{argsort}(\bar{\boldsymbol{\pi}})$

Algorithm 5 Matrix Borda (MB) Algorithm ((Rajkumar & Agarwal, 2014))

Input: Pairwise comparison matrix $\bar{\mathbf{P}} \in [0, 1]^{n \times n}$ satisfying the following conditions:

- (i) for all $i \neq j$: $\bar{p}_{ij} + \bar{p}_{ji} = 1$ or $\bar{p}_{ij} = \bar{p}_{ji} = 0$
- (ii) for every i : $\bar{p}_{ii} = 0$

For $i = 1$ to n : $\bar{f}_i = \frac{1}{n} \sum_{k=1}^n \bar{p}_{ki}$

Output: Permutation $\bar{\sigma}_{\text{MB}} \in \text{argsort}(\bar{\mathbf{f}})$

B. Supplement to Section 4 (Performance of Rank Centrality, Matrix Borda and SVM-RankAggregation Algorithms under Preferences with Cycles) and Example 3

Additional details for Examples 1–3. For completeness, we first give here the permutations output by different algorithms for the preference matrices given in Examples 1, 2 and 3. For the PM algorithm, the value of the parameter c is chosen as the maximum of the values prescribed by Theorem 10 and 11.

Algorithm 6 SVM-RankAggregation (SVM-RA) Algorithm ((Rajkumar & Agarwal, 2014))

Input: Pairwise comparison matrix $\overline{\mathbf{P}} \in [0, 1]^{n \times n}$ satisfying the following conditions:

- (i) for all $i \neq j$: $\overline{p}_{ij} + \overline{p}_{ji} = 1$ or $\overline{p}_{ij} = \overline{p}_{ji} = 0$
- (ii) for every i : $\overline{p}_{ii} = 0$

Construct $\overline{\mathbf{P}}$ -induced dataset $S_{\overline{\mathbf{P}}}$ (see Definition 12)

If $S_{\overline{\mathbf{P}}}$ is linearly separable by hyperplane through origin, then

train hard-margin linear SVM on $S_{\overline{\mathbf{P}}}$; obtain $\overline{\alpha} \in \mathbb{R}^n$

else

train soft-margin linear SVM (with any suitable value for regularization parameter) on $S_{\overline{\mathbf{P}}}$; obtain $\overline{\alpha} \in \mathbb{R}^n$

For $i = 1$ to n : $\overline{f}_i = \sum_{k=1}^n \overline{\alpha}_k \overline{p}_{ki}$

Output: Permutation $\overline{\sigma}_{\text{SVM-RA}} \in \text{argsort}(\overline{\mathbf{f}})$

For the matrix in Example 1, we have the following:

$$\begin{aligned} \sigma_{\text{RC}} &= (2\ 1\ 6\ 3\ 4\ 7\ 5) \\ \sigma_{\text{MB}} &= (2\ 1\ 6\ 3\ 4\ 7\ 5) \\ \sigma_{\text{SVM-RA}} &= (7\ 6\ 2\ 1\ 5\ 4\ 3) \\ \sigma_{\text{MC}} &= (1\ 2\ 6\ 3\ 4\ 7\ 5) \\ \sigma_{\text{UM}} &= (1\ 2\ 3\ 4\ 5\ 6\ 7) \\ \sigma_{\text{PM}} &= (1\ 2\ 6\ 7\ 3\ 4\ 5) \end{aligned}$$

For the matrix in Example 2, we have the following:

$$\begin{aligned} \sigma_{\text{RC}} &= (4\ 3\ 1\ 2\ 6\ 5) \\ \sigma_{\text{MB}} &= (4\ 1\ 2\ 3\ 6\ 5) \\ \sigma_{\text{SVM-RA}} &= (4\ 2\ 6\ 1\ 3\ 5) \\ \sigma_{\text{MC}} &= (1\ 2\ 3\ 4\ 6\ 5) \\ \sigma_{\text{UM}} &= (1\ 2\ 3\ 4\ 5\ 6) \\ \sigma_{\text{PM}} &= (1\ 2\ 3\ 6\ 5\ 4) \end{aligned}$$

For the matrix in Example 3, we have the following:

$$\begin{aligned} \sigma_{\text{RC}} &= (4\ 6\ 5\ 7\ 1\ 2\ 3\ 8) \\ \sigma_{\text{MB}} &= (4\ 6\ 5\ 7\ 1\ 2\ 3\ 8) \\ \sigma_{\text{SVM-RA}} &= (4\ 6\ 5\ 2\ 1\ 7\ 3\ 8) \\ \sigma_{\text{MC}} &= (2\ 4\ 6\ 5\ 1\ 7\ 3\ 8) \\ \sigma_{\text{UM}} &= (5\ 2\ 4\ 6\ 7\ 3\ 1\ 8) \\ \sigma_{\text{PM}} &= (5\ 4\ 6\ 2\ 7\ 3\ 1\ 8) \end{aligned}$$

We next give additional examples showing that the MB algorithm can fail to rank the Condorcet winner/Copeland set/Markov set at the top even in $\mathcal{P}^{\text{DAG}} \setminus \mathcal{P}^{\text{LN}}$, and the RC algorithm can fail to do so even in $\mathcal{P}^{\text{LN}} \setminus \mathcal{P}^{\text{BTL}}$.

Example 6 (MB algorithm can fail to rank the Condorcet winner/Copeland set/Markov set at the top even in $\mathcal{P}^{\text{DAG}} \setminus \mathcal{P}^{\text{LN}}$). Let $n = 3$, and consider

$$\mathbf{P} = \begin{bmatrix} 0 & 0.4 & 0.4 \\ 0.6 & 0 & 0.1 \\ 0.6 & 0.9 & 0 \end{bmatrix}$$

Here $\text{CW}(\mathbf{P}) = 1$, $\text{CO}(\mathbf{P}) = \text{MA}(\mathbf{P}) = \{1\}$, but the permutation σ_{MB} produced by running MB on \mathbf{P} is $\sigma_{\text{MB}} = (2\ 1\ 3)$.

Example 7 (RC algorithm can fail to rank the Condorcet winner/Copeland set/Markov set at the top even in $\mathcal{P}^{\text{LN}} \setminus \mathcal{P}^{\text{BTL}}$).

Let $n = 8$, and consider

$$\mathbf{P} = \begin{bmatrix} 0 & 0.49 & 0.3095 & 0.3 & 0.3 & 0.3 & 0.3 & 0.3 \\ 0.51 & 0 & 0.1 & 0.49 & 0.3 & 0.3 & 0.3 & 0.3 \\ 0.6905 & 0.9 & 0 & 0.2 & 0.1 & 0.3 & 0.3 & 0.3 \\ 0.7 & 0.51 & 0.8 & 0 & 0.1 & 0.3 & 0.3 & 0.3 \\ 0.7 & 0.7 & 0.9 & 0.9 & 0 & 0.3 & 0.3 & 0.3 \\ 0.7 & 0.7 & 0.7 & 0.7 & 0.7 & 0 & 0.3 & 0.3 \\ 0.7 & 0.7 & 0.7 & 0.7 & 0.7 & 0.7 & 0 & 0.3 \\ 0.7 & 0.7 & 0.7 & 0.7 & 0.7 & 0.7 & 0.7 & 0 \end{bmatrix}$$

Here $\text{CW}(\mathbf{P}) = 1$, $\text{CO}(\mathbf{P}) = \text{MA}(\mathbf{P}) = \{1\}$, but the permutation σ_{RC} produced by running the RC algorithm on \mathbf{P} is $\sigma_{\text{RC}} = (2\ 1\ 3\ 4\ 6\ 5\ 7\ 8)$.

C. Proofs of Results in Sections 5–6

The overall strategy followed by our proofs is largely similar to that of (Rajkumar & Agarwal, 2014): namely, that the empirical pairwise comparison matrix $\hat{\mathbf{P}}$ concentrates around the true pairwise preference matrix \mathbf{P} , and that when $\hat{\mathbf{P}}$ becomes sufficiently close to \mathbf{P} , the specific algorithm satisfies the desired property. However the details differ considerably depending on the algorithm and property of interest. The proofs for the Matrix Copeland (MC) and Unweighted Markov (UM) algorithms are relatively straightforward; the proofs for the Parametrized Markov (PM) algorithm, on the other hand, require additional tools. In particular, as noted in the main text, the proof of Theorem 9 involves reasoning about a new property of preference matrices that we term the *restricted low-noise* (RLN) property; the proofs of Theorems 10–11 make use of the Cho-Meyer perturbation bound for Markov chains. Details follow.

C.1. Proof of Theorem 5

Proof. Let m satisfy the given assumption. Then by Lemma 4, we have with probability at least $1 - \delta$, the following event holds:

$$|p_{ij} - \hat{p}_{ij}| < \frac{\gamma_{\min}}{2} \text{ for all } i, j.$$

Under this event, we have for all i, j : $p_{ij} > \frac{1}{2} \iff \hat{p}_{ij} > \frac{1}{2}$, and therefore $G_{\mathbf{P}} = G_{\hat{\mathbf{P}}}$, thus giving for all $i \in [n]$:

$$\hat{f}_i = \frac{1}{n} \sum_{j=1}^n \mathbf{1}(\hat{p}_{ij} < \frac{1}{2}) = \frac{1}{n} \sum_{j=1}^n \mathbf{1}(p_{ij} < \frac{1}{2}) = \frac{d(i)}{n}.$$

The first claim follows. To prove the second claim, note that if $\mathbf{P} \in \mathcal{P}^{\text{DAG}}$, we have for any i, j , $i \succ_{\mathbf{P}} j \implies d(i) > d(j)$. The claim follows. \square

C.2. Proof of Theorem 6

Proof. Let m satisfy the given assumption. Then by Lemma 4, we have with probability at least $1 - \delta$, the following event holds:

$$|p_{ij} - \hat{p}_{ij}| < \frac{\gamma_{\text{TC}}}{2} \text{ for all } i, j.$$

Under this event, we have for all $i \in \text{TC}(\mathbf{P}), j \notin \text{TC}(\mathbf{P}), \hat{p}_{ij} < \frac{1}{2}$ (and therefore the edge (i, j) is present $G_{\hat{\mathbf{P}}}$), thus giving

$$\text{for all } i \in \text{TC}(\mathbf{P}): \hat{f}_i = \frac{1}{n} \sum_{j=1}^n \mathbf{1}(\hat{p}_{ij} < \frac{1}{2}) \geq \frac{n - |\text{TC}(\mathbf{P})|}{n} \text{ and} \quad (2)$$

$$\text{for all } j \notin \text{TC}(\mathbf{P}): \hat{f}_j = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(\hat{p}_{ji} < \frac{1}{2}) \leq \frac{n - |\text{TC}(\mathbf{P})| - 1}{n} \text{ (since } \hat{p}_{jj} = 0) \quad (3)$$

Thus under the above event, we have $\hat{f}_i > \hat{f}_j$ for all $i \in \text{TC}(\mathbf{P}), j \notin \text{TC}(\mathbf{P})$. The claim follows. \square

C.3. Proof of Theorem 7

Proof. Let m satisfy the given assumption. Then by Lemma 4, we have with probability at least $1 - \delta$, the following event holds:

$$|p_{ij} - \hat{p}_{ij}| < \frac{\gamma_{\min}}{2} \text{ for all } i, j. \quad (4)$$

Under this event, we have for all $i, j : p_{ij} > \frac{1}{2} \iff \hat{p}_{ij} > \frac{1}{2}$, and therefore $G_{\mathbf{P}} = G_{\hat{\mathbf{P}}}$; it is easy to verify that in this case, the Markov chain $\check{\mathbf{H}}$ constructed by the UM algorithm is the same as the Markov chain $\tilde{\mathbf{P}}$ used to define the Markov set $\text{MA}(\mathbf{P})$, and therefore their corresponding stationary distributions $\hat{\pi}$ and π are also the same. The claim follows. \square

C.4. Proof of Theorem 8

Proof. Let m satisfy the given assumption. Then by Lemma 4, we have with probability at least $1 - \delta$, the following event holds:

$$|p_{ij} - \hat{p}_{ij}| < \frac{\gamma_{\min}}{2} \text{ for all } i, j. \quad (5)$$

Under this event, we have for all $i, j : \hat{p}_{ij} > \frac{1}{2} \iff p_{ij} > \frac{1}{2}$. In particular, we have for all $i \in \text{TC}(\mathbf{P}), j \notin \text{TC}(\mathbf{P})$, $\hat{p}_{ij} < \frac{1}{2}$, and thus the Markov chain $\check{\mathbf{H}}$ constructed by the UM algorithm satisfies

$$\check{h}_{ij} = \begin{cases} 0 & \text{if } i \in \text{TC}(\mathbf{P}), j \notin \text{TC}(\mathbf{P}) \\ \frac{1}{n} & \text{if } i \notin \text{TC}(\mathbf{P}), j \in \text{TC}(\mathbf{P}). \end{cases} \quad (6)$$

This implies in particular that the stationary distribution $\hat{\pi}$ of $\check{\mathbf{H}}$ satisfies

$$\hat{\pi}_i = 0 \text{ for all } i \notin \text{TC}(\mathbf{P}).$$

Moreover, for all $i, j \in \text{TC}(\mathbf{P})$, we have $\hat{p}_{ij} > \frac{1}{2} \iff p_{ij} > \frac{1}{2}$ as above, and therefore the items in $\text{TC}(\mathbf{P})$ form a recurrent class in the above Markov chain; this implies that

$$\hat{\pi}_i > 0 \text{ for all } i \in \text{TC}(\mathbf{P}).$$

Thus under the above event, we have $\hat{\pi}_i > \hat{\pi}_j$ for all $i \in \text{TC}(\mathbf{P}), j \notin \text{TC}(\mathbf{P})$. The claim follows. \square

C.5. Auxiliary Results Needed for Proof of Theorem 9

Theorem 9 claims that for $\mathbf{P} \in \mathcal{P}^{\text{BTL}}$, the PM algorithm using any $1 \leq c < \infty$ (for sufficiently large sample size, with high probability) recovers an optimal permutation w.r.t. PD error. As noted earlier, we will prove the result of Theorem 9 for a slightly larger set of preference matrices than \mathcal{P}^{BTL} , namely for all \mathbf{P} satisfying the *restricted low-noise* (RLN) property, defined as follows:

$$\mathcal{P}_n^{\text{RLN}} = \left\{ \mathbf{P} \in \mathcal{P}^n : \forall i \neq j \neq k : i \succ_{\mathbf{P}} j \implies p_{kj} < p_{ki} \right\}.$$

We will show below the following three results:

- (i) $\mathcal{P}^{\text{BTL}} \subseteq \mathcal{P}^{\text{RLN}}$ (Proposition 13).
- (ii) For any $1 \leq c < \infty$, \mathcal{P}^{RLN} is closed under g_c , i.e. $\mathbf{P} \in \mathcal{P}^{\text{RLN}} \implies \mathbf{H}^c = g_c(\mathbf{P}) \in \mathcal{P}^{\text{RLN}}$ (Proposition 14).
- (iii) For any $\mathbf{P} \in \mathcal{P}^{\text{RLN}}$, the permutation σ_{RC} produced by running the RC algorithm on \mathbf{P} satisfies $\sigma_{\text{RC}} \in \operatorname{argmin}_{\sigma \in \mathcal{S}_n} \operatorname{er}_{\mathbf{P}}^{\text{PD}}[\sigma]$ (Theorem 15).

These results will imply that for any $\mathbf{P} \in \mathcal{P}^{\text{RLN}}$ (and therefore in particular for any $\mathbf{P} \in \mathcal{P}^{\text{BTL}}$) and any $1 \leq c < \infty$, the permutation σ_{PM} produced by running the PM algorithm on \mathbf{P} (i.e. the RC algorithm on $\mathbf{H}^c = g_c(\mathbf{P})$) satisfies

$$\sigma_{\text{PM}} \in \operatorname{argmin}_{\sigma \in \mathcal{S}_n} \operatorname{er}_{\mathbf{H}^c}^{\text{PD}}[\sigma].$$

Moreover, since $p_{ij} > \frac{1}{2} \iff g_c(p_{ij}) > \frac{1}{2}$, we have $\operatorname{argmin}_{\sigma \in S_n} \operatorname{er}_{\mathbf{H}^c}^{\text{PD}}[\sigma] = \operatorname{argmin}_{\sigma \in S_n} \operatorname{er}_{\mathbf{P}}^{\text{PD}}[\sigma]$, which further gives that for \mathbf{P} and c as above,

$$\sigma_{\text{PM}} \in \operatorname{argmin}_{\sigma \in S_n} \operatorname{er}_{\mathbf{P}}^{\text{PD}}[\sigma].$$

We will then be able to argue that for sufficiently large sample size m , when $\widehat{\mathbf{P}}$ is sufficiently close to \mathbf{P} , the permutation $\widehat{\sigma}_{\text{PM}}$ returned by running the PM algorithm on $\widehat{\mathbf{P}}$ will also minimize the PD error w.r.t. \mathbf{P} , which will prove Theorem 9.

Proposition 13. $\mathcal{P}_n^{\text{BTL}} \subseteq \mathcal{P}_n^{\text{RLN}}$.

Proof. We have,

$$\begin{aligned} \mathbf{P} \in \mathcal{P}_n^{\text{BTL}} &\implies \text{there exists } \mathbf{w} \in \mathbb{R}_+^n : \text{ for all } i \neq j, p_{ij} = \frac{w_j}{w_i + w_j} \\ &\implies \text{for any } i, j, \text{ if } p_{ij} < p_{ji}, \text{ then } w_i > w_j \\ &\implies \text{for any } i, j, \text{ for any } k \neq \{i, j\}, p_{ik} = \frac{w_k}{w_i + w_k} > p_{jk} = \frac{w_k}{w_j + w_k} \\ &\implies \mathbf{P} \in \mathcal{P}_n^{\text{RLN}}. \end{aligned}$$

This proves the claim. □

Proposition 14. Let $\mathbf{P} \in \mathcal{P}_n^{\text{RLN}}$. Let $1 \leq c < \infty$ and define \mathbf{H}^c as $h_{ij}^c = g_c(p_{ij})$ for all $i \neq j$ and $h_{ii}^c = 0$ for all i . Then $\mathbf{H}^c \in \mathcal{P}_n^{\text{RLN}}$.

Proof. Let $i, j \in [n]$. We have,

$$\begin{aligned} i \succ_{\mathbf{H}^c} j &\implies h_{ij}^c < h_{ji}^c \\ &\implies \frac{p_{ij}^c}{p_{ij}^c + p_{ji}^c} < \frac{p_{ji}^c}{p_{ji}^c + p_{ij}^c} \\ &\implies p_{ij} < p_{ji} \\ &\implies p_{ki} > p_{kj} \text{ for all } k \neq \{i, j\} \quad (\text{since } \mathbf{P} \in \mathcal{P}_n^{\text{RLN}}) \\ &\implies p_{ki} + 1 - p_{jk}p_{ik} > p_{kj} + 1 - p_{jk}p_{ik} \text{ for all } k \neq \{i, j\} \\ &\implies p_{ki}p_{jk} > p_{ik}p_{kj} \text{ for all } k \neq \{i, j\} \\ &\implies \frac{p_{ik}}{p_{ki}} < \frac{p_{kj}}{p_{jk}} \text{ for all } k \neq \{i, j\} \\ &\implies \left(\frac{p_{ik}}{p_{ki}}\right)^c < \left(\frac{p_{kj}}{p_{jk}}\right)^c \text{ for all } c \geq 1, \text{ for all } k \neq \{i, j\} \\ &\implies g_c(p_{ki}) > g_c(p_{kj}) \text{ for all } c \geq 1, \text{ for all } k \neq \{i, j\} \\ &\implies h_{ki}^c > h_{jk}^c \text{ for all } c \geq 1, \text{ for all } k \neq \{i, j\} \end{aligned}$$

Thus we have $\mathbf{H}^c \in \mathcal{P}_n^{\text{RLN}}$. □

Theorem 15. Let $\mathbf{P} \in \mathcal{P}_n^{\text{RLN}}$. Then the permutation σ_{RC} produced by running the Rank Centrality algorithm on \mathbf{P} satisfies

$$\sigma_{\text{RC}} \in \operatorname{argmin}_{\sigma \in S_n} \operatorname{er}_{\mathbf{P}}^{\text{PD}}[\sigma].$$

Proof. Let $\check{\mathbf{P}}$ be the Markov chain constructed from \mathbf{P} by the Rank Centrality algorithm and let π be its stationary distribution. Fix any i, j such that $i \succ_{\mathbf{P}} j$. Assume for the sake of contradiction that $\pi_i \leq \pi_j$. We have from the stationary

distribution equations,

$$\begin{aligned}
 \pi_i &= \sum_{k=1}^n \pi_k \check{p}_{ki} \\
 &= \sum_{k \neq \{i,j\}} \pi_k \check{p}_{ki} + \pi_i \check{p}_{ii} + \pi_j \check{p}_{ji} \\
 &= \sum_{k \neq \{i,j\}} \pi_k \frac{p_{ki}}{n} + \pi_i \left(1 - \sum_{k \neq i} \check{p}_{ik}\right) + \pi_j \frac{p_{ji}}{n} \\
 &= \sum_{k \neq \{i,j\}} \pi_k \frac{p_{ki}}{n} + \pi_i \left(\frac{1}{n} + \sum_{k \neq i} \frac{p_{ki}}{n}\right) + \pi_j \frac{p_{ji}}{n}
 \end{aligned}$$

Similarly we have

$$\pi_j = \sum_{k \neq \{i,j\}} \pi_k \frac{p_{kj}}{n} + \pi_j \left(\frac{1}{n} + \sum_{k \neq j} \frac{p_{kj}}{n}\right) + \pi_i \frac{p_{ij}}{n}$$

Notice that the equations for π_i and π_j are of the form

$$\begin{aligned}
 \pi_i &= a_1 + b_1 \pi_i + c_1 \pi_j \\
 \pi_j &= a_2 + b_2 \pi_j + c_2 \pi_i
 \end{aligned}$$

where

$$\begin{aligned}
 a_1 &= \sum_{k \neq \{i,j\}} \pi_k \frac{p_{ki}}{n}, \quad b_1 = \frac{1}{n} + \sum_{k \neq i} \frac{p_{ki}}{n}, \quad c_1 = \frac{p_{ji}}{n} \\
 a_2 &= \sum_{k \neq \{i,j\}} \pi_k \frac{p_{kj}}{n}, \quad b_2 = \frac{1}{n} + \sum_{k \neq j} \frac{p_{kj}}{n}, \quad c_2 = \frac{p_{ij}}{n}.
 \end{aligned}$$

Since $\mathbf{P} \in \mathcal{P}^{\text{RLN}}$ and $i \succ_{\mathbf{P}} j$, we have that $a_1 \geq a_2, b_1 > b_2$ and $c_1 > c_2$. Observe that $0 < b_1, b_2 \leq 1$. We claim that neither b_1 nor b_2 can be exactly equal to 1. To see this, first consider b_2 . We have $b_2 = 1$ only if $p_{kj} = 1$ for all $k \neq j$. But as $i \succ_{\mathbf{P}} j$, we have $p_{ij} < p_{ji}$ and so this cannot happen. Thus $b_2 < 1$. Next consider b_1 . We have $b_1 = 1$ only if $p_{ki} = 1$ for all $k \neq i$. When this happens, it is easy to see that $\text{CW}(\mathbf{P}) = i$ and so $\pi_i = 1$ and $\pi_j = 0$ for all $j \neq i$. This contradicts our assumption that $\pi_i \leq \pi_j$. So in what follows, assume $0 < b_1, b_2 < 1$. In that case, we have

$$\begin{aligned}
 \pi_i &= \frac{a_1}{1-b_1} + \frac{c_1}{1-b_1} \pi_j \\
 &> \frac{a_2}{1-b_2} + \frac{c_2}{1-b_2} \pi_i \quad (\text{as } \mathbf{P} \in \mathcal{P}^{\text{RLN}} \text{ and by our assumption that } \pi_i \leq \pi_j) \\
 &= \pi_j
 \end{aligned}$$

which contradicts the assumption that $\pi_i \leq \pi_j$. Thus, we must have that

$$\text{for any } i, j, \quad i \succ_{\mathbf{P}} j \implies \pi_i > \pi_j.$$

The claim follows. □

C.6. Proof of Theorem 9

Proof. Let m satisfy the given assumption. Then by Lemma 4, we have with probability at least $1 - \delta$, the following event holds:

$$|p_{ij} - \hat{p}_{ij}| < \frac{\min(\gamma_{\min}, r_{\min})}{2} \text{ for all } i, j. \tag{7}$$

Under this event, we have

$$\begin{aligned}
 \text{for all } i, j: \quad \widehat{p}_{ij} > \frac{1}{2} &\implies p_{ij} > \frac{1}{2} \quad (\text{since } |\widehat{p}_{ij} - p_{ij}| < \frac{\gamma_{\min}}{2} \text{ for all } i, j) \\
 &\implies p_{ki} > p_{kj} \text{ for all } k \neq \{i, j\} \quad (\text{since } \mathbf{P} \in \mathcal{P}^{\text{BTL}} \subseteq \mathcal{P}^{\text{RLN}} \text{ from Proposition 13}) \\
 &\implies \widehat{p}_{ki} > \widehat{p}_{kj} \text{ for all } k \neq \{i, j\} \quad (\text{since } |\widehat{p}_{km} - p_{km}| < \frac{r_{\min}}{2} \text{ for all } k, m).
 \end{aligned}$$

Thus under the above event, we have $\widehat{\mathbf{P}} \in \mathcal{P}^{\text{RLN}}$, and therefore by Proposition 14 we have that the matrix $\widehat{\mathbf{H}}^c$ constructed by the PM algorithm satisfies $\widehat{\mathbf{H}}^c \in \mathcal{P}^{\text{RLN}}$, and by Theorem 15 that

$$\widehat{\sigma}_{\text{PM}} \in \operatorname{argmin}_{\sigma \in \mathcal{S}_n} \operatorname{er}_{\widehat{\mathbf{H}}^c}^{\text{PD}}[\sigma]. \quad (8)$$

Since under the above event we also have

$$\text{for all } i, j: \quad p_{ij} > \frac{1}{2} \iff \widehat{p}_{ij} > \frac{1}{2} \iff \widehat{h}_{ij}^c > \frac{1}{2},$$

it follows that under this event, $\operatorname{argmin}_{\sigma \in \mathcal{S}_n} \operatorname{er}_{\widehat{\mathbf{H}}^c}^{\text{PD}}[\sigma] = \operatorname{argmin}_{\sigma \in \mathcal{S}_n} \operatorname{er}_{\mathbf{P}}^{\text{PD}}[\sigma]$ and therefore

$$\widehat{\sigma}_{\text{PM}} \in \operatorname{argmin}_{\sigma \in \mathcal{S}_n} \operatorname{er}_{\mathbf{P}}^{\text{PD}}[\sigma].$$

This proves the result. □

C.7. Auxiliary Results Needed for Proofs of Theorems 10–11

The first result we will need for proving Theorems 10–11 is the following perturbation bound for Markov chains:

Theorem 16 ((Cho & Meyer, 2001; Seneta, 1988)). *Let \mathbf{R} and \mathbf{R}' be the transition probability matrices of finite Markov chains having unique stationary distributions $\boldsymbol{\pi}$ and $\boldsymbol{\pi}'$ respectively. Then*

$$\|\boldsymbol{\pi} - \boldsymbol{\pi}'\|_1 \leq \frac{\|\mathbf{R} - \mathbf{R}'\|_\infty}{1 - \tau(\mathbf{R})},$$

where $1 - \tau(\mathbf{R}) = \min_{i \neq j} \sum_{k=1}^n \min\{r_{ik}, r_{jk}\}$.

We will use the above theorem to show the following two results:

- (i) For any \mathbf{P} and sufficiently large c , the permutation σ_{PM} produced by running the PM algorithm on \mathbf{P} satisfies $\sigma_{\text{PM}}(i) < \sigma_{\text{PM}}(j)$ for all $i \in \text{MA}(\mathbf{P}), j \notin \text{MA}(\mathbf{P})$ (Theorem 19).
- (ii) For any \mathbf{P} and sufficiently large c , the permutation σ_{PM} produced by running the PM algorithm on \mathbf{P} satisfies $\sigma_{\text{PM}}(i) < \sigma_{\text{PM}}(j)$ for all $i \in \text{TC}(\mathbf{P}), j \notin \text{TC}(\mathbf{P})$ (Theorem 20).

In proving these results, we will make use of Theorem 16 to bound the difference between the stationary probability vectors $\boldsymbol{\pi}$ and $\boldsymbol{\pi}^c$ of the Markov chains associated with the matrices $\mathbf{H}^c = g_c(\mathbf{P})$ and $\mathbf{H} = g_\infty(\mathbf{P})$. For this, we will find the following two lemmas useful:

Lemma 17. *Let $\mathbf{P} \in \mathcal{P}_n$. Let $\mathbf{H} = g_\infty(\mathbf{P})$, with $h_{ij} = \mathbf{1}(p_{ij} > 1/2)$, and let $\check{\mathbf{H}}$ be the Markov chain transition matrix corresponding to \mathbf{H} as constructed by the Rank Centrality algorithm applied to \mathbf{H} . Then*

$$1 - \tau(\check{\mathbf{H}}) \geq \frac{1}{n}.$$

Proof. We have,

$$\begin{aligned}
 1 - \tau(\check{\mathbf{H}}) &= \min_{i \neq j} \sum_{k=1}^n \min\{\check{h}_{ik}, \check{h}_{jk}\} \\
 &= \min_{i \neq j} \left(\sum_{k \neq \{i, j\}} \min\{\check{h}_{ik}, \check{h}_{jk}\} + \min\{\check{h}_{ii}, \check{h}_{ji}\} + \min\{\check{h}_{ij}, \check{h}_{jj}\} \right) \\
 &= \min_{i \neq j} \left(\sum_{k \neq \{i, j\}} \min\{\check{h}_{ik}, \check{h}_{jk}\} + \check{h}_{ji} + \check{h}_{ij} \right) \text{ (since } \check{h}_{ii} \geq \check{h}_{ij}, \check{h}_{ji} \text{ for all } i \neq j) \\
 &= \min_{i \neq j} \left(\sum_{k \neq \{i, j\}} \min\{\check{h}_{ik}, \check{h}_{jk}\} + \check{h}_{ji} + \check{h}_{ij} \right) \\
 &= \min_{i \neq j} \left(\sum_{k \neq \{i, j\}} \min\{\check{h}_{ik}, \check{h}_{jk}\} + \frac{1}{n} \right) \text{ (since } \check{h}_{ji} + \check{h}_{ij} = \frac{1}{n} \text{ for all } i \neq j) \\
 &\geq \frac{1}{n}.
 \end{aligned}$$

□

Lemma 18. Let $\mathbf{P} \in \mathcal{P}_n$ and $1 \leq c < \infty$. Let $\mathbf{H}^c = g_c(\mathbf{P})$ and $\mathbf{H} = g_\infty(\mathbf{P})$, and let $\check{\mathbf{H}}^c, \check{\mathbf{H}}$ be the Markov chain transition matrices corresponding to \mathbf{H}^c and \mathbf{H} as constructed by the Rank Centrality algorithm applied to \mathbf{H}^c and \mathbf{H} . Let $\alpha_{\min} = \min_{i \neq j: i \succ_{\mathbf{P}} j} \frac{p_{ji}}{p_{ij}}$. Let $\epsilon > 0$. If $c \geq \frac{\ln(2/\epsilon)}{\ln(\alpha_{\min})}$, then

$$\|\check{\mathbf{H}}^c - \check{\mathbf{H}}\|_\infty < \epsilon$$

Proof. Consider any $i \neq j$. If $p_{ij} < \frac{1}{2}$, we have

$$|\check{h}_{ij}^c - \check{h}_{ij}| = \frac{|g_c(p_{ij}) - \mathbf{1}(p_{ij} > \frac{1}{2})|}{n} = \frac{g_c(p_{ij})}{n} = \frac{1}{n(1 + (\frac{p_{ji}}{p_{ij}})^c)}.$$

Similarly if $p_{ij} > \frac{1}{2}$, we have

$$|\check{h}_{ij}^c - \check{h}_{ij}| = \frac{|g_c(p_{ij}) - \mathbf{1}(p_{ij} > \frac{1}{2})|}{n} = \frac{|g_c(p_{ij}) - 1|}{n} = \frac{1}{n(1 + (\frac{p_{ji}}{p_{ij}})^c)}.$$

In both cases, we have

$$|\check{h}_{ij}^c - \check{h}_{ij}| \leq \frac{1}{n(1 + \alpha_{\min}^c)}.$$

Let c satisfy the given assumption. Then we have for all $i \neq j$,

$$|\check{h}_{ij}^c - \check{h}_{ij}| \leq \frac{\epsilon}{2n}.$$

Moreover, for all i , we have

$$|\check{h}_{ii}^c - \check{h}_{ii}| = \left| \sum_{k \neq i} (\check{h}_{ik}^c - \check{h}_{ik}) \right| \leq \sum_{k \neq i} |\check{h}_{ik}^c - \check{h}_{ik}| \leq \frac{\epsilon}{2}.$$

Thus

$$\|\check{\mathbf{H}}^c - \check{\mathbf{H}}\|_\infty = \max_i \sum_k |\check{h}_{ik}^c - \check{h}_{ik}| \leq \max_i \left(|\check{h}_{ii}^c - \check{h}_{ii}| + \sum_{k \neq i} |\check{h}_{ik}^c - \check{h}_{ik}| \right) \leq \frac{\epsilon}{2} + \frac{(n-1)\epsilon}{2n} < \epsilon.$$

□

Theorem 19. Let $\mathbf{P} \in \mathcal{P}_n$. Let $\mathbf{H} = g_\infty(\mathbf{P})$, and let $\check{\mathbf{H}}$ be the Markov chain transition matrix corresponding to \mathbf{H} as constructed by the Rank Centrality algorithm applied to \mathbf{H} . Let $\boldsymbol{\pi}$ be the stationary probability vector of $\check{\mathbf{H}}$. Let $\beta_2(\mathbf{P}) = \max_{i \in [n]} \pi_i - \max_{k \notin \text{MA}(\mathbf{P})} \pi_k$ and $\alpha_{\min} = \min_{i,j:i \succ_{\mathbf{P}} j} \frac{p_{ji}}{p_{ij}}$. If $c \geq \frac{\ln(4n/\beta_2(\mathbf{P}))}{\ln(\alpha_{\min})}$, then the permutation σ_{PM} produced by running the PM algorithm on \mathbf{P} satisfies

$$\sigma_{\text{PM}}(i) < \sigma_{\text{PM}}(j) \text{ for all } i \in \text{MA}(\mathbf{P}), j \notin \text{MA}(\mathbf{P}).$$

Proof. Let c satisfy the given assumption. Let $\check{\mathbf{H}}^c$ be the Markov chain transition matrix corresponding to $\mathbf{H}^c = g_c(\mathbf{P})$ as constructed by the Rank Centrality algorithm applied to \mathbf{H}^c . Let $\boldsymbol{\pi}^c$ be the stationary probability vector of $\check{\mathbf{H}}^c$. By Lemma 18, we have

$$\|\check{\mathbf{H}}^c - \check{\mathbf{H}}\|_\infty \leq \frac{\beta_2(\mathbf{P})}{2n}.$$

By Theorem 16 and Lemma 17, it then follows that

$$\|\boldsymbol{\pi}^c - \boldsymbol{\pi}\|_\infty \leq \|\boldsymbol{\pi}^c - \boldsymbol{\pi}\|_1 \leq \frac{\beta_2(\mathbf{P})}{2}.$$

By definition of $\beta_2(\mathbf{P})$ and $\text{MA}(\mathbf{P})$, this implies in particular that

$$\pi_i^c > \pi_j^c \text{ for all } i \in \text{MA}(\mathbf{P}), j \notin \text{MA}(\mathbf{P}).$$

The claim follows. □

Theorem 20. Let $\mathbf{P} \in \mathcal{P}_n$. Let $\alpha_{\min} = \min_{i,j:i \succ_{\mathbf{P}} j} \frac{p_{ji}}{p_{ij}}$. If $c \geq \frac{(n+1)\ln(4n)}{\ln(\alpha_{\min})}$, then the permutation σ_{PM} produced by the running the PM algorithm on \mathbf{P} satisfies

$$\sigma_{\text{PM}}(i) < \sigma_{\text{PM}}(j) \text{ for all } i \in \text{TC}(\mathbf{P}), j \notin \text{TC}(\mathbf{P}).$$

Proof. Let c satisfy the given assumption. Let $\check{\mathbf{H}}^c$ and $\check{\mathbf{H}}$ be the Markov chain transition matrices corresponding to $\mathbf{H}^c = g_c(\mathbf{P})$ and $\mathbf{H} = g_\infty(\mathbf{P})$ as constructed by the Rank Centrality algorithm applied to \mathbf{H}^c and \mathbf{H} . Let $\boldsymbol{\pi}^c, \boldsymbol{\pi}$ be the stationary probability vectors of $\check{\mathbf{H}}^c, \check{\mathbf{H}}$. By Lemma 18, we have

$$\|\check{\mathbf{H}}^c - \check{\mathbf{H}}\|_\infty \leq \frac{1}{2n^{n+1}}.$$

By Theorem 16 and Lemma 17, it then follows that

$$\|\boldsymbol{\pi}^c - \boldsymbol{\pi}\|_\infty \leq \|\boldsymbol{\pi}^c - \boldsymbol{\pi}\|_1 \leq \frac{1}{2n^n}.$$

By definition of \mathbf{H} , we have that $\pi_i > 0$ if and only if $i \in \text{TC}(\mathbf{P})$. Now, define

$$\pi_{\max}^{\text{TC}} = \max_{i \in \text{TC}(\mathbf{P})} \pi_i$$

$$\pi_{\min}^{\text{TC}} = \min_{i \in \text{TC}(\mathbf{P})} \pi_i,$$

and let

$$m_0 \in \operatorname{argmin}_{i \in \text{TC}(\mathbf{P})} \pi_i.$$

If $|\text{TC}(\mathbf{P})| = 1$, then clearly m_0 is the unique Condorcet winner and $\pi_{m_0} = 1$, with $\pi_j = 0$ for $j \neq m_0$; in this case

$$\pi_{\min}^{\text{TC}} = 1 \geq \frac{1}{n^n}.$$

If $|\text{TC}(\mathbf{P})| > 1$, then it must be the case that $m_0 \succ_{\mathbf{P}} m_1$ for some $m_1 \in \text{TC}(\mathbf{P})$ (otherwise $m_0 \notin \text{TC}(\mathbf{P})$). From the stationary equations, we have that

$$\pi_{\min}^{\text{TC}} = \pi_{m_0} = \sum_{k=1}^n \pi_k \check{h}_{km_0} \geq \pi_{m_1} \check{h}_{m_1 m_0} \geq \frac{\pi_{m_1}}{n} \quad (\text{from definition of } \check{\mathbf{H}} \text{ and since } m_0 \succ_{\mathbf{P}} m_1).$$

If $|\text{TC}(\mathbf{P})| = 2$, then $\pi_{m_1} = \pi_{\max} \geq \frac{1}{n}$, which again gives

$$\pi_{\min}^{\text{TC}} \geq \frac{1}{n^2} \geq \frac{1}{n^n}.$$

One can similarly show that regardless of the size of $\text{TC}(\mathbf{P})$, we have

$$\pi_{\min}^{\text{TC}} \geq \frac{1}{n^n}.$$

Thus the above result gives us

$$\|\boldsymbol{\pi}^c - \boldsymbol{\pi}\|_{\infty} \leq \frac{\pi_{\min}^{\text{TC}}}{2}.$$

By definition of π_{\min}^{TC} , this implies in particular that

$$\pi_i^c > \pi_j^c \text{ for all } i \in \text{TC}(\mathbf{P}), j \notin \text{TC}(\mathbf{P}).$$

The claim follows. □

C.8. Proof of Theorem 10

Proof. Let c and m satisfy the given assumptions. Applying Theorem 19 to $\widehat{\mathbf{P}}$, we have that $\widehat{\sigma}_{\text{PM}}$ satisfies

$$\widehat{\sigma}_{\text{PM}}(i) < \widehat{\sigma}_{\text{PM}}(j) \text{ for all } i \in \text{MA}(\widehat{\mathbf{P}}), j \notin \text{MA}(\widehat{\mathbf{P}}).$$

Now by Lemma 4, we have with probability at least $1 - \delta$, the following event holds:

$$|p_{ij} - \widehat{p}_{ij}| < \frac{\gamma_{\min}}{2} \text{ for all } i, j. \tag{9}$$

Under this event, we have

$$\text{for all } i, j: \quad p_{ij} > \frac{1}{2} \iff \widehat{p}_{ij} > \frac{1}{2}. \tag{10}$$

Thus under the above event, we have $\text{MA}(\widehat{\mathbf{P}}) = \text{MA}(\mathbf{P})$. The claim follows. □

C.9. Proof of Theorem 11

Proof. Let c and m satisfy the given assumptions. Applying Theorem 20 to $\widehat{\mathbf{P}}$, we have that $\widehat{\sigma}_{\text{PM}}$ satisfies

$$\widehat{\sigma}_{\text{PM}}(i) < \widehat{\sigma}_{\text{PM}}(j) \text{ for all } i \in \text{TC}(\widehat{\mathbf{P}}), j \notin \text{TC}(\widehat{\mathbf{P}}).$$

Now by Lemma 4, we have with probability at least $1 - \delta$, the following event holds:

$$|p_{ij} - \widehat{p}_{ij}| < \frac{\gamma_{\min}}{2} \text{ for all } i, j. \tag{11}$$

Under this event, we have

$$\text{for all } i, j: \quad p_{ij} > \frac{1}{2} \iff \widehat{p}_{ij} > \frac{1}{2}. \tag{12}$$

Thus under the above event, we have $\text{TC}(\widehat{\mathbf{P}}) = \text{TC}(\mathbf{P})$. The claim follows. □