



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Rank Centrality: Ranking from Pairwise Comparisons

Sahand Negahban, Sewoong Oh, Devavrat Shah

To cite this article:

Sahand Negahban, Sewoong Oh, Devavrat Shah (2017) Rank Centrality: Ranking from Pairwise Comparisons. Operations Research 65(1):266-287. <https://doi.org/10.1287/opre.2016.1534>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2016, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Rank Centrality: Ranking from Pairwise Comparisons

Sahand Negahban,^a Sewoong Oh,^b Devavrat Shah^c

^a Department of Statistics, Yale University, New Haven, Connecticut 06510; ^b Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana–Champaign, Urbana, Illinois 61801; ^c Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, Cambridge, Massachusetts 02139

Contact: sahand.negahban@yale.edu (SN); swoh@illinois.edu (SO); devavrat@mit.edu (DS)

Received: January 15, 2014

Revised: December 2, 2014, November 12, 2015

Accepted: May 6, 2016

Published Online in Articles in Advance:
October 18, 2016

Subject Classifications: analysis of algorithms;
games/group decisions: voting/committees;
statistics: estimation

Area of Review: Games, Information, and
Networks

<https://doi.org/10.1287/opre.2016.1534>

Copyright: © 2017 INFORMS

Abstract. The question of aggregating pairwise comparisons to obtain a global ranking over a collection of objects has been of interest for a very long time: be it ranking of online gamers (e.g., MSR’s TrueSkill system) and chess players, aggregating social opinions, or deciding which product to sell based on transactions. In most settings, in addition to obtaining a ranking, finding ‘scores’ for each object (e.g., player’s rating) is of interest for understanding the intensity of the preferences.

In this paper, we propose RANK CENTRALITY, an iterative rank aggregation algorithm for discovering scores for objects (or items) from pairwise comparisons. The algorithm has a natural *random walk* interpretation over the graph of objects with an edge present between a pair of objects if they are compared; the score, which we call Rank Centrality, of an object turns out to be its stationary probability under this random walk.

To study the efficacy of the algorithm, we consider the popular Bradley-Terry-Luce (BTL) model (equivalent to the Multinomial Logit (MNL) for pairwise comparisons) in which each object has an associated score that determines the probabilistic outcomes of pairwise comparisons between objects. In terms of the pairwise marginal probabilities, which is the main subject of this paper, the MNL model and the BTL model are identical. We bound the finite sample error rates between the scores assumed by the BTL model and those estimated by our algorithm. In particular, the number of samples required to learn the score well with high probability depends on the structure of the comparison graph. When the Laplacian of the comparison graph has a strictly positive spectral gap, e.g., each item is compared to a subset of randomly chosen items, this leads to dependence on the number of samples that is nearly order optimal.

Experimental evaluations on synthetic data sets generated according to the BTL model show that our algorithm performs as well as the maximum likelihood estimator for that model and outperforms other popular ranking algorithms.

Funding: This research is supported in part by Army Research Office MURI [Award W911NF-11-1-0036], National Science Foundation [Grants CMMI-1462158, MES-1450848, CNS-1527754, and CCF-1553452].

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/opre.2016.1534>.

Keywords: rank aggregation • rank centrality • Markov chain • random walk

1. Introduction

Rank aggregation is an important task in a wide range of learning and social contexts arising in recommendation systems, information retrieval, and sports and competitions. Given n items, we wish to infer relevancy scores or an ordering on the items based on partial orderings provided through many (possibly contradictory) samples. Frequently, the available data that are presented to us is in the form of a comparison: player A defeats player B ; book A is purchased when books A and B are displayed (a bigger collection of books implies multiple pairwise comparisons); movie A is liked more compared to movie B . From such partial preferences in the form of comparisons, we frequently wish to deduce not only the order of the underlying objects, but also the scores associated with the objects

themselves so as to deduce the intensity of the resulting preference order.

For example, the Microsoft TrueSkill engine assigns scores to online gamers based on the outcomes of (pairwise) games between players. Indeed, it assumes that each player has inherent “skill” and the outcomes of the games are used to learn these skill parameters, which in turn lead to scores associated with each player. In most such settings, similar model-based approaches are employed.

In this paper, we have set out with the following goal: develop an algorithm for the above stated problem in which (a) is computationally simple, (b) works with available (comparison) data only, and (c) when data is generated as per a reasonable model, then the algorithm should do as well as the best model aware

algorithm. The main result of this paper is an affirmative answer to these questions.

Related work. Most rating-based systems rely on users to provide explicit numeric scores for their interests. While these assumptions have led to a flurry of theoretical research for item recommendations based on matrix completion (cf. Candès and Recht 2009, Keshtavan et al. 2010, Negahban and Wainwright 2012), arguably numeric scores provided by individual users are generally inconsistent. Furthermore, in a number of learning contexts as illustrated above, explicit scores are not available.

These observations have led to the need to develop methods that can aggregate such forms of ordering information into relevance ratings. In general, however, designing consistent aggregation methods can be challenging due in part to possible contradictions between individual preferences. For example, if we consider items A , B , and C , one user might prefer A to B , while another prefers B to C , and a third user prefers C to A . Such problems have been well studied starting with (and potentially even before) Condorcet (1785). In the celebrated work by Arrow (1963), existence of a rank aggregation algorithm with reasonable sets of properties (or axioms) was shown to be impossible.

In this paper, we are interested in a more restrictive setting: we have outcomes of pairwise comparisons between pairs of items, rather than a complete ordering as considered in Arrow (1963). Based on those pairwise comparisons, we want to obtain a ranking of items along with a score for each item indicating the intensity of the preference. One reasonable way to think about our setting is to imagine that there is a distribution over orderings or rankings or permutations of items (also known as the *discrete choice model* in the literature on social choice) and every time a pair of items is compared, the outcome is generated as per this underlying distribution. Examples of popular distributions over permutations include the Plackett-Luce model (Luce 1959, Plackett 1975) and the Mallows model (Mallows 1957). With this, our question becomes even harder than the setting considered by Arrow (1963) as, in that work, effectively the entire distribution over permutations was already known!

Indeed, such hurdles have not stopped the scientific community as well as practical designers from designing such systems. Chess rating systems and the more recent MSR TrueSkill Ranking system are prime examples. Our work falls precisely into this realm: design algorithms that work well in practice, makes sense in general, and perhaps more importantly, have attractive theoretical properties under common comparative judgment models.

An important and landmark model in this class is called the Plackett-Luce model, which is also known as the multinomial logit (MNL) model (cf. McFadden

1973) in the operations research and social science literature. A special case of the Plackett-Luce model applied to pairwise comparisons is known as the Bradley-Terry-Luce (BTL) model (Bradley and Terry 1955, Luce 1959). It has been the backbone of many practical system designs including pricing in the airline industry, e.g., see Talluri and Van Ryzin (2005). Adler et al. (1994) used such models to design adaptive algorithms that select the winner from a small number of rounds. Interestingly enough, the (near-)optimal performance of their adaptive algorithm for winner selection is matched by our nonadaptive algorithm for assigning scores to obtain global rankings of all players.

We propose a new rank aggregation algorithm,¹ which we call RANK CENTRALITY, that builds on a long line of research in using eigenvectors of certain matrices to find global rankings of items, which dates back to Seeley (1949). This line of research is referred to as *spectral ranking* and for an extensive survey we refer to Vigna (2009). Given pairwise comparisons of items from a single individual on all possible choices of pairs, Wei (1952) introduced a ranking algorithm based on the leading eigenvector of the matrix representing the comparisons outcome. A slight generalization accounting for data from multiple decision makers was proposed by Kendall (1955). Keener (1993), and more recent work by Dwork et al. (2001), proposed several variations of spectral algorithms for ranking from pairwise comparisons. We propose RANK CENTRALITY for ranking from pairwise comparisons by using the leading eigenvector of a particular matrix formed by constructing a Markov chain corresponding to a random walk on a graph. Although it appears to be similar to the existing spectral ranking approaches, the precise form of the algorithm proposed is distinct and this precise form *does matter*: the empirical results using synthetic data presented in Section 3.3 make this clear. In summary, building on the classical field of spectral ranking, we propose a novel spectral ranking algorithm and provide a firm theoretical grounding by showing that it is a provably near-optimal estimator for a popular discrete choice model, i.e., the BTL model formally defined in Section 2.1.

Numerous spectral ranking algorithms have been proposed in the past, one of the most popular examples being PageRank (Brin and Page 1998). However, almost invariably, the question of when one should choose to use a particular spectral ranking algorithm is left open. One notable exception is the work of Altman and Tennenholtz (2005), which provides a set of axioms satisfied by the PageRank algorithm and proves that PageRank is the only rank aggregation algorithm that satisfies those particular axioms. Hence, it provides a guideline for deciding when PageRank should be used, i.e., in applications where the specific set of axioms

make sense. In a similar spirit, Rank Centrality is a spectral ranking algorithm with a theoretical justification suggesting that it should be used in applications where the BTL or MNL model makes sense (in the remainder of this manuscript, we shall use the BTL model to represent the BTL and MNL models).

There has been significant work on rankings from pairwise comparison in the last several years. A popular model is a distribution over permutations known as the Mallows model, which assigns probability to observed rankings according to the Kendall- τ distance to a true ranking. Since the maximum likelihood estimation is provably difficult, Dwork et al. (2001b) studied this problem (also known as the Kemeny optimization) when full rankings are observed and provided a 2-approximation algorithm. This was later improved by Ailon et al. (2008) and also generalized to partial rankings (Ailon 2010). Recently, Lu and Boutilier (2011) proposed an expectation-maximization approach with novel sampling schemes to learn the Mallows model from pairwise comparisons. These distance-based approaches aim to provide good approximation algorithms for the provably difficult problem of minimizing the Kendall- τ distance and some variations of it (e.g., Farnoud et al. 2012).

Learning to rank from pairwise comparisons has also been studied in applications where one might observe more than just the ordinal outcome of pairwise comparisons. Additional data on cardinal preferences such as the margin of victory (the difference between the winning team's score and the losing team's score) in a football match has led to score-based methods for ranking, where the goal is to find scores for each team such that the difference of the scores is consistent with the observed margins of victory (Hochbaum 2006, Gleich and Lim 2011, Jiang et al. 2011). More recently, Volkovs and Zemel (2012) proposed a unified model that generalizes both the BTL model and the cardinal preferences. These approaches add to the traditional approaches based on some notion of distance, such as the Kendall- τ distance, and probabilistic models, such as the BTL model.

Another probabilistic model directly parameterizes the distribution of pairwise comparisons for all the pairs and asks the question of whether existing pairwise ranking algorithms are consistent or not (Duchi et al. 2010, Rajkumar and Agarwal 2014). It is shown that many existing algorithms do not meet the proposed "consistency" criteria and new regret/optimization-based algorithms are presented.

The algorithm proposed by Ammar and Shah (2011) can be viewed as a natural adaption of Borda count based on pairwise comparison data. They establish it to be equivalent to Borda count based on entire distribution when perfect pairwise marginals are available, i.e., a large sample limit. In Braverman and Mossel

(2008), the authors present an algorithm that produces an ordering based on $O(n \log n)$ pairwise comparisons on adaptively selected pairs. They assume that there is an underlying true ranking and one observes noisy comparison results. Each time a pair is queried, we are given the true ordering of the pair with probability $1/2 + \gamma$ for some $\gamma > 0$, which does not depend on the items being compared.

Our contributions. In this paper, we introduce RANK CENTRALITY, an iterative algorithm that takes the noisy comparison answers between a subset of all possible pairs of items as input and produces scores for each item as the output. The proposed algorithm has a nice intuitive explanation. Consider a graph with nodes/vertices corresponding to the items of interest (e.g., players). Construct a random walk on this graph where at each time, the random walk is likely to go from vertex i to vertex j if items i and j were ever compared; and if so, the likelihood of going from i to j depends on how often i lost to j . That is, the random walk is more likely to move to a neighbor who has more "wins." How frequently this walk visits a particular node in the long run, or equivalently the stationary distribution, is the score of the corresponding item. Thus, effectively this algorithm captures preference of the given item versus all of the others, not just immediate neighbors: the global effect induced by transitivity of comparisons is captured through the stationary distribution.

Such an interpretation of the stationary distribution of a Markov chain or a random walk has been an effective measure of relative importance of a node in wide class of graph problems, popularly known as the *Network Centrality*; cf. (Newman 2010). Notable examples of such network centralities include the random surfer model on the web graph for the version of the PageRank (Brin and Page 1998), which computes the relative importance of a web page, a model of a random crawler in a peer-to-peer file-sharing network to assign trust value to each peer in EigenTrust (Kamvar et al. 2003), and a random walk interpretation of Rumor Centrality that assigns likelihood to each node for being the source of information (or rumor) spread in a network graph based on the footprint of infection under the susceptible-infected model Shah and Zaman (2011, 2016).

The computation of the stationary distribution of the Markov chain boils down to "power iteration" using transition matrix leading to a nice iterative algorithm. To establish rigorous properties of the algorithm, we analyze its performance under the BTL model described in Section 2.1.

Formally, we establish the following result: given n items, when comparisons between randomly chosen $\omega(n \log n)$ pairs of items are produced as per an (unknown) underlying BTL model, RANK CENTRALITY

learns the true score up to an arbitrary accuracy with high probability as $n \rightarrow \infty$. It should be noted that $\Omega(n \log n)$ is a necessary number of (random) comparisons for any algorithm to even produce a consistent ranking with high probability since with fewer edges (comparisons) the resulting random graph will be disconnected with positive probability. In that sense, RANK CENTRALITY is nearly order optimal.

In general, the comparisons may not be available between randomly chosen pairs. Let $G = ([n], E)$ denote the graph of comparisons between these n objects with an edge $(i, j) \in E$ if and only if objects i and j are compared. In this setting, we establish that with $O(\xi^{-2} n \text{poly}(\log n))$ comparisons, RANK CENTRALITY learns the true score of the underlying BTL model up to an arbitrarily small error with high probability. Here, ξ is the spectral gap for the Laplacian of G and this is how the graph structure of comparisons plays a role. Indeed, as a special case when comparisons are chosen at random, the induced graph is Erdős-Rényi for which ξ is strictly positive, independent of n , with high probability, leading to the (order) optimal performance of the algorithm as stated earlier.

To understand the performance of RANK CENTRALITY compared to the other options, we perform an experimental study. It shows that the performance of RANK CENTRALITY is identical to the ML estimation of the BTL model. Furthermore, it outperforms other popular choices. In summary, RANK CENTRALITY (a) is computationally simple, (b) always produces a solution using available data, and (c) has near-optimal performance with respect to a reasonable generative model.

Some remarks about our analytic technique. Our analysis boils down to studying the induced stationary distribution of the random walk or Markov chain corresponding to the algorithm. Like most such scenarios, the only hope to obtain meaningful results for such “random noisy” Markov chain is to relate it to a stationary distribution of a *known* Markov chain. Through recent concentration of measure results for random matrices and comparison technique using Dirichlet forms for characterizing the spectrum of reversible/self-adjoint operators, along with the known expansion property of the random graph, we obtain the eventual result. Indeed, it is the consequence of such existing powerful results that lead to near-optimal analytic results for the random comparison model and characterization of the algorithm’s performance for general setting.

As an important comparison, we provide analysis of sample complexity required by the maximum likelihood estimator (MLE) using the state-of-the-art analytic techniques; cf. Negahban and Wainwright (2012). Subsequent to our work, Hajek et al. (2014) extended our analysis of MLE and established that

MLE also achieves near-optimal performance guarantees (up to a logarithmic factor). Our numerical experiments suggests something even stronger, the resulting error is effectively identical for both MLE and RANK CENTRALITY.

Organization. The remainder of the paper is organized as follows. In Section 2, we describe the model, problem statement, and the rank centrality algorithm. Section 3 describes the main results—the key theoretical properties of rank centrality as well as its empirical performance in the context of two real data sets from NASCAR and One Day International (ODI) cricket. We provide comparison of the Rank Centrality with the maximum likelihood estimator using the existing analytic techniques in the same section. We derive the Cramer-Rao lower bound on the square error for estimating parameters by any algorithm—across a range of parameters, the performance of Rank Centrality and MLE matches the lower bound implied by Cramer-Rao bound as explained in Section 3 as well. Finally, Section 4 details proofs of all results. We discuss and conclude in Section 5.

Notation. In the remainder of this paper, we use C, C' , etc. to denote absolute constants, and their value might change from line to line. We use A^T to denote the transpose of a matrix. The Euclidean norm of a vector is denoted by $\|x\| = \sqrt{\sum_i x_i^2}$, and the operator norm of a linear operator is denoted by $\|A\|_2 = \max_x (x^T A x / (x^T x))$. When we say *with high probability*, we mean that the probability of a sequence of events $\{\mathcal{A}_n\}_{n=1}^\infty$ goes to one as n grows: $\lim_{n \rightarrow \infty} \mathbb{P}(\mathcal{A}_n) = 1$. Also define $[n] = \{1, 2, \dots, n\}$ to be the set of all integers from 1 to n .

2. Model, Problem Statement, and Algorithm

2.1. Model

In this section, we discuss a model of comparisons between various items. This model will be used to analyze the Rank Centrality algorithm.

Bradley-Terry-Luce model for comparative judgment.

When comparing pairs of items from n items of interest, represented as $[n] = \{1, \dots, n\}$, the Bradley-Terry-Luce model assumes that there is a weight or score $w_i \in \mathbb{R}_+ \equiv \{x \in \mathbb{R} : x > 0\}$ associated with each item $i \in [n]$. The outcome of a comparison for a pair of items i and j is determined only by the corresponding weights w_i and w_j . Let Y_{ij}^l denote the outcome of the l th comparison of the pair i and j , such that $Y_{ij}^l = 1$ if j is preferred over i and 0 otherwise. Then, according to the BTL model,

$$Y_{ij}^l = \begin{cases} 1 & \text{with probability } \frac{w_j}{w_i + w_j}, \\ 0 & \text{otherwise.} \end{cases}$$

Furthermore, conditioned on the score vector $w = (w_1, \dots, w_n)^T$, it is assumed that the random variables Y_{ij}^l 's are independent of one another for all i, j , and l .

Since the BTL model is invariant under the scaling of the scores, an n -dimensional representation of the scores is not unique. Indeed, under the BTL model, a score vector $w \in \mathbb{R}_+^n$ is the equivalence class $[w] = \{w' \in \mathbb{R}_+^n \mid w' = aw, \text{ for some } a > 0\}$. The outcome of a comparison only depends on the equivalence class of the score vector.

To get a unique representation, we represent each equivalence class by its projection onto the standard orthogonal simplex such that $\sum_i w_i = 1$. This representation naturally defines a distance between two equivalent classes as the Euclidean distance between two projections:

$$d(w, w') = \left\| \frac{1}{\langle w, \mathbb{1} \rangle} w - \frac{1}{\langle w', \mathbb{1} \rangle} w' \right\|.$$

Our main result provides an upper bound on the (normalized) distance between the estimated score vector and the true underlying score vector.

Bradley-Terry-Luce is equal to pairwise marginals of multinomial logit (MNL)/Plackett-Luce. We take a brief detour to remind the reader that the BTL model is identical to the MNL model in the sense that the pairwise distributions between objects induced under BTL are identical to that under MNL. Consider an equivalent way to describe an MNL model. Each object i has an associated score $w_i > 0$. A random ordering over all n objects is drawn as follows: iteratively fill the ordered positions $1, \dots, n$ by choosing object $i(k)$ for position k , among the remaining objects (not chosen in the first $1, \dots, k-1$ positions) with probability proportional to its weight $w_{i(k)}$. It can be easily verified that in the random ordering of n objects generated as per this process, i is ranked higher than j with probability $w_i / (w_i + w_j)$.

Sampling model. We also assume that we perform a fixed k number of comparisons for all pairs i and j that are considered (e.g., a best of k series). This assumption is mainly to simplify notations, and the analysis as well as the algorithm easily generalizes to the case when we might have a different number of comparisons for different pairs. Given observations of pairwise comparisons among n items according to this sampling model, we define a *comparisons graph* $G = ([n], E, A)$ as a graph of n items where two items are connected if we have comparisons data on that pair and A denotes the weights on each of the edges in E .

2.2. Rank Centrality

In our setting, we will assume that a_{ij} represents the fraction of times object j has been preferred to object i ,

for example the fraction of times chess player j has defeated player i . Given the notation above, we have that $a_{ij} = (1/k) \sum_{l=1}^k Y_{ij}^l$. Consider a random walk on a weighted directed graph $G = ([n], E, A)$, where a pair $(i, j) \in E$ if and only if the pair has been compared. The weight edges are defined based on the outcome of the comparisons: $A_{ij} = a_{ij} / (a_{ij} + a_{ji})$ and $A_{ji} = a_{ji} / (a_{ij} + a_{ji})$ (note that $a_{ij} + a_{ji} = 1$ in our setting). We let $A_{ij} = 0$ if the pair has not been compared. Note that by the strong law of large numbers, as the number $k \rightarrow \infty$ the quantity A_{ij} converges to $w_j / (w_i + w_j)$ almost surely.

A random walk can be represented by a time-independent transition matrix P , where $P_{ij} = \mathbb{P}(X_{t+1} = j \mid X_t = i)$. By definition, the entries of a transition matrix are nonnegative and satisfy $\sum_j P_{ij} = 1$. One way to define a valid transition matrix of a random walk on G is to scale all the edge weights by $1/d_{\max}$, where we define d_{\max} as the maximum out-degree of a node. This rescaling ensures that each row-sum is at most one. Finally, to ensure that each row-sum is exactly one, we add a self-loop to each node. Concretely,

$$P_{ij} = \begin{cases} \frac{1}{d_{\max}} A_{ij} & \text{if } i \neq j, \\ 1 - \frac{1}{d_{\max}} \sum_{k \neq i} A_{ik} & \text{if } i = j. \end{cases} \quad (1)$$

The choice to construct our random walk as above is not arbitrary. In an ideal setting with infinite samples ($k \rightarrow \infty$) per comparison the transition matrix P would define a reversible Markov chain under the BTL model. Recall that a Markov chain is reversible if it satisfies the *detailed balance equation*: there exists $v \in \mathbb{R}_+^n$ such that $v_i P_{ij} = v_j P_{ji}$ for all i, j ; and in that case, $\pi \in \mathbb{R}_+^n$ defined as $\pi_i = v_i / (\sum_j v_j)$ is its unique stationary distribution. In the ideal setting (say $k \rightarrow \infty$), we will have $P_{ij} = \tilde{P}_{ij} \equiv (1/d_{\max}) w_j / (w_i + w_j)$. That is, the random walk will move from state i to state j with probability proportional to the chance that item j is preferred to item i . In such a setting, it is clear that $v = w$ satisfies the reversibility conditions. Therefore, under these ideal conditions it immediately follows that the vector $w / \sum_i w_i$ acts as a valid stationary distribution for the Markov chain defined by \tilde{P} , the ideal matrix. Hence, as long as the graph G is connected and at least one node has a self-loop, then we are guaranteed that our graph has a unique stationary distribution proportional to w . If the Markov chain is reversible then we may apply the spectral analysis of self-adjoint operators, which is crucial in the analysis of the behavior of the method.

In our setting, the matrix P is a noisy version (due to finite sample error) of the ideal matrix \tilde{P} discussed above. Therefore, it naturally suggests the following algorithm as a surrogate. We estimate the probability distribution obtained by applying matrix P repeated starting from any initial condition. Precisely, let $p_t(i) =$

$\mathbb{P}(X_t = i)$ denote the distribution of the random walk at time t with $p_0 = (p_0(i)) \in \mathbb{R}_+^n$ as an arbitrary starting distribution on $[n]$. Then,

$$p_{t+1}^T = p_t^T P. \quad (2)$$

When the transition matrix has a unique left largest eigenvector, then starting from any initial distribution p_0 , the limiting distribution π is unique. This stationary distribution π is the top left eigenvector of P , which makes computing π a simple eigenvector computation. Formally, we state the algorithm, which assigns numerical scores to each node, which we shall call RANK CENTRALITY:

RANK CENTRALITY

Input: $G = ([n], E, A)$

Output: rank $\{\pi(i)\}_{i \in [n]}$

- 1: Compute the transition matrix P according to (1);
 - 2: Compute the stationary distribution π (as the limit of (2)).
-

The stationary distribution of the random walk is a fixed point of the following equation:

$$\pi(i) = \sum_j \pi(j) \frac{A_{ji}}{\sum_t A_{it}}.$$

This suggests an alternative intuitive justification: an object receives a high rank if it has been preferred to other high ranking objects or if it has been preferred to many objects.

One key question remains: does P have a well-defined unique stationary distribution? Since the Markov chain has a finite state space, there is always a stationary distribution or solution of the above stated fixed-point equations. However, it may not be unique if the Markov chain P is not irreducible. The irreducibility follows easily when the graph is connected and for all edges $(i, j) \in E$, $a_{ij} > 0$, $a_{ji} > 0$. Interestingly enough, we show that the iterative algorithm produces a meaningful solution with near-optimal sample complexity as stated in Theorem 2 when the pairs of objects that are compared are chosen at random.

3. Main Results

The main result of this paper derives sufficient conditions under which the proposed iterative algorithm finds a solution that is close to the true solution (under the BTL model) for a general model with arbitrary connected comparison graph G . This result is stated as Theorem 1. In words, the result implies that to learn the true score correctly as per our algorithm, it is sufficient to have the number of comparisons scaling as $O(\xi^{-2} n \text{poly}(\log n))$, where ξ is the spectral gap of the

Laplacian of the graph G . This result explicitly identifies the role played by the graph structure in the ability of the algorithm to learn the true scores.

In the special case, when the pairs of objects to be compared are chosen at random, that is the induced G is an Erdős-Rényi random graph, the spectral gap ξ can be lower bounded by a constant with high probability and hence the resulting number of comparisons required scales as $O(n \text{poly}(\log n))$. This is effectively the optimal sample complexity.

The bounds are presented as the rescaled Euclidean norm between our estimate π and the underlying stationary distribution of \tilde{P} . This error metric provides us with a means to quantify the relative certainty in guessing if one item is preferred over another.

After presenting our main theoretical result, we describe illustrative simulation results. We also present application of the algorithm in the context of two real data sets: results of NASCAR race for ranking drivers, and results of ODI Cricket for ranking teams. We shall discuss relation between Rank Centrality, the maximum likelihood estimator, and the information theoretic lower bound to conclude that both MLE and Rank Centrality are near optimal when the pairs are chosen according to the Erdős-Rényi random graph.

3.1. Rank Centrality: Error Bound for General Graphs

Recall that in the general setting, each pair of objects or items are chosen for comparisons as per the comparisons graph $G([n], E)$. For each such pair, we have k comparisons available. The result below characterizes the performance of Rank Centrality for such a general setting.

Before we state the result, we present a few necessary notations. Let d_i denote the degree of node i in G ; let the max-degree be denoted by $d_{\max} \equiv \max_i d_i$ and min-degree be denoted by $d_{\min} \equiv \min_i d_i$; let $\kappa \equiv d_{\max}/d_{\min}$. The *random walk normalized Laplacian matrix* of the graph G is defined as $L = D^{-1}B$, where D is the diagonal matrix with $D_{ii} = d_i$ and B is the adjacency matrix with $B_{ij} = B_{ji} = 1$ if $(i, j) \in E$ and 0 otherwise. This normalized Laplacian, defined thus, can be thought of as a transition matrix of a reversible random walk on graph G : from each node i , jump to one of its neighbors j with equal probability. Given this, it is well known that the random walk normalized Laplacian of the graph has real eigenvalues denoted as

$$-1 \leq \lambda_n(L) \leq \dots \leq \lambda_1(L) = 1. \quad (3)$$

We shall denote the *spectral gap* of the Laplacian as

$$\xi \equiv 1 - \lambda_{\max}(L),$$

where

$$\lambda_{\max}(L) \equiv \max\{\lambda_2(L), -\lambda_n(L)\}. \quad (4)$$

There is one-to-one correspondence between the eigenvalues of the random walk normalized Laplacian L and the standard (symmetric) normalized Laplacian $\| - D^{-1/2} B D^{-1/2}$. Now we state the result establishing the performance of Rank Centrality.

Theorem 1. *Given n objects and a connected comparison graph $G = ([n], E)$, let each pair $(i, j) \in E$ be compared for k times with outcomes produced as per a BTL model with parameters w_1, \dots, w_n . Then, for some positive constant $C \geq 8$ and when $k \geq 4C^2(b^5 \kappa^2 / d_{\max} \xi^2) \log n$, the following bound on the normalized error holds with probability at least $1 - 4n^{-C/8}$:*

$$\frac{\|\pi - \tilde{\pi}\|}{\|\tilde{\pi}\|} \leq \frac{C b^{5/2} \kappa}{\xi} \sqrt{\frac{\log n}{k d_{\max}}},$$

where $\tilde{\pi}(i) = w_i / \sum_{\ell} w_{\ell}$, $b \equiv \max_{i,j} w_i / w_j$, and $\kappa \equiv d_{\max} / d_{\min}$.

3.2. Rank Centrality: Error Bound for Random Graphs

Now we consider the special case when the comparison graph G is an Erdős-Rényi random graph with pair (i, j) being compared with probability d/n . When d is polylogarithmic in n , we provide a strong performance guarantee. Specifically, the result stated below suggests that with $O(n \text{ poly}(\log n))$ comparisons, Rank Centrality manages to learn the true scores with high probability.

Theorem 2. *Given n objects, let the comparison graph $G = ([n], E)$ be generated by selecting each pair (i, j) to be in E with probability d/n independently of everything else. Each such chosen pair of objects is compared k times with the outcomes of comparisons produced as per a BTL model with parameters w_1, \dots, w_n . Then, if $d \geq 10C^2 \log n$ and $k d \geq 128C^2 b^5 \log n$, the following bound on the error rate holds with probability at least $1 - 10n^{-C/8}$:*

$$\frac{\|\pi - \tilde{\pi}\|}{\|\tilde{\pi}\|} \leq 8C b^{5/2} \sqrt{\frac{\log n}{k d}},$$

where $\tilde{\pi}(i) = w_i / \sum_{\ell} w_{\ell}$ and $b \equiv \max_{i,j} w_i / w_j$.

Remarks. Some remarks are in order. First, Theorem 2 immediately implies that as long as kd grows super linear in $\log n$, then the error goes to 0. Furthermore, in the context that the number of items n goes to ∞ as long as we choose $d = \Omega(\log n)$ and $kd = \omega(\log n)$, the relative error goes to 0 as $n \rightarrow \infty$ with high probability. That is, with $\omega(n \log n)$ total samples, the relative error goes to 0 with high probability. It is well known that for Erdős-Rényi graphs, the induced graph G is connected with high probability only when $d = \Omega(\log n)$, i.e., when total number of pairs sampled scales as $\Omega(n \log n)$. Thus, RANK CENTRALITY is nearly order optimal in this setting.

Second, the b parameter should be treated as constant. It is the *dynamic* range in which we are trying to resolve the uncertainty between scores. We are considering a regime that there exists some uncertainty in the samples. Otherwise, if the weight of a single item were an order n greater than the weights of other items, then it would effectively be preferred with certainty. Hence, we would remove it from the items under consideration.

Third, for a general graph, Theorem 1 implies that by choice of $kd_{\max} = O(\kappa^2 \xi^{-2} \log n)$, Rank Centrality learns a score vector close to the true scores with high probability. That is, effectively the Rank Centrality algorithm requires $O(n \kappa^2 \xi^{-2} \text{poly}(\log n))$ comparisons to learn scores well. Ignoring κ , the graph structure plays a role through ξ^{-2} , the squared inverse of the spectral gap of Laplacian of G , indicating the performance of Rank Centrality. A reversible natural random walk on G , whose transition matrix is the Laplacian, has its mixing time scaling as ξ^{-2} (precisely, relaxation time). In that sense, the mixing time of natural random walk on G ends up playing an important role in the ability of Rank Centrality to learn the true scores. Hence, if one has the option to choose which pairs to compare, our analysis in Theorem 1 suggests that one should choose pairs such that the resulting graph has large spectral gap. Spectral gap of the comparisons graph also plays an important role in Osting et al. (2013), where the goal is to choose pairs to compare under a different model where cardinal preferences (as opposed to ordinal preferences) are observed.

Finally, if we wish to obtain a relative accuracy of ϵ with probability at least $1 - \delta$ for a fixed number of items n , then our results also show that we require $kd \geq [512 b^5 / \epsilon^2] \cdot \max(\log^2(10/\delta) / \log n, \log n)$.

3.3. Experimental Results

Under the BTL model, define an error metric of an estimated ordering σ as the weighted sum of pairs (i, j) whose ordering is incorrect:

$$D_w(\sigma) = \left\{ \frac{1}{2n \|w\|^2} \sum_{i < j} (w_i - w_j)^2 \cdot \mathbb{I}((w_i - w_j)(\sigma_i - \sigma_j) > 0) \right\}^{1/2},$$

where $\mathbb{I}(\cdot)$ is an indicator function. This is a more natural error metric compared to the Kemeny distance, which is an unweighted version of the above sum, since $D_w(\cdot)$ is less sensitive to errors between pairs with similar weights. Further, assuming without loss of generality that w is normalized such that $\sum_i w_i = 1$, the next lemma connects the error in $D_w(\cdot)$ to the bound provided in Theorem 2. Hence, the same upper bound holds for D_w error. A proof of this lemma is provided in the online appendix.

Lemma 1. Let σ be an ordering of n items induced by a scoring π . Then,

$$D_w(\sigma) \leq \frac{\|w - \pi\|}{\|w\|}.$$

Synthetic Data. To begin with, we generate data synthetically as per a BTL model for a specific choice of scores. For a given n and b , the scores are chosen such that the ratio between two consecutive scores are fixed to be $b^{1/n}$, i.e., $w_1 = b^{(1-n)/2n}$, $w_2 = b^{(3-n)/2n}$, $w_3 = b^{(5-n)/2n}$, etc. A representative result is depicted in Figure 1: for fixed $n = 400$ and a fixed $b = 10$, it shows how the error scales when varying two key parameters—varying the number of comparisons per pair with fixed $d = 10 \log n$ (on left), and varying the sampling probability with fixed $k = 32$ (on right). This figure compares performance of Rank Centrality with variety of other algorithms. Next, we provide a brief description of various algorithms that we shall compare with.

Regularized Rank Centrality. When there are items that have been compared only a few times, the scores to those items might be sensitive to the randomness in the outcome of the comparisons, or even worse the resulting comparisons graph might not be connected. To make the random walk irreducible and get a ranking that is more robust against comparisons noise in those edges with only a few comparisons, one can add regularization to Rank Centrality. A reasonable way to add regularization is to consider the transition probability P_{ij} as the prediction of the event that j beats i , given data (a_{ij}, a_{ji}) . The Rank Centrality, in nonregularized settings, uses the Haldane prior of $\text{Beta}(0, 0)$, which gives $P_{ij} \propto a_{ij} / (a_{ij} + a_{ji})$. To add regularization, one can use different priors, for example $\text{Beta}(\varepsilon, \varepsilon)$, which gives

$$P_{ij} = \frac{1}{d_{\max}} \frac{a_{ij} + \varepsilon}{a_{ij} + a_{ji} + 2\varepsilon}. \quad (5)$$

When the prior is unknown, a reasonable choice in practice is $\varepsilon = 1$.

Maximum Likelihood Estimator (MLE). The ML estimator directly maximizes the likelihood assuming the BTL model (Ford 1957). If we reparameterize the problem so that $\theta_i = \log(w_i)$, then we obtain our estimates $\hat{\theta}$ by solving the convex program

$$\hat{\theta} \in \arg \min_{\theta} \sum_{(i,j) \in E} \sum_{l=1}^k \{ \log(1 + \exp(\theta_j - \theta_i)) - Y_{ij}^l (\theta_j - \theta_i) \}, \quad (6)$$

which is a pairwise logistic regression model. The MLE is known to be consistent (Ford 1957). The finite sample analysis of MLE is provided in Section 3.5.

For comparison with Regularized Rank Centrality, we provide regularized MLE or regularized Logistic Regression:

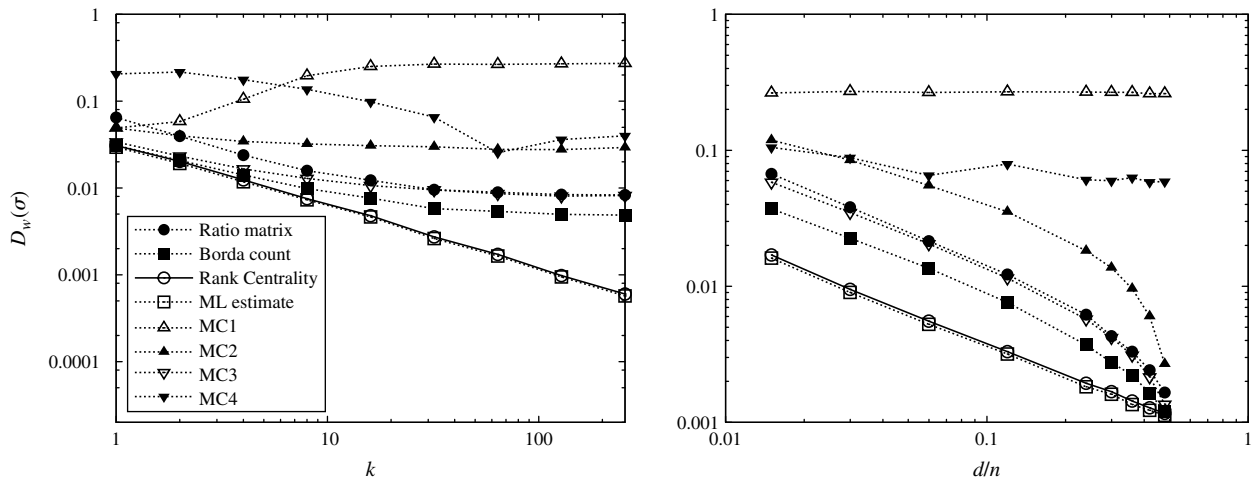
$$\arg \min_{\theta} \left[\sum_{(i,j) \in E} \sum_{l=1}^k \{ \log(1 + \exp(\theta_j - \theta_i)) - Y_{ij}^l (\theta_j - \theta_i) \} + \frac{1}{2} \lambda \|\theta\|^2 \right]. \quad (7)$$

Borda Count. The (generalized) Borda Count method, analyzed recently by Ammar and Shah (2011), scores an item by counting the number of wins divided by the total number of comparisons:

$$s(i) = \frac{\text{No. of times item } i \text{ has won}}{\text{No. of times item } i \text{ has been compared}}.$$

This can be thought of as an extension of the standard Borda Count for aggregating full rankings (de Borda 1781), which is widely used in psychology (David 1963, Kendall and Smith 1940, Mosteller 1951). If we break the full rankings into pairwise comparisons and apply the pairwise version of the Borda Count from (Ammar

Figure 1. Average error $D_w(\sigma)$ of various rank aggregation algorithms averaged over 20 instances.



Notes. In the figure on the left, d and n are fixed while k is increased. The figure on the right keeps $k = 32$ fixed, and lets d increase.

and Shah 2011), then it produces the same ranking as the standard Borda Count applied to the original full rankings. This is different from how HodgeRank from Jiang et al. (2011) generalizes Borda count, which does not normalize the scores by the number of comparisons.

Spectral Ranking Algorithms. Rank Centrality can be classified as part of the spectral ranking algorithms, which assign scores to the items according to the leading eigenvector of a matrix that represents the data. Different choices of the matrix based on data can lead to different algorithms. A few prominent examples are *Ratio matrix* in Saaty (2003) and those in Dwork et al. (2001). In the Ratio Matrix algorithm, a matrix $M \in \mathbb{R}^{n \times n}$ with $M_{ij} = a_{ij}/a_{ji}$ is constructed (and $M_{ii} = 1$), and the scores for the items are assigned as per the top eigenvector of this ratio matrix. Dwork et al. (2001) introduced four spectral ranking algorithms called MC1, MC2, MC3, and MC4. They are all based on a random walk very similar (but distinct) to that of Rank Centrality. These algorithms use the stationary distributions of the following Markov chains, respectively, translated to account for the pairwise comparisons data: $P_{ij}^{(\text{MC1})} = 1/|\{\ell: a_{i\ell} > 0\}|$, $P_{ij}^{(\text{MC2})} = a_{ij}/\sum_{\ell \neq i} a_{i\ell}$,

$$P_{ij}^{(\text{MC3})} = \begin{cases} a_{ij}/\deg(i) & \text{if } i \neq j, \\ 1 - \sum_{\ell \neq i} a_{i\ell}/\deg(i) & \text{if } i = j, \end{cases}$$

$$P_{ij}^{(\text{MC4})} = \begin{cases} 1/n & \text{if } a_{ij} \geq a_{ji}, \\ 0 & \text{if } a_{ij} < a_{ji}, \\ 1 - \sum_{\ell \neq i} |\{\ell: a_{i\ell} \geq a_{\ell i}\}|/n & \text{if } i = j, \end{cases}$$

where $\deg(i)$ is the number of items that item i has been compared to.

We make note of the following observations from Figure 1. First, the error achieved by our Rank Centrality is comparable to that of ML estimator, and vanishes at the rate of $1/\sqrt{k}$ as predicted by our main result. Moreover, as predicted by our bounds, the error scales as $1/\sqrt{d}$. Second, for fixed d , both the Borda Count and Ratio Matrix algorithms have strictly positive error even if we take $k \rightarrow \infty$. This exhibits that these are inherently inefficient algorithms. Third, despite strong similarity between Rank Centrality and the Markov chain based algorithms of Dwork et al. (2001), the careful choice of the transition matrix of Rank Centrality makes a noticeable difference as shown in the figure—like Borda Count and Ratio Matrix, for fixed d, n , despite k increasing the error remains finite (and at times gets worse!).

Real Data Sets. Next we show that Rank Centrality is more robust to randomly missing data compared to existing spectral ranking approaches on real data sets, which are not necessarily derived from the BTL model.

Data Set 1: Washington Post. This is the public data set collected from an online polling on Washington Post² from December 2010 to January 2011. Using allourideas³ platform developed by Salganik and Levy (2012), they asked who had the worst year in Washington, where each user was asked to compare a series of randomly selected pairs of political entities. There are 67 political entities in the data set, and the resulting graph is a complete graph on these 67 nodes. We used Rank Centrality and other algorithms to aggregate this data. We use this data set primarily to check the “robustness” of algorithms rather than understanding their ability to identify ground truth, as by design it is not available.

Now each algorithm gives different ground truth rankings given the full set of data. This ground truth is compared to a ranking we get from only a subset of the data, which is generated by sampling each edge with a given sampling rate and revealing only the data on those sampled edges. We want to measure how much each algorithm is affected by eliminating edges from the complete graph. Let σ_{GT} be the ranking we get by applying our choice of rank aggregation algorithm to the complete data set, and σ_{Sample} be the ranking we get from sampled data set. To measure the resulting error in the ranking, we use the following metric:


$$D_{L_1}(\sigma_{\text{GT}}, \sigma_{\text{Sample}}) = \frac{1}{n} \sum_i |\sigma_{\text{GT}}(i) - \sigma_{\text{Sample}}(i)|.$$

Figure 2 illustrates that Rank Centrality, ML estimator and MC2 are less sensitive to sampling the data set, compared to Borda Count, MC1, MC3, and MC4. Hence they are more robust when available comparisons data are limited.

Data Set 2: NASCAR 2002. Table 1 shows ranking of drivers from NASCAR 2002 season racing results. Hunter (2004) used this data set for studying rank-aggregation algorithms, and we use the data set, at Guiver and Snelson (2009). The data set has 87 different drivers who competed in a total of 36 races in which 43 drivers were racing at each race. Some of the drivers raced in all 36 races, whereas some drivers only participated in one. To break the racing results into parity comparisons and to be able to run the comparison-based algorithm, like Hunter (2004), Guiver and Snelson (2009), we eliminated four drivers who finished last in every race they participated. Therefore, for the data set we used, there are a total of 83 drivers.

Table 1 shows the top 10 and bottom 10 drivers according to their average place, and their ranking from Rank Centrality and ML estimator. The unregularized Rank Centrality can over fit the data by placing P. J. Jones and Scott Pruett in the first and second places. They have high average place, but they only participated in one race. In contrast, the regularized

Figure 2. Experimental results on a real data set shows that Rank Centrality, ML estimator, and MC2 are less sensitive to having limited data.



version places them lower and gives the top ranking to those players with more races. Similarly, Morgan Shepherd is placed last in the regularized version, because he had consistently low performance in five races. Similarly, the ML estimator with regularization gives the top (and bottom) rankings to those players with more races.

Table 1. ϵ -regularized Rank Centrality for the top 10 and bottom 10 2002 NASCAR drivers, as ranked by average place.

Driver	Races	Av. place	Rank Centrality				ML estimator	
			$\epsilon = 0$		$\epsilon = 3$		$\lambda = 0.01$	
			π	Rank	π	Rank	e^θ	Rank
P. J. Jones	1	4.00	0.1837	1	0.0181	11	0.0124	23
Scott Pruett	1	4.00	0.0877	2	0.0176	12	0.0124	24
Mark Martin	36	12.17	0.0302	5	0.0220	2	0.0203	1
Tony Stewart	36	12.61	0.0485	3	0.0219	1	0.0199	2
Rusty Wallace	36	13.17	0.0271	6	0.0209	3	0.0193	3
Jimmie Johnson	36	13.50	0.0211	12	0.0199	5	0.0189	4
Sterling Marlin	29	13.86	0.0187	14	0.0189	10	0.0177	8
Mike Bliss	1	14.00	0.0225	10	0.0148	18	0.0121	27
Jeff Gordon	36	14.06	0.0196	13	0.0193	8	0.0184	5
Kurt Busch	36	14.06	0.0253	7	0.0200	4	0.0184	6
⋮								
Carl Long	2	40.50	0.0004	77	0.0087	68	0.0106	59
Christian Fittipaldi	1	41.00	0.0001	83	0.0105	49	0.0111	40
Hideo Fukuyama	2	41.00	0.0004	76	0.0088	67	0.0106	60
Jason Small	1	41.00	0.0002	80	0.0105	48	0.0111	41
Morgan Shepherd	5	41.20	0.0002	78	0.0059	83	0.0092	75
Kirk Shelmerdine	2	41.50	0.0002	81	0.0084	70	0.0105	61
Austin Cameron	1	42.00	0.0005	75	0.0107	44	0.0111	43
Dave Marcis	1	42.00	0.0012	71	0.0105	47	0.0111	44
Dick Trickle	3	42.00	0.0001	82	0.0071	77	0.0100	65
Joe Varde	1	42.00	0.0002	79	0.0110	43	0.0111	42

Data Set 3: ODI Cricket. Table 2 shows ranking of international cricket teams from the 2012 season of the ODI cricket match, where 16 teams played a total of 362 games. Like NASCAR data set, in Table 2, teams with a smaller number of matches, such as Scotland and Ireland, are moved toward the middle with regularization, and New Zealand is moved toward the end. Notice that regularized or not, the ranking from Rank Centrality is different from the simple ranking from average place or winning ratio, because we give more score for winning against stronger opponents. The regularized ML estimator produces similar ranking as the regularized Rank Centrality. These data on ODI cricket match are publicly available, for example, from <http://www.cricmetric.com/blog/>.

3.4. Information-Theoretic Lower Bound

In previous sections, we presented the achievable error rate based on a particular low-complexity algorithm. In this section, we ask how this bound compares to the fundamental limit under the BTL model.

Our result in Theorem 2 provides an upper bound on the achievable error rate between estimated scores and the true underlying scores. We provide a constructive argument to lower bound the minimax error rate over a class of BTL models. Concretely, we consider the scores coming from a simplex with bounded dynamic range defined as

$$\mathcal{S}_b \equiv \left\{ \tilde{\pi} \in \mathbb{R}^n \mid \sum_{i \in [n]} \tilde{\pi}_i = 1, \max_{i,j} \frac{\tilde{\pi}_i}{\tilde{\pi}_j} \leq b \right\}.$$

Table 2. Applying ε -regularized Rank Centrality to One Day International (ODI) cricket match results from 2012.

Team	Matches	Win ratio	deg	Rank Centrality				ML estimator	
				$\varepsilon = 0$		$\varepsilon = 1$		$\lambda = 0.01$	
				π	Rank	π	Rank	e^θ	Rank
South Africa	43	0.6744	11	0.1794	2	0.0943	2	0.0924	2
India	76	0.6382	11	0.1317	4	0.0911	3	0.0923	3
Australia	72	0.6319	13	0.1798	1	0.0900	4	0.0881	4
England	60	0.6000	10	0.1526	3	0.0957	1	0.0927	1
Scotland	15	0.6000	7	0.0029	12	0.0620	7	0.0627	7
Sri Lanka	78	0.5577	12	0.1243	5	0.0801	5	0.0768	5
Pakistan	65	0.5385	13	0.0762	6	0.0715	6	0.0755	6
Ireland	32	0.5316	13	0.0124	11	0.0561	8	0.0539	9
Afghanistan	20	0.5000	7	0.0005	15	0.0435	13	0.0472	12
West Indies	55	0.4091	12	0.0396	7	0.0546	9	0.0592	8
New Zealand	50	0.3800	10	0.0354	8	0.0466	12	0.0514	10
Bangladesh	51	0.3333	11	0.0320	9	0.0500	10	0.0492	11
Netherlands	24	0.3333	10	0.0017	13	0.0432	14	0.0427	14
Zimbabwe	40	0.3250	11	0.0307	10	0.0481	11	0.0439	13
Canada	22	0.2273	11	0.0003	16	0.0365	16	0.0364	15
Kenya	21	0.1905	10	0.0007	14	0.0367	15	0.0356	16

Notes. The degree of a team in the comparisons graph is the number of teams it has played against.

We constrain the scores to be on the simplex, because we represent the scores by its projection onto the standard simplex as explained in Section 2.1. Then, we can prove the following lower bound on the minimax error rate.

Theorem 3. Consider a minimax scenario where we first choose an algorithm \mathcal{A} that estimates the BTL weights, say $\pi^{\mathcal{A}}$, from given observations and for this particular algorithm \mathcal{A} , nature chooses the worst-case true BTL weights $\tilde{\pi}$. Let \mathcal{S}_b denote the space of all BTL score vectors $\tilde{\pi}$ with dynamic range at most b as defined above. Then

$$\inf_{\mathcal{A}} \sup_{\tilde{\pi} \in \mathcal{S}_b} \frac{\mathbb{E}[\|\pi^{\mathcal{A}} - \tilde{\pi}\|]}{\|\tilde{\pi}\|} \geq \frac{b-1}{240\sqrt{10}(b+1)} \frac{1}{\sqrt{kd}}, \quad (8)$$

where the infimum ranges over all estimation algorithms \mathcal{A} that are measurable functions over the observations. Here a pair of items is chosen to be compared with probability d/n , and for this chosen pair k comparison observations are generated as per the underlying BTL model.

By definition the dynamic range is always at least one. When $b = 1$, we can trivially achieve a minimax rate of zero. Since the infimum ranges over all measurable functions, it includes a trivial estimator that always outputs $(1/n)\mathbb{1}$ regardless of the observations, and this estimator achieves zero error when $b = 1$. In the regime where the dynamic range b is bounded away from one and bounded above by a constant, Theorem 3 establishes that the upper bound obtained in Theorem 2 is minimax-optimal up to factors logarithmic in the number of items n .

3.5. MLE: Error Bounds Using a State-of-the-Art Method

It is well known that the maximum-likelihood estimate of a set of parameters is asymptotically normal with mean 0 and covariance equal to the inverse Fisher information of the set of parameters. In this section we wish to show the behavior of the estimates obtained through the logistic regression-based approach for estimating the parameters $\theta_i^* = \log w_i$ in a finite sample setting.

Model. Recall that the logistic regression-based method reparameterizes the model so that given items i and j the probability that i defeats j is

$$P(i \text{ defeats } j) = \frac{\exp(\theta_i^* - \theta_j^*)}{1 + \exp(\theta_i^* - \theta_j^*)}.$$

To ensure identifiability we also assume that $\sum_i \theta_i^* = 0$, so that we also enforce the constraint $\sum \hat{\theta}_i = 0$. We also recall that we let $b = w_{\max}/w_{\min}$. Similarly, we let $\tilde{b} := \theta_{\max}^* - \theta_{\min}^*$ and enforce the constraint that $\hat{\theta}_{\max} - \hat{\theta}_{\min} \leq \tilde{b}'$ where $\tilde{b} \leq \tilde{b}'$. For simplicity we assume that $\tilde{b}' = \tilde{b}$.

Finally, recall that we are given m i.i.d. observations. We take $l \in \{1, 2, \dots, n\}$ and let v_l to be the outcome of the l th comparison. Furthermore, if during the l th competition item i competed against item j we take $x_l = e_i - e_j$ where e_i is the standard basis vector with entries that are all zero except for the i th entry, which equals one. Note that in this context the ordering of the competition does matter. Finally, we define the inner product between two vectors $x, y \in \mathbb{R}^n$ to be $\langle x, y \rangle =$

$\sum_{i=1}^n x_i y_i$. Therefore, under the BTL model with parameters θ^* we have that

$$v_l = \begin{cases} 1 & \text{with probability } \exp \langle x_l, \theta^* \rangle / (1 + \exp \langle x_l, \theta^* \rangle), \\ 0 & \text{otherwise.} \end{cases}$$

Now the estimation procedure is of the form

$$\hat{\theta} = \arg \min_{\theta} \mathcal{L}_m(\theta, v, x)$$

where

$$\mathcal{L}_m(\theta, v, x) = \frac{1}{m} \sum_{l=1}^n \{ \log(1 + \exp \langle x_l, \theta \rangle) - v_l \langle x_l, \theta \rangle \}. \quad (9)$$

Results. Before proceeding we recall that $\|\theta^*\|_2 \leq \tilde{b}\sqrt{n}$. With that in mind we have the following theorem.

Theorem 4. Suppose that we have $m > 12n \log n$ observations of the form (i, j, y) where i and j are drawn uniformly at random from $[n]$ and y is Bernoulli with parameter $\exp(\theta_i^* - \theta_j^*) / (1 + \exp(\theta_i^* - \theta_j^*))$. Then, we have with probability at least $1 - 2/n$

$$\|\hat{\theta} - \theta^*\| \leq 6 \frac{(1+b)^2}{b} \sqrt{\frac{n^2 \log n}{m}}.$$

With the assumption that $\|\theta^*\|_\infty \leq \tilde{b}$, we have $\|\theta^*\| \leq \tilde{b}\sqrt{n}$.

3.6. Cramér-Rao Lower Bound

The Fisher information matrix (FIM) encodes the amount of information that the observed measurements carry about the parameter of interest. The Cramér-Rao bound we derive in this section provides a lower bound on the expected squared Euclidean norm $\mathbb{E}[\|\tilde{\pi} - \pi\|^2]$ of any unbiased estimator and is directly related to the (inverse of) Fisher information matrix.

Denote the log-likelihood function as

$$\begin{aligned} \ell(\tilde{\pi} | a) &= \sum_{(i,j) \in E} \log f(a_{ij}, a_{ji} | \tilde{\pi}), \quad \text{where} \\ f(a_{ij}, a_{ji} | \tilde{\pi}) &= \left(\frac{\tilde{\pi}_j}{\tilde{\pi}_i + \tilde{\pi}_j} \right)^{k_{ij} a_{ij}} \left(\frac{\tilde{\pi}_i}{\tilde{\pi}_i + \tilde{\pi}_j} \right)^{k_{ji} a_{ji}}, \end{aligned}$$

and k_{ij} is the number of times the pair (i, j) was compared. The Fisher information matrix with the BTL weights $\tilde{\pi}$ is defined as $F(\tilde{\pi}) \in \mathbb{R}^{n \times n}$ with

$$\begin{aligned} F(\tilde{\pi})_{ij} &= \mathbb{E}_a \left[-\frac{\partial^2 \ell(\tilde{\pi} | a)}{\partial \tilde{\pi}_i \partial \tilde{\pi}_j} \right] \\ &= \begin{cases} \sum_{i' \in \partial i} \frac{k_{ii'}}{(\tilde{\pi}_i + \tilde{\pi}_{i'})^2} \frac{\tilde{\pi}_{i'}}{\tilde{\pi}_i} & \text{if } i = j, \\ -\frac{k_{ij}}{(\tilde{\pi}_i + \tilde{\pi}_j)^2} & \text{if } (i, j) \in E, \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

This follows from the fact that

$$\begin{aligned} \frac{\partial \ell(\tilde{\pi} | a)}{\partial \tilde{\pi}_i} &= \sum_{i' \in \partial i} \frac{-k_{ii'}(a_{ii'} + a_{i'i})}{\tilde{\pi}_i + \tilde{\pi}_{i'}} + \frac{k_{ii'} a_{i'i}}{\tilde{\pi}_i}, \quad \text{and} \\ \frac{\partial^2 \ell(\tilde{\pi} | a)}{\partial \tilde{\pi}_i \partial \tilde{\pi}_j} &= \begin{cases} \sum_{i' \in \partial i} k_{ii'} \left(\frac{1}{(\tilde{\pi}_i + \tilde{\pi}_{i'})^2} - \frac{a_{i'i}}{(\tilde{\pi}_i)^2} \right) & \text{if } i = j, \\ \frac{k_{ij}}{(\tilde{\pi}_i + \tilde{\pi}_j)^2} & \text{if } (i, j) \in E, \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Let π denote our estimate of the weights. Applying the CRB (Rao 1945), we get the following lower bound for all unbiased estimators π :

$$E[\|\pi - \tilde{\pi}\|^2] \geq \text{Trace}(F(\tilde{\pi})^{-1}).$$

This bound depends on $\tilde{\pi}$ and the graph structure. Although a closed-form expression is difficult to get and Rank Centrality as well as the ML estimate is biased, we compare our numerical experiments with a numerically computed CRB on the same graph and the same weights $\tilde{\pi}$.

3.6.1. Numerical Comparisons. In Figure 3, the average normalized root mean squared error (RMSE) is shown as a function of various model parameters. We fixed the control parameters as $k = 32$, $n = 400$, $d = 60$, and $b = 10$ with pairs assigned according to Erdős-Renyi graph $G(n, d/n)$. Each point in the figure is averaged over 20 random instances \mathcal{S} . Let $\tilde{\pi}^{(i)}$ be the resulting estimate at the i th experiment, then


$$\text{RMSE} = \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \frac{\|\pi^{(i)} - \tilde{\pi}\|}{\|\tilde{\pi}\|}. \quad (10)$$

For all ranges of model parameters k , d , and b , RMSE achieved using Rank Centrality is almost indistinguishable from that of the ML estimate and also the CRB.

CRB provides a lower bound on the expected mean squared error for unbiased estimators. Although we are plotting average root mean squared error, as opposed to average mean squared error, we do not expect any estimator to achieve RMSE better than the CRB as long as there is a concentration.

The ML estimator in (7) with $\lambda = 0$ finds an estimate $\pi = e^{\hat{\theta}}$ that maximizes the log likelihood, and in general the ML estimate does not coincide with the minimum mean squared error estimator. From the figure we see that it in fact achieves the minimum mean squared error and matches the CRB.

What is perhaps surprising is that for all the parameters that we experimented with, the RMSE achieved by Rank Centrality is almost indistinguishable with that of ML estimate and the CRB. Thus, coupled with the minimax lower bounds, one cannot do better than Rank Centrality under the BTL model.

Figure 3. (Color online) Comparisons of Rank Centrality, the ML estimator, and the Cramér-Rao bound.

Notes. All three lines are almost indistinguishable for all ranges of model parameters.

3.7. Discussion of Results

In this section we review the results that we have established above. In Theorem 1 we establish upper bounds on the error when samples are drawn from an arbitrary graph and when each edge is compared k times. This bound depends on the spectral gap of the underlying graph, which shows that graphs with a larger spectral gap achieve smaller estimation error. For the case of Erdős-Renyi graphs, Theorem 2 provides an upper bound on the error achieved by Rank Centrality. In Theorem 3 we prove that the bound is near optimal, up to logarithmic factors, in an information theoretic sense. That is, no method, regardless of computational power can achieve better performance on the same statistical model. For a tighter analysis of the optimality of Rank Centrality, we provide numerical experiments under the BTL model and compare it to the Cramer Rao lower bound established in Section 3.6. Comparisons with the Cramer-Rao bound in Figure 3 suggests that the error achieved by Rank Centrality is indistinguishable from the fundamental Cramer-Rao lower bound, and hence exactly optimal for a certain class of estimators.

For completeness, we further provide an analysis of the error achieved by the MLE in Theorem 4. Building upon our analysis, Hajek et al. (2014) shows that MLE is near order optimal, just like RANK CENTRALITY.

Finally, we compare the computational cost of Rank Centrality versus the MLE. While it is difficult to make an exact, theoretical, comparison, we nevertheless compare their computational cost by means of popular implementations on a common computation platform. For RANK CENTRALITY, the implementation is based on using `eigs` function MATLAB. For MLE, the implementation is based on the basic first-order method. In a collection of experiments (with varying problem parameters), RANK CENTRALITY converges an order of magnitude faster than the MLE. It should be noted that the first-order method has tunable step size and our implementation did not attempt to optimize this selection when varying problem parameters. Finally, MLE can be viewed as a standard logistic

regression. Therefore, the `lm` function of R-package can be used to solve for MLE. Again, in the same computation environment, the resulting MLE is an order of magnitude slower compared to the MATLAB implementation of RANK CENTRALITY, but faster than the first-order method.

4. Proofs

We may now present proofs of Theorems 1 and 2. We first present a proof of convergence for general graphs in Theorem 1. This result follows from Lemma 2 that we state below, which shows that our algorithm enjoys convergence properties that result in useful upper bounds. The lemma is made general and uses standard techniques of spectral theory. The main difficulty arises in establishing that the Markov chain P satisfies certain properties that we will discuss subsequently. Given the proof for the general graph, Theorem 2 follows by showing that in the case of Erdős-Renyi graphs, certain spectral properties are satisfied with high probability.

The next set of proofs involve the information-theoretic lower bound stated in Theorem 3 and the proof of Theorem 4 establishing the finite sample error analysis of MLE.

4.1. Proof of Theorem 1: General Graph

In this section, we characterize the error rate achieved by our ranking algorithm. Given the random Markov chain P , where the randomness comes from the outcome of the comparisons, we will show that it does not deviate too much from its expectation \tilde{P} , where we recall that \tilde{P} is defined as

$$\tilde{P}_{ij} = \begin{cases} \frac{1}{d_{\max}} \frac{w_j}{w_i + w_j} & \text{if } i \neq j, \\ 1 - \frac{1}{d_{\max}} \sum_{\ell \neq i} \frac{w_\ell}{w_i + w_\ell} & \text{if } i = j, \end{cases}$$

for all $(i, j) \in E$ and $\tilde{P}_{ij} = 0$ otherwise.

Recall from the discussion following Equation (1) that the transition matrix P used in our ranking algorithm has been carefully chosen such that the

corresponding expected transition matrix \tilde{P} has two important properties. First, the stationary distribution of \tilde{P} , which we denote with $\tilde{\pi}$ is proportional to the weight vectors w . Furthermore, when the graph is connected and has self loops (which at least one exists), this Markov chain is irreducible and aperiodic so that the stationary distribution is unique. The next important property of \tilde{P} is that it is reversible— $\tilde{\pi}(i)\tilde{P}_{ij} = \tilde{\pi}(j)\tilde{P}_{ji}$. This observation implies that the operator \tilde{P} is symmetric in an appropriately defined inner product space. The symmetry of the operator \tilde{P} will be crucial in applying ideas from spectral analysis to prove our main results.

Let Δ denote the fluctuation of the transition matrix around its mean, such that $\Delta \equiv P - \tilde{P}$. The following lemma bounds the deviation of the Markov chain after t steps in terms of two important quantities: the spectral radius of the fluctuation $\|\Delta\|_2$ and the spectral gap $1 - \lambda_{\max}(\tilde{P})$, where

$$\lambda_{\max}(\tilde{P}) \equiv \max\{\lambda_2(\tilde{P}), -\lambda_n(\tilde{P})\}.$$

Since $\lambda(\tilde{P})$'s are sorted, $\lambda_{\max}(\tilde{P})$ is the second largest eigenvalue in absolute value.

Lemma 2. For any Markov chain $P = \tilde{P} + \Delta$ with a reversible Markov chain \tilde{P} , let p_t be the distribution of the Markov chain P when started with initial distribution p_0 . Then,

$$\frac{\|p_t - \tilde{\pi}\|}{\|\tilde{\pi}\|} \leq \rho^t \frac{\|p_0 - \tilde{\pi}\|}{\|\tilde{\pi}\|} \sqrt{\frac{\tilde{\pi}_{\max}}{\tilde{\pi}_{\min}}} + \frac{1}{1 - \rho} \|\Delta\|_2 \sqrt{\frac{\tilde{\pi}_{\max}}{\tilde{\pi}_{\min}}}, \quad (11)$$

where $\tilde{\pi}$ is the stationary distribution of \tilde{P} , $\tilde{\pi}_{\min} = \min_i \tilde{\pi}(i)$, $\tilde{\pi}_{\max} = \max_i \tilde{\pi}(i)$, and $\rho = \lambda_{\max}(\tilde{P}) + \|\Delta\|_2 \sqrt{\tilde{\pi}_{\max}/\tilde{\pi}_{\min}}$.

The above result provides a general mechanism for establishing error bounds between an estimated stationary distribution π and the desired stationary distribution $\tilde{\pi}$. It is worth noting that the result only requires control on the quantities $\|\Delta\|_2$ and $1 - \rho$. We may now state two technical lemmas that provide control on the quantities $\|\Delta\|_2$ and $1 - \rho$, respectively.

Lemma 3. For some constant $C \geq 8$, the error matrix $\Delta = P - \tilde{P}$ satisfies

$$\|\Delta\|_2 \leq C \sqrt{\frac{\log n}{k d_{\max}}}$$

with probability at least $1 - 4n^{-C/8}$.

The next lemma provides our desired bound on $1 - \rho$.

Lemma 4. If $\|\Delta\|_2 \leq C \sqrt{\log n / (k d_{\max})}$ and $k \geq 4C^2 b^5 \cdot d_{\max} \log n (1/d_{\min} \xi)^2$, then

$$1 - \rho \geq \frac{\xi d_{\min}}{b^2 d_{\max}}.$$

Proof of Theorem 1. With the above stated Lemmas, we shall proceed with the proof of Theorem 1. When there is a positive spectral gap such that $\rho < 1$, the first term in (11) vanishes as t grows. The rest of the first term is bounded and independent of t . Formally, we have

$$\tilde{\pi}_{\max}/\tilde{\pi}_{\min} \leq b, \quad \|\tilde{\pi}\| \geq 1/\sqrt{n}, \quad \text{and} \quad \|p_0 - \tilde{\pi}\| \leq 2,$$

by the assumption that $\max_{i,j} w_i/w_j \leq b$ and the fact that $\tilde{\pi}(i) = w_i/(\sum_j w_j)$. Hence, the error between the distribution at the t th iteration p^t and the true stationary distribution $\tilde{\pi}$ is dominated by the second term in Equation (11). Substituting the bounds in Lemmas 3 and 4, the dominant second term in Equation (11) is bounded by

$$\lim_{t \rightarrow \infty} \frac{\|p_t - \tilde{\pi}\|}{\|\tilde{\pi}\|} \leq \frac{C b^{5/2}}{\xi d_{\min}} \sqrt{\frac{d_{\max} \log n}{k}}$$

with probability of at least $1 - 4n^{-C/8}$. In fact, we only need $t = \Omega(\log n + \log b + \log(d_{\max} \log n / (d_{\min}^2 k \xi^2)))$ to ensure that the above bound holds up to a constant factor. This finishes the proof of Theorem 1. Notice that in order for this result to hold, we need $k \geq 4C^2 b^5 d_{\max} \log n (1/d_{\min} \xi)^2$ for Lemma 4.

4.1.1. Proof of Lemma 2. Due to the reversibility of \tilde{P} , we can view it as a self-adjoint operator on an appropriately defined inner product space. This observation allows us to apply the well-understood spectral analysis of self-adjoint operators. To that end, define an inner product space $L^2(\tilde{\pi})$ as a space of n -dimensional vectors, \mathbb{R}^n , endowed with

$$\langle a, b \rangle_{\tilde{\pi}} = \sum_{i=1}^n a_i \tilde{\pi}_i b_i.$$

Similarly, we define $\|a\|_{\tilde{\pi}} = \sqrt{\langle a, a \rangle_{\tilde{\pi}}}$ as the 2-norm in $L^2(\tilde{\pi})$. An operator (matrix) A is self-adjoint with respect to $L^2(\pi)$ if $\langle u, Av \rangle_{\tilde{\pi}} = \langle Au, v \rangle_{\tilde{\pi}}$ for all $u, v \in \mathbb{R}^n$. For a self-adjoint operator A in $L^2(\tilde{\pi})$, we define $\|A\|_{\tilde{\pi},2} = \max_a \|Aa\|_{\tilde{\pi}} / \|a\|_{\tilde{\pi}}$ as the operator norm. These norms are related to the corresponding norms in the Euclidean space through the following inequalities:

$$\sqrt{\tilde{\pi}_{\min}} \|a\| \leq \|a\|_{\tilde{\pi}} \leq \sqrt{\tilde{\pi}_{\max}} \|a\|, \quad (12)$$

$$\sqrt{\frac{\tilde{\pi}_{\min}}{\tilde{\pi}_{\max}}} \|A\|_2 \leq \|A\|_{\tilde{\pi},2} \leq \sqrt{\frac{\tilde{\pi}_{\max}}{\tilde{\pi}_{\min}}} \|A\|_2. \quad (13)$$

It is easy to check that, a reversible Markov chain \tilde{P} is self-adjoint in $L^2(\tilde{\pi})$ because of the *detailed-balanced condition*, where $\tilde{\pi}$ is the unique stationary distribution of \tilde{P} .

Consider a symmetrized version of \tilde{P} , defined as $S = \tilde{\Pi}^{1/2} \tilde{P} \tilde{\Pi}^{-1/2}$, where $\tilde{\Pi}$ is a diagonal matrix with $\tilde{\Pi}_{ii} = \tilde{\pi}(i)$. Again, reversibility of \tilde{P} makes S symmetric. It

can be verified that \tilde{P} and S have the same set of eigenvalues. By the Perron-Frobenius theorem, the eigenvalues are in $[-1, 1]$ with largest being equal to 1. Let them be denoted as $1 = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \geq -1$, and let $\lambda_{\max} = \max\{|\lambda_n|, \lambda_2\}$. Let u_i be the left eigenvector of S corresponding to λ_i for $1 \leq i \leq n$. Then the i th left eigenvector of \tilde{P} is given by $v_i = \tilde{\Gamma}^{1/2} u_i$. Since the first left eigenvector of \tilde{P} is the stationary distribution, i.e., $v_1 = \tilde{\pi}$, we have that $u_1(i) = \tilde{\pi}(i)^{1/2}$ or $\tilde{\Gamma}^{-1/2} u_1 = \mathbb{1}$. Finally, define rank-1 projection of S as $S_1 = \lambda_1 u_1 u_1^T = u_1 u_1^T$ and let $\tilde{P}_1 = \tilde{\Gamma}^{-1/2} S_1 \tilde{\Gamma}^{1/2}$.

Our interest is in Markov chain $P = \tilde{P} + \Delta$ and iterates obtained from it $p_t^T = p_{t-1}^T P$. Then,

$$p_t^T - \tilde{\pi}^T = (p_{t-1} - \tilde{\pi})^T (\tilde{P} + \Delta) + \tilde{\pi}^T \Delta. \quad (14)$$

Using the fact that $(p_\ell - \tilde{\pi})^T \tilde{\Gamma}^{-1/2} u_1 = (p_\ell - \tilde{\pi})^T \mathbb{1} = 0$ for any probability distribution p_ℓ , we get $(p_\ell - \tilde{\pi})^T \tilde{P}_1 = (p_\ell - \tilde{\pi})^T \tilde{\Gamma}^{-1/2} u_1 u_1^T \tilde{\Gamma}^{1/2} = 0$. Then, from (14) we get

$$p_t^T - \tilde{\pi}^T = (p_{t-1} - \tilde{\pi})^T (\tilde{P} - \tilde{P}_1 + \Delta) + \tilde{\pi}^T \Delta.$$

By definition of \tilde{P}_1 , it follows that $\|\tilde{P} - \tilde{P}_1\|_{\tilde{\pi}, 2} = \|S - S_1\|_2 = \lambda_{\max}$. Let $\rho = \lambda_{\max} + \|\Delta\|_{\tilde{\pi}, 2}$, then

$$\begin{aligned} \|p_t - \tilde{\pi}\|_{\tilde{\pi}} &\leq \|p_{t-1} - \tilde{\pi}\|_{\tilde{\pi}} (\|\tilde{P} - \tilde{P}_1\|_{\tilde{\pi}, 2} + \|\Delta\|_{\tilde{\pi}, 2}) + \|\tilde{\pi}^T \Delta\|_{\tilde{\pi}} \\ &\leq \rho^t \|p_0 - \tilde{\pi}\|_{\tilde{\pi}} + \sum_{\ell=0}^{t-1} \rho^{t-1-\ell} \|\tilde{\pi}^T \Delta\|_{\tilde{\pi}}. \end{aligned}$$

Dividing each side by $\|\tilde{\pi}\|$ and applying the bounds in (12) and (13), we get

$$\frac{\|p_t - \tilde{\pi}\|}{\|\tilde{\pi}\|} \leq \rho^t \sqrt{\frac{\tilde{\pi}_{\max}}{\tilde{\pi}_{\min}}} \frac{\|p_0 - \tilde{\pi}\|}{\|\tilde{\pi}\|} + \sum_{\ell=0}^{t-1} \rho^{t-1-\ell} \sqrt{\frac{\tilde{\pi}_{\max}}{\tilde{\pi}_{\min}}} \frac{\|\tilde{\pi}^T \Delta\|}{\|\tilde{\pi}\|}.$$

This finishes the proof of the desired claim.

4.1.2. Proof of Lemma 3. Our interest is in bounding $\|\Delta\|_2$. Now $\Delta = P - \tilde{P}$ so that for $1 \leq i, j \leq n$,

$$\Delta_{ij} = \frac{1}{kd_{\max}} C_{ij}, \quad (15)$$

where C_{ij} is distributed as per $B(k, p_{ij}) - kp_{ij}$ if $(i, j) \in E$ and $C_{ij} = 0$ otherwise. Here $B(k, p_{ij})$ is a binomial random variable with parameter k and $p_{ij} \equiv w_j/(w_i + w_j)$. It should be noted that $C_{ij} + C_{ji} = 0$ and C_{ij} are independent across all the pairs with $i < j$. For $1 \leq i \leq n$

$$\begin{aligned} \Delta_{ii} &= P_{ii} - \tilde{P}_{ii} = \left(1 - \sum_{j \neq i} P_{ij}\right) - \left(1 - \sum_{j \neq i} \tilde{P}_{ij}\right) \\ &= \sum_{j \neq i} \tilde{P}_{ij} - P_{ij} = - \sum_{j \neq i} \Delta_{ij}. \end{aligned} \quad (16)$$

Given the above dependence between diagonal and off-diagonal entries, we shall bound $\|\Delta\|_2$ as follows: let D be the diagonal matrix with $D_{ii} = \Delta_{ii}$ for $1 \leq i \leq n$ and $\bar{\Delta} = \Delta - D$. Then,

$$\|\Delta\|_2 = \|D + \bar{\Delta}\|_2 \leq \|D\|_2 + \|\bar{\Delta}\|_2. \quad (17)$$

We shall establish the bound of $O(\sqrt{\log n/(kd_{\max})})$ for both $\|D\|_2$ and $\|\bar{\Delta}\|_2$ to establish the Lemma 3.

Bounding $\|D\|_2$. Since D is a diagonal matrix, $\|D\|_2 = \max_i |D_{ii}| = \max_i |\Delta_{ii}|$. For a given fixed i , as per (15)–(16), $kd_{\max} \Delta_{ii}$ can be expressed as the summation of at most kd_{\max} independent, zero-mean random variables taking values in the range of at most 1. Therefore, by an application of the Azuma-Hoeffding's inequality, it follows that

$$\mathbb{P}(kd_{\max} |\Delta_{ii}| > t) \leq 2 \exp\left(-\frac{t^2}{2kd_{\max}}\right). \quad (18)$$

By selection of $t = C\sqrt{kd_{\max} \log n}$ for an appropriately large constant, it follows from the above display that

$$\begin{aligned} \mathbb{P}\left(\|D\|_2 \geq C\sqrt{\frac{\log n}{kd_{\max}}}\right) &\leq \sum_{i=1}^n \mathbb{P}\left(|\Delta_{ii}| > C\sqrt{\frac{\log n}{kd_{\max}}}\right) \\ &\leq 2n^{-C^2/2+1}. \end{aligned} \quad (19) \quad (20)$$

Bounding $\|\bar{\Delta}\|_2$ when $d_{\max} \leq \log n$. Toward this goal, we shall make use of the following standard inequality: for any square matrix M ,

$$\|M\|_2 \leq \sqrt{\|M\|_1 \|M\|_{\infty}}, \quad (21)$$

where $\|M\|_1 = \max_i \sum_j |M_{ij}|$ and $\|M\|_{\infty} = \|M^T\|_1$. In words, $\|M\|_2^2$ is bounded above by product of the maximal row-sum and column-sum of absolute values of M . Since Δ_{ij} and Δ_{ji} are identically distributed and entries along each row (and hence each column) are independent, it is sufficient to obtain a high probability bound ($\geq 1 - 1/\text{poly}(n)$) for maximal row-sum of absolute values of $\bar{\Delta}$; exactly the same bound will apply for column-sum; and using union bound the desired result will follow.

To that end, consider the sum of the absolute values of the i th row-sum of $\bar{\Delta}$ and for simplicity let us denote it by R_i . Then,

$$R_i = \frac{1}{kd_{\max}} \sum_{j \neq i} |C_{ij}|, \quad (22)$$

where recall that $C_{ij} = X_{ij} - kp_{ij}$ with X_{ij} an independent binomial random variable with parameters k, p_{ij} . Therefore, for any $s > 0$,

$$\begin{aligned} \mathbb{P}(R_i > s) &= \mathbb{P}\left(\sum_{j \in \partial i} |C_{ij}| > kd_{\max} s\right) \\ &\leq \sum_{j \in \partial i} \sum_{\xi_j \in \{-1, +1\}} \mathbb{P}\left(\sum_j \xi_j C_{i,j} > kd_{\max} s\right), \\ &\quad \text{by the union bound} \\ &\leq \sum_{j \in \partial i} \sum_{\xi_j \in \{-1, +1\}} \exp\left(\frac{-2k^2 d_{\max}^2 s^2}{d_i k}\right), \end{aligned}$$

where the last inequality follows from Hoeffding's bound and the fact that $X_{ij} = \sum_{j=1}^k (y_{ij} - p_{ij})$, where y_{ij} are Bernoulli random variables with mean p_{ij} . Now,

the number of terms in the sum is 2^{d_i} , the summand is constant, and $d_i \leq d_{\max}$. Thus, the last inequality is upper bounded by

$$\sum_{j \in \partial i} \sum_{\xi_j \in \{-1, +1\}} \exp\left(\frac{-2k^2 d_{\max}^2 s^2}{d_i k^2}\right) \leq \exp(-2kd_{\max} s^2 + d_i \ln 2).$$

By an application of the union bound

$$\begin{aligned} \mathbb{P}(\|\bar{\Delta}\|_2 \geq s) &\leq 2n \mathbb{P}(R_i \geq s) \\ &\leq 2n \exp(-2kd_{\max} s^2 + d_{\max} \ln 2). \end{aligned}$$

Now, if we set $s = (C/2)\sqrt{(\log n + d_{\max} \ln 2)/(kd_{\max})}$ we have that

$$\mathbb{P}\left(\|\bar{\Delta}\|_2 \geq C/2 \sqrt{\frac{\log n + d_{\max} \ln 2}{kd_{\max}}}\right) \leq 2n^{-(C^2/2-1)}.$$

Finally, using the assumption that $d_{\max} \leq \log n$ yields

$$\|\bar{\Delta}\|_2 \leq C \sqrt{\frac{\log n}{kd_{\max}}}$$

with probability at least $1 - 2n^{-C^2/2+1}$.

Bounding $\|\bar{\Delta}\|_2$ When $d_{\max} \geq \log n$. Toward this goal, we shall make use of the recent results on the concentration of the sum of independent random matrices. For completeness, we recall the following result (Tropp 2012).

Lemma 5 (Theorem 6.2 (Tropp 2012)). *Consider a finite sequence $\{\tilde{Z}^{ij}\}_{i < j}$ of independent random self-adjoint matrices with dimensions $n \times n$. Assume that*

$$\mathbb{E}[\tilde{Z}^{ij}] = 0 \quad \text{and} \quad \mathbb{E}(\tilde{Z}^{ij})^p \leq \frac{p!}{2} R^{p-2} (\tilde{A}^{ij})^2, \quad \text{for } p = 2, 3, 4, \dots$$

where $A \leq B$ if and only if $B - A$ is a positive semidefinite matrix.

Define $\tilde{\sigma}^2 \equiv \|\sum_{i < j} (\tilde{A}^{ij})^2\|_2$. Then, for all $t \geq 0$,

$$\mathbb{P}\left(\left\|\sum_{i < j} \tilde{Z}^{ij}\right\|_2 \geq t\right) \leq 2n \exp\left\{\frac{-t^2/2}{\tilde{\sigma}^2 + Rt}\right\}.$$

We wish to prove concentration results on $\bar{\Delta} = \Delta - D = \sum_{i < j} Z^{ij}$, where

$$Z^{ij} = (e_i e_j^T - e_j e_i^T)(P_{ij} - \bar{P}_{ij}) \quad \text{for } (i, j) \in E,$$

and $Z^{ij} = 0$ if i and j are not connected. The Z^{ij} 's as defined are zero-mean and independent, however, they are not self-adjoint. Nevertheless, we can symmetrize it by applying the dilation ideas presented in the paper (Tropp 2012):

$$\tilde{Z}^{ij} \equiv \begin{pmatrix} 0 & Z^{ij} \\ (Z^{ij})^T & 0 \end{pmatrix}.$$

Now we can apply the above lemma to these self-adjoint, independent, and zero-mean random matrices.

To find R and \tilde{A}^{ij} 's that satisfy the conditions of the lemma, first consider a set of matrices $\{A^{ij}\}_{i < j}$ such that $\tilde{Z}^{ij} = \Delta_{ij} A^{ij}$ and

$$A^{ij} = \begin{pmatrix} 0 & e_i e_j^T - e_j e_i^T \\ e_j e_i^T - e_i e_j^T & 0 \end{pmatrix},$$

if $(i, j) \in E$ and zero otherwise. In the following, we show that the condition on the p th moment is satisfied with $R = 1/\sqrt{kd_{\max}^2}$ and $(\tilde{A}^{ij})^2 = (1/(kd_{\max}^2))(A^{ij})^2$ such that

$$\mathbb{E}[(\tilde{Z}^{ij})^p] \leq \frac{p!}{2} \left(\frac{1}{\sqrt{kd_{\max}^2}}\right)^{p-2} \frac{1}{kd_{\max}^2} (A^{ij})^2. \quad (23)$$

We can also show that $\tilde{\sigma}^2 \equiv \|\sum_{i < j} (\tilde{A}^{ij})^2\|_2 = 1/(kd_{\max})$, since

$$\begin{aligned} \sum_{i < j} (\tilde{A}^{ij})^2 &= \sum_{i < j} \frac{1}{kd_{\max}^2} \mathbb{I}_{((i, j) \in E)} \begin{pmatrix} e_i e_i^T + e_j e_j^T & 0 \\ 0 & e_i e_i^T + e_j e_j^T \end{pmatrix} \\ &= \frac{1}{kd_{\max}^2} \sum_{i=1}^n d_i \begin{pmatrix} e_i e_i^T & 0 \\ 0 & e_i e_i^T \end{pmatrix}, \end{aligned}$$

where $\mathbb{I}_{(\cdot)}$ is the indicator function. Using $d_i \leq d_{\max}$ and the structure of matrices in the summation in the last term, it can be easily verified that the $\|\cdot\|_2$ norm of the resulting matrix is at most $1/kd_{\max}$. Now we can apply the results of Lemma 5 to obtain a bound on $\|\sum_{i < j} Z^{ij}\|_2 = \|\sum_{i < j} \tilde{Z}^{ij}\|_2$:

$$\mathbb{P}\left(\left\|\sum_{i < j} Z^{ij}\right\| \geq t\right) \leq 2n \exp\left(\frac{-t^2/2}{(1/kd_{\max}) + (t/\sqrt{kd_{\max}^2})}\right).$$

Under our assumption that $d_{\max} \geq \log n$ and choosing $t = C\sqrt{\log n/(kd_{\max})}$, the tail probability is bounded by

$$2n \exp\{- (C^2 \log n / 2) / (1 + C)\}.$$

Hence, we get the desired bound that $\|\Delta - D\|_2 \leq C\sqrt{\log n/(kd_{\max})}$ with probability at least $1 - 2n^{-C/4+1}$, where we have used the fact that $C \geq 8$.

Now we are left to prove that the condition (23) holds. A quick calculation shows that

$$(A^{ij})^p = \begin{cases} (A^{ij})^2 & \text{for } p \text{ even,} \\ A^{ij} & \text{for } p \text{ odd.} \end{cases} \quad (24)$$

Furthermore, we can verify that the eigenvalues of A^{ij} are either 1 or -1 . Hence, $(A^{ij})^p \leq (A^{ij})^2$ for all $p \geq 1$. Thus, given the fact that $\tilde{Z}^{ij} = \Delta_{ij} A^{ij}$ we have that $\mathbb{E}[(\tilde{Z}^{ij})^p] = \mathbb{E}[\Delta_{ij}^p (A^{ij})^p] \leq |\mathbb{E}[\Delta_{ij}^p]| (A^{ij})^2$ for all p . This fact follows since for any constant $c \in \mathbb{R}$, $cA^{ij} \leq |c|(A^{ij})^2$ and

$c(A^{ij})^2 \leq |c|(A^{ij})^2$. Hence, coupling these observation with the identities presented in Equation (24) we have

$$\mathbb{E}[(\tilde{Z}^{ij})^p] \leq \mathbb{E}[|\Delta_{ij}|^p](A^{ij})^2,$$

where we used Jensen's inequality for $|\mathbb{E}[\Delta_{ij}^p]| \leq \mathbb{E}[|\Delta_{ij}|^p]$.

Next, it remains to construct a bound on $\mathbb{E}[|\Delta_{ij}|^p]$:

$$\mathbb{E}[|\Delta_{ij}|^p] \leq \frac{p!}{2} \left(\frac{1}{\sqrt{k d_{\max}^2}} \right)^p. \quad (25)$$

From (15), we have $\Delta_{ij} = P_{ij} - \tilde{P}_{ij} = (1/(k d_{\max})) C_{ij}$. Therefore,

$$\mathbb{E}[|\Delta_{ij}|^p] = (1/k d_{\max})^p \mathbb{E}[|C_{ij}|^p].$$

Applying Azuma-Hoeffding's inequality to C_{ij} , we have that

$$\mathbb{P}\left(\frac{1}{k d_{\max}} |C_{ij}| \geq t\right) \leq 2 \exp(-2t^2 d_{\max}^2 k).$$

That is, $(1/(k d_{\max})) C_{ij}$ is a sub-Gaussian random variable. And therefore, it follows that for $p \geq 2$,

$$\mathbb{E}\left[\left|\frac{1}{k d_{\max}} C_{ij}\right|^p\right] \leq \frac{p!}{2} \left(\frac{1}{\sqrt{k d_{\max}^2}} \right)^p.$$

This proves the desired bound in (25).

4.1.3. Proof of Lemma 4. By Lemma 3, we have for some $C \geq 8$ that

$$\begin{aligned} 1 - \rho &= 1 - \lambda_{\max}(\tilde{P}) - \|\Delta\|_2 \sqrt{b} \\ &\geq 1 - \lambda_{\max}(\tilde{P}) - C \sqrt{b \log n / (k d_{\max})} \end{aligned}$$

with probability at least $1 - 4n^{-C/8}$. In this section we prove that there is a positive gap: $(d_{\min}/2 b^2 d_{\max}) \xi$. We will first prove that

$$1 - \lambda_{\max}(\tilde{P}) \geq \frac{\xi d_{\min}}{b^2 d_{\max}}. \quad (26)$$

This implies that we have the desired eigen-gap for $k \geq 4C^2 b^5 d_{\max} \log n (1/d_{\min} \xi)^2$ such that $C \sqrt{b \log n / (k d_{\max})} \leq (d_{\min}/2 b^2 d_{\max}) \xi$.

To prove (26), we use comparison theorems (Diaconis and Saloff-Coste 1993), which bound the spectral gap of the Markov chain \tilde{P} of interest using a few comparison inequalities related to a more tractable Markov chain, which is the simple random walk on the graph. We define the transition matrix of the simple random walk on the graph G as

$$Q_{ij} = \frac{1}{d_i} \quad \text{for } (i, j) \in E,$$

and the stationary distribution of this Markov chain is $\mu(i) = d_i / \sum_j d_j$. Further, since the detailed balance

equation is satisfied, Q is a reversible Markov chain. Formally, $\mu(i)Q_{ij} = 1/\sum_\ell d_\ell = \mu(j)Q_{ji}$ for all $(i, j) \in E$.

The following key lemma is a special case of a more general result (Diaconis and Saloff-Coste 1993) proved for two arbitrary reversible Markov chains, which are not necessarily defined on the same graph. For completeness, we provide a proof of this lemma later in this section, following a technique similar to the one in Boyd et al. (2005) used to prove a similar result for a special case when the stationary distribution is uniform.

Lemma 6. Let Q, μ and $\tilde{P}, \tilde{\mu}$ be reversible Markov chains on a finite set $[n]$ representing random walks on a graph $G = ([n], E)$, i.e., $\tilde{P}(i, j) = 0$ and $Q(i, j) = 0$ if $(i, j) \notin E$. For $\alpha \equiv \min_{(i, j) \in E} \{\tilde{\mu}(i)\tilde{P}_{ij}/\mu(i)Q_{ij}\}$ and $\beta \equiv \max_i \{\tilde{\mu}(i)/\mu(i)\}$,

$$\frac{1 - \lambda_{\max}(\tilde{P})}{1 - \lambda_{\max}(Q)} \geq \frac{\alpha}{\beta}. \quad (27)$$

By assumption, we have $\xi \equiv 1 - \lambda_{\max}(Q)$. To prove that there is a positive spectral gap for the random walk of interest as in (26), we are left to bound α and β . We have $\mu(i)Q_{ij} = 1/\sum_\ell d_\ell \leq 1/|E|$ and $\mu(i) \geq (d_i/|E|)$. Also, by assumption that $\max_{i,j} w_i/w_j \leq b$, we have $\tilde{\mu}(i)\tilde{P}_{ij} = w_i w_j / (d_{\max}(w_i + w_j) \sum_\ell w_\ell) \geq 1/(b n d_{\max})$ and $\tilde{\mu}(i) = w_i / \sum_\ell w_\ell \leq b/n$. Then, $\alpha = \min_{(i, j) \in E} \{\tilde{\mu}(i)\tilde{P}_{ij}/\mu(i)Q_{ij}\} \geq |E|/(n b d_{\max})$ and $\beta = \max_i \{\tilde{\mu}(i)/\mu(i)\} \leq b|E|/n d_{\min}$. Hence, $\alpha/\beta \geq d_{\min}/(d_{\max} b^2)$ and this finishes the proof of the bound in (26).

4.1.4. Proof of Lemma 6. Since $1 - \lambda_{\max} = \min\{1 - \lambda_2, 1 + \lambda_n\}$, we will first show that $1 - \lambda_2(Q) \leq (\beta/\alpha) \cdot (1 - \lambda_2(\tilde{P}))$ and $1 + \lambda_n(Q) \leq (\beta/\alpha)(1 + \lambda_n(\tilde{P}))$. The desired bound in (27) immediately follows from the fact that $\min\{a, b\} \leq \min\{a', b'\}$ if $a \leq b$ and $a' \leq b'$.

A reversible Markov chain Q is self-adjoint in $L_2(\mu)$. Then, the second largest eigenvalue $\lambda_2(Q)$ can be represented by the Dirichlet form \mathcal{E} defined as

$$\begin{aligned} \mathcal{E}^Q(\phi, \phi) &\equiv \langle (I - Q)\phi, \phi \rangle_\mu \\ &= \frac{1}{2} \sum_{i,j} (\phi(i) - \phi(j))^2 \mu(i)Q(i, j). \end{aligned}$$

For $\lambda_n(Q)$, we use

$$\begin{aligned} \mathcal{F}^Q(\phi, \phi) &\equiv \langle (I + Q)\phi, \phi \rangle_\mu \\ &= \frac{1}{2} \sum_{i,j} (\phi(i) + \phi(j))^2 \mu(i)Q(i, j). \end{aligned}$$

Following the usual variational characterization of the eigenvalues (see, for instance, Horn and Johnson 1985, p. 176) gives

$$1 - \lambda_2(Q) = \min_{\phi \perp \mathbb{1}} \frac{\mathcal{E}^Q(\phi, \phi)}{\langle \phi, \phi \rangle_\mu}, \quad (28)$$

$$1 + \lambda_n(Q) = \min_{\phi} \frac{\mathcal{F}^Q(\phi, \phi)}{\langle \phi, \phi \rangle_\mu}. \quad (29)$$

By the definitions of α and β , we have $\tilde{\pi}(i, j) \geq \alpha \mu(i) Q(i, j)$ and $\tilde{\pi}(i) \leq \beta \mu(i)$ for all i and j , which implies

$$\begin{aligned}\mathcal{E}^{\tilde{P}}(\phi, \phi) &\geq \alpha \mathcal{E}^Q(\phi, \phi), \\ \mathcal{F}^{\tilde{P}}(\phi, \phi) &\geq \alpha \mathcal{F}^Q(\phi, \phi), \\ \langle \phi, \phi \rangle_{\tilde{\pi}} &\leq \beta \langle \phi, \phi \rangle_{\mu}.\end{aligned}$$

Together with (28), this implies $1 - \lambda_2(Q) \leq (\beta/\alpha) \cdot (1 - \lambda_2(\tilde{P}))$ and $1 + \lambda_n(Q) \leq (\beta/\alpha)(1 + \lambda_n(\tilde{P}))$. This finishes the proof of the desired bound.

4.2. Proof of Theorem 2: Random Sampling

Given the proof of Theorem 1 in the previous section, we only need to prove that for an Erdős-Renyi graph with average degree $d \geq C' \log n$ the following are true:

$$(1/2)d \leq d_i \leq (3/2)d, \quad (30)$$

$$1/2 \leq \xi. \quad (31)$$

Then, it follows that $\kappa \leq 3$ and $(1/2)d \leq d_{\min} \leq d_{\max} \leq (3/2)d$. By Theorem 1, it follows that

$$\frac{\|\pi - \tilde{\pi}\|}{\|\tilde{\pi}\|} \leq 6Cb^{5/2} \sqrt{\frac{\log n}{kd}},$$

with probability at least $1 - 4n^{-C/8}$ for some positive constant $C \geq 8$ and for $kd \geq 288C^2b^5 \log n$.

We can apply standard concentration inequalities to establish Equation (30). Applying Chernoff's inequality, we get $\mathbb{P}[|d_i - d| > (1/2)d] \leq 2e^{-d/16}$. Hence, for $d \geq C' \log n$, Equation (30) is true with probability at least $1 - 2n^{-C'/16}$.

Finally, we finish the proof with a result on the lower bound of the spectral gap $\xi = 1 - \lambda_{\max}(D^{-1}B)$.

Lemma 7. Consider a random graph G drawn from the Erdős-Renyi distribution $G(n, d/n)$. Then if $d \geq 10C^2 \log n$, we have $\xi \geq 1/2$ with probability at least $1 - n^{-Cn/(n-d)/8}$.

The proof of this result can be found in the online appendix.

4.3. Proof of Theorem 3: Information-Theoretic Lower Bound

In this section, we prove Theorem 3 using an information-theoretic method that allows us to reduce the stochastic inference problem into a multiway hypothesis testing problem.

This estimation problem can be reduced to the following hypothesis testing problem. Consider a set $\{\tilde{\pi}^{(1)}, \dots, \tilde{\pi}^{(M(\delta))}\}$ of $M(\delta)$ vectors on the standard orthogonal simplex that are separated by δ , such that $\|\tilde{\pi}^{(\ell_1)} - \tilde{\pi}^{(\ell_2)}\| \geq \delta$ for all $\ell_1 \neq \ell_2$. To simplify the notations, we are going to use M as a shorthand for $M(\delta)$. Suppose we choose an index $L \in \{1, \dots, M\}$ uniformly

at random. Then, we are given noisy outcomes of pairwise comparisons with $w = \tilde{\pi}^{(L)}$ from the BTL model. We use X to denote this set of observations. Let π be the estimation produced by an algorithm using the noisy observations. Given this, the best estimation of the “index” is \hat{L} , where $\hat{L} = \arg \min_{\ell \in [M]} \|\pi - \tilde{\pi}^{(\ell)}\|$.

By construction of our packing set, when we make a mistake in the hypothesis testing, our estimate is at least $\delta/2$ away from the true weight $\tilde{\pi}^{(L)}$. Precisely, $\hat{L} \neq L$ implies that $\|\pi - \tilde{\pi}^{(L)}\| \geq \delta/2$. Then,

$$\begin{aligned}\mathbb{E}[\|\pi - \tilde{\pi}^{(L)}\|] &\geq \frac{\delta}{2} \mathbb{P}(\hat{L} \neq L) \\ &\geq \frac{\delta}{2} \left\{ 1 - \frac{I(\hat{L}; L) + \log 2}{\log M} \right\},\end{aligned} \quad (32)$$

where $I(\cdot; \cdot)$ denotes the mutual information between two random variables and the second inequality follows from Fano's inequality.

These random vectors form a Markov chain $L - \tilde{\pi}^{(L)} - X - \pi - \hat{L}$, where $X - Y - Z$ indicates that X and Z are conditionally independent given Y . Let $\mathbb{P}_{L, X}(\ell, x)$ denote the joint probability function, and $\mathbb{P}_{X|L}(x|\ell)$, $\mathbb{P}_L(\ell)$ and $\mathbb{P}_X(x)$ denote the conditional and marginal probability functions. Then, by data processing inequality for a Markov chain, we get

$$\begin{aligned}I(L; \hat{L}) &\leq I(L; X) \\ &= \mathbb{E}_{L, X} \left[\log \left(\frac{\mathbb{P}_{L, X}(L, X)}{\mathbb{P}_L(L) \mathbb{P}_X(X)} \right) \right] \\ &= \frac{1}{M} \sum_{\ell \in [M]} \mathbb{E}_X \left[\log \left(\frac{\mathbb{P}_{X|L}(X|\ell)}{\mathbb{P}_X(X)} \right) \right] \\ &= \frac{1}{M} \sum_{\ell \in [M]} \mathbb{E}_X \left[\log \left(\frac{\mathbb{P}_{X|L}(X|\ell)}{\sum_{\ell_2 \in [M]} \mathbb{P}_{X|L}(X|\ell_2) \mathbb{P}(\ell_2)} \right) \right] \\ &\leq \frac{1}{M} \sum_{\ell \in [M]} \sum_{\ell_2 \in [M]} \mathbb{P}(\ell_2) \mathbb{E}_X \left[\log \left(\frac{\mathbb{P}_{X|L}(X|\ell)}{\mathbb{P}_{X|L}(X|\ell_2)} \right) \right] \\ &= \frac{1}{M^2} \sum_{\ell_1, \ell_2} D_{\text{KL}}(\mathbb{P}_{X|L}(X|\ell_1) \| \mathbb{P}_{X|L}(X|\ell_2)),\end{aligned} \quad (33)$$

where $D_{\text{KL}}(\cdot \| \cdot)$ is the Kullback-Leibler (KL) divergence and the inequality follows from the concavity of logarithm and Jensen's inequality.

The KL divergence between the observations coming from two different BTL models depends on how we sample the comparisons. We are sampling each pair of items for comparison with probability d/n , and we are comparing each of these sampled pairs k times. Let X_{ij} denote the outcome of k comparisons for a sampled pair of items (i, j) . To simplify notations, we drop the subscript $X|L$ whenever it is clear from the context. Then,

$$\begin{aligned}D_{\text{KL}}(\mathbb{P}(X|\ell_1) \| \mathbb{P}(X|\ell_2)) &= \frac{d}{n} \sum_{1 \leq i < j \leq n} D_{\text{KL}}(\mathbb{P}(X_{ij}|\ell_1) \| \mathbb{P}(X_{ij}|\ell_2)) \\ &\leq 2n^2kd \|(\tilde{\pi}^{(\ell_1)} - \tilde{\pi}^{(\ell_2)})\|^2,\end{aligned} \quad (34)$$

where in the last inequality we used the fact that

$$\begin{aligned} D_{\text{KL}}(\mathbb{P}(X_{ij} | \ell_1) \| \mathbb{P}(X_{ij} | \ell_2)) \\ \leq \frac{k(\tilde{\pi}_j^{(\ell_2)}(\tilde{\pi}_i^{(\ell_1)} - \tilde{\pi}_i^{(\ell_2)})^2 + \tilde{\pi}_i^{(\ell_2)}(\tilde{\pi}_j^{(\ell_1)} - \tilde{\pi}_j^{(\ell_2)})^2)}{\tilde{\pi}_i^{(\ell_2)}\tilde{\pi}_j^{(\ell_2)}(\tilde{\pi}_i^{(\ell_1)} + \tilde{\pi}_j^{(\ell_1)})} \\ \leq 2kn^2((\tilde{\pi}_i^{(\ell_1)} - \tilde{\pi}_i^{(\ell_2)})^2 + (\tilde{\pi}_j^{(\ell_1)} - \tilde{\pi}_j^{(\ell_2)})^2), \end{aligned}$$

for k independent trials of Bernoulli random variables, and $\tilde{\pi}_i^{(\ell)} \geq 1/(2n)$ for all i and ℓ , which follows from our construction of the packing set in Lemma 8 and our choice of δ .

The remainder of the proof relies on the following key technical lemma, on the construction of a suitable packing set that has enough number of entries that are reasonably separated. This is proved in Section 4.3.1.

Lemma 8. For $n \geq 90$ and for any positive $\delta \leq 1/2\sqrt{10n}$, there exists a set of n -dimensional vectors $\{\tilde{\pi}^{(1)}, \dots, \tilde{\pi}^{(M)}\}$ with cardinality $M = e^{n/128}$ such that $\sum_i \tilde{\pi}_i^{(\ell)} = 1$ and

$$\frac{1 - 2\delta\sqrt{10n}}{n} \leq \tilde{\pi}_i^{(\ell)} \leq \frac{1 + 2\delta\sqrt{10n}}{n},$$

for all $i \in [n]$ and $\ell \in [M]$, and

$$\delta \leq \|\tilde{\pi}^{(\ell_1)} - \tilde{\pi}^{(\ell_2)}\| \leq \sqrt{13}\delta,$$

for all $\ell_1 \neq \ell_2$.

Substituting this bound in Equations (34), (33), and (32), we get

$$\begin{aligned} \max_{\ell \in [M]} \mathbb{E}[\|\pi - \tilde{\pi}^{(\ell)}\|] &\geq \mathbb{E}[\|\pi - \tilde{\pi}^{(L)}\|] \\ &\geq \frac{\delta}{2} \left\{ 1 - \frac{3,328n^2kd\delta^2 + 128\log 2}{n} \right\}. \end{aligned}$$

Choosing $\delta = (b-1)/(30\sqrt{10}(b+1)\sqrt{kdn})$, we know that $3,328n^2kd\delta^2 + 128\log 2 \leq (1/2)n$ for all b and all $n \geq 682$. This implies that

$$\max_{\ell \in [M]} \mathbb{E}[\|\pi - \tilde{\pi}^{(\ell)}\|] \geq \frac{b-1}{120(b+1)\sqrt{10kdn}}.$$

From Lemma 8, it follows that $\|\tilde{\pi}^{(\ell)}\| \leq 2/\sqrt{n}$ for all ℓ . Then, scaling the bound by $1/\|\tilde{\pi}^{(\ell)}\|$, the normalized minimax rate is lower bounded by $(b-1)/(240(b+1)\sqrt{10kd})$. Also, for this choice of δ , the dynamic range is at most b . From Lemma 8, the dynamic range is upper bounded by

$$\max_{\ell, i, j} \frac{\tilde{\pi}_i^{(\ell)}}{\tilde{\pi}_j^{(\ell)}} \leq \frac{1 + 2\delta\sqrt{10n}}{1 - 2\delta\sqrt{10n}}.$$

This is monotonically increasing in δ for $\delta < 1/(2\sqrt{10n})$. Hence, for $\delta \leq (b-1)/((b+1)2\sqrt{10n})$, which is always true for our choice of δ , the dynamic range is upper bounded by b . This finishes the proof of the desired bound on normalized minimax error rates for general b .

4.3.1. Proof of Lemma 8. We show that a random construction succeeds in generating a set of M vectors on the standard orthogonal simplex satisfying the conditions with a strictly positive probability. Let $M = e^{n/128}$ and for each $\ell \in [M]$, we construct independent random vectors $\tilde{\pi}^{(\ell)}$ according to the following procedure. For a positive α to be specified later, we first draw n random variables uniformly from $[(1 - \alpha\delta\sqrt{n})/n, (1 + \alpha\delta\sqrt{n})/n]$. Let $Y^{(\ell)} = [Y_1^{(\ell)}, \dots, Y_n^{(\ell)}]$ denote this random vector in n dimensions. Then we project this onto the n -dimensional simplex by setting

$$\tilde{\pi}^{(\ell)} = Y^{(\ell)} + (1/n - \bar{Y}^{(\ell)})\mathbb{1},$$

where $\bar{Y}^{(\ell)} = (1/n) \sum_i Y_i^{(\ell)}$. By construction, the resulting vector is on the standard orthogonal simplex: $\sum_i \tilde{\pi}_i^{(\ell)} = 1$. Also, applying Hoeffding's inequality for $\bar{Y}^{(\ell)}$, we get that

$$\mathbb{P}\left(\left|\bar{Y}^{(\ell)} - \frac{1}{n}\right| > \frac{\alpha\delta}{\sqrt{n}}\right) \leq 2e^{-n/2}.$$

By union bound, this holds uniformly for all ℓ with probability at least $1 - 2e^{-63n/128}$. In particular, this implies that

$$\frac{1 - 2\alpha\delta\sqrt{n}}{n} \leq \tilde{\pi}_i^{(\ell)} \leq \frac{1 + 2\alpha\delta\sqrt{n}}{n}, \quad (35)$$

for all $i \in [n]$ and $\ell \in [M]$.

Next, we use standard concentration results to bound the distance between two vectors:

$$\|\tilde{\pi}^{(\ell_1)} - \tilde{\pi}^{(\ell_2)}\|^2 = \|Y^{(\ell_1)} - Y^{(\ell_2)}\|^2 - n(\bar{Y}^{(\ell_1)} - \bar{Y}^{(\ell_2)})^2.$$

Applying Hoeffding's inequality for the first term, we get $\mathbb{P}(|\sum_i (Y_i^{(\ell_1)} - Y_i^{(\ell_2)})^2 - (2/3)\alpha^2\delta^2| \geq (1/2)\alpha^2\delta^2) \leq 2e^{-n/32}$. Similarly for the second term, we can show that $\mathbb{P}(|\sum_i (Y_i^{(\ell_1)} - Y_i^{(\ell_2)})| \geq (1/4)\alpha\delta\sqrt{n}) \leq 2e^{-n/32}$. Substituting these bounds, we get

$$\frac{1}{10}\alpha^2\delta^2 \leq \|\tilde{\pi}^{(\ell_1)} - \tilde{\pi}^{(\ell_2)}\|^2 \leq \frac{13}{10}\alpha^2\delta^2, \quad (36)$$

with probability at least $1 - 4e^{-n/32}$. Applying union bound over $\binom{M}{2} \leq e^{n/64}$ pairs of vectors, we get that the lower and upper bound holds for all pairs $\ell_1 \neq \ell_2$ with probability at least $1 - 4e^{-n/64}$.

The probability that both conditions (35) and (36) are satisfied is at least $1 - 4e^{-n/64} - 2e^{-63n/128}$. For $n \geq 90$, the probability of success is strictly positive. Hence, we know that there exists at least one set of vectors that satisfy the conditions. Setting $\alpha = \sqrt{10}$, we have constructed a set that satisfies all the conditions.

4.4. Proof of Theorem 4: Finite Sample Analysis of MLE

The proof of this theorem will follow in two parts. First we will show that if the gradient of the loss $\nabla \mathcal{L}_m$ evaluated at θ^* is small, then the error between θ^*

and $\hat{\theta}$ is also small. To that end we begin with a simple inequality:

$$\mathcal{L}_m(\hat{\theta}) \leq \mathcal{L}_m(\theta^*).$$

Let $\Delta = \hat{\theta} - \theta^*$. We can add and subtract $\langle \nabla \mathcal{L}_m(\theta^*), \Delta \rangle$ from the above equation to obtain

$$\mathcal{L}_m(\theta^* + \Delta) - \mathcal{L}_m(\theta^*) - \langle \nabla \mathcal{L}_m(\theta^*), \Delta \rangle \leq \langle \nabla \mathcal{L}_m(\theta^*), \Delta \rangle.$$

Now assume $\|\nabla \mathcal{L}_m(\theta^*)\|_2 \leq c$. By the Cauchy-Schwartz inequality we have that

$$\mathcal{L}_m(\theta^* + \Delta) - \mathcal{L}_m(\theta^*) - \langle \nabla \mathcal{L}_m(\theta^*), \Delta \rangle \leq c \|\Delta\|_2.$$

Therefore, we if we prove that

$$\mathcal{L}_m(\theta^* + \Delta) - \mathcal{L}_m(\theta^*) - \langle \nabla \mathcal{L}_m(\theta^*), \Delta \rangle \geq \frac{\mu}{2} \|\Delta\|_2^2, \quad (37)$$

then we immediately have that $\|\Delta\|_2 \leq 2c/\mu$. We now proceed to establish the above inequality.

4.4.1. Proof of Equation (37). By Taylor's theorem and the definition of \mathcal{L}_m from Equation (9) for some $v \in [0, 1]$ we have

$$\begin{aligned} \mathcal{L}_m(\theta^* + \Delta) - \mathcal{L}_m(\theta^*) - \langle \nabla \mathcal{L}_m(\theta^*), \Delta \rangle \\ = \frac{1}{2m} \sum_{l=1}^m \frac{\exp(\langle \theta^*, x_l \rangle + v \langle \theta^*, x_l \rangle)}{(1 + \exp(\langle \theta^*, x_l \rangle + v \langle \theta^*, x_l \rangle))^2} (\langle \Delta, x_l \rangle)^2. \end{aligned}$$

Now, by assumption $\sum_i \theta_i^* = \sum_i \hat{\theta}_i = 0$; and $\theta_{\max}^* - \theta_{\min}^*$ and $\hat{\theta}_{\max} - \hat{\theta}_{\min} \leq \log(b)$ so that $|\langle \theta^*, x_l \rangle + v \langle \theta^*, x_l \rangle| \leq \log(b)$. Therefore,

$$\begin{aligned} \mathcal{L}_m(\theta^* + \Delta) - \mathcal{L}_m(\theta^*) - \langle \nabla \mathcal{L}_m(\theta^*), \Delta \rangle \\ \geq \frac{1}{2m} \sum_{l=1}^m \frac{b}{(1+b)^2} (\langle \Delta, x_l \rangle)^2. \end{aligned}$$

Thus, what remains is to establish a lower bound on

$$\frac{1}{m} \sum_{l=1}^m (\langle \Delta, x_l \rangle)^2.$$

We appeal to the following lemma for the lower bound.

Lemma 9. Given $m > 12n \log n$ i.i.d. samples y_l, x_l we have that

$$\frac{1}{m} \sum_{l=1}^m (\langle \Delta, x_l \rangle)^2 \geq \frac{1}{3n} \|\Delta\|_2^2$$

with probability at least $1 - 1/n$.

Finally, we present the following lemma that establishes an upper bound on $\|\nabla \mathcal{L}_m(\theta^*)\|_2$.

Lemma 10. Given m observations (v_l, x_l) we have that

$$\|\nabla \mathcal{L}_m(\theta^*)\|_2 \leq 2\sqrt{\frac{\log n}{m}}$$

with probability at least $1 - 1/n$.

Therefore, putting everything together we have that

$$\|\Delta\|_2 \leq \frac{6(1+b)^2}{b\sqrt{n^2 \log n / m}},$$

which establishes the desired result.

4.4.2. Proof of Lemma 9. To prove this lemma we note that

$$\frac{1}{m} \sum_{l=1}^m (\langle \Delta, x_l \rangle)^2 = \frac{1}{m} \sum_{l=1}^m \Delta^T x_l x_l^T \Delta.$$

Thus, it is sufficient to prove a lower bound on $\lambda_{\min}((1/m) \sum_{l=1}^m x_l x_l^T)$. To do so we may again appeal to recent results on random matrix theory Tropp (2012).

Lemma 11 [Theorem 1.4 (Tropp 2012)]. Consider a finite sequence $\{X_k\}$ of independent, random, self-adjoint matrices with dimensions d . Assume that each random matrix satisfies $\mathbb{E} X_k = 0$ and $\lambda_{\max}(X_k) \leq R$ almost surely. Then, for all $t \geq 0$,

$$\begin{aligned} \mathbb{P}\left\{\lambda_{\max}\left(\sum_k X_k\right) \geq t\right\} &\leq d \cdot \exp\left(\frac{-t^2/2}{\sigma^2 + Rt/3}\right), \\ \text{where } \sigma^2 &:= \left\|\sum_k \mathbb{E}(X_k^2)\right\|, \end{aligned} \quad (38)$$

and $\|X\|$ for a matrix X represents the operator norm of X or its largest singular value.

To apply the above lemma we let

$$X_l = x_l x_l^T - \frac{2}{n(I - \mathbb{1}\mathbb{1}^T/n)}.$$

Therefore, the X_l are zero-mean, i.i.d., and symmetric. Furthermore, $\|X_l\| \leq 2$ and $\mathbb{E} X_l^2 = 4/n(I - \mathbb{1}\mathbb{1}^T/n) - 4/n^2(I - \mathbb{1}\mathbb{1}^T/n)$. Therefore, applying the above lemma to both X_l and $-X_l$ yields the inequality

$$\mathbb{P}\left\{\left\|\sum_l X_l/m\right\| \geq t\right\} \leq 2n \exp\left(\frac{-t^2/2}{4/(nm) + 2t/(3m)}\right).$$

Thus, with probability at least $1 - 1/n$,

$$\left\|\frac{1}{m} \sum_l X_l\right\| \leq \max\left(4\sqrt{\frac{2 \log n}{nm}}, 8/3 \frac{\log n}{m}\right).$$

Hence, as long as $12n \log n < m$, then

$$\left\|\frac{1}{m} \sum_l X_l\right\| \leq 4\sqrt{\frac{2 \log n}{nm}},$$

with probability at least $1 - 1/n$.

With the above result in hand we now have that

$$\left\|\frac{1}{m} \sum_{l=1}^m x_l x_l^T - \frac{2}{n}(I - \mathbb{1}\mathbb{1}^T/n)\right\| \leq 4\sqrt{\frac{2 \log n}{nm}}.$$

Therefore,

$$\frac{1}{m} \sum_{l=1}^m \Delta^T x_l x_l^T \Delta \geq \frac{2}{n} \|\Delta\|_2^2 \left(1 - 2\sqrt{\frac{2n \log n}{m}}\right),$$

where we have used the fact that $\Delta = \hat{\theta} - \theta^*$ and $\sum_i \hat{\theta}_i = \sum_i \theta_i^* = 0$. Recalling that, $m > 12n \log n$ the above inequality can be lower bounded by $(1/(3n))\|\Delta\|_2^2$, establishing the desired result.

4.4.3. Proof of Lemma 10. To establish this result, we will proceed by showing each individual element of $\nabla \mathcal{L}_m$ is upper bounded by $2\sqrt{\log n/(nm)}$ with high probability. Recall that

$$\nabla \mathcal{L}_m = \frac{1}{m} \sum_{l=1}^m x_l (\mathbb{E}[X_l | x_l] - X_l).$$

Consequently, focusing on a single component $\nabla \mathcal{L}_{mk}$ we have that

$$\nabla \mathcal{L}_{mk} = \frac{1}{m} \sum_{l=1}^m (x_l)_k (\mathbb{E}[X_l | x_l] - X_l).$$

Thus, the k th component of $\nabla \mathcal{L}_m$ is the average over m independent mean zero random variables that are upper bounded by 1 and that each have variance upper bounded by $1/n$. Therefore, an application of Bernstein's inequality yields

$$\mathbb{P}(|\nabla \mathcal{L}_{mk}| \geq t) \leq 2 \exp\left(\frac{-t^2}{2/(nm) + 2t/(3m)}\right).$$

Therefore,

$$\begin{aligned} \mathbb{P}(\|\nabla \mathcal{L}_m\|_\infty \geq t) &\leq n \mathbb{P}(|\nabla \mathcal{L}_{mk}| \geq t) \\ &\leq 2n \exp\left(\frac{-t^2}{2/(nm) + 2t/(3m)}\right). \end{aligned}$$

Using arguments similar to those to establish the results in Section 4.4.1 we have that with probability at least $1 - 2/n$

$$\|\nabla \mathcal{L}_m\|_\infty \leq 2\sqrt{\frac{\log n}{nm}},$$

as desired.

5. Conclusion

The main contribution of this paper is the design and analysis of RANK CENTRALITY: an iterative algorithm for rank aggregation using pairwise comparisons. We established the efficacy of the algorithm by analyzing its performance when data are generated as per the popular BTL or MNL model. We have obtained an analytic bound on the finite sample error rates between the scores assumed by the BTL model and those estimated by our algorithm. As shown, these lead to near-optimal dependence on the number of samples required to learn the scores well by our algorithm under random selection of pairs for comparison. More generally, the comparison graph structure plays a crucial role in the performance of the algorithm.

For a tighter analysis of the optimality of Rank Centrality, we provide numerical experiments under the BTL model and compare it to the Cramer Rao lower bound. Comparisons with the Cramer-Rao bound in Figure 3 suggests that the error achieved by Rank

Centrality is indistinguishable from the fundamental Cramer-Rao lower bound, and thus suggest its stronger optimality properties compared to what we can establish.

For completeness, we further provided an analysis of the error achieved by the MLE. Building upon our analysis, Hajek et al. (2014) shows that the MLE is near order optimal, just like RANK CENTRALITY. It is worth noting, however, that empirically the computational cost of RANK CENTRALITY seems much better than that of finding the MLE.

Endnotes

¹ Similar algorithms, based on the comparison data matrix have been proposed in the literature. As discussed in detail in Section 3.3, they are all different from RANK CENTRALITY.

² <http://www.washingtonpost.com/wp-srv/interactivity/worst-year-voting.html>.

³ <http://www.allourideas.org>.

References

- Adler M, Gemmell P, Harchol-Balter M, Karp RM, Kenyon C (1994) Selection in the presence of noise: The design of playoff systems. *Proc. 5th Annual ACM-SIAM Sympos. Discrete Algorithms, SODA '94* (SIAM, Philadelphia), 564–572.
- Ailon N (2010) Aggregation of partial rankings, p-ratings and top-m lists. *Algorithmica* 57(2):284–300.
- Ailon N, Charikar M, Newman A (2008) Aggregating inconsistent information: Ranking and clustering. *J. ACM (JACM)* 55(5):23.
- Altman A, Tennenholtz M (2005) Ranking systems: The pagerank axioms. Riedl J, Kearns MJ, Reiter MK, eds. *Proc. 6th ACM Conf. Electronic Commerce, EC '05* (ACM, New York), 1–8.
- Ammar A, Shah D (2011) Ranking: Compare, don't score. *49th Annual Allerton Conf. Comm., Control, and Comput., Allerton '11* (IEEE, Piscataway, NJ), 776–783.
- Arrow KJ (1963) *Social Choice and Individual Values* (Yale University Press, New Haven, CT).
- Boyd S, Ghosh A, Prabhakar B, Shah D (2005) Mixing times for random walks on geometric random graphs. Demetrescu C, Sedgewick R, Tamassia R, eds. *Proc. Second Workshop Analytic Algorithms Combinatorics, ANALCO '05* (SIAM, Philadelphia), 240–249.
- Bradley RA, Terry ME (1955) Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika* 39(3/4): 324–345.
- Braverman M, Mossel E (2008) Noisy sorting without resampling. Teng S-H, ed. *Proc. 19th Annual ACM-SIAM Sympos. Discrete Algorithms, SODA '08* (SIAM, Philadelphia), 268–276.
- Brin S, Page L (1998) The anatomy of a large-scale hypertextual web search engine. *Seventh Internat. World-Wide Web Conf., WWW7* (Elsevier, Amsterdam), 107–117.
- Candès EJ, Recht B (2009) Exact matrix completion via convex optimization. *Foundations of Computational Math.* 9(6):717–772.
- Condorcet M (1785) *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix* (l'Imprimerie Royale, Paris).
- David HA (1963) *The Method of Paired Comparisons* (Charles Griffin and Company, London).
- de Borda J-C (1781) Mémoire sur les élections au scrutin. *Mémoires de l'Académie royale des sciences*, 657–664.
- Diaconis P, Saloff-Coste L (1993) Comparison theorems for reversible Markov chains. *Ann. Appl. Probab.* 3(3):696–730.
- Duchi JC, Mackey L, Jordan MI (2010) On the consistency of ranking algorithms. Fürnkranz J, Joachims T, eds. *Proc. 27th Internat. Conf. Machine Learning, ICML '10* (Omnipress, Madison, WI), 327–334.

- Dwork C, Kumar R, Naor M, Sivakumar D (2001) Rank aggregation methods for the web. Shen VY, Saito N, Lyu MR, Zurko ME, eds. *Proc. 10th Internat. World Wide Web Conf., WWW '10* (ACM, New York), 613–622.
- Dwork C, Kumar R, Naor M, Sivakumar D (2001b) Rank aggregation methods for the web. *Proc. 10th Internat. World Wide Web Conf.* (ACM, New York), 613–622.
- Farnoud F, Touri B, Milenkovic O (2012) Novel distance measures for vote aggregation. Preprint, arXiv:1203.6371.
- Ford Jr LR (1957) Solution of a ranking problem from binary comparisons. *Amer. Math. Monthly* 64(8):28–33.
- Gleich DF, Lim L (2011) Rank aggregation via nuclear norm minimization. Apté C, Ghosh J, Smyth P, eds. *Proc. 17th ACM SIGKDD Internat. Conf. Knowledge Discovery and Data Mining* (ACM, New York), 60–68.
- Guiver J, Snelson E (2009) Bayesian inference for Plackett-Luce ranking models. Pohorecký J, Danyluk A, Bottou L, Littman ML, eds. *Proc. 26th Annual Internat. Conf. Machine Learn., ICML '09* (ACM, New York), 377–384. [Available at <http://sites.stat.psu.edu/~dhunter/code/btmatlab/>]
- Hajek B, Oh S, Xu J (2014) Minimax-optimal inference from partial rankings. Ghahramani Z, Welling M, Cortes C, Lawrence ND, Weinberger KQ, eds. *Adv. Neural Inform. Processing Systems (NIPS)*, 1475–1483.
- Hochbaum DS (2006) Ranking sports teams and the inverse equal paths problem. Spirakis P, Mavronicolas M, Kontogiannis S, eds. *Internet and Network Economics* (Springer, Berlin), 307–318.
- Horn RA, Johnson CR (1985) *Matrix Analysis* (Cambridge University Press, Cambridge, UK).
- Hunter DR (2004) MM algorithms for generalized Bradley-Terry models. *Ann. Statist.* 32(1):384–406.
- Jiang X, Lim L, Yao Y, Ye Y (2011) Statistical ranking and combinatorial Hodge theory. *Math. Programming* 127(1):203–244.
- Kamvar SD, Schlosser MT, Garcia-Molina H (2003) The eigentrust algorithm for reputation management in P2P networks. Hencsey G, White B, Chen Y-F, Kovács L, Lawrence S, eds. *Proc. 12th Internat. Conf. World Wide Web, WWW '03* (ACM, New York), 640–651.
- Keener JP (1993) The Perron-Frobenius theorem and the ranking of football teams. *SIAM Rev.* 35(1):80–93.
- Kendall MG (1955) Further contributions to the theory of paired comparisons. *Biometrics* 11(1):43–62.
- Kendall MG, Smith BB (1940) On the method of paired comparisons. *Biometrika* 31(3):324–345.
- Keshavan RH, Montanari A, Oh S (2010) Matrix completion from noisy entries. *J. Machine Learn. Res.* 11(1):2057–2078.
- Lu T, Boutilier C (2011) Learning Mallows models with pairwise preferences. *Proc. 28th Internat. Conf. Machine Learn. ICML '11* (Omnipress, Madison, WI), 145–152.
- Luce DR (1959) *Individual Choice Behavior* (John Wiley & Sons, New York).
- Mallows CL (1957) Non-null ranking models. I. *Biometrika* 44(1): 114–130.
- McFadden D (1973) Conditional logit analysis of qualitative choice behavior. *Frontiers Econometrics* 1(1):105–142.
- Mosteller F (1951) Remarks on the method of paired comparisons: I. The least squares solution assuming equal standard deviations and equal correlations. *Psychometrika* 16(1):3–9.
- Negahban S, Wainwright MJ (2012) Restricted strong convexity and (weighted) matrix completion: Optimal bounds with noise. *J. Machine Learn. Res.* 13(1):1665–1697.
- Newman MEJ (2010) *Networks: An Introduction* (Oxford University Press, Oxford, UK).
- Osting B, Brune C, Osher S (2013) Enhanced statistical rankings via targeted data collection. *Proc. 30th Internat. Conf. Machine Learn., ICML '13* (JMLR.org), 489–497.
- Plackett RL (1975) The analysis of permutations. *Appl. Statist.* 24(2): 193–202.
- Rajkumar A, Agarwal S (2014) A statistical convergence perspective of algorithms for rank aggregation from pairwise data. *Proc. 31st Internat. Conf. Machine Learn., ICML '14* (JMLR.org), 118–126.
- Rao CR (1945) Information and accuracy attainable in the estimation of statistical parameters. *Bull. Calcutta Math. Soc.* 37(3):81–91.
- Saaty TL (2003) Decision making with the AHP: Why is the principal eigenvector necessary. *Eur. J. Oper. Res.* 145(1):85–91.
- Salganik MJ, Levy KEC (2012) Wiki surveys: Open and quantifiable social data collection. Preprint, arXiv:1202.0500.
- Seeley JR (1949) The net of reciprocal influence. *Canadian J. Psych.* 3(4):234–240.
- Shah D, Zaman T (2011) Rumors in a network: Who's the culprit? *IEEE Trans. Inform. Theory* 57(8):5163–5181.
- Shah D, Zaman T (2016) Finding rumor sources on random trees. *Oper. Res.* 64(3):736–755.
- Talluri KT, Van Ryzin G (2005) *The Theory and Practice of Revenue Management* (Springer, New York).
- Tropp J (2012) User-friendly tail bounds for sums of random matrices. *Foundations of Computational Math.* 12(4):389–434.
- Vigna S (2009) Spectral ranking. Preprint, arXiv: 0912.0238.
- Volkovs MN, Zemel RS (2012) A flexible generative model for preference aggregation. Mille A, Gandon FL, Misselis J, Rabinovich M, Staab S, eds. *Proc. 21st Internat. Conf. World Wide Web, WWW '12* (ACM, New York), 479–488.
- Wei TH (1952) The algebraic foundations of ranking theory. Ph.D. thesis, University of Cambridge, Cambridge, UK.

Sahnd Negahban is an assistant professor in the Statistics Department at Yale. His research interest is to develop theoretically sound methods, which are both computationally and statistically efficient, for extracting information from large data sets.

Sewoong Oh is an assistant professor of Industrial and Enterprise Systems Engineering at the University of Illinois at Urbana-Champaign. He was coawarded the Kenneth C. Sevick outstanding student paper award at the SIGMETRICS 2010 and the best paper award at the SIGMETRICS 2015. He was awarded the National Science Foundation faculty early career award in 2016. His current research interests are in developing efficient algorithms with provable guarantees for extracting useful information from social platforms such as ranking, crowdsourcing, and recommendation systems.

Devavrat Shah is a Jamieson Career Development Associate Professor with the department of electrical engineering and computer science at MIT. He is a member of the Laboratory for Information and Decision Systems and Operations Research Center. His research focus is on the theory of large complex networks, which include network algorithms, stochastic networks, network information theory, and large-scale statistical inference. He was the winner of the George B. Dantzig best dissertation award from INFORMS in 2005. He has been coawarded best paper awards at the IEEE INFOCOM '04, ACM SIGMETRICS/Performance '06; and he has supervised best student paper awards at Neural Information Processing Systems '08, ACM SIGMETRICS/Performance '09 and Management Science and Operations Management Paper competition '10. He was awarded the first ACM SIGMETRICS Rising Star Award 2008 for his work on network scheduling algorithms. He received the 2010 Erlang Prize from INFORMS, which is given to a young researcher for outstanding contributions to applied probability.