

Two-Sample Testing on Ranked Preference Data and the Role of Modeling Assumptions

Charvi Rastogi

*Machine Learning Department
Carnegie Mellon University
Pittsburgh, PA 15213, USA*

CRASTOGI@CS.CMU.EDU

Sivaraman Balakrishnan

*Department of Statistics
Carnegie Mellon University
Pittsburgh, PA 15213, USA*

SIVA@STAT.CMU.EDU

Nihar B. Shah

*Machine Learning Department and Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213, USA*

NIHARS@CS.CMU.EDU

Aarti Singh

*Machine Learning Department
Carnegie Mellon University
Pittsburgh, PA 15213, USA*

AARTI@CS.CMU.EDU

Editor: Shivani Agarwal, Francis Bach

Abstract

A number of applications require two-sample testing on ranked preference data. For instance, in crowdsourcing, there is a long-standing question of whether pairwise-comparison data provided by people is distributed identically to ratings-converted-to-comparisons. Other applications include sports data analysis and peer grading. In this paper, we design two-sample tests for pairwise-comparison data and ranking data. For our two-sample test for pairwise-comparison data, we establish an upper bound on the sample complexity required to correctly test whether the distributions of the two sets of samples are identical. Our test requires essentially no assumptions on the distributions. We then prove complementary lower bounds showing that our results are tight (in the minimax sense) up to constant factors. We investigate the role of modeling assumptions by proving lower bounds for a range of pairwise-comparison models (WST, MST, SST, parameter-based such as BTL and Thurstone). We also provide tests and associated sample complexity bounds for partial (or total) ranking data. Furthermore, we empirically evaluate our results via extensive simulations as well as three real-world data sets consisting of pairwise-comparisons and rankings. By applying our two-sample test on real-world pairwise-comparison data, we conclude that ratings and rankings provided by people are indeed distributed differently.

Keywords: Two-sample testing, pairwise-comparisons, ranking, minimax theory, social choice models

1. Introduction

Data in the form of pairwise-comparisons, or more generally partial or total rankings, arises in a wide variety of settings. For instance, when eliciting data from people (say, in crowdsourcing), there is a long-standing debate over the difference between two methods of data collection: asking people to compare pairs of items or asking people to provide numeric scores to the items. A natural question here is whether people implicitly generate pairwise-comparisons using a fundamentally different mechanism than first forming numeric scores and then converting them to a comparison. Thus, we are interested in testing if the data obtained from pairwise-comparisons is distributed identically to if the numeric scores were converted to pairwise-comparisons (Raman and Joachims, 2014; Shah et al., 2016). As another example consider sports and online games, where a match between two players or two teams is a pairwise-comparison between them (Herbrich et al., 2007; Hvattum and Arntzen, 2010; Van Der Maas and Wagenmakers, 2005). Again, a natural question that arises here is whether the relative performance of the teams has changed significantly across a certain period of time (e.g., to design an appropriate rating system (Cattelan et al., 2013)). A third example is peer grading where students are asked to compare pairs of homeworks (Shah et al., 2013) or rank a batch of homeworks (Lamon et al., 2016; Raman and Joachims, 2014). A question of interest here is whether a certain group of students (female/senior/...) grade very differently as compared to another group (male/junior/...) (Shah et al., 2018). Additionally, consumer preferences as pairwise-comparisons or partial (or total) rankings can be used to investigate whether a certain group (married/old/...) make significantly different choices about purchasing products as opposed to another group (single/young/...) (Cavagnaro and Davis-Stober, 2014; Regenwetter et al., 2011).

Each of the aforementioned problems involves two-sample testing, that is, testing whether the distribution of the data from two populations is identical or not. With this motivation, in this paper we consider the problem of two-sample testing on preference data in the form of pairwise-comparisons and, more generally, partial and total rankings. First, we focus our efforts on preference data in the form of pairwise-comparisons. Specifically, consider a collection of items (e.g., teams in a sports league). The data we consider comprises comparisons between pairs of these items, where the outcome of a comparison involves one of the items beating the other. In the two-sample testing problem, we have access to two sets of such pairwise-comparisons, obtained from two different sources (e.g., the current season in a sports league forming one set of pairwise-comparisons and the previous season forming a second set). The goal is to test whether the underlying distributions (winning probabilities) in the two sets of data are identical or different. Similarly, when the data comprises of partial or total rankings over a collection of items from two different sources, our goal is to test whether the distributions over total rankings for the two sources are identical or not. Specifically, we consider the case where a partial ranking is defined as a total ranking over some subset of the collection of items.

Contributions. We now outline the contributions of this paper; the theoretical contributions for the pairwise-comparison setting are also summarized in Table 1.

- First, we present a test for two-sample testing with pairwise-comparison data and associated upper bounds on its minimax sample complexity. Our test makes essentially no assumptions on the outcome probabilities of the pairwise-comparisons.

- Second, we prove information-theoretic lower bounds on the critical testing radius for this problem. Our bounds show that our test is minimax optimal for this problem.
- As a third contribution, we investigate the role of modeling assumptions: What if one could assume one of the popular models (e.g., BTL, Thurstone, parameter-based, SST, MST, WST) for pairwise-comparison outcomes? We show that our test is minimax optimal under WST and MST models. We also provide an information-theoretic lower bound under the SST and parameter-based models. Conditioned on the planted clique hardness conjecture, we prove a computational lower bound for the SST model with a single observation per pair of items, which matches the sample complexity upper bound attained by our test, up to logarithmic factors.
- Fourth, we conduct experiments on two real-world pairwise-comparison data sets. Our test detects a statistically significant difference between the distributions of directly-elicited pairwise-comparisons and converting numeric scores to comparison data. On the other hand, from the data available for four European football leagues over two seasons, our test does not detect any statistically significant difference between the relative performance of teams across two consecutive seasons.
- Finally, we present algorithms for two-sample testing on partial (or total) ranking data for two partial ranking models—namely, the Plackett-Luce model and a more general marginal probability based model. We provide upper bounds on sample complexity for the test for the Plackett-Luce model controlling both the Type I and Type II error. Moreover, our test for the marginal probability based model controls the Type I error. We apply our test to a real-world data set on sushi preferences. Our test finds a statistically significant difference in sushi preferences across sections of different demographics based on age, gender and region of residence.

A shorter version of this paper (Rastogi et al., 2020) was presented at the IEEE International Symposium on Information Theory (ISIT) 2020.

Related literature. The problem of two-sample testing on ranked preference data is at the intersection of two rich areas of research—two-sample testing and analyzing ranked preference data.

The problem of two-sample testing has a long history in statistics, and classical tests include the t-test and Pearson’s χ^2 test (see for instance Lehmann and Romano, 2005 and references therein). More recently, non-parametric tests (Gretton et al., 2012a,b; Rosenbaum, 2005; Kim et al., 2020b; Szekely and Rizzo, 2004) have gained popularity but these can perform poorly in structured, high-dimensional settings. The minimax perspective on hypothesis testing which we adopt in this work originates in the work of Ingster (1994) (and was developed further in Ingster (1997); Ingster and Suslina (2003); Ingster (1994)). Several recent works have studied the minimax rate for two-sample testing for high-dimensional multinomials (Balakrishnan and Wasserman, 2018, 2019; Chan et al., 2014; Valiant and Valiant, 2017; Valiant, 2011), and testing for sparsity in regression (Carpentier et al., 2018; Collier et al., 2017), we build on some of these ideas in our work. We also note the work of Mania et al. (2018) who propose a kernel-based two-sample test for distributions over total rankings.

Model (\mathcal{M})	Upper Bound	Lower Bound	Computational Lower Bound
Model-free	for $k > 1$, $\epsilon_{\mathcal{M}}^2 \leq c \frac{1}{kd}$ (Thm. 1)	$\epsilon_{\mathcal{M}}^2 > c \frac{\mathbb{I}(k > 1)}{kd} + \frac{\mathbb{I}(k = 1)}{4}$ (Prop. 4)	$\epsilon_{\mathcal{M}}^2 > c \frac{\mathbb{I}(k > 1)}{kd} + \frac{\mathbb{I}(k = 1)}{4}$
WST and MST	$\epsilon_{\mathcal{M}}^2 \leq c \frac{1}{kd}$	$\epsilon_{\mathcal{M}}^2 > c \frac{1}{kd}$ (Thm. 3)	$\epsilon_{\mathcal{M}}^2 > c \frac{1}{kd}$
SST	$\epsilon_{\mathcal{M}}^2 \leq c \frac{1}{kd}$	$\epsilon_{\mathcal{M}}^2 > c \frac{1}{kd^{3/2}}$	for $k = 1$, $\epsilon_{\mathcal{M}}^2 > \frac{1}{kd(\log \log(d))^2}$ (Thm. 6)
Parameter-based	$\epsilon_{\mathcal{M}}^2 \leq c \frac{1}{kd}$	$\epsilon_{\mathcal{M}}^2 > c \frac{1}{kd^{3/2}}$ (Thm. 5)	$\epsilon_{\mathcal{M}}^2 > c \frac{1}{kd^{3/2}}$

Table 1: This table summarizes our results for two-sample testing of pairwise-comparison data (introduced formally in Equation 1), for common pairwise-comparison models. Here, d denotes the number of items, and we obtain k samples (comparisons) per pair of items from each of the two populations. In this work, we provide upper and lower bounds on the critical testing radius $\epsilon_{\mathcal{M}}$, defined in (3). The upper bound in Theorem 1 is due to the test in Algorithm 1 which is computationally efficient. We note that the constant c varies from result to result.

The analysis of pairwise-comparison data is a rich field of study, dating back at least 90 years to the seminal work of Louis Thurstone (1927) and subsequently Bradley and Terry (1952) and Luce (1959). Along with this, Plackett (1975) and Luce (1959) worked on the now-well-known Plackett-Luce model for partial and total rankings. In the past two decades, motivated by crowdsourcing and other applications, there is significant interest in studying such data in a high-dimensional setting, that is, where the number of items d is not a fixed constant. A number of papers (Shah et al., 2016; Chen and Suh, 2015; Negahban et al., 2012; Rajkumar and Agarwal, 2014; Szörényi et al., 2015; Guiver and Snelson, 2009; Maystre and Grossglauser, 2015, and references therein) in this space analyze parameter-based models such as the BTL and the Thurstone models for pairwise-comparison data and the Plackett-Luce model for partial (or total) ranking data. Here the goal is usually to estimate the parameters of the model or the underlying ranking of the items. The papers (Ailon, 2012; Braverman and Mossel, 2008; Chatterjee and Mukherjee, 2019; Chen et al., 2018; Falahatgar et al., 2017; Rajkumar et al., 2015, and references therein) also study ranking from pairwise-comparisons, under some different assumptions.

Of particular interest is the paper by Aldous (2017) which uses the BTL model to make match predictions in sports, and also poses the question of analyzing the average change in the performance of teams over time. While this paper suggests some simple statistics to test for change, designing principled tests is left as an open problem. To this end, we provide a two-sample test without any assumptions and with rigorous guarantees, and also use it subsequently to conduct such a test on real-world data.

A recent line of work (Heckel et al., 2019; Shah et al., 2017; Shah and Wainwright, 2018) focuses on the role of the modeling assumptions in estimation and ranking from pairwise-

comparisons. We study the role of modeling assumptions from the perspective of two-sample testing and prove performance guarantees for some pairwise-comparison models.

Organization. The remainder of this paper is organized as follows. In Section 2, we formally describe the problem setup and provide some background on the minimax perspective on hypothesis testing. We also provide a detailed description of the pairwise-comparison models studied in this work. In Section 3 we present our minimax optimal test for pairwise-comparison data and present the body of our main technical results for the pairwise-comparison setting with brief proof sketches and defer technical aspects of the proofs to Section 6. Then, in Section 4 we extend our results for the two-sample testing problem on partial (or total) ranking data. We describe two partial ranking models and provide testing algorithms and associated sample complexity bounds. The corresponding proofs are in Section 6. In Section 5, we present our findings from implementing our testing algorithms on three real-world data sets. Furthermore, we present results of simulations on synthetic data which validate our theoretical findings. We conclude with a discussion in Section 7.

2. Background and problem formulation for pairwise-comparison setting

In this section, we provide a more formal statement of the problem of two-sample testing using pairwise-comparison data along with background on hypothesis testing and the associated definition of risk, and various types of ranking models.

2.1 Problem statement

Our focus in this paper is on the two-sample testing problem where the two sets of samples come from two potentially different populations. Here, we describe the model of the data we consider in our work. Specifically, consider a collection of d items. The two sets of samples comprise outcomes of comparisons between various pairs of these items. In the first set of samples, the outcomes are governed by an unknown matrix $P \in [0, 1]^{d \times d}$. The $(i, j)^{\text{th}}$ entry of matrix P is denoted as p_{ij} , and any comparison between items i and j results in i beating j with probability p_{ij} , independent of other outcomes. We assume there are no ties. Analogously, the second set of samples comprises outcomes of pairwise-comparisons between the d items governed by a (possibly different) unknown matrix $Q \in [0, 1]^{d \times d}$, wherein item i beats item j with probability q_{ij} , the $(i, j)^{\text{th}}$ entry of matrix Q . For any pair (i, j) of items, we let k_{ij}^p and k_{ij}^q denote the number of times a pair of items (i, j) is compared in the first and second set of samples respectively. Let X_{ij} denote the number of times item $i \in [d]$ beats item $j \in [d]$ in the first set of samples, and let Y_{ij} denote the analogous quantity in the second set of samples. It follows that X_{ij} and Y_{ij} are Binomial random variables independently distributed as $X_{ij} \sim \text{Bin}(k_{ij}^p, p_{ij})$ and $Y_{ij} \sim \text{Bin}(k_{ij}^q, q_{ij})$. We adopt the convention of setting $X_{ij} = 0$ when $k_{ij}^p = 0$, and $Y_{ij} = 0$ when $k_{ij}^q = 0$, and $k_{ii}^p = k_{ii}^q = 0$.

Our results apply to both the symmetric and asymmetric settings of pairwise-comparisons:

Symmetric setting: The literature on the analysis of pairwise-comparison data frequently considers a symmetric setting where “ i vs. j ” and “ j vs. i ” have an identical meaning. Our results apply to this setting, for which we impose the additional constraints that $p_{ji} = 1 - p_{ij}$

and $q_{ji} = 1 - q_{ij}$ for all $(i, j) \in [d]^2$. In addition, for every $1 \leq i \leq j \leq d$, we set $k_{ji}^p = k_{ji}^q = 0$ (and hence $X_{ji} = Y_{ji} = 0$), and let $k_{ij}^p, k_{ij}^q, X_{ij}$ and Y_{ij} represent the comparisons between the pair of items (i, j) .

Asymmetric setting: Our results also apply to an asymmetric setting where “ i vs. j ” may have a different meaning as compared to “ j vs. i ”. For instance, in a setting of sports where “ i vs. j ” could indicate i as the home team and j as the visiting team. This setting does not impose the restrictions described in the symmetric setting above.

Hypothesis test. Consider any class \mathcal{M} of pairwise-comparison probability matrices, and any parameter $\epsilon > 0$. Then, the goal is to test the hypotheses

$$\begin{aligned} H_0 &: P = Q \\ H_1 &: \frac{1}{d} \|P - Q\|_{\mathbb{F}} \geq \epsilon, \end{aligned} \tag{1}$$

where $P, Q \in \mathcal{M}$.

2.2 Hypothesis testing and risk

We now provide a brief background on hypothesis tests and associated terminology. In hypothesis testing, the Type I error is defined as the probability of rejecting the null hypothesis H_0 when the null hypothesis H_0 is actually true, an upper bound on the Type I error is denoted by α ; the Type II error is defined as the probability of failing to reject the null when the alternate hypothesis H_1 is actually true, an upper bound on Type II error is denoted by β . The performance of the testing algorithm is evaluated by measuring its Type I error and its power, which is defined as one minus the Type II error.

Consider the hypothesis testing problem defined in (1). We define a test ϕ as $\phi : \{k_{ij}^p, k_{ij}^q, X_{ij}, Y_{ij}\}_{(i,j) \in [d]^2} \mapsto \{0, 1\}$. Let \mathbb{P}_0 and \mathbb{P}_1 denote the distribution of the input variables under the null and under the alternate respectively. Here, we assume that the variables k_{ij}^p and k_{ij}^q are fixed for all $(i, j) \in [d]^2$. Let \mathcal{M}_0 and $\mathcal{M}_1(\epsilon)$ denote the set of matrix pairs (P, Q) that satisfy the null condition and the alternate condition in (1) respectively. Then, we define the minimax risk (Ingster, 1994, 1997; Ingster and Suslina, 2003) as

$$\mathcal{R}_{\mathcal{M}} = \inf_{\phi} \left\{ \sup_{(P,Q) \in \mathcal{M}_0} \mathbb{P}_0(\phi = 1) + \sup_{(P,Q) \in \mathcal{M}_1(\epsilon)} \mathbb{P}_1(\phi = 0) \right\}, \tag{2}$$

where the infimum is over all $\{0, 1\}$ -valued tests ϕ . It is common to study the minimax risk via a coarse lens by studying instead the critical radius or the minimax separation. The critical radius is the smallest value ϵ for which a hypothesis test has non-trivial power to distinguish the null from the alternate. Formally, we define the critical radius as

$$\epsilon_{\mathcal{M}} = \inf\{\epsilon : \mathcal{R}_{\mathcal{M}} \leq 1/3\}. \tag{3}$$

The constant $1/3$ is arbitrary; we could use any specified constant in $(0, 1)$. In this paper, we focus on providing tight bounds on the critical radius.

2.3 A range of pairwise-comparison models

A model for the pairwise-comparison probabilities is a set of matrices in $[0, 1]^{d \times d}$. In the context of our problem setting, assuming a model means that the matrices P and Q are guaranteed to be drawn from this set. In this paper, the proposed test and the associated guarantees do not make any assumptions on the pairwise-comparison probability matrices P and Q . In other words, we allow P and Q to be any arbitrary matrices in $[0, 1]^{d \times d}$. However, there are a number of models which are popular in the literature on pairwise-comparisons, and we provide a brief overview of them here. We analyze the role of these modeling assumptions in our two-sample testing problem. In what follows, we let $M \in [0, 1]^{d \times d}$ denote a generic pairwise-comparison probability matrix, with M_{ij} representing the probability that item $i \in [d]$ beats item $j \in [d]$. The models impose conditions on the matrix M .

- *Parameter-based models:* A parameter-based model is associated with some known, non-decreasing function $f : \mathbb{R} \rightarrow [0, 1]$ such that $f(\theta) = 1 - f(-\theta) \quad \forall \theta \in \mathbb{R}$. We refer to any such function f as being “valid”. The parameter-based model associated to a given valid function f is given by

$$M_{ij} = f(w_i - w_j) \quad \text{for all pairs } (i, j), \quad (4)$$

for some unknown vector $w \in \mathbb{R}^d$ that represents the notional qualities of the d items. It is typically assumed that the vector w satisfies the conditions $\sum_{i \in [d]} w_i = 0$ and that $\|w\|_\infty$ is bounded above by a known constant.

- *Bradley-Terry-Luce (BTL) model:* This is a specific parameter-based model with $f(\theta) = \frac{1}{1 + e^{-\theta}}$.
- *Thurstone model:* This is a specific parameter-based model with $f(\theta) = \Phi(\theta)$, where Φ is the standard Gaussian CDF.
- *Strong stochastic transitivity (SST):* The model assumes that the set of items $[d]$ is endowed with an unknown total ordering π , where $\pi(i) < \pi(j)$ implies that item i is preferred to item j . A matrix $M \in [0, 1]^{d \times d}$ is said to follow the SST model if it satisfies the shifted-skew-symmetry condition $M_{ij} = 1 - M_{ji}$ for every pair $i, j \in [d]$ and the condition

$$M_{i\ell} \geq M_{j\ell} \quad \text{for every } i, j \in [d] \text{ such that } \pi(i) < \pi(j) \text{ and for every } \ell \in [d]. \quad (5)$$

- *Moderate stochastic transitivity (MST):* The model assumes that the set of items $[d]$ is endowed with an unknown total ordering π . A matrix $M \in [0, 1]^{d \times d}$ is said to follow the MST model if it satisfies $M_{ij} = 1 - M_{ji}$ for every pair $i, j \in [d]$ and the condition

$$M_{i\ell} \geq \min\{M_{ij}, M_{j\ell}\} \quad \text{for every } i, j, \ell \in [d] \text{ such that } \pi(i) < \pi(j) < \pi(\ell). \quad (6)$$

- *Weak stochastic transitivity (WST):* The model assumes that the set of items $[d]$ is endowed with an unknown total ordering π . A matrix $M \in [0, 1]^{d \times d}$ is said to follow the WST model if it satisfies $M_{ij} = 1 - M_{ji}$ for every pair $i, j \in [d]$ and the condition

$$M_{ij} \geq \frac{1}{2} \quad \text{for every } i, j \in [d] \text{ such that } \pi(i) < \pi(j). \quad (7)$$

Model hierarchy: There is a hierarchy between these models, that is, {BTL, Thurstone} \subset parameter-based \subset SST \subset MST \subset WST \subset model-free.

3. Main results for pairwise-comparison setting

We now present our main theoretical results for pairwise-comparison data.

3.1 Test and guarantees

Our first result provides an algorithm for two-sample testing in the problem (1), and associated upper bounds on its sample complexity. Importantly, we do not make any modeling assumptions on the probability matrices P and Q . First we consider a per-pair fixed-design setup in Theorem 1 where for every pair of items (i, j) , the sample sizes k_{ij}^p, k_{ij}^q are equal to k . Following that, in Corollary 2, we consider a random-design setup wherein for every pair of items (i, j) , the sample sizes k_{ij}^p, k_{ij}^q are drawn i.i.d. from some distribution \mathcal{D} supported over non-negative integers.

Input: Samples X_{ij}, Y_{ij} denoting the number of times item i beat item j in the observed k_{ij}^p, k_{ij}^q pairwise-comparisons from populations denoted by probability matrices P, Q respectively.

Test Statistic:

$$T = \sum_{i=1}^d \sum_{j=1}^d \mathbb{I}_{ij} \frac{k_{ij}^q (k_{ij}^q - 1) (X_{ij}^2 - X_{ij}) + k_{ij}^p (k_{ij}^p - 1) (Y_{ij}^2 - Y_{ij}) - 2(k_{ij}^p - 1)(k_{ij}^q - 1) X_{ij} Y_{ij}}{(k_{ij}^p - 1)(k_{ij}^q - 1)(k_{ij}^p + k_{ij}^q)} \quad (8)$$

where $\mathbb{I}_{ij} = \mathbb{I}(k_{ij}^p > 1) \times \mathbb{I}(k_{ij}^q > 1)$.

Output: If $T \geq 11d$, where $11d$ is the threshold, then reject the null.

Algorithm 1: Two-sample test with pairwise-comparisons for model-free setting

Our test is presented in Algorithm 1. The test statistic (8) is designed such that it has an expected value of zero under the null and a large expected value under the alternate. That is, the test statistic (8) is designed symmetrically across P and Q such that it has expected value zero, if $P = Q$. Similarly, to reject the null with high probability, under the alternate $\|P - Q\|_F \geq \epsilon d$, the test statistic (8) is designed to increase as $\|P - Q\|_F$ increases. In fact, we will later show in Section 6.2 that the test statistic increases quadratically with $\|P - Q\|_F$. The following theorem characterizes the performance of this test, thereby establishing an upper bound on the sample complexity of this two-sample testing problem in a random-design setting.

Theorem 1 *Consider the testing problem in (1) with \mathcal{M} as the class of all pairwise probability matrices. Suppose the number of (per pair) comparisons between the two populations is fixed, $k_{ij}^p = k_{ij}^q = k$ (for all $i \neq j$ in the asymmetric setting and all $i < j$ in the symmetric setting).*

There is a constant $c > 0$ such that for any $\epsilon > 0$, if $k > 1$ and $\epsilon^2 \geq c \frac{1}{kd}$, then the sum of Type I error and Type II error of Algorithm 1 is at most $\frac{1}{3}$.

The proof is provided in Section 6.2. Theorem 1 provides a guarantee of correctly distinguishing between the null and the alternate with probability at least $\frac{2}{3}$. The value $\frac{2}{3}$ is closely tied to the specific threshold used in the test above. More generally, for any specified constant $\nu \in (0, 1)$, the test achieves a Type I error at most ν by setting the threshold as $d\sqrt{\frac{24(1-\nu)}{\nu}}$. Similarly, for any specified constant $\nu \in (0, 1)$, the test achieves a Type II error at most ν , if $\epsilon^2 \geq \frac{\nu_1}{kd}$, wherein $\nu_1 > 0$ is a constant that depends on ν . The power of our test approaches 1 at the rate $\frac{c}{kd\epsilon^2}$, where c is some constant, that is the power of our test increases as ϵ increases.

Moreover, if the sample complexity is increased by some factor R , then running Algorithm 1 on R independent instances of the data and taking the majority answer results in error probability that decreases exponentially with R as $(\exp(-2R))$, while the sample complexity increases only linearly in R . One can thus have a very small probability of error of, for instance, d^{-50} with $k = \tilde{O}\left(\frac{1}{d\epsilon^2}\right)$. Later, in Proposition 4, we show that under the fixed k condition, $k_{ij}^p = k_{ij}^q = k$, we have that $k > 1$ is necessary for our two-sample testing problem. It is also interesting to note that the estimation rate to test the hypotheses in (1) is $k = O\left(\frac{\log(d)}{\epsilon^2}\right)$ while the rate for our testing algorithm is $k = O\left(\frac{1}{d\epsilon^2}\right)$.

Now, we consider the random-design setup wherein for every pair of items (i, j) , the sample sizes k_{ij}^p, k_{ij}^q are drawn i.i.d. from some distribution \mathcal{D} supported over non-negative integers. Let μ and σ denote the mean and standard deviation of distribution \mathcal{D} respectively, and let $p_1 := \Pr_{Z \sim \mathcal{D}}(Z = 1)$. We assume that \mathcal{D} has a finite mean and that

$$\mu \geq c_1 p_1; \quad \mu \geq c_2 \sigma, \tag{9}$$

for some constants $c_1 > 1$ and $c_2 > 1$. Many commonly occurring distributions obey these properties, for instance, Binomial distribution, Poisson distribution, geometric distribution and discrete uniform distribution, with appropriately chosen parameters.

Corollary 2 *Consider the testing problem in (1) with \mathcal{M} as the class of all pairwise probability matrices. Suppose the number of comparisons in the two populations k_{ij}^p, k_{ij}^q are drawn i.i.d. from some distribution \mathcal{D} that satisfies (9) (for all $i \neq j$ in the asymmetric setting and all $i < j$ in the symmetric setting). There is a constant $c > 0$ such that if $\epsilon^2 \geq c \max\{\frac{1}{\mu d}, \frac{1}{d^2}\}$, then the sum of Type I error and Type II error of Algorithm 1 is at most $\frac{1}{3}$.*

The proof of Corollary 2 is in Section 6.1. In Corollary 2, we see that the even under the random-design setup, our test achieves the same testing rate as in the per-pair fixed-design setup considered in Theorem 1, for $\mu \leq d$.

We now evaluate the performance of Algorithm 1 when k_{ij}^p, k_{ij}^q are drawn i.i.d. from one of the following commonly occurring distributions. Consider any arbitrary matrices P and Q . We specialise Corollary 2 to these distributions by stating the sample complexity that

guarantees that the probability of error is at most $\frac{1}{3}$ in the two-sample testing problem (1), wherein constant c may depend on c_1, c_2 for each distribution. Note that, as in Corollary 2, we assume $\epsilon^2 d^2 \geq c'$ where c' is some positive constant

- Binomial distribution $(k_{ij}^p, k_{ij}^q \stackrel{\text{iid}}{\sim} \text{Bin}(n, a))$: Sufficient condition $n \geq c \max\{\frac{1}{ad\epsilon^2}, \frac{1}{a}\}$.
- Poisson distribution $(k_{ij}^p, k_{ij}^q \stackrel{\text{iid}}{\sim} \text{Poisson}(\lambda))$: Sufficient condition $\lambda \geq c \max\{\frac{1}{d\epsilon^2}, 1\}$.
- Geometric Distribution $(k_{ij}^p, k_{ij}^q \stackrel{\text{iid}}{\sim} \text{Geometric}(a))$: Sufficient condition $\frac{1}{a} \geq c \max\{\frac{1}{d\epsilon^2}, 1\}$.
- Discrete Uniform Distribution $(k_{ij}^p, k_{ij}^q \stackrel{\text{iid}}{\sim} \text{Unif}(0, n))$: Sufficient condition $n \geq c \max\{\frac{1}{d\epsilon^2}, 1\}$.

Next, we note that a sharper but non-explicit threshold in Algorithm 1 can be obtained using the permutation test method to control the Type I error. We detail this approach in Algorithm 2.

Input : Samples X_{ij}, Y_{ij} denoting the number of times item i beat item j in the observed k_{ij}^p, k_{ij}^q pairwise-comparisons from populations denoted by probability matrices P, Q respectively. Significance level $\alpha \in (0, 1)$. Iteration count γ .

(1) Compute the test statistic T defined in (8).

(2) For $\ell \leftarrow 1$ to γ :

(i) Repeat this step independently for all $i \neq j$ in the asymmetric setting and for all $i < j$ in the symmetric setting. Collect the $(k_{ij}^p + k_{ij}^q)$ samples together and reassign k_{ij}^p of the samples chosen uniformly at random to P and the rest to Q . Compute the new values of X_{ij} and Y_{ij} based on this reassignment.

(ii) Using the new values of X_{ij} and Y_{ij} , recompute the test statistic in (8). Denote the computed test statistic as T_ℓ .

Output : Reject the null if $p = \sum_{\ell=1}^{\gamma} \frac{1}{\gamma} \mathbb{1}(T_\ell - T) < \alpha$.

Algorithm 2: Permutation test with pairwise-comparisons for model-free setting.

More generally, the results in Theorem 1 and Corollary 2 (and the following converse results in Theorem 3 and Proposition 4) also apply to the two-sample testing problem of comparing two Bernoulli matrices (or vectors) P and Q , wherein each entry of the matrices (or vectors) is a Bernoulli parameter. In this problem, we want to test whether two Bernoulli matrices are identical or not, and we have access to some observations of some (or all) of the underlying Bernoulli random variables.

We conclude this section with a proof sketch for Theorem 1; the complete proof is provided in Section 6.1 and 6.2.

Proof Sketch for Theorem 1. The test statistic T is designed to ensure that $\mathbb{E}_{H_0}[T] = 0$ for any arbitrary pairwise probability matrices P, Q such that $P = Q$, and for any values of $\{k_{ij}^p, k_{ij}^q\}_{1 \leq i, j \leq d}$. We lower bound the expected value of T under the alternate hypothesis as $\mathbb{E}_{H_1}[T] \geq ckd^2\epsilon^2$ (Lemma 10). Next, we show that the variance of T is upper bounded under the null by $24d^2$ and under the alternate by $24d^2 + 4kd^2\epsilon^2$ (Lemma 11). These lemmas

allow us to choose a suitable threshold value of $11d$. Finally, using Chebyshev's inequality comparing the square of expectation with the variance, we obtain the desired upper bound on the sample complexity with guarantees on both Type I and Type II errors.

3.2 Converse results and the role of modeling assumptions

In this section we look at the role of modeling assumptions on the pairwise-comparison probability matrices in the two-sample testing problem in (1).

Lower bound for MST, WST, and model-free classes. Having established an upper bound on the rate of two-sample testing without modeling assumptions on the pairwise-comparison probability matrices P, Q , we show matching lower bounds that hold under the MST class. The WST and model-free classes are both supersets of MST, and hence the following guarantees automatically apply to them as well.

Theorem 3 *Consider the testing problem in (1) with \mathcal{M} as the class of matrices described by the MST model. Suppose we have k comparisons for each pair (i, j) from each population. There exists a constant $c > 0$, such that the critical radius $\epsilon_{\mathcal{M}}$ is lower bounded as $\epsilon_{\mathcal{M}}^2 > \frac{c}{kd}$.*

The lower bound on the rate matches the rate derived for Algorithm 1 in Theorem 1, thereby establishing the minimax optimality of our algorithm (up to constant factors). The MST class is a subset of the WST model class. This proves that Algorithm 1 is simultaneously minimax optimal under the MST and WST modeling assumptions in addition to the model-free setting. We provide a proof sketch for Theorem 3 in Section 3.2.1; the complete proof is in Section 6.3.3.

Necessity of $\mu > p_1$. Recall that the upper bound derived in Theorem 1 under the model-free setting holds under the assumption that $k > 1$ and, similarly, Corollary 2 holds under the assumption that $\mu \geq c_1 p_1$ with $c_1 > 1$, as stated in (9). We now state a negative result for the case $\mu \leq p_1$, which implies that $k_{ij}^p, k_{ij}^q \leq 1 \forall (i, j)$ under the random-design setup and $k \leq 1$ under the per-pair fixed-design setup.

Proposition 4 *Consider the testing problem in (1) with \mathcal{M} as the class of all pairwise probability matrices. Suppose we have at most one comparison for each pair (i, j) from each population (for all $i \neq j$ in the asymmetric setting and all $i < j$ in the symmetric setting). Then, for any value of $\epsilon \leq \frac{1}{2}$, the minimax risk defined in (2) is at least $\frac{1}{2}$, thus, $\epsilon_{\mathcal{M}}^2 \geq \frac{1}{4}$.*

We provide some intuition for this result here. If $k_{ij}^p = k_{ij}^q \leq 1 \forall (i, j)$, then at best one has access to first order information of each entry of P and Q , that is, one has access to only $\Pr(X_{ij} = 1), \Pr(Y_{ij} = 1), \Pr(X_{ij} = 1, Y_{ij} = 1)$ for each pair (i, j) . This observation allows us to construct a case wherein the null and the alternate cannot be distinguished from each other by any test, due to the inaccessibility of higher order information of the underlying Bernoulli random variables. The complete proof is provided in Section 6.3.2.

Lower bound for parameter-based class. We now prove an information-theoretic lower bound for our two-sample testing problem wherein the probability matrices follow the parameter-based model.

Theorem 5 Consider the testing problem in (1). Consider any arbitrary non-decreasing function $f : \mathbb{R} \rightarrow [0, 1]$ such that $f(\theta) = 1 - f(-\theta) \quad \forall \theta \in \mathbb{R}$, with \mathcal{M} as the parameter-based class of probability matrices associated to the given function. Suppose we have k comparisons for each pair (i, j) from each population. There exists a constant $c > 0$, such that the critical radius $\epsilon_{\mathcal{M}}$ is lower bounded as $\epsilon_{\mathcal{M}}^2 > \frac{c}{kd^{3/2}}$.

This lower bound also applies to probability matrices in the SST class described in (5). We provide a brief proof sketch in Section 3.2.1; the complete proof is in Section 6.3.4. Observe that the lower bound on testing rate obtained in Theorem 5 does not match the testing rate obtained in Theorem 1. In the next part of this section, we provide a computational lower bound in place of a statistical lower bound to bridge this gap.

Computational lower bound for SST class. Given the polynomial gap between Theorem 1 and Theorem 5, it is natural to wonder whether there is another polynomial-time testing algorithm for testing under the SST and/or parameter-based modeling assumption. We answer this question in the negative, for the SST model and single observation model ($k = 1$), conditionally on the average-case hardness of the planted clique problem (Jerrum, 1992; Kučera, 1995). In informal terms, the planted clique conjecture asserts that there is no polynomial-time algorithm that can detect the presence of a planted clique of size $\kappa = o(\sqrt{d})$ in an Erdős-Rényi random graph with d nodes. We construct SST matrices that are similar to matrices in the planted clique problem and as a direct consequence of the planted clique conjecture, we have the following result.

Theorem 6 Consider the testing problem in (1) with \mathcal{M} as the class of matrices described by the SST model. Suppose the planted clique conjecture holds. Suppose we have one comparison for each pair (i, j) from each population. Then there exists a constant $c > 0$ such that for polynomial-time testing algorithms the critical radius $\epsilon_{\mathcal{M}}$ is lower bounded as $\epsilon_{\mathcal{M}}^2 > \frac{c}{d(\log \log(d))^2}$.

Thus, for $k = 1$, the computational lower bound on the testing rate for the SST model matches the rate derived for Algorithm 1 (up to logarithmic factors). The proof of Theorem 6 is provided in Section 6.3.5. We devote the rest of this section to a sketch of the proofs of Theorem 3 and Theorem 5.

3.2.1 PROOF SKETCHES FOR THEOREM 3 AND THEOREM 5

To prove the information-theoretic lower bound under the different modeling assumptions, we construct a null and alternate belonging to the corresponding class of probability matrices. The bulk of our technical effort is devoted to upper bounding the chi-square divergence between the probability measure under the null and the alternate. We then invoke Le Cam's lower bound for testing to obtain a lower bound on the minimax risk which gives us the information-theoretic lower bound. We now look at the constructions for the two modeling assumptions.

Lower bound construction for MST class (Section 6.3.3). We construct a null and alternate such that under the null $P = Q = [\frac{1}{2}]^{d \times d}$ and under the alternate $P = [\frac{1}{2}]^{d \times d}$ and $Q \in \Theta$ with $\frac{1}{d} \|P - Q\|_F = \epsilon$. For this, we define a parameter $\eta \in [0, \frac{1}{2}]$ and then define Θ as a set of

matrices in which the upper right quadrant has exactly one entry equal to $\frac{1}{2} + \eta$ in each row and each column and the remaining entries above the diagonal are $\frac{1}{2}$. The entries below the diagonal follow from the shifted-skew-symmetry condition. We consider the alternate where Q is chosen uniformly at random from the set Θ of probability matrices in MST class.

Lower bound construction for parameter-based class (Section 6.3.4). The construction is same as the construction given above except we define a different set Θ of probability matrices. According to the parameter-based model, the matrices P and Q depend on the vectors $w_p \in \mathbb{R}^d$ and $w_q \in \mathbb{R}^d$ respectively. Now, for simplicity in this sketch, suppose that d is even. We set $w_p = [0, \dots, 0]$, which fixes $p_{ij} = \frac{1}{2} \forall (i, j)$. Consider a collection of vectors each with half the entries as δ and the other half as $-\delta$, thereby ensuring that $\sum_{i \in [d]} w_i = 0$. We set δ to ensure that each of the probability matrices induced by this collection of vectors obey $\frac{1}{d} \|P - Q\|_F = \epsilon$. We then consider the setting where Q is chosen uniformly at random from the set of pairwise-comparison probability matrices induced by the collection of values of w_Q .

4. Two-sample testing with partial or total ranking data

In this section, we extend our work from the previous sections to two-sample testing for ranking data. We focus on the two-sample testing problem where the two sets of samples from two potentially different populations comprise of partial or total rankings over various subsets of d items. Specifically, we consider the case where a partial ranking is defined as a total ranking over some subset of d items. Let λ_P and λ_Q be two unknown probability distributions over the set of all d -length rankings. We observe two sets of partial or total rankings, one set from each of two populations. The partial rankings in the first set are assumed to be drawn i.i.d. according to λ_P , and the partial rankings in second set are drawn i.i.d. according to λ_Q . Each sample obtained is a ranking over a subset of items of size ranging from 2 to d . Henceforth, we use the term total ranking to specify a ranking over all d items. We assume there are no ties.

Hypothesis test Our goal is to test the hypothesis,

$$\begin{aligned} H_0 : \lambda_P &= \lambda_Q \\ H_1 : \lambda_P &\neq \lambda_Q. \end{aligned} \tag{10}$$

In the sequel, we consider this hypothesis testing problem under certain modeling assumptions on λ_P and λ_Q .

4.1 Models

We now describe two partial ranking models that we analyse subsequently.

Marginal probability based model This is a non-parametric partial ranking model that is entirely specified by the probability distribution over all total rankings, given by λ_P in the first population and λ_Q in the second population. The distribution λ defines the partial ranking model for the corresponding population as follows. Let S_d denote the set of all total rankings over the d items. Consider some subset of items $\Omega \subseteq [d]$ of size $m \in \{2, \dots, d\}$, and let τ_Ω be a ranking of the items in this set. Then, we define a set of all total rankings that

obey the partial ranking τ_Ω as

$$S(\tau_\Omega) = \{\tau \in S_d : \tau(\tau_\Omega^{-1}(1)) < \tau(\tau_\Omega^{-1}(2)) < \dots < \tau(\tau_\Omega^{-1}(m))\}. \quad (11)$$

The marginal probability based partial ranking model gives the probability of a partial ranking τ_Ω as

$$\mathbb{P}(\tau_\Omega) = \sum_{\tau \in S(\tau_\Omega)} \lambda(\tau), \quad (12)$$

where λ represents λ_P or λ_Q for the corresponding population. This model defines the probability of a partial ranking similarly to the non-parametric choice model described in Farias et al. (2013). In fact, their choice model defined over sets of size 2 is the same as our model over partial rankings of size 2. Our model has the desired property that given a partial ranking over the set Ω containing item i and item j , the marginal probability that item i is ranked higher than item j , denoted by $\mathbb{P}(i \succ j)$, does not depend on other items in set Ω . Subsequently, the marginal probability is expressed as

$$\begin{aligned} \mathbb{P}(i \succ j | \Omega) &= \sum_{\tau \in S(\tau_\Omega), \tau(i) < \tau(j)} \lambda(\tau) \\ &= \sum_{\tau \in S_d, \tau(i) < \tau(j)} \lambda(\tau) \end{aligned}$$

Now, for the two populations we define the marginal probability of pairwise-comparisons over items (i, j) as p_{ij} and q_{ij} for all pairs (i, j) with $i < j$. Note that this model has the property that $p_{ij} = 1 - p_{ji}$ and $q_{ij} = 1 - q_{ji}$ for all (i, j) . We also note that the Plackett-Luce model described next is a subset of this model.

Plackett-Luce model This model introduced by Luce (1959) and Plackett (1975) is a commonly used parametric model for ranking data. In this model, each item has a notional quality parameter $w_i \in \mathbb{R}$, $\forall i \in [d]$. Under the Plackett-Luce model, the partial rankings in each population are generated according to the corresponding underlying quality parameters, namely $w_{i \in [d]}^P$ and $w_{i \in [d]}^Q$. The weight parameters completely define the probability distribution λ over the set of all total rankings. In this model, a partial (or total) ranking τ is generated in a sequential manner where each item in a ranking is viewed as chosen from the set of items ranked lower. The probability of choosing an item i from any set $S \subseteq [d]$ is given by $\frac{\exp(w_i)}{\sum_{i' \in S} \exp(w_{i'})}$. To explain the sequential generation procedure, we show an example here,

$$\mathbb{P}(i_1 \succ i_2 \succ \dots \succ i_\ell) = \prod_{j=1}^{\ell} \frac{\exp(w_{i_j})}{\sum_{j'=j}^{\ell} \exp(w_{i_{j'}})}.$$

An important property of the Plackett-Luce model is that the marginal probability that item i is ranked higher than item j , $\mathbb{P}(i \succ j)$ does not depend on the other items in the ranking, in fact, $\mathbb{P}(i \succ j) = \frac{\exp(w_i)}{(\exp(w_i) + \exp(w_j))}$. For each pair (i, j) , we denote the marginal probability $\mathbb{P}(i \succ j)$ corresponding to the parameters $w_{i \in [d]}^P$ as p_{ij} . Similarly, we denote the marginal probability $\mathbb{P}(i \succ j)$ corresponding to the parameters $w_{i \in [d]}^Q$ as q_{ij} . These pairwise marginal

probabilities p_{ij} and q_{ij} are collected in pairwise-comparison probability matrices P and Q respectively.

Finally, with this notation, we specialise the hypothesis testing problem in (10) for the two partial ranking models described above, in terms of the pairwise probability matrices P and Q . For any given parameter $\epsilon > 0$, we define the two-sample testing problem as

$$\begin{aligned} H_0 &: P = Q \\ H_1 &: \frac{1}{d} \|P - Q\|_{\mathbb{F}} \geq \epsilon. \end{aligned} \tag{13}$$

We note that under the Plackett-Luce model, the null condition in (10) is equivalent to the null condition in (13). Moreover, under the Plackett-Luce model, difference in two probability distributions λ_P and λ_Q is captured by difference in the pairwise probability matrices P and Q . Thus, we specialise the alternate condition in (10) in terms of scaled Frobenius distance between the pairwise probability matrices, P and Q , denoted by the parameter ϵ , to get the alternate condition in (13). Furthermore, under the marginal probability based model, the null condition in (10) implies $P = Q$ whereas the converse is not true. That is, there exist pairs of probability distributions over the set of all d -length rankings λ_P and λ_Q , that follow the marginal probability based model with $\lambda_P \neq \lambda_Q$, such that their corresponding pairwise probability matrices P and Q are equal. Thus, under the marginal probability based model, by conducting a test for the hypothesis testing problem in (13) that controls the Type I error at level α , we get control over Type I error at level α for the hypothesis testing problem in (10).

We are now ready to describe our main results for two-sample testing with partial (or total) ranking data.

4.2 Main results

Our testing algorithms for ranking data build upon the test statistic in Algorithm 1. To test for difference in probability distributions λ_P and λ_Q , we first use a rank breaking method to convert the data into pairwise-comparisons, on which we apply the test statistic in (8). Given a rank breaking method, denoted by R , and rankings from the two populations, S_{P_i} and S_{Q_i} for $i \in [N]$, then the rank breaking algorithm yields pairwise-comparison data as $R(S_{P_{i \in [N]}}) = \{k_{ij}^p, X_{ij}\}_{(i,j) \in [d]^2}$ and, similarly, $R(S_{Q_{i \in [N]}}) = \{k_{ij}^q, Y_{ij}\}_{(i,j) \in [d]^2}$. Here, $k_{ij}^p, k_{ij}^q, X_{ij}, Y_{ij}$ represent the pairwise-comparison data as defined in Section 2.1. Now, we describe three rank breaking methods that we subsequently use in our testing algorithms, Algorithm 3 and Algorithm 4.

1. **Random disjoint:** In this method, denoted by $R_{\mathbb{R}}$, given a set of N partial (or total) rankings, we randomly break each ranking up into pairwise-comparisons such that no item is in more than one pair. In this method, each m -length ranking yields $\lfloor \frac{m}{2} \rfloor$ pairwise-comparisons.
2. **Deterministic disjoint:** We use this rank breaking method, denoted by $R_{\mathbb{D}}$, when we have N total rankings. In this method, we deterministically break each ranking into pairwise-comparisons so that no item is in more than one pair. So, we get $\lfloor \frac{d}{2} \rfloor$ pairwise-comparisons from each total ranking. First, we want the number of samples

to be divisible by d , so we throw away $(N \bmod d)$ rankings chosen randomly. Then arbitrarily without looking at the data, partition the remaining rankings into $\lfloor \frac{N}{d} \rfloor$ disjoint subsets each containing d rankings. Within each subset, we convert the d rankings into $d \lfloor \frac{d}{2} \rfloor$ pairwise-comparisons deterministically such that we observe at least one comparison between each pair $(i, j) \in [d]$ with $i < j$. We keep exactly one pairwise-comparison for each pair in a subset. In this manner, we get to observe exactly $\lfloor \frac{N}{d} \rfloor$ comparisons between each pair of items.

3. **Complete:** In this method, denoted by R_C , given a set of N partial (or total) rankings, we break each ranking into all possible pairwise-comparisons for that ranking. In this method, each m -length ranking yields $\binom{m}{2}$ pairwise-comparisons.

Now, equipped with the rank breaking methods, we describe our first result which provides an algorithm (Algorithm 3) for the two-sample testing problem in (13) for the Plackett-Luce model, and associated upper bounds on its sample complexity.

Input : Two sets S_P and S_Q of m -length partial rankings, where $2 \leq m \leq d$. The two sets of partial rankings, S_P and S_Q correspond to pairwise probability matrices P and Q respectively, according to the Plackett-Luce model. Rank breaking method, $R \in \{R_R, R_D, R_C\}$.

(1) Using the rank breaking method get

$$\{k_{ij}^p, X_{ij}\}_{(i,j) \in [d]^2, i < j} \leftarrow R(S_P); \quad \{k_{ij}^q, Y_{ij}\}_{(i,j) \in [d]^2, i < j} \leftarrow R(S_Q).$$

(2) Execute Algorithm 1.

Algorithm 3: Two-sample testing with partial ranking data for Plackett-Luce model.

We note that both Algorithm 3 and Algorithm 4, defined in this section, can be used with any of the three rank breaking methods described. The subsequent guarantees provided depend on the rank breaking method used, as we see in Theorem 7 and Theorem 8.

In our results for two-sample testing under the Plackett-Luce modeling assumption, we consider two cases. In the first case, for some $m \in \{2, \dots, d-1\}$, each sample is a ranking of some m items chosen uniformly at random from the set of d items. In the second case, the samples comprise of total rankings, that is, $m = d$. The following two theorems characterize the performance of Algorithm 3 thereby establishing an upper bound on the sample complexity of the two-sample testing problem defined in (13). In these theorems we use the disjoint rank breaking methods so that the pairwise-comparisons created from a ranking are independent.

Theorem 7 *Consider the testing problem in (13) where pairwise probability matrices P and Q follow the Plackett-Luce model. Suppose we have N samples, where for some $m \in \{2, \dots, d-1\}$, each sample is a ranking of some m items chosen uniformly at random from the set of d items. Then, there are positive constants c, c_0, c_1 and c_2 such that if $N \geq c \frac{d^2 \log(d)}{m} \lceil \frac{c_0}{d\epsilon^2} \rceil$ and $\epsilon \geq c_1 d^{-c_2}$, then Algorithm 3 with the “Random disjoint” rank breaking method will correctly distinguish between $P = Q$ and $\frac{1}{d} \|P - Q\|_F \geq \epsilon$, with probability at least $\frac{2}{3}$.*

The proof of Theorem 7 is provided in Section 6.4. The lower bound assumption on ϵ in Theorem 7 is to ensure that the sufficient number of pairwise comparisons needed after applying the random disjoint rank breaking algorithm, is not very large. Theorem 7 is a combined result of the random disjoint rank breaking algorithm and the result in Theorem 1. When we have total rankings, Algorithm 3 with the “Deterministic disjoint” rank breaking method yields an improvement in the sample complexity by a logarithmic factor. We state this formally in the following theorem.

Theorem 8 *Consider the testing problem in (13) where pairwise probability matrices P and Q follow the Plackett-Luce model. Suppose we have N samples of total rankings from each population. Then, there are positive constants c, c_1 and c_2 such that if $N \geq 2d \lceil \frac{c}{d\epsilon^2} \rceil$ and $\epsilon \geq c_1 d^{-c_2}$, then Algorithm 3 with the “Deterministic disjoint” rank breaking method will correctly distinguish between $P = Q$ and $\frac{1}{d} \|P - Q\|_F \geq \epsilon$, with probability at least $\frac{2}{3}$.*

The proof of Theorem 8 is provided in Section 6.5. These two results provide an upper bound on the sample complexity when using partial (and total) rankings for the two-sample testing problem in (13) under the Plackett-Luce model. In Theorem 7 and Theorem 8 the lower bound of $\frac{2}{3}$ on probability of success is tied to the specific threshold used in Algorithm 1 in the same manner as described for Theorem 1. Specifically, for any constant $\nu \in (0, 1)$, Algorithm 3 can achieve a probability of error at most ν with the same order of sample complexity as mentioned in Theorem 7 and Theorem 8.

Algorithm 3 addresses the problem of two-sample testing under the Plackett-Luce model. Now, we provide a permutation test based algorithm for the more general, non-parametric model, namely, marginal probability based model. The permutation test method described in Algorithm 4 gives a sharper (implicit) threshold than that in Algorithm 3. Note that Algorithm 4 doesn’t require any assumptions on the length of the partial-ranking data, the partial-ranking data in each population can be of varying lengths. Moreover, as we will see in Theorem 9, the Type I error guarantee of Algorithm 4 holds even if the pairwise-comparisons created from the rank breaking method are dependent, hence the guarantee does not depend on the choice of the rank breaking method.

The key difference between the permutation testing algorithm for pairwise-comparison data, described in Section 3.1, and the permutation testing algorithm for partial ranking data, described in Algorithm 4, is the shuffling step. In our partial ranking based setup, each ranking sample is obtained independent of all else while the pairwise-comparisons obtained from a rank are not necessarily independent of each other. Hence, in the partial ranking based permutation testing algorithm (Algorithm 4), we re-distribute ranking samples between

the two populations and not the pairwise-comparisons.

Input : Two sets of partial rankings S_P and S_Q from two populations corresponding to the probability distributions λ_P and λ_Q . Significance level $\alpha \in (0, 1)$. Rank breaking method, $R \in \{R_R, R_D, R_C\}$. Iteration count γ .

(1) Using the rank breaking method get

$$\{k_{ij}^p, X_{ij}\}_{(i,j) \in [d]^2, i < j} \leftarrow R(S_P); \quad \{k_{ij}^q, Y_{ij}\}_{(i,j) \in [d]^2, i < j} \leftarrow R(S_Q).$$

(2) Compute the test statistic T defined in (8).

(3) {Repeat γ times} Put the samples in S_P and S_Q together and reassign the samples at random such that the number of samples assigned to each population is the same as before. Repeat Step 1 and Step 2. Denote the computed test statistic as T_ℓ for the ℓ^{th} iteration.

Output : Reject the null if $p = \sum_{\ell=1}^{\gamma} \frac{1}{\gamma} \mathbb{1}(T_\ell - T) < \alpha$.

Algorithm 4: Two-sample testing algorithm with partial ranking data for marginal probability based model.

We now show that Algorithm 4 controls the Type I error of the two-sample testing problem in (10) under the more general, marginal probability based partial ranking model. This result relies on considerably weaker assumptions than Theorem 7 and Theorem 8. In particular, we do not assume that each ranking is of the same length. We only assume that the (sub)set of items ranked in each sample from each population is sampled independently from the same distribution. Specifically, let there be any probability distribution over all non-empty subsets of $[d]$. Then, the set of items ranked in each sample for each population is sampled i.i.d. from this distribution. Moreover, the number of samples from the two populations need not be equal.

Theorem 9 *Consider any probability distributions λ_P and λ_Q and the two-sample testing problem in (10). Suppose we have partial ranking data from each population such that the sets of items ranked in each sample in each population is sampled i.i.d. from any probability distribution over all non-empty subsets of $[d]$. Suppose the partial ranking data follows the marginal probability based model. Then, for any significance level $\alpha \in (0, 1)$ the permutation testing method of Algorithm 4 has Type I error at most α .*

The proof of Theorem 9 is provided in Section 6.6. Recall that the Plackett-Luce model is a special case of the marginal probability based model, and hence as a direct corollary, the guarantees for Algorithm 4 established in Theorem 9 also apply to the Plackett-Luce model.

5. Experiments

In this section, we present results from experiments on simulated and real-world data sets, to gain a further understanding of the problem of two-sample testing on pairwise-comparison data.




Figure 1: Power of the testing algorithm versus the scaling factor of the sample size parameter $n = \frac{1}{ad\epsilon^2}$. We use Algorithm 2 which uses the test statistic in (8) with the permutation testing method. The test is conducted at a significance level of 0.05 (indicated by the horizontal line at $y = 0.05$). Unless specified otherwise, the parameters are fixed as $d = 20, \epsilon = 0.05, a = 1$.

5.1 Pairwise-comparison data

We now describe real-world experiments and synthetic simulations we conduct for two-sample testing on pairwise-comparison data. In these experiments, we use the test statistic we designed in Algorithm 1 along with the permutation testing method as described in Algorithm 2 to obtain an implicit value of the threshold and control Type I error.

5.1.1 SYNTHETIC SIMULATIONS

We conduct two sets of experiments via synthetic simulations. The first set of experiments empirically evaluates the dependence of the power of our test (Algorithm 1) with respect to individual problem parameters. In each of the simulations, we set the significance level to be 0.05. Specifically, given the problem parameters n, a, d and ϵ , we consider the random-design setting with $k_{ij}^p, k_{ij}^q \stackrel{\text{iid}}{\sim} \text{Bin}(n, a)$. We consider the asymmetric and model-free setting, fix $P = [\frac{1}{2}]^{d \times d}$ and set $Q = P + \Delta$ where Δ is sampled uniformly at random from the set of all matrices in $[-\frac{1}{2}, \frac{1}{2}]^{d \times d}$ with $\frac{1}{d} \|\Delta\|_F = \epsilon$. In Figure 1a, b and c, we vary the parameter d, ϵ and a respectively, keeping the other parameters fixed. Recall that our results in Theorem 1 and Corollary 2 predict the sample complexity as $n = \Theta(\frac{1}{ad\epsilon^2})$. To test this theoretical prediction, we set the sample size n (on the x-axis) as $n = \frac{1}{ad\epsilon^2}$, and plot the power of the test on the y-axis. Each plot point in Figure 1 is obtained by averaging over 400 iterations of the experiment, and the threshold for the test is obtained by running the permutation test method over 5000 iterations. Observe that, interestingly in each figure, the curves across all values of the varied parameters nearly coincide, thereby validating the sample complexity predicted by our theoretical results.

The second set of experiments empirically investigates the role of the underlying pairwise-comparison models in two-sample testing with our test (Algorithm 1). We consider the random-design setup in the symmetric setting with $k_{ij}^p, k_{ij}^q \stackrel{\text{iid}}{\sim} \text{Bin}(n, a) \forall i < j$. We generate the matrices P and Q in three ways: model-free, BTL and SST. In the model-free setting, we

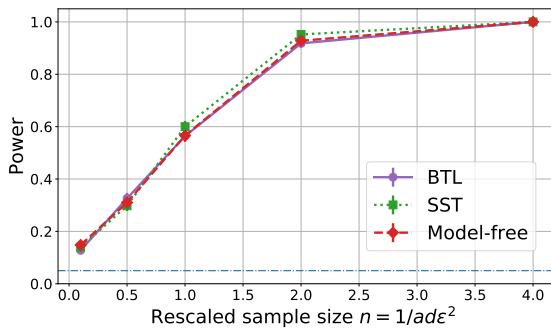


Figure 2: Power of our test (Algorithm 2) under three different models for pairwise-comparisons: BTL, SST and the model-free setting. The parameters of the problem are set as $d = 20$, $\epsilon^2 = 0.05$, $a = 1$ and the test is conducted at a significance level of 0.05 (indicated by the horizontal line at $y = 0.05$).

generate P and Q in a manner similar to the first set of simulations above, with the additional constraints $\Delta_{ji} = 1 - \Delta_{ij} \forall i \leq j$. For the BTL and SST models, we fix $P = [\frac{1}{2}]^{d \times d}$. For the BTL model, we choose w_p according to the construction in Section 6.3.4 to obtain Q . For the SST model, we set $Q = P + \Delta$, where matrix Δ is generated as follows. We generate Δ by arranging $\binom{d}{2}$ random variables uniformly distributed over $[0, 1]$, in a row-wise decreasing and column-wise increasing order in the upper triangle matrix ($\Delta_{ij} = 1 - \Delta_{ji}$) and normalizing to make $\frac{1}{d} \|\Delta\|_F = \epsilon$. This construction ensures that Δ lies in the SST class, and since matrix P is a constant matrix, Q is also guaranteed to lie in the SST class. The results of the simulations are shown in Figure 2. The results show that in the settings simulated, the power of the testing algorithm is identical in all the models considered. This leaves an open question whether there exists a tighter information-theoretic lower bound for the SST and the parameter-based model that matches the upper bound derived for the test in Algorithm 1 or if there exists a test statistic for these models with a better rate.

5.1.2 REAL-WORLD DATA

In this section, we describe the results of our experiments on two real-world data sets. In these experiments, we use Algorithm 2 to obtain a p -value for the experiment.

Ordinal versus cardinal An important question in the field of crowdsourcing and data-elicitation from people is whether pairwise-comparisons provided by people (ordinal responses) are distributed similarly to if they provide ratings (cardinal responses) which are then converted to pairwise-comparisons (Shah et al., 2016; Raman and Joachims, 2014). In this section, we use the permutation based two-sample test described in Algorithm 2 to address this question. We use the data set from Shah et al. (2016) comprising six different experiments on the Amazon Mechanical Turk crowdsourcing platform. In each experiment, workers are asked to either provide ratings for the set of items in that experiment (age for photo given, number of spelling mistakes in a paragraph, distance between two cities, relevance of web-based search results, quality of taglines for a product, frequency of a piano sound clip) or provide pairwise-comparisons. The number of items in each experiment ranged from 10 to 25. For each of the six experiments, there were 100 workers, and each worker was assigned to either the ordinal or the cardinal version of the task uniformly at random.

The first set of samples corresponds to the elicited ordinal responses and the second set of samples are obtained by converting the elicited ratings to ordinal data. We have a total of 2017 ordinal responses and 1671 cardinal-converted-to-ordinal responses. More details about the data set and experiment are provided in the appendix.

Using Algorithm 2, we test for difference in the two resulting distributions for the entire data set ($d = 74$). We observe that *the test rejects the null with a p-value of 0.003, thereby concluding a statistically significant difference between the ordinal and the cardinal-converted-to-ordinal data.*

European football leagues In the second data set, we investigate whether the relative performances of the teams (in four European football leagues: English Premier League, Bundesliga, Ligue 1, La Liga) changed significantly from the 2016-17 season to the 2017-18 season. The leagues are designed such that each pair of teams plays twice in a season (one home, one away game), so we have at most two pairwise-comparisons per pair within a league (we do not consider the games that end in a draw). Each league has 15-17 common teams across two consecutive seasons. This gives a total of 801 and 788 pairwise-comparisons in the 2016-17 and 2017-18 seasons respectively. More details about the experiment are provided in the appendix.

Using the test statistic of Algorithm 1 with permutation testing, we test for a difference in the two resulting distributions for the entire data set ($d = 67$). We observe that *the test fails to reject the null with a p-value of 0.971, that is, it does not recognize any significant difference between the relative performance of the European football teams in 2017-18 season and the 2016-17 season from the data available.* Running the test for each league individually also fails to reject the null.

5.2 Partial and total ranking data

We now describe the experiments we conducted on real-world data for two-sample testing on partial (and total) ranking data. In these experiments, we use the test statistic (8) along with the permutation testing method, as explained in Algorithm 4.

For our experiments, we use the ‘‘Sushi preference data set’’ Kamishima (2003), in which subjects rank different types of sushi according to their preferences. The data set contains two sets of ranking data. In the first set, the subjects are asked to provide a total ranking over 10 items (popular types of sushi). In this set, all subjects are asked to rank the same 10 objects. This set contains 5000 such total rankings.

In the second set of ranking data, a total of 100 types of sushi are considered. We first describe how the 100 types are chosen. The authors in Kamishima (2003) surveyed menu data from 25 sushi restaurants found on the internet. For each type of sushi sold at the restaurant, they counted the number of restaurants that offered the item. From these counts, they derived the probabilities that each item would be supplied. By eliminating unfamiliar or low frequency items, they came up with a list of 100 items. Each subject in this set is asked to rank a subset of 10 items randomly selected from the 100 items, according to the probability distribution described above. This set contains responses from 5000 subjects.

In addition, this data set contains demographic information about all the subjects, including their

- (a) Gender {Male, Female}

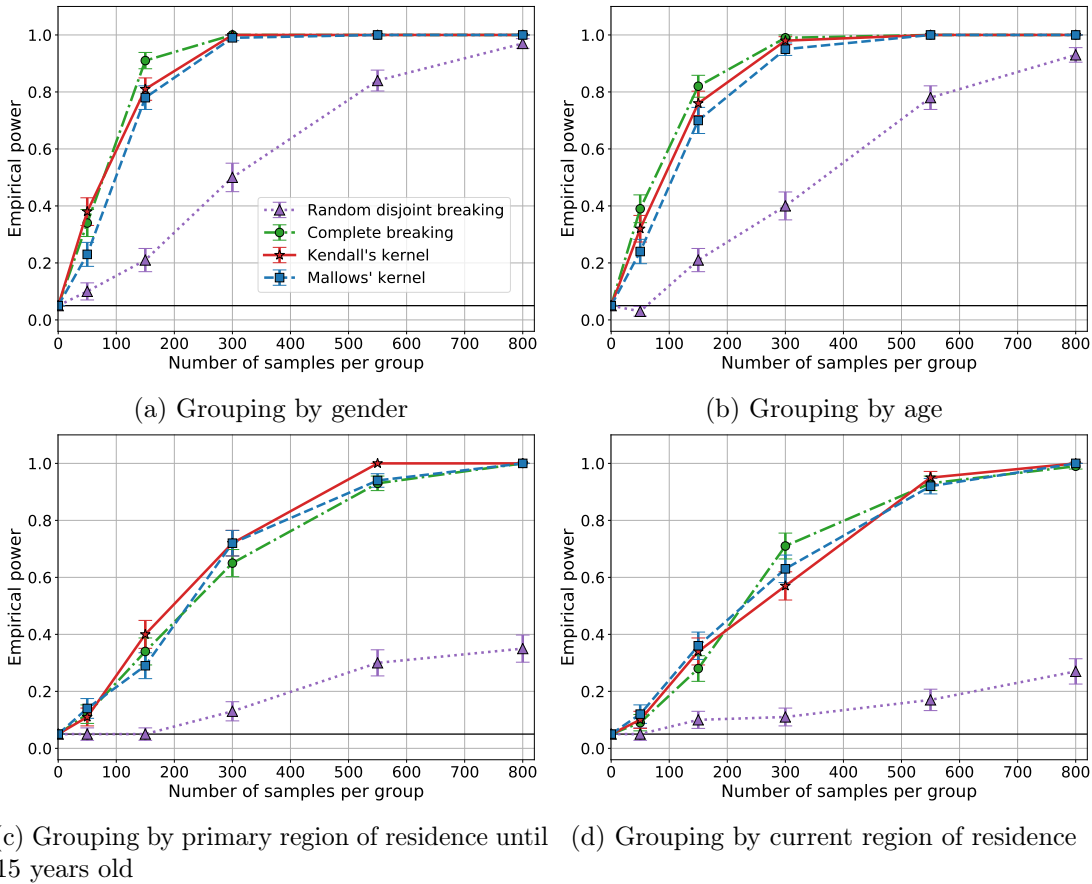


Figure 3: Empirical power of our test statistic T with the permutation testing method described in Algorithm 4 in testing for difference in sushi preference from the first set of responses with $d = 10$. The responses obtained comprise of total rankings from each subject. Test results are shown for differences between demographic division based on the information available. Two different rank breaking methods are used for our algorithm, namely, “Random Disjoint” and “Complete”. We also show the results for kernel-based two-sample testing with Kendall’s kernel and Mallows’ kernel as in Mania et al. (2018). The x-axis shows the number of samples (total rankings) from each group used to conduct the test and the y-axis shows the empirical power of our test. The test is conducted at a significance level of 0.05 (indicated by the horizontal line at $y = 0.05$). Empirical power is computed as an average over 100 trials.

- (b) Age {Above 30, Below 30}
- (c) Current region of residence {East, West}
- (d) Primary region of residence until 15 years old {East, West}.

Using our testing algorithm, we test for a difference in preferences across the two sections within each demographic mentioned above. In the first set of experiments, we implement the permutation testing method with our test statistic (8) on the first set of ranking data with $d = 10$ for each demographic division. We show the results in Figure 3 for two rank-breaking methods, namely, “Random Disjoint” and “Complete”. In addition, we show the results of two

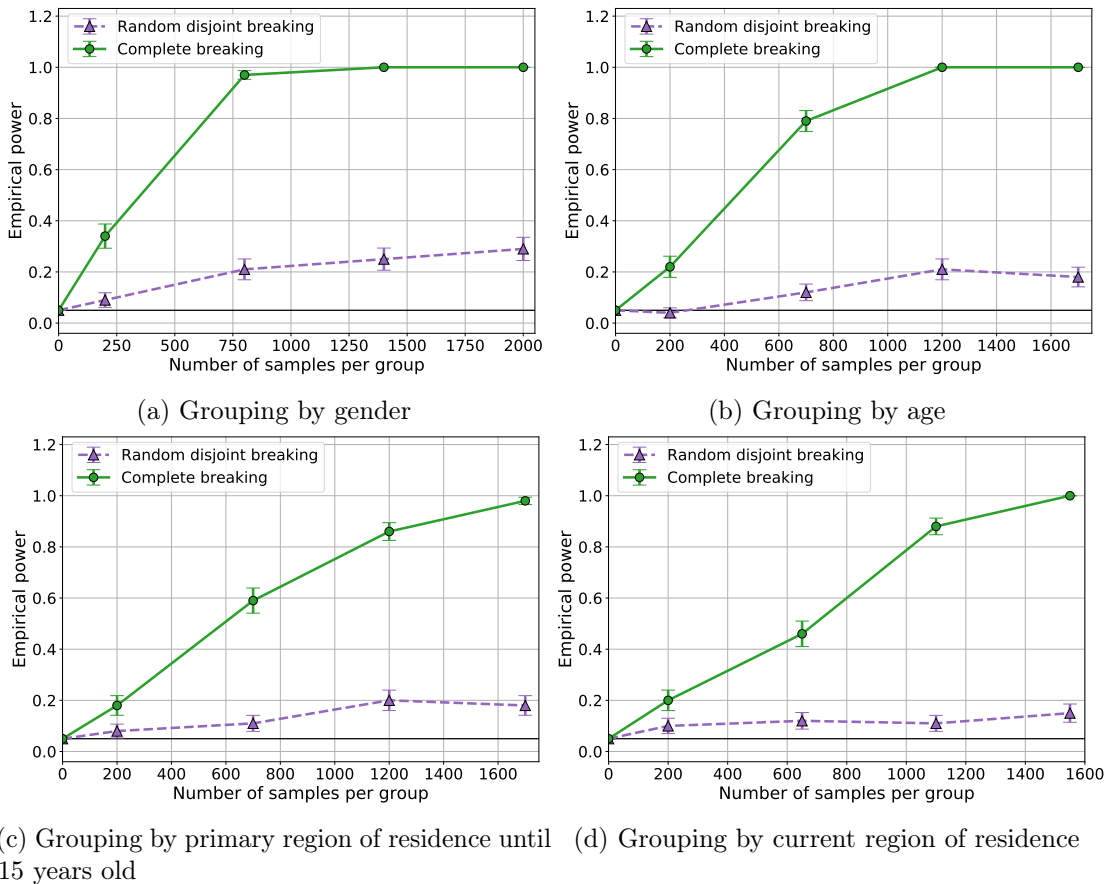


Figure 4: Empirical power of our test statistic T with the permutation testing method described in Algorithm 4 in testing for difference in sushi preference from the first set of responses with $d = 100$. Test results are shown for differences between demographic division based on the information available. Two different rank breaking methods are used, namely, “Random Disjoint” and “Complete”. The x-axis shows the number of samples (total rankings) from each sub-group used to conduct the test and the y-axis shows the empirical power of our test. The test is conducted at a significance level of 0.05 (indicated by the horizontal line at $y = 0.05$). Empirical power is computed as an average over 100 iterations.

kernel-based two-sample testing methods for total rankings designed in Mania et al. (2018), namely, Kendall’s kernel and Mallows’ kernel. We randomly sub-sampled n samples from each sub-group of subjects and used 200 permutations to determine the rejection threshold for the permutation test. In these experiments, we show the empirical power of our testing method, which is the fraction of times our test rejected the null in a total of 100 trials. We show all the results of using this method on the sushi data set in Figure 3. *Across each demographic division, our test detects a statistically significant difference in distribution over sushi preferences for the 10 types of sushi included.* Moreover, our testing algorithm with “Complete” rank breaking method performs competitively with the kernel-based two-sample testing methods introduced in Mania et al. (2018).

We note that since our testing algorithms also work with partial ranking data, our testing algorithms are much more general than the testing algorithms in Mania et al. (2018), as we demonstrate in our next set of experiments on the second sushi preference data set. We use Algorithm 4 to test if the preferences of the subjects in the second sushi data set also varies across the different demographics. Recall that this data set has $d = 100$ items in total but each ranking only ranks a subset of 10 items. The other details of the experiment are the same as the previous experiment. The results are shown in Figure 4. Again, across each demographic, our test detects a statistically significant difference in distribution over sushi preferences for the 100 types of sushi included.

6. Proofs

This section is devoted to the proofs of our main results. In Section 6.1 and Section 6.2 we prove the positive results from Section 3.1, and in Section 6.3 we prove the converse results from Section 3.2. Lastly, Sections 6.4-6.6 are devoted to proofs of results under the partial (or total) ranking setting mentioned in Section 4.2.

Throughout these and other proofs, we use the notation c, c', c_0, c_1 and so on to denote positive constants whose values may change from line to line.

6.1 Proof of Corollary 2

In this section we present the complete proof of Corollary 2. We first present the proof for the random-design setup described in Corollary 2 and then specialise the proof in Section 6.2 to the per-pair fixed-design setup in Theorem 1. To prove our result, we analyse the expected value and the variance of the test statistic T in Algorithm 1 in the following two lemmas. Recall that under the random-design setup k_{ij}^p, k_{ij}^q are distributed independently and identically according to some distribution \mathcal{D} that satisfies the conditions in (9).

Lemma 10 *For T as defined in Algorithm 1, with $k_{ij}^p, k_{ij}^q \stackrel{\text{iid}}{\sim} \mathcal{D}$, under the null $\mathbb{E}_{H_0}[T] = 0$ and under the alternate,*

$$\mathbb{E}_{H_1}[T] \geq c\mu\|P - Q\|_F^2.$$

The proof of Lemma 10 is provided in Section 6.1.1. Now, with a view to applying Chebyshev's concentration inequality, we bound the variance of T .

Lemma 11 *For T as defined in Algorithm 1, with $k_{ij}^p, k_{ij}^q \stackrel{\text{iid}}{\sim} \mathcal{D}$, where \mathcal{D} obeys the conditions described in (9), under the null*

$$\text{Var}_{H_0}[T] \leq 24d^2,$$

and under the alternate,

$$\text{Var}_{H_1}[T] \leq 24d^2 + 8\mu\|P - Q\|_F^2 + c'\mu^2\|P - Q\|_F^2$$

where $c' > 0$ is a constant.

The proof for Lemma 11 is provided in Section 6.1.2. We now have to control Type I error and Type II error. Using one-sided Chebyshev's inequality for the test statistic T , which has $\mathbb{E}_{H_0}[T] = 0$, we derive an upper bound on Type I error as follows

$$\mathbb{P}_{H_0}(T \geq t) \leq \frac{\text{Var}_{H_0}[T]}{\text{Var}_{H_0}[T] + t^2}. \quad (14)$$

Observe that if $t = 11d$ then the Type I error is upper bounded by $\frac{1}{6}$. In addition, if Type I error is required to be at most ν , then we set the threshold equal to $d\sqrt{\frac{24(1-\nu)}{\nu}}$. We now move to controlling the Type II error of the testing algorithm. We again invoke Chebyshev's inequality as follows

$$\mathbb{P}_{H_1}(T < t) \leq \frac{\text{Var}_{H_1}[T]}{\text{Var}_{H_1}[T] + (\mathbb{E}_{H_1}[T] - t)^2}. \quad (15)$$

To guarantee that Type II error is at most $\frac{1}{6}$, we substitute the bounds on $\mathbb{E}_{H_1}[T]$, $\text{Var}_{H_0}[T]$, $\text{Var}_{H_1}[T]$ from Lemma 10 and Lemma 11 in (15) to get the sufficient condition

$$\begin{aligned} 5(24d^2 + 8\mu\|P - Q\|_{\mathbb{F}}^2 + c'\mu^2\|P - Q\|_{\mathbb{F}}^2) &\leq (c\mu\|P - Q\|_{\mathbb{F}}^2 - 11d)^2 \\ 40\mu\|P - Q\|_{\mathbb{F}}^2 + 22cd\mu\|P - Q\|_{\mathbb{F}}^2 + 5c'\mu^2\|P - Q\|_{\mathbb{F}}^2 &\leq c^2\mu^2\|P - Q\|_{\mathbb{F}}^4 + d^2. \end{aligned}$$

This condition yields

$$40 + 22cd + 5c'\mu \leq c^2\mu\|P - Q\|_{\mathbb{F}}^2. \quad (16)$$

Recall that under the alternate $\frac{1}{d}\|P - Q\|_{\mathbb{F}} \geq \epsilon$. According to the final condition derived here (16), under the regime $\mu > d$, we have control over total probability of error if $\epsilon^2 d^2 \geq c'$ for some constant $c' > 0$. Under the regime $\mu \leq d$, the condition (16) simplifies as

$$\epsilon^2 \geq \frac{c''}{\mu d}, \quad (17)$$

where $c'' > 0$ is some constant. This gives the sufficient condition to control total probability of error (sum of Type I error and Type II error) to be at most $\frac{1}{3}$ under the setting where $k_{ij}^p, k_{ij}^q \stackrel{\text{iid}}{\sim} \mathcal{D}$.

6.1.1 PROOF OF LEMMA 10

We now prove the bounds on the expected value of the test statistic defined in Algorithm 1. Recall that for each (i, j) , given k_{ij}^p, k_{ij}^q , we have $X_{ij} \sim \text{Bin}(k_{ij}^p, p_{ij})$ and $Y_{ij} \sim \text{Bin}(k_{ij}^q, q_{ij})$. Also, $k_{ij}^p, k_{ij}^q \stackrel{\text{iid}}{\sim} \mathcal{D}$ wherein $\mathbb{E}[k_{ij}^p] = \mu$, $\text{Var}[k_{ij}^p] = \sigma^2$, $\Pr(k_{ij}^p = 1) = p_1$ and \mathcal{D} obeys (9). We denote the vector of k_{ij}^p and k_{ij}^q for all (i, j) by \mathbf{k}^p and \mathbf{k}^q respectively. Now, the conditional expectation of T is expressed as

$$\mathbb{E}[T | \mathbf{k}^p, \mathbf{k}^q] = \sum_{i=1}^d \sum_{j=1}^d \frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{k_{ij}^p + k_{ij}^q} (p_{ij} - q_{ij})^2. \quad (18)$$

Using the law of total expectation, we have

$$\begin{aligned}\mathbb{E}[T] &= \mathbb{E}[\mathbb{E}[T | \mathbf{k}^p, \mathbf{k}^q]] \\ &= \sum_{i=1}^d \sum_{j=1}^d \mathbb{E} \left[\frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{k_{ij}^p + k_{ij}^q} (p_{ij} - q_{ij})^2 \right] \\ &= \mathbb{E} \left[\frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{k_{ij}^p + k_{ij}^q} \right] \|P - Q\|_F^2\end{aligned}$$

Clearly, $\mathbb{E}_{H_0}[T] = 0$. To find a lower bound for $\mathbb{E}_{H_1}[T]$, we first note that

$$\begin{aligned}\mathbb{E} \left[\frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{k_{ij}^p + k_{ij}^q} \right] &= \mathbb{E} \left[\frac{\mathbb{I}_{ij}^0 k_{ij}^p k_{ij}^q}{k_{ij}^p + k_{ij}^q} \right] - 2 \sum_{k \in [d]} \mathbb{P}(k_{ij}^p = 1, k_{ij}^q = k) \frac{k}{k+1} + \frac{1}{2} \mathbb{P}(k_{ij}^p = 1, k_{ij}^q = 1) \\ &\geq \mathbb{E} \left[\frac{\mathbb{I}_{ij}^0 k_{ij}^p k_{ij}^q}{k_{ij}^p + k_{ij}^q} \right] - 2p_1.\end{aligned}\tag{19}$$

where $\mathbb{I}_{ij}^0 = \mathbb{I}(k_{ij}^p > 0) \times \mathbb{I}(k_{ij}^q > 0)$. Furthermore, we see that for any event E ,

$$\mathbb{E} \left[\frac{\mathbb{I}_{ij}^0 k_{ij}^p k_{ij}^q}{k_{ij}^p + k_{ij}^q} \right] \geq \mathbb{E} \left[\frac{\mathbb{I}_{ij}^0 k_{ij}^p k_{ij}^q}{k_{ij}^p + k_{ij}^q} \mid E \right] \Pr(E).\tag{20}$$

We define the event E as

$$\begin{aligned}\mu - c\sigma &\leq k_{ij}^p \leq \mu + c\sigma, \text{ and} \\ \mu - c\sigma &\leq k_{ij}^q \leq \mu + c\sigma\end{aligned}\tag{21}$$

with some constant $c > 1$ such that $\mu - c\sigma > 0$. Using Chebyshev's inequality, we get that $\Pr(E) \geq (1 - \frac{1}{c^2})^2$. Finally, we combine (19), (20) and (21), to get

$$\mathbb{E} \left[\frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{k_{ij}^p + k_{ij}^q} \right] \geq \frac{(\mu - c\sigma)^2}{2(\mu + c\sigma)} \left(1 - \frac{1}{c^2}\right)^2 - 2p_1.\tag{22}$$

Since \mathcal{D} obeys the conditions in (9), we have $\mu \geq c_1 p_1$ and $\mu \geq c_2 \sigma$. Therefore, there is a constant $c > 0$ that depends on c_1, c_2 , such that $\mathbb{E}[T] \geq c\mu \|P - Q\|_F^2$. This proves Lemma 10.

6.1.2 PROOF OF LEMMA 11

To analyse the variance of the test statistic T , we note that pairwise-comparisons for each pair are obtained independently. This allows us to compute the variance for each pair (i, j) separately, as variance of sum is equal to the sum of variances. The following analysis of the variance of the test statistic T applies under both the null and the alternate. The law of total variance states that

$$\text{Var}[T] = \mathbb{E}[\text{Var}[T | \mathbf{k}^p, \mathbf{k}^q]] + \text{Var}[\mathbb{E}[T | \mathbf{k}^p, \mathbf{k}^q]].\tag{23}$$

We evaluated the term $\text{Var}[T|\mathbf{k}^p, \mathbf{k}^q]$, present in the expression above, in Wolfram Mathematica. We show the output here,

$$\begin{aligned}
 \text{Var}[T|\mathbf{k}^p, \mathbf{k}^q] &\leq \sum_{i=1}^d \sum_{j=1}^d \frac{2\mathbb{I}_{ij}k_{ij}^p(k_{ij}^p-1)k_{ij}^q(k_{ij}^q-1)}{(k_{ij}^p-1)^2(k_{ij}^q-1)^2(k_{ij}^p+k_{ij}^q)^2} \left(k_{ij}^q(k_{ij}^q-1)p_{ij}^4(3-2k_{ij}^p) \right. \\
 &\quad + 2p_{ij}^3k_{ij}^q(k_{ij}^q-1)(-2+2q_{ij}k_{ij}^p-2q_{ij}+k_{ij}^p) \\
 &\quad + 2p_{ij}q_{ij}(k_{ij}^p-1)(k_{ij}^q-1)(1+2q_{ij}^2k_{ij}^p-q_{ij}-2q_{ij}k_{ij}^p+q_{ij}k_{ij}^q) \\
 &\quad + p_{ij}^2(k_{ij}^q-1)(2q_{ij}(k_{ij}^p-1)(k_{ij}^p-1-2k_{ij}^q)+k_{ij}^p-2q_{ij}^2(k_{ij}^p-1)(k_{ij}^p+k_{ij}^q-1)) \\
 &\quad \left. - q_{ij}^2(q_{ij}-1)k_{ij}^p(k_{ij}^p-1)(1-3q_{ij}+2q_{ij}k_{ij}^q) \right) \\
 &\leq \sum_{i=1}^d \sum_{j=1}^d \frac{8\mathbb{I}_{ij}}{(k_{ij}^p+k_{ij}^q)^2} \left(k_{ij}^p(k_{ij}^p-1)(k_{ij}^q-1)(2p_{ij}(p_{ij}-q_{ij})^2) \right. \\
 &\quad + k_{ij}^q(k_{ij}^q-1)(k_{ij}^p-1)(2q_{ij}(p_{ij}-q_{ij})^2) \\
 &\quad + 2p_{ij}q_{ij}(k_{ij}^p-1)(k_{ij}^q-1)(1-p_{ij})(1-q_{ij}) \\
 &\quad \left. + p_{ij}^2k_{ij}^q(k_{ij}^q-1)(1-p_{ij})^2 + q_{ij}^2k_{ij}^p(k_{ij}^p-1)(1-q_{ij})^2 \right). \tag{24}
 \end{aligned}$$

Applying the trivial upper bound $p_{ij} \leq 1, q_{ij} \leq 1 \forall (i, j)$, we get

$$\text{Var}[T|\mathbf{k}^p, \mathbf{k}^q] \leq \sum_{i=1}^d \sum_{j=1}^d 8\mathbb{I}_{ij} \left(\frac{k_{ij}^p k_{ij}^q}{(k_{ij}^p + k_{ij}^q)} (p_{ij} - q_{ij})^2 + 3 \right) \tag{25}$$

Following this, we evaluate the first term on the right hand side of (23) as

$$\mathbb{E}[\text{Var}[T|\mathbf{k}^p, \mathbf{k}^q]] \leq 24d^2 + 8\|P - Q\|_{\mathbb{F}}^2 \mathbb{E} \left[\frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{k_{ij}^p + k_{ij}^q} \right]. \tag{26}$$

To further simplify the upper bound in (26), we observe that

$$\mathbb{E} \left[\frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{k_{ij}^p + k_{ij}^q} \right] \leq \frac{1}{2} \mathbb{E}[\max\{k_{ij}^p, k_{ij}^q\}]. \tag{27}$$

We exploit the independence of k_{ij}^p, k_{ij}^q to get the CDF of $\max\{k_{ij}^p, k_{ij}^q\}$ as

$$\mathbb{P}(\max\{k_{ij}^p, k_{ij}^q\} \leq x) = \mathbb{P}(k_{ij}^p \leq x) \mathbb{P}(k_{ij}^q \leq x).$$

Through the CDF, we derive the PDF as

$$\begin{aligned}
 \mathbb{P}(\max\{k_{ij}^p, k_{ij}^q\} = x) &= \mathbb{P}(\max\{k_{ij}^p, k_{ij}^q\} \leq x) - \mathbb{P}(\max\{k_{ij}^p, k_{ij}^q\} \leq x-1) \\
 &= \mathbb{P}(k_{ij}^p \leq x)^2 - \mathbb{P}(k_{ij}^p \leq x-1)^2 \\
 &= \mathbb{P}(k_{ij}^p = x)(\mathbb{P}(k_{ij}^p \leq x) + \mathbb{P}(k_{ij}^p \leq x-1)) \\
 &\leq 2\mathbb{P}(k_{ij}^p = x) \tag{28}
 \end{aligned}$$

We substitute this inequality in (27) to get

$$\mathbb{E} \left[\frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{k_{ij}^p + k_{ij}^q} \right] \leq \mu. \quad (29)$$

As a result, following from (26), we have

$$\mathbb{E}[\text{Var}[T|\mathbf{k}^p, \mathbf{k}^q]] \leq 24d^2 + 8\mu \|P - Q\|_{\mathbb{F}}^2. \quad (30)$$

Now, the remaining (second) term on the right hand side of (23) is

$$\begin{aligned} \text{Var}[\mathbb{E}[T|\mathbf{k}^p, \mathbf{k}^q]] &= \sum_{i=1}^d \sum_{j=1}^d \text{Var} \left[\frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{(k_{ij}^p + k_{ij}^q)} \right] (p_{ij} - q_{ij})^4 \\ &\leq \text{Var} \left[\frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{(k_{ij}^p + k_{ij}^q)} \right] \sum_{i=1}^d \sum_{j=1}^d (p_{ij} - q_{ij})^2 \end{aligned} \quad (31)$$

$$\leq \text{Var} \left[\frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{(k_{ij}^p + k_{ij}^q)} \right] \|P - Q\|_{\mathbb{F}}^2. \quad (32)$$

To bound the variance term in the previous equation, we see that

$$\begin{aligned} \text{Var} \left[\frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{(k_{ij}^p + k_{ij}^q)} \right] &= \mathbb{E} \left[\left(\frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{(k_{ij}^p + k_{ij}^q)} \right)^2 \right] - \mathbb{E} \left(\left[\frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{(k_{ij}^p + k_{ij}^q)} \right] \right)^2 \\ &\stackrel{(a)}{\leq} \frac{1}{4} \mathbb{E} \left[\left(\max\{k_{ij}^p, k_{ij}^q\} \right)^2 \right] - c\mu^2 \\ &\stackrel{(b)}{\leq} \frac{1}{2} \mathbb{E}[(k_{ij}^p)^2] - c\mu^2 \\ &= \frac{1}{2}(\mu^2 + \sigma^2) - c\mu^2 \\ &\stackrel{(c)}{\leq} \left(\frac{1}{2} + \frac{1}{2c^2} - c \right) \mu^2 = c'\mu^2, \end{aligned} \quad (33)$$

where inequality (a) follows from (22), inequality (b) follows similarly to the result in (28), and inequality (c) is a result of (9). Thus, the upper bound in (32) becomes

$$\text{Var}[\mathbb{E}[T|\mathbf{k}^p, \mathbf{k}^q]] \leq c'\mu^2 \|P - Q\|_{\mathbb{F}}^2. \quad (34)$$

Finally, we put together the terms in (23) by combining (30) and (34) to get the desired upper bound on the variance of the test statistic under the alternate hypothesis, which is

$$\text{Var}[T] \leq 24d^2 + 8\mu \|P - Q\|_{\mathbb{F}}^2 + c'\mu^2 \|P - Q\|_{\mathbb{F}}^2. \quad (35)$$

Additionally, to obtain the upper bound on the variance of the test statistic under the null, we substitute $\|P - Q\|_{\mathbb{F}}$ as zero in (35). This completes the proof of Lemma 11.

6.2 Proof of Theorem 1

In this proof, we first specialise the statements of Lemma 10 and Lemma 11 to the per-pair fixed-design setup where $k_{ij}^p = k_{ij}^q = k \forall (i, j) \in [d]$, for some positive integer $k > 1$. Under this setting, following from (18), we have

$$\mathbb{E}[T] = \sum_{i=1}^d \sum_{j=1}^d \frac{1}{2} \mathbb{I}(k > 1) k (p_{ij} - q_{ij})^2 = \frac{1}{2} k \|P - Q\|_{\mathbb{F}}^2. \quad (36)$$

Similarly, we note that in (23) we have that $\text{Var}[\mathbb{E}[T|\mathbf{k}^p, \mathbf{k}^q]] = 0$, which in combination with (26) implies that

$$\text{Var}[T] \leq 24d^2 + 4k \|P - Q\|_{\mathbb{F}}^2. \quad (37)$$

Now, invoking Chebyshev's inequality as described in (14) and (15) to control Type I and Type II error at level $\frac{1}{6}$, we set the threshold as $11d$ to get the sufficient condition as

$$\epsilon^2 \geq \frac{c}{kd} \quad (38)$$

for some positive constant c . This proves Theorem 1.

6.3 Proof of converse results

In this section we prove all the claims made in Section 3.2. We begin with some background.

6.3.1 PRELIMINARIES FOR PROOF OF LOWER BOUND

We begin by briefly introducing the lower bound technique applied in Theorem 3 and Theorem 5. The main objective of the proof is to construct a set of null and alternate such that the minimax risk of testing defined in (2) is lower bounded by a constant. To lower bound the minimax risk, we analyse the χ^2 distance between the resulting distributions of the null and the alternate. We construct the null and alternate as follows. Let $P_0 = [\frac{1}{2}]^{d \times d}$. Under the null, we fix $P = Q = P_0$ and under the alternate, $P = P_0, Q \in \Theta$ where Θ is a set of matrices from the model class \mathcal{M} to be defined subsequently. We assume a uniform probability measure over Θ . The set Θ is chosen such that $\frac{1}{d} \|P_0 - Q\|_{\mathbb{F}} = \epsilon$ for all $Q \in \Theta$.

In our problem setup, we observe matrices X, Y wherein each element is the outcome of k observations of the corresponding Bernoulli random variable. For each pair (i, j) , we have $X_{ij} \sim \text{Bin}(k, p_{ij}), Y_{ij} \sim \text{Bin}(k, q_{ij})$. For simplicity of notation, we will denote the matrix distribution corresponding to the pairwise-comparison probability matrix P_0 by \mathbb{P}_0 , that is, $X \sim \mathbb{P}_0$ when $P = P_0$, and $Y \sim \mathbb{P}_0$ when $Q = P_0$. For the case where $Y_{ij} \sim \text{Bin}(k, q_{ij})$ and $Q \sim \text{Unif}(\Theta)$, we denote the resulting matrix distribution as $Y \sim \mathbb{P}_{\Theta}$. We now have all the parts required to derive the χ^2 divergence between the null and the alternate defined in this section.

The χ^2 divergence between the distribution of X, Y under the null and the distribution of X, Y under the alternate is given by

$$\begin{aligned} \chi^2((X, Y)_{H_0}, (X, Y)_{H_1}) &= \chi^2(X_{H_0}, X_{H_1}) + \chi^2(Y_{H_0}, Y_{H_1}) + \chi^2(X_{H_0}, X_{H_1})\chi^2(Y_{H_0}, Y_{H_1}) \\ &= \chi^2(\mathbb{P}_0, \mathbb{P}_0) + \chi^2(\mathbb{P}_0, \mathbb{P}_\Theta) + \chi^2(\mathbb{P}_0, \mathbb{P}_0)\chi^2(\mathbb{P}_0, \mathbb{P}_\Theta) \\ &= \chi^2(\mathbb{P}_0, \mathbb{P}_\Theta). \end{aligned} \quad (39)$$

This reduces our two-sample testing problem into a goodness of fit testing problem for the given model class, where the null distribution is given by \mathbb{P}_0 and the alternate distribution is given by \mathbb{P}_Θ . This goodness of fit testing problem is written as

$$\begin{aligned} H_0 : P &= P_0 \\ H_1 : P &\sim \text{Unif}(\Theta), \end{aligned} \quad (40)$$

where $P_0 = \left[\frac{1}{2}\right]^{d \times d}$.

Continuing with the reduction in (39) and (40), Le Cam's method for testing states that the minimax risk (2) for the hypothesis testing problem in (40), is lower bounded as (Lemma 3 in Collier et al. (2017))

$$\mathcal{R}_{\mathcal{M}} \geq \frac{1}{2} \left(1 - \sqrt{\chi^2(\mathbb{P}_0, \mathbb{P}_\Theta)}\right). \quad (41)$$

Therefore, if the χ^2 divergence is upper bounded by some constant $c < 1$, then no algorithm can correctly distinguish between the null and the alternate with probability of error less than $\frac{1}{2}(1 - \sqrt{c})$. Consequently, by deriving the value of ϵ corresponding to $c = \frac{1}{9}$, we will get the desired lower bound on the critical radius defined in (3) for the two-sample testing problem in (1).

We now delve into the technical part of the proof in which we derive the χ^2 divergence between \mathbb{P}_0 and \mathbb{P}_Θ . For a probability distribution \mathbb{P}_0 and a mixture probability measure \mathbb{P}_Θ , we know (from Lemma 7 in Carpentier et al. (2018)) that

$$\chi^2(\mathbb{P}_0, \mathbb{P}_\Theta) = \mathbb{E}_{(Q, Q') \sim \text{Unif}(\Theta)} \left(\int \frac{d\mathbb{P}_Q d\mathbb{P}_{Q'}}{d\mathbb{P}_0} \right) - 1. \quad (42)$$

Here $\mathbb{E}_{(Q, Q') \sim \text{Unif}(\Theta)}$ denotes the expectation with respect to the distribution of the pair (Q, Q') where Q and Q' are sampled independently and uniformly at random from the set Θ (with replacement). According to the null and alternate construction described in the beginning of this section, recall that $X \sim \mathbb{P}_0$ implies that $X_{ij} \sim \text{Bin}(k, \frac{1}{2}) \forall (i, j)$. Similarly $X \sim \mathbb{P}_Q$ implies that $X_{ij} \sim \text{Bin}(k, q_{ij}) \forall (i, j)$. With this information, we simplify the χ^2 divergence as

$$\chi^2(\mathbb{P}_0, \mathbb{P}_\Theta) = \mathbb{E}_{(Q, Q') \sim \text{Unif}(\Theta)} \sum_{v \in V} \left(\prod_{i=1}^d \prod_{j=1}^d \frac{\binom{k}{v_{ij}} q_{ij}^{v_{ij}} (1 - q_{ij})^{k - v_{ij}} \binom{k}{v'_{ij}} (q'_{ij})^{v'_{ij}} (1 - q'_{ij})^{k - v'_{ij}}}{\binom{k}{v_{ij}} \left(\frac{1}{2}\right)^k} \right) - 1. \quad (43)$$

where $V \in \mathbb{R}^{d(d-1)}$ is the set of all possible vectors with each element belonging to the set $\{0, 1, \dots, k\}$. There are $(k+1)^{d(d-1)}$ such vectors. We further simplify the summation over V as

$$\chi^2(\mathbb{P}_0, \mathbb{P}_\Theta) = \mathbb{E}_{(Q, Q') \sim \text{Unif}(\Theta)} \prod_{i=1}^d \prod_{j=1}^d \left(\sum_{\ell=0}^k \frac{\binom{k}{\ell} q_{ij}^\ell (1 - q_{ij})^{k-\ell} (q'_{ij})^\ell (1 - q'_{ij})^{k-\ell}}{\left(\frac{1}{2}\right)^k} \right) - 1. \quad (44)$$

This gives us the χ^2 divergence for the construction defined in terms of the elements of the matrices in the set Θ . Later, we will see that the set Θ designed for the different modeling assumptions considered (namely, MST and parameter-based) consist solely of matrices with entries from the set $\{\frac{1}{2} - \eta, \frac{1}{2}, \frac{1}{2} + \eta\}$. This information enables us to further simplify the expression for $\chi^2(\mathbb{P}_0, \mathbb{P}_\Theta)$.

Consider a pair of matrices (Q, Q') sampled uniformly at random from the set Θ . Let an agreement be defined as the occurrence of $\frac{1}{2} + \eta$ (or $\frac{1}{2} - \eta$) in the same position in Q and Q' and a disagreement is defined as the occurrence of $\frac{1}{2} + \eta$ in Q or Q' in the same position as $\frac{1}{2} - \eta$ in Q' or Q respectively. Next, we define two statistics b_1 and b_2 that quantify the number of agreements and disagreements, respectively, in the matrix pair (Q, Q') as shown here

$$\begin{aligned} b_1(Q, Q') &= \sum_{i=1}^d \sum_{j=1}^d \left[\mathbb{1}_{\{q_{ij}=q'_{ij}=\frac{1}{2}+\eta\}} + \mathbb{1}_{\{q_{ij}=q'_{ij}=\frac{1}{2}-\eta\}} \right], \\ b_2(Q, Q') &= \sum_{i=1}^d \sum_{j=1}^d \left[\mathbb{1}_{\{q_{ij}=\frac{1}{2}+\eta\}} \mathbb{1}_{\{q'_{ij}=\frac{1}{2}-\eta\}} + \mathbb{1}_{\{q_{ij}=\frac{1}{2}-\eta\}} \mathbb{1}_{\{q'_{ij}=\frac{1}{2}+\eta\}} \right]. \end{aligned} \quad (45)$$

Using these definitions, we state the following Lemma to analyse $\chi^2(\mathbb{P}_0, \mathbb{P}_\Theta)$ in (44).

Lemma 12 *Consider two pairwise-comparison probability matrices Q and Q' with $q_{ij} \in \{\frac{1}{2} - \eta, \frac{1}{2}, \frac{1}{2} + \eta\}$ and $q'_{ij} \in \{\frac{1}{2} - \eta, \frac{1}{2}, \frac{1}{2} + \eta\}$. Suppose $b_1(Q, Q') = b_1$ and $b_2(Q, Q') = b_2$. Then, we have*

$$\prod_{i=1}^d \prod_{j=1}^d \left(\sum_{\ell=0}^k \frac{\binom{k}{\ell} q_{ij}^\ell (1 - q_{ij})^{k-\ell} (q'_{ij})^\ell (1 - q'_{ij})^{k-\ell}}{\left(\frac{1}{2}\right)^k} \right) \leq (1 + 4\eta^2)^{k(b_1 - b_2)}. \quad (46)$$

The proof is provided at the end of this subsection. Using Lemma 12 and (44), we get an upper bound on the χ^2 divergence as

$$\chi^2(\mathbb{P}_0, \mathbb{P}_\Theta) \leq \mathbb{E}_{(Q, Q') \sim \text{Unif}(\Theta)} \left[(1 + 4\eta^2)^{k(b_1 - b_2)} \right] - 1. \quad (47)$$

We conclude the background section on the converse results here. We will use the equations discussed in this section to derive the lower bound for the different modeling assumptions in Theorem 3 and Theorem 5 in their respective proofs.

Proof of Lemma 12 Let

$$G(Q, Q') = \prod_{i=1}^d \prod_{j=1}^d \left(\sum_{\ell=0}^k \frac{\binom{k}{\ell} q_{ij}^{\ell} (1 - q_{ij})^{k-\ell} (q'_{ij})^{\ell} (1 - q'_{ij})^{k-\ell}}{\left(\frac{1}{2}\right)^k} \right) - 1. \quad (48)$$

Let $g(q_{ij}, q'_{ij})$ denote the summation in the equation above, that is

$$g(q_{ij}, q'_{ij}) = 2^k \sum_{\ell=0}^k \binom{k}{\ell} q_{ij}^{\ell} (1 - q_{ij})^{k-\ell} (q'_{ij})^{\ell} (1 - q'_{ij})^{k-\ell}. \quad (49)$$

Notice that if $q_{ij} = \frac{1}{2}$ or $q'_{ij} = \frac{1}{2}$ then $g(q_{ij}, q'_{ij}) = 1$. Additionally, $g(\frac{1}{2} + \eta, \frac{1}{2} + \eta) = g(\frac{1}{2} - \eta, \frac{1}{2} - \eta)$ and

$$\begin{aligned} g\left(\frac{1}{2} + \eta, \frac{1}{2} + \eta\right) &= 2^k \sum_{\ell=0}^k \left(\frac{1}{2} + \eta\right)^{2\ell} \sum_{\ell=0}^k \left(\frac{1}{2} - \eta\right)^{2k-2\ell} \\ &= 2^k \left(\frac{1}{2} + 2\eta^2\right)^k \sum_{\ell=0}^k \binom{k}{\ell} \left(\frac{1}{2} + \frac{\eta}{\frac{1}{2} + 2\eta^2}\right)^{\ell} \left(\frac{1}{2} + \frac{\eta}{\frac{1}{2} - 2\eta^2}\right)^{k-\ell} \\ &= (1 + 4\eta^2)^k. \end{aligned} \quad (50)$$

Also, note that $g(\frac{1}{2} + \eta, \frac{1}{2} - \eta) = g(\frac{1}{2} - \eta, \frac{1}{2} + \eta)$ and

$$\begin{aligned} g\left(\frac{1}{2} - \eta, \frac{1}{2} + \eta\right) &= 2^k \sum_{\ell=0}^k \binom{k}{\ell} \left(\frac{1}{2} - \eta\right)^{\ell} \left(\frac{1}{2} + \eta\right)^{k-\ell} \left(\frac{1}{2} + \eta\right)^{\ell} \left(\frac{1}{2} - \eta\right)^{k-\ell} \\ &= \left(\frac{1}{2} - 2\eta^2\right)^k \sum_{\ell=0}^k \binom{k}{\ell} \\ &= (1 - 4\eta^2)^k. \end{aligned} \quad (51)$$

Therefore, if the pair of matrices Q, Q' have b_1 agreements and b_2 disagreements, then using (50), (51) we get

$$\begin{aligned} G(Q, Q') &= g\left(\frac{1}{2} + \eta, \frac{1}{2} + \eta\right)^{b_1} g\left(\frac{1}{2} + \eta, \frac{1}{2} - \eta\right)^{b_2} \\ &= (1 + 4\eta^2)^{kb_1} (1 - 4\eta^2)^{kb_2} \\ &\leq (1 + 4\eta^2)^{k(b_1 - b_2)}. \end{aligned}$$

This proves Lemma 12.

6.3.2 PROOF OF PROPOSITION 4

In this section, we provide a construction of the null and the alternate in (1) under the model-free assumption that proves the statement of Proposition 4. To this end, let P_0 be a pairwise probability matrix with the $(i, j)^{th}$ element denoted by p_{ij} for all $i, j \in [d]$. We

will provide more details about P_0 subsequently. Consider the case where under the null $P = Q = P_0$ and under the alternate $P = P_0$ and $Q \sim \text{Bernoulli}(P_0)$. Under this notation, we have that under the alternate $q_{ij} \sim \text{Bernoulli}(p_{ij})$. We choose P_0 such that for all $i, j \in [d]$ we have $0 \leq p_{ij} \leq \frac{1}{2}$. With this, we argue that under the alternate construction, for any realization of Q , we have

$$\|P_0 - Q\|_F^2 \geq \sum_{i=1}^d \sum_{j=1}^d p_{ij}^2.$$

In this manner, by choosing an appropriate P_0 , we construct the alternate for any given ϵ which satisfy the conditions in the two-sample testing problem in (1). Note that the maximum value of $\|P_0 - Q\|_F^2$ attainable in this construction is when $p_{ij} = \frac{1}{2}$ for all $(i, j) \in [d]$. In this setting, $\|P_0 - Q\|_F^2 = \frac{d^2}{4}$, for all realizations of Q . Thus, in our construction, the parameter ϵ is at most $\frac{1}{2}$.

Now, in Proposition 4, we consider the case where we have one pairwise-comparison for each pair in each population, that is, $k_{ij}^p = k_{ij}^q = 1 \forall i, j \in [d]$. Recall that the observed matrices corresponding to the two populations are denoted by X and Y which are distributed as $X \sim \text{Bernoulli}(P)$ and $Y \sim \text{Bernoulli}(Q)$. Now, for our construction, we see that X is distributed identically under the null and the alternate as $X \sim \text{Bernoulli}(P_0)$, and Y is distributed as $Y \sim \text{Bernoulli}(P_0)$ under the null and $Y \sim \text{Bernoulli}(Q)$ under the alternate. Thus, to distinguish between the null and the alternate, we must be able to distinguish between the product distribution $\mathbb{P}_0 := \text{Bernoulli}(P_0) \times \text{Bernoulli}(P_0)$ and the product distribution $\mathbb{P}_1 := \text{Bernoulli}(P_0) \times \text{Bernoulli}(Q)$ where $Q \sim \text{Bernoulli}(P_0)$.

For the setting with one comparison per pair, we have access to only first order statistics for matrices X and Y . Since the Bernoulli parameters for all pairs (i, j) are independently chosen under the model-free setting, we look at the first order statistics of any pair (i, j) , which are given by $\Pr(X_{ij} = 1), \Pr(Y_{ij} = 1), \Pr(X_{ij} = 1, Y_{ij} = 1)$. Now, observe that under both the distributions \mathbb{P}_0 and \mathbb{P}_1 we have that

$$\Pr(X_{ij} = 1) = p_{ij}; \quad \Pr(Y_{ij} = 1) = p_{ij}; \quad \Pr(X_{ij} = 1, Y_{ij} = 1) = p_{ij}^2.$$

Since the first order statistics under both distributions \mathbb{P}_0 and \mathbb{P}_1 are identical, we conclude that no algorithm can distinguish between these distributions with a probability of error less than half. In turn, the minimax risk defined in (2) is at least half. This proves Proposition 4.

6.3.3 PROOF OF THEOREM 3

In this section, we establish a lower bound on the critical radius (3) for the two-sample testing problem defined in (1) under the assumption of the MST class as stated in Theorem 3. First, we provide a construction for the null and the alternate in Section 6.3.1. In this construction, we set $P = Q = P_0$ under the null and $P = P_0, Q \sim \text{Unif}(\Theta)$ under the alternate where Θ is a set of matrices belonging to the MST class. To complete the description of the construction, we now describe the set Θ for the MST class of pairwise-comparison probability matrices. The probability matrices in Θ correspond to a fixed ranking of items. Each matrix in Θ is such that the upper right quadrant has exactly one element in each row and each column equal to $\frac{1}{2} + \eta$ for some $\eta \in (0, \frac{1}{2})$. The rest of the elements

above the diagonal are half. The elements below the diagonal follow from the shifted-skew-symmetry condition imposed on MST probability matrices. It can be verified that all matrices $Q \in \Theta$ lie in the MST class. Note that the set Θ has $(d/2)!$ matrices. Since each matrix has a total of d elements equal to $\frac{1}{2} \pm \eta$, we get that $\frac{1}{d^2} \|P_0 - Q\|_{\mathbb{F}}^2 = \epsilon^2 = \frac{\eta^2}{d}$. This implies $\epsilon^2 \leq \frac{1}{4d}$.

Now, to derive bounds on the minimax risk according to (41), we analyse the χ^2 divergence between \mathbb{P}_0 and \mathbb{P}_Θ . From the analysis of $\chi^2(\mathbb{P}_0, \mathbb{P}_\Theta)$ in Section 6.3.1, specifically (47), we have that

$$\chi^2(\mathbb{P}_0, \mathbb{P}_\Theta) \leq \mathbb{E}_{(Q, Q') \sim \text{Unif}(\Theta)} \left[(1 + 4\eta^2)^{k(b_1 - b_2)} \right] - 1$$

where b_1 and b_2 are the number of agreements and disagreements between the matrices Q, Q' , as defined in (45). Now, to compute the χ^2 divergence, we want to find the probability that two matrices picked uniformly at random from Θ have i agreements in the upper right quadrant (the total number of agreements is $2i$ due to shifted-skew-symmetry). Given a matrix Q from set Θ , for i agreements, we want to choose a matrix $Q' \in \Theta$ such that exactly i of the perturbed elements share the same position as in Q . There are $\binom{d/2}{i}$ ways of choosing the i elements. Now that the i elements have their position fixed, we have to find the number of ways we can rearrange the remaining $\frac{d}{2} - i$ elements such that none of them share a position with the remaining perturbed elements in Q . This problem is the same as reshuffling and matching envelopes with letters such that no letter matches with originally intended envelope. The number of ways to rearrange a set of i objects in such a manner is given by $i! \left(\frac{1}{2!} - \frac{1}{3!} + \dots + (-1)^i \frac{1}{i!} \right)$. Thus the number of ways of rearrangement for $\frac{d}{2} - i$ items is upper bounded by $\frac{1}{2}(d/2 - i)!$. Thus, the probability of $2i$ agreements is upper bounded as

$$\mathbb{P}(b_1 = 2i) \leq \frac{(d/2)!}{(d/2 - i)!i!} \frac{(d/2 - i)!}{2(d/2)!} \leq \frac{1}{2(i)!}. \quad (52)$$

Then, we further simplify (47) as

$$\begin{aligned} \chi^2(\mathbb{P}_0, \mathbb{P}_\Theta) &\leq \mathbb{E}_{(Q, Q') \sim \text{Unif}(\Theta)} \left[(1 + 4\eta^2)^{k(b_1 - b_2)} \right] - 1 \\ &\leq \mathbb{E}_{(Q, Q') \sim \text{Unif}(\Theta)} \left[(1 + 4\eta^2)^{kb_1} \right] - 1 \\ &\leq \sum_{i=0}^{d/2} \mathbb{P}(b_1 = 2i) (1 + 4\eta^2)^{2ki} - 1. \end{aligned} \quad (53)$$

Notice that if we choose $k = \frac{c}{4\eta^2}$ with some constant $c \in (0, 1)$ then we have that

$$\begin{aligned} (1 + 4\eta^2)^k &= \sum_{\ell=0}^k \binom{k}{\ell} (4\eta^2)^\ell \\ &\leq 1 + \sum_{\ell=1}^k (4\eta^2 k)^\ell \\ &\leq 1 + \sum_{\ell=1}^k c^\ell \\ &\leq 1 + c', \end{aligned}$$

where c' is some positive constant. Using this, we show that the χ^2 divergence in (53) is upper bounded by a constant for $k \leq \frac{c}{\eta^2}$, as follows

$$\begin{aligned} \chi^2(\mathbb{P}_0, \mathbb{P}_\Theta) &\leq \sum_{i=0}^{d/2} \mathbb{P}(b_1 = 2i) (1 + 4\eta^2)^{2ki} - 1 \\ &\leq \sum_{i=0}^{d/2} \mathbb{P}(b_1 = 2i) (1 + c')^{2i} - 1 \\ &= \sum_{i=0}^{d/2} \mathbb{P}(b_1 = 2i) - 1 + \sum_{i=0}^{d/2} \mathbb{P}(b_1 = 2i) ((1 + c')^{2i} - 1) \\ &\stackrel{(a)}{\leq} \sum_{i=0}^{d/2} \frac{1}{2(i!)} ((1 + c')^{2i} - 1) \\ &\leq \frac{1}{2} \left(\exp((1 + c')^2) - \exp(1) + \sum_{i=d/2}^{\infty} \frac{1}{i!} \right) \\ &\leq c'', \end{aligned}$$

where c'' is some positive constant. The inequality (a) follows from (52). This proves that there exists a constant c , such that if $k \leq \frac{c}{\eta^2} = \frac{c}{d\epsilon^2}$, then the χ^2 divergence is upper bounded by $\frac{1}{9}$. According to (41), this implies that the minimax risk is at least $\frac{1}{3}$. This establishes the lower bound on the critical testing radius for two-sample testing under the MST modeling assumption as $\epsilon_{\mathcal{M}}^2 > c \frac{1}{kd}$ and proves Theorem 3.

6.3.4 PROOF OF THEOREM 5

Consider any arbitrary non-decreasing function $f : \mathbb{R} \rightarrow [0, 1]$ such that $f(\theta) = 1 - f(-\theta) \quad \forall \theta \in \mathbb{R}$. In order to prove the lower bound on testing radius stated in Theorem 5, we construct a set of matrices Θ based on the parameter-based pairwise-comparison probability model described in (4) associated to the given function f . Observe that $f(0) = \frac{1}{2}$. Recall

that under the parameter-based model the sum of all weights is fixed as $\sum_{i=1}^d w_i = 0$. We use the weight parameter to define the construction for the lower bound.

Recall the null and alternate construction described in Section 6.3.1 to prove the lower bound. Accordingly, we set $P = Q = P_0$ under the null and $P = P_0, Q \sim \text{Unif}(\Theta)$ under the alternate where Θ is a set of matrices belonging to the parameter-based class. The weights $w_{P_0} = [0, \dots, 0] \in \mathbb{R}^d$ correspond to the pairwise probability matrix $P_0 = [\frac{1}{2}]^{d \times d}$. Now for creating the set Θ , consider a collection of weight vectors w_Θ each with half the entries as δ and the other half as $-\delta$, thereby ensuring that $\sum_{i \in [d]} w_i = 0$. We set δ to ensure that each of the probability matrices induced by this collection of vectors obey $\frac{1}{d} \|P_0 - Q\|_F = \epsilon$. We define the set of matrices Θ as the set of pairwise-comparison probability matrices induced by the collection of values of w_Θ . Clearly, there are $\binom{d}{d/2}$ matrices in Θ . A similar argument holds for odd d wherein $\frac{d-1}{2}$ elements of the weight vector are δ and $\frac{d-1}{2}$ elements are $-\delta$. Since f is monotonic, $f(-2\delta) \leq f(0) \leq f(2\delta)$ and we have that $f(2\delta) = 1 - f(-2\delta)$, we define $f(-2\delta) = \frac{1}{2} - \eta$ and $f(2\delta) = \frac{1}{2} + \eta$ for some $0 < \eta \leq \frac{1}{2}$.

Similar to the proof of Theorem 3, we use (47) to bound the χ^2 divergence between the null and the alternate constructed. We first note that if we sample two matrices Q and Q' uniformly at random (with replacement) from Θ , then if the number of agreements is $\frac{i^2}{2}$ then the number of disagreements is equal to $\frac{(d-i)^2}{2}$. The probability of $\frac{i^2}{2}$ agreements is given by

$$\mathbb{P} \left(b_1 = \frac{i^2}{2}, b_2 = \frac{(d-i)^2}{2} \right) = \frac{\binom{d/2}{i/2} \binom{d/2}{i/2}}{\binom{d}{d/2}}.$$

Following from (44), the χ^2 divergence is

$$\begin{aligned} \chi^2(\mathbb{P}_0, \mathbb{P}_\Theta) &= \mathbb{E}_{(Q, Q') \sim \text{Unif}(\Theta)} \left[(1 + 4\eta^2)^{k(b_1 - b_2)} \right] - 1 \\ &\leq \sum_{i=0}^d \frac{\binom{d/2}{i/2} \binom{d/2}{i/2}}{\binom{d}{d/2}} (1 + 4\eta^2)^{k(i^2 - (d-i)^2)/2} - 1. \end{aligned}$$

For ease of presentation, we replace $\frac{d}{2}$ by z and $\frac{i}{2}$ by ℓ , to get

$$\begin{aligned} \chi^2(\mathbb{P}_0, \mathbb{P}_\Theta) &\leq \sum_{\ell=0}^z \frac{\binom{z}{\ell} \binom{z}{\ell}}{\binom{2z}{2z}} (1 + 4\eta^2)^{2k(2\ell z - z^2)} - 1 \\ &\leq \sum_{\ell=0}^z \left(\frac{1}{2} \right)^z \binom{z}{\ell} (1 + 4\eta^2)^{2k(2\ell z - z^2)} - 1 \\ &\leq \sum_{\ell=0}^z \left(\frac{1}{2} \right)^z \binom{z}{\ell} \exp(8\eta^2 k(2\ell z - z^2)) - 1. \end{aligned}$$

Here, we see that the summation in the final expression is equal to the expectation of $\exp(8\eta^2 k(2\ell z - z^2))$ over the random variable ℓ where $\ell \sim \text{Bin}(z, \frac{1}{2})$. So,

$$\begin{aligned} \chi^2(\mathbb{P}_0, \mathbb{P}_\Theta) &\leq \mathbb{E}_\ell \left[\exp(8\eta^2 k(2\ell z - z^2)) \right] - 1 \\ &\leq \sum_{i=0}^{\infty} \frac{(8\eta^2 k z)^i}{i!} \mathbb{E}_\ell \left[(2\ell - z)^i \right] - 1, \end{aligned} \tag{54}$$

where $\mathbb{E}_\ell[(2\ell - z)^i]$ is the scaled centered i^{th} moment of $\text{Bin}(z, \frac{1}{2})$. To get the expression for the centered moments, we first find the moment generating function of the random variable $\ell' = 2\ell - z$, as

$$\begin{aligned} \mathbb{E} \left[e^{(2\ell - z)t} \right] &= e^{-zt} \mathbb{E} \left[e^{2\ell t} \right] \\ &= e^{-zt} \sum_{\ell=0}^z \binom{z}{\ell} \left(\frac{1}{2} e^{2t} \right)^\ell \left(\frac{1}{2} \right)^{z-\ell} \\ &= e^{-zt} \left(\frac{1}{2} + \frac{1}{2} e^{2t} \right)^z \\ &= \left(\frac{e^{-t} + e^t}{2} \right)^z \\ &= (\cosh t)^z. \end{aligned}$$

Then, we have

$$\mathbb{E} [(2\ell - z)^i] = \left. \frac{d^i (\cosh t)^z}{dt^i} \right|_{t=0},$$

which is the i^{th} derivative of $(\cosh t)^z$ evaluated at $t = 0$. This leads to the fact that for odd i , $\mathbb{E}[(2\ell - z)^i] = 0$ and for even i , $\mathbb{E}[(2\ell - z)^i] \leq \frac{i!}{(i/2)!} z^{i/2}$. Using this with (54), we get

$$\chi^2(\mathbb{P}_0, \mathbb{P}_\Theta) \leq \sum_{i=1}^{\infty} c_i (64\eta^4 k^2 z^3)^i \quad (55)$$

where $c_i = \frac{1}{(i/2)!}$, that is, c_i is decreasing as i increases.

Thus, we see that if $k \leq \frac{c}{\eta^2 d^{3/2}}$ then there is a small enough c such that the χ^2 divergence is upper bounded by $\frac{1}{9}$. In this construction, we have $\epsilon^2 = \eta^2/2$, therefore, using (41), the lower bound for two-sample testing under the parameter-based modeling assumption is given as $\epsilon^2 = \Omega\left(\frac{1}{kd^{3/2}}\right)$. This proves Theorem 5.

6.3.5 PROOF OF THEOREM 6

To prove Theorem 6, we use the conjectured average-case hardness of the planted clique problem. In informal terms, the planted clique conjecture asserts that it is hard to detect the presence of a planted clique in an Erdős-Rényi random graph. In order to state it more precisely, let $G(d, \kappa)$ be a random graph on d vertices constructed in one of the following two ways:

H_0 : Every edge is included in $G(d, \kappa)$ independently with probability $\frac{1}{2}$

H_1 : Every edge is included in $G(d, \kappa)$ independently with probability $\frac{1}{2}$. In addition, a set of κ vertices is chosen uniformly at random and all edges with both endpoints in the chosen set are added to G .

The planted clique conjecture then asserts that when $\kappa = o(\sqrt{d})$, then there is no polynomial-time algorithm that can correctly distinguish between H_0 and H_1 with an error

probability that is strictly bounded below $\frac{1}{2}$. We complete the proof by identifying a subclass of SST matrices and showing that any testing algorithm that can distinguish between the subclass of SST matrices and the all half matrix, can also be used to detect a planted clique in an Erdős-Rényi random graph.

Consider the null with $P = Q = [\frac{1}{2}]^{d \times d}$ and the alternate such that $P = [\frac{1}{2}]^{d \times d}$ and Q is chosen uniformly at random from set Θ . The set of probability matrices Θ contains all $(d \times d)$ matrices with the upper left and lower right quadrant equal to all half, the upper right quadrant is all half except a $(\kappa \times \kappa)$ planted clique (i.e., a (κ, κ) submatrix with all entries equal to one). Then we have $\epsilon^2 = \kappa^2/2d^2$. The bottom left quadrant follows from the skew symmetry property. Recall that we observe one sample per pair of items ($i > j$). This testing problem is reduced to a goodness-of-fit testing problem as shown in (39) and (40).

Consider the set of $(\frac{d}{2} \times \frac{d}{2})$ matrices comprising the top right $(\frac{d}{2} \times \frac{d}{2})$ sub-matrix of every matrix in Θ . We claim that this set is identical to the set of all possible matrices in the planted clique problem with $\frac{d}{2}$ vertices and a planted clique of size κ . Indeed, the null contains the all-half matrix corresponding to the absence of a planted clique, and the alternate contains all symmetric matrices that have all entries equal to half except for a (κ, κ) all-ones sub-matrix corresponding to the planted clique. We choose the parameter $\kappa = \frac{\sqrt{d}}{\log \log(d)}$ so that any constant multiple of it will be within the hardness regime of planted clique (for sufficiently large values of d). Now, we leverage the planted clique hardness conjecture to state that the null in our construction cannot be distinguished from the alternate by any polynomial-time algorithm with probability of error less than $\frac{1}{2}$. This implies that for polynomial-time testing it is necessary that $\epsilon^2 \geq \frac{c}{d(\log \log(d))^2}$. This proves Theorem 6.

6.4 Proof of Theorem 7

Bounding the Type I error. In this proof, we first bound the Type I error and subsequently bound the Type II error. To bound the probability of error of Algorithm 3, we study the distribution of the test statistic T under the null and the alternate. Algorithm 3 uses the test statistic defined in (8). To understand the distribution of the test statistic, we first look at the distribution of X_{ij} and Y_{ij} .

In Algorithm 3, we break the partial ranks into disjoint pairwise-comparisons. Under the Plackett-Luce model disjoint pairwise-comparisons obtained from the same partial ranking are mutually independent. Additionally, since the Plackett-Luce model obeys the property of independence of irrelevant alternatives, the probability of observing item i being ranked ahead of item j in a partial ranking is independent of the other items being ranked in that partial ranking. Thus, for any pair of items (i, j) , the probability of i beating j , conditioned on the event that the pair (i, j) was observed, is always equal to p_{ij} for the population corresponding to pairwise probability matrix P and q_{ij} for the population corresponding to pairwise probability matrix Q . This holds true irrespective of which other items are involved in that partial (or total) ranking. With this in mind, we identify the distribution of X_{ij} conditioned on k_{ij}^p as $X_{ij} | k_{ij}^p \sim \text{Bin}(k_{ij}^p, p_{ij})$. Similarly, we have $Y_{ij} | k_{ij}^q \sim \text{Bin}(k_{ij}^q, q_{ij})$. Let $\mathbf{k}^p, \mathbf{k}^q$ denote the vector of k_{ij}^p, k_{ij}^q for all $(i, j), i < j$. The conditional expectation of T is

$$\mathbb{E}[T | \mathbf{k}^p, \mathbf{k}^q] = \sum_{i=1}^{j-1} \sum_{j=1}^d \frac{\mathbb{I}_{ij} k_{ij}^p k_{ij}^q}{k_{ij}^p + k_{ij}^q} (p_{ij} - q_{ij})^2.$$

Under the null we have $p_{ij} = q_{ij}$. Clearly, using the law of total expectation, we see that $\mathbb{E}_{H_0}[T] = \mathbb{E}[\mathbb{E}[T | \mathbf{k}^p, \mathbf{k}^q]] = 0$. We now upper bound the variance of T under the null. Recall from (23) and (24) that

$$\begin{aligned} \text{Var}_{H_0}[T] &\leq \sum_{i=1}^{j-1} \sum_{j=1}^d 8\mathbb{I}_{ij} \left(\frac{k_{ij}^p k_{ij}^q}{(k_{ij}^p + k_{ij}^q)} (p_{ij} - q_{ij})^2 + 3 \right) + \text{Var}[\mathbb{E}[T | \mathbf{k}^p, \mathbf{k}^q]] \\ &\leq 24d^2. \end{aligned} \tag{56}$$

Now, we have the information to bound the Type I error. To get a bound on the Type I error with threshold t , we use the one sided Chebyshev's inequality,

$$\mathbb{P}_{H_0}(T \geq t) \leq \frac{\text{Var}_{H_0}[T]}{\text{Var}_{H_0}[T] + t^2}. \tag{57}$$

Using the bound in (56) and (57), we observe that if $t = 11d$ then the Type I error is at most $\frac{1}{6}$. This concludes the proof that Algorithm 3 controls the Type I error of the test (13) at level $\frac{1}{6}$.

Bounding the Type 2 error. We now analyse the Type II error of Algorithm 3, that is, the probability of our algorithm failing to reject the null, under the alternate. We consider two cases depending on whether the pairwise-comparison data created through the rank breaking method has at least k pairwise-comparisons per pair (i, j) , $i < j$, or not, for some positive integer k . We will define k later in the proof. Let the case where the pairwise-comparisons created in each population have at least k comparisons of each pair be denoted by C_1 and let the associated Type II error be denoted by β_1 . Let the Type II error associated with the remaining case be denoted by β_2 . Our objective is to provide an upper bound on the total Type II error which is $\beta = \mathbb{P}(C_1)\beta_1 + (1 - \mathbb{P}(C_1))\beta_2$.

First, we derive a bound on $\mathbb{P}(C_1)$. To start, we note that the probability of observing a specific pair from a total ranking is $\frac{1}{d}$ if d is odd and $\frac{1}{d-1}$ if d is even. Recall that for a given m , each sample is a ranking of some m items chosen uniformly at random from the set of d items. Under this setting, we see that the probability that ‘‘Random disjoint’’ rank breaking yields a specific pairwise-comparison from a m -length partial ranking is $\frac{m}{d(d-1)}$ if m is even and $\frac{m-1}{d(d-1)}$ if m is odd. Henceforth, in this proof, we assume that m is even. The proof follows similarly for odd m . Thus, the number of pairwise-comparisons observed of any pair (i, j) is a binomial random variable with Bernoulli parameter $\frac{m}{d(d-1)}$. Consequently, if we have N samples from each population, then for the population corresponding to the pairwise probability matrix P , for all pairs (i, j) we have $k_{ij}^p \sim \text{Bin}(N, \frac{m}{d(d-1)})$. Similarly for the population corresponding to pairwise probability matrix Q , for all pairs (i, j) we have $k_{ij}^q \sim \text{Bin}(N, \frac{m}{d(d-1)})$. Now, we are equipped to compute the probability of case C_1 . We divide the samples available in each population into k sections of equal sizes. Let the samples in each population be indexed from 1 to N then we assign the first $\lfloor \frac{N}{k} \rfloor$ into the first section and so on. Now, we know that the probability of observing a pair (i, j) at least once in one such section is given by $1 - (1 - \frac{m}{d(d-1)})^{\lfloor \frac{N}{k} \rfloor}$. Using this, we get the following union bound,

$$\mathbb{P}(k_{ij}^p \geq k) \geq 1 - k \left(1 - \frac{m}{d(d-1)} \right)^{\frac{N}{k}}.$$

The same inequality holds for k_{ij}^q for all (i, j) . Then, the probability that all pairs of items had at least k pairwise-comparisons in both populations is lower bounded as

$$\begin{aligned} \mathbb{P}(C_1) &\geq 1 - \frac{2kd(d-1)}{2} \left(1 - \frac{m}{d^2}\right)^{\frac{N}{k}} \\ &\geq 1 - kd^2 \exp\left(-\frac{Nm}{kd^2}\right) \end{aligned}$$

We see that, if $N = ckd^2 \log(d)/m$ for some positive constant c , then $\mathbb{P}(C_1) \geq 1 - \frac{k}{d^{c-2}}$.

Conditioned on the case C_1 , we invoke Theorem 1 to control the Type II error. Recall that Theorem 1 asserts that there is a constant $c_0 > 0$ such that if we have k pairwise-comparisons of each pair (i, j) from each population where $k \geq \max\{c_0 \frac{1}{d\epsilon^2}, 2\}$ then the Type II error of Algorithm 1 for the testing problem (13) is upper bounded by $\frac{1}{12}$. To apply this result to Algorithm 3 conditioned on case C_1 , we keep k pairwise-comparisons for each pair where $k = 2\lceil c_0 \frac{1}{d\epsilon^2} \rceil$. Observe that, since we assume $\epsilon \geq c_1 d^{-c_2}$ for some positive constants c_1 and c_2 , we get the inequality $k \leq c'_1 d^{c'_2}$ for some positive constants c'_1 and c'_2 . Under this inequality, we get that there exist positive constants c, c'_1 and c'_2 such that $\mathbb{P}(C_1) > 11/12$.

Next, we observe that the Type II error conditioned on the complement of C_1 is at most 1. Therefore, the total probability of failing to reject the null under the alternate is given by

$$\begin{aligned} \beta &= \mathbb{P}(C_1)\beta_1 + (1 - \mathbb{P}(C_1))\beta_2 \\ &\leq \frac{1}{12} + \frac{1}{12} = \frac{1}{6}. \end{aligned}$$

This concludes the proof that for some constant $C > 0$, if $N \geq C \frac{d^2 \log(d)}{m} \lceil \frac{c_0}{d\epsilon^2} \rceil$, then the probability of error of Algorithm 3 is at most $\frac{1}{3}$.

6.5 Proof of Theorem 8

The proof of Theorem 8 follows similarly to the proof of Theorem 7. Both theorems establish the performance of Algorithm 3 for the two-sample testing problem stated in (13). The difference lies in the assumption on the partial ranking data available and consequently in the rank breaking algorithm used. In Theorem 8, we assume we have total ranking data that is then deterministically converted to pairwise-comparisons in the following manner. We have a total of N total rankings available from each population. We divide these into subsets each containing d rankings as described in the ‘‘Deterministic disjoint’’ rank breaking method. Notice that we can break the ranking data available in a section into pairwise-comparisons such that we observe each unique pair of items at least one time. We prove this statement at the end of the section. We repeat this breaking technique for all subsets. Consequently, we get $k = \lfloor \frac{N}{d} \rfloor \geq 2\lceil c \frac{1}{d\epsilon^2} \rceil$ pairwise-comparisons for all pairs (i, j) from each population. With this in mind, we apply Theorem 1 to obtain the desired result.

Finally, to complete the proof we show that it is indeed possible to break d total rankings such that we observe each of $\binom{d}{2}$ unique pairs at least once. We use a mathematical induction

based argument. As a first step we observe that our hypothesis is true for $d = 2$. In the inductive step, we assume that the hypothesis is true for all natural numbers $d \in \{2, \dots, r\}$. Now, we wish to prove the hypothesis is true for $d = r + 1$. First, consider the case where r is even. We divide the set of r items into two groups with $r/2$ items in each. From the inductive step we know that our hypothesis is true for $d = r/2$. Consequently, we get $2^{\binom{r/2}{2}}$ unique pairs from within the two groups which use $r/2$ total rankings. Next, we arrange the items in group one in a list against the items in group two and make pairs by choosing the items in i^{th} position in both the lists. This gives the breaking for one total ranking. We do this $r/2$ times, each time cyclically shifting the first list by one item. This step gives $r^2/4$ unique pairs that are different from the pairs obtained in the previous step and uses $r/2$ total rankings. This proves our hypothesis for $d = r$ for even r . To prove our hypothesis for odd r , we prove our hypothesis for $r + 1$ which is even using the same method described in the previous step. We complete our proof by noting that if the hypothesis is true for even r then it must be true for $r - 1$. This concludes our proof of Theorem 8

6.6 Proof of Theorem 9

The idea of the proof is to show that under the null hypothesis, a ranking sample sourced from the first population is mutually independent of and identically distributed as a ranking sample sourced from the second population. If this statement is true, then shuffling the population labels of ranking data, does not alter the distribution of the test statistic (8). This in turn controls the Type I error of the permutation test method.

Under the null, for some distribution λ over all total rankings, we have that $\lambda_P = \lambda_Q = \lambda$. This implies that under the marginal probability based model, the probability of any given partial ranking over a set of items is the same for both the populations. Specifically, conditioned on the set of items being ranked, each partial ranking in each population is sampled independently and identically, according to the distribution λ . Recall that the set of items being ranked in each population is sampled independently and identically from some distribution over all non-empty subsets in $[d]$. Consequently, each ranking sample is independent of all other ranking samples obtained from the two populations. Moreover, using the law of total probability over all the non-empty subsets in $[d]$, we get that each ranking sample obtained in each population is identically distributed. With this, we conclude that shuffling the population labels of ranking data does not alter the distribution of the test statistic. Thus, for a permutation test with γ iterations, the p -value of the test is distributed uniformly over $\{0, 1/\gamma, 2/\gamma, \dots, 1\}$. Hence, for any given significance level $\alpha \in (0, 1)$, by applying a threshold of α on the p -value of the test, we are guaranteed to have Type I error at most α .

7. Discussion and open problems

We conclude with a discussion focused on open problems in this area. We provide algorithms for two-sample testing on pairwise-comparison and ranking data to distinguish between two potentially different populations in terms of their underlying distributions. Through our analysis, we see that our testing algorithm for pairwise-comparison data is simultaneously minimax optimal under the model-free setting as well as the MST and WST model. There is a gap between the testing rate of our algorithm and our information-theoretic lower bound

for the SST and parameter-based models, and closing this gap is an open problem of interest. In addition, our lower bound does not consider the random sampling regime we address in Corollary 2, thus, obtaining a lower bound in this regime is another open problem of interest. Second, in the future, our work may help in studying two-sample testing problems pertaining to more general aspects of data from people such as function evaluations (Xu et al., 2020; Noothigattu et al., 2018), issues of calibration (Wang and Shah, 2019), and strategic behavior (Xu et al., 2019). Thirdly, in practice we use the permutation method to calibrate our tests, ensuring valid Type I error control even when the distribution of the test statistic is difficult to characterize analytically (for instance, in the setting with partial rank data). Understanding the power of tests calibrated via the permutation method is an active area of research (Kim et al., 2020a) and it would be interesting to understand this in the context of the tests developed in our work. Finally, the literature on analyzing pairwise-comparison data builds heavily on probability models from social choice theory (some are described in this work). A natural related question, that has received some recent attention in Seshadri and Ugander (2019), is the design of goodness-of-fit hypothesis tests to test whether given pairwise-comparison data obeys certain modeling assumptions.

Acknowledgments

This work was supported in part by NSF grants DMS 1713003, CCF 1763734, 1755656 and CAREER 1942124.

Appendix A. Additional details of experiments

We now provide more details about the experiments described in Section 5.1.2.

Ordinal versus cardinal

The data set from Shah et al. (2016) used in the “Ordinal versus cardinal” experiment comprises of six different experiments on Amazon Mechanical Turk crowdsourcing platform. We describe each experiment briefly here.

- Photo age: There are 10 objects in this experiment wherein each object is a photograph of a different face. The worker is either shown pairs of photos together and asked to identify the older of the two or they provide the numeric age for each photo. There are a total of 225 ordinal responses and 275 cardinal-converted-to-ordinal responses.
- Spelling mistakes: There are 8 objects in this experiment wherein each object is a paragraph of text in English possibly with some spelling mistakes. The worker is either shown pairs of paragraphs and asked to identify the paragraph with more spelling mistakes or they are asked to provide the count of spelling mistakes for all 8 paragraphs. There are a total of 184 ordinal responses and 204 cardinal-converted-to-ordinal responses.
- Distances between cities: There are 16 objects in this experiment wherein each object is a pair of cities (no two objects share a common city). The worker is either shown two pairs of cities at a time and asked to identify the pair that is farther from each other, or they are asked to estimate the distances for the 16 pairs of cities. There are a total of 408 ordinal responses and 392 cardinal-converted-to-ordinal responses.
- Search results: There are 20 objects in this experiment wherein each object is the result of an internet based search query of the word "internet". The worker is either asked to compare pairs of results based on their relevance or they are shown all the results and asked to rate the relevance of each result on a scale of 0-100. There are a total of 630 ordinal responses and 370 cardinal-converted-to-ordinal responses.
- Taglines: There are 10 objects in this experiment wherein each object is a tagline for a product described to the worker. The worker is either asked to compare the quality of pairs of taglines or they are asked to provide ratings for each tagline on a scale of 0-10. There are a total of 305 ordinal responses and 195 cardinal-converted-to-ordinal responses.
- Piano : There are 10 objects in this experiment wherein each object is a sound clip of a piano key played at a certain frequency. The worker is either given pairs of sound clips and asked to identify the clip with the higher frequency or they are asked to estimate the frequency of the 10 clips. There are a total of 265 ordinal responses and 235 cardinal-converted-to-ordinal responses.

In our main experiment, we combine the data from all the experiments described above and test for statistically significant difference between the underlying distributions for ordinal responses and ordinal-converted-to-cardinal responses. We also test for difference in each individual experiment (which however have smaller sample sizes), the results are provided in

Table 2. We observe that the qualitatively more subjective experiments (photo age, search results, taglines) have a lower p -value, which indicates that the ordinal responses are more different from cardinal-converted-to-ordinal responses in a more subjective setting.

Experiment	Combined	Age	Spellings	Distances	Search results	Taglines	Piano
p -value	0.003	0.001	0.657	0.75	0.187	0.0829	0.514

Table 2: p -value of two-sample test comparing the distribution of ordinal responses and the distribution of cardinal-converted-to-ordinal responses in the experiments described above

European football leagues

In this experiment, we obtain the match scores for four different European football leagues (English Premier League, Bundesliga, Ligue 1, La Liga) across two seasons (2016-2017, 2017-2018). There were 17 teams that played the two seasons in each of EPL, La Liga, Ligue 1 and 16 teams in Bundesliga. To test for statistically significant difference between the relative performance of the participating teams in the two consecutive seasons we combined the data from all four leagues. We also tested for difference in each individual league, the results are displayed in Table 3. From the 2016-2017 season we have 202 pairwise-comparisons in EPL, 170 pairwise-comparisons in Bundesliga, 215 pairwise-comparisons in La Liga, and 201 pairwise-comparisons in Ligue 1. From the 2017-2018 season we have 214 pairwise-comparisons in EPL, 178 pairwise-comparisons in Bundesliga, 208 pairwise-comparisons in La Liga, and 201 pairwise-comparisons in Ligue 1. From the number of comparisons available our test does not detect any significant difference between the relative performance of teams in European football leagues over two consecutive seasons.

League	Combined	EPL	Bundesliga	La Liga	Ligue 1
p -value	0.971	0.998	0.691	0.67	0.787

Table 3: p -value of two-sample test comparing relative performance of teams in a football league over two consecutive seasons.

References

- Nir Ailon. An active learning algorithm for ranking from pairwise preferences with an almost optimal query complexity. *Journal of Machine Learning Research*, 13:137–164, 2012.
- David Aldous. Elo ratings and the sports model: A neglected topic in applied probability? *Statistical Science*, 32(4):616–629, 2017.
- Sivaraman Balakrishnan and Larry Wasserman. Hypothesis testing for high-dimensional multinomials: A selective review. *Ann. Appl. Stat.*, 12(2):727–749, 2018.
- Sivaraman Balakrishnan and Larry Wasserman. Hypothesis testing for densities and high-dimensional multinomials: Sharp local minimax rates. *Ann. Statist.*, 47(4):1893–1927, 08 2019.

- Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*, pages 324–345, 1952.
- Mark Braverman and Elchanan Mossel. Noisy sorting without resampling. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 268–276, 2008.
- Alexandra Carpentier, Olivier Collier, Laëtitia Comminges, Alexandre B Tsybakov, and Yuhao Wang. Minimax rate of testing in sparse linear regression. *arXiv preprint arXiv:1804.06494*, 2018.
- Manuela Cattelan, Cristiano Varin, and David Firth. Dynamic Bradley–Terry modelling of sports tournaments. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 62(1):135–150, 2013.
- Daniel R Cavagnaro and Clinton P Davis-Stober. Transitive in our preferences, but transitive in different ways: An analysis of choice variability. *Decision*, 1(2):102, 2014.
- Siu On Chan, Ilias Diakonikolas, Paul Valiant, and Gregory Valiant. Optimal algorithms for testing closeness of discrete distributions. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1193–1203. SIAM, 2014.
- S. Chatterjee and S. Mukherjee. Estimation in tournaments and graphs under monotonicity constraints. *IEEE Transactions on Information Theory*, 65(6):3525–3539, 2019.
- Xi Chen, Sivakanth Gopi, Jieming Mao, and Jon Schneider. Optimal instance adaptive algorithm for the top- k ranking problem. *IEEE Transactions on Information Theory*, 64(9):6139–6160, 2018.
- Yuxin Chen and Changho Suh. Spectral MLE: Top- k rank aggregation from pairwise comparisons. In *International Conference on Machine Learning*, pages 371–380, 2015.
- Olivier Collier, Laëtitia Comminges, and Alexandre B Tsybakov. Minimax estimation of linear and quadratic functionals on sparsity classes. *The Annals of Statistics*, 45(3):923–958, 2017.
- Moein Falahatgar, Alon Orlitsky, Venkatadheeraj Pichapati, and Ananda Theertha Suresh. Maximum selection and ranking under noisy comparisons. In *34th International Conference on Machine Learning*, pages 1088–1096, 2017.
- Vivek F Farias, Srikanth Jagabathula, and Devavrat Shah. A nonparametric approach to modeling choice with limited data. *Management science*, 59(2):305–322, 2013.
- A. Gretton, K. Borgwardt, M. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13:723–773, March 2012a.
- Arthur Gretton, Dino Sejdinovic, Heiko Strathmann, Sivaraman Balakrishnan, Massimiliano Pontil, Kenji Fukumizu, and Bharath K. Sriperumbudur. Optimal kernel choice for large-scale two-sample tests. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25, pages 1205–1213. Curran Associates, Inc., 2012b.

- John Guiver and Edward Snelson. Bayesian inference for plackett-luce ranking models. In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09*, page 377–384, New York, NY, USA, 2009. Association for Computing Machinery.
- Reinhard Heckel, Nihar B Shah, Kannan Ramchandran, and Martin J Wainwright. Active ranking from pairwise comparisons and when parametric assumptions do not help. *The Annals of Statistics*, 47(6):3099–3126, 2019.
- Ralf Herbrich, Tom Minka, and Thore Graepel. Trueskill: A Bayesian skill rating system. In *Advances in neural information processing systems*, pages 569–576, 2007.
- Lars Magnus Hvattum and Halvard Arntzen. Using Elo ratings for match result prediction in association football. *International Journal of forecasting*, 26(3):460–470, 2010.
- Yuri I. Ingster. Minimax detection of a signal in ℓ_p metrics. *Journal of Mathematical Sciences*, 68(4):503–515, 1994.
- Yuri I. Ingster. Adaptive chi-square tests. *Zapiski Nauchnykh Seminarov POMI*, 244:150–166, 1997.
- Yuri I. Ingster and Irina A. Suslina. *Nonparametric Goodness-of-Fit Testing Under Gaussian Models*. Lecture Notes in Statistics. Springer, 2003.
- Mark Jerrum. Large cliques elude the metropolis process. *Random Struct. Algorithms*, 3: 347–360, 1992.
- Toshihiro Kamishima. Nantonac collaborative filtering: recommendation based on order responses. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 583–588, 2003.
- Ilmun Kim, Sivaraman Balakrishnan, and Larry Wasserman. Minimax optimality of permutation tests, 2020a.
- Ilmun Kim, Sivaraman Balakrishnan, and Larry Wasserman. Robust multivariate nonparametric tests via projection-averaging. *To appear in The Annals of Statistics*, 2020b.
- Luděk Kučera. Expected complexity of graph partitioning problems. *Discrete Appl. Math.*, 57(2-3):193–212, February 1995. ISSN 0166-218X.
- Alec Lamon, Dave Comroe, Peter Fader, Daniel McCarthy, Rob Ditto, and Don Huesman. Making WHOOPPEE: A collaborative approach to creating the modern student peer assessment ecosystem. In *EDUCAUSE*, 2016.
- E. L. Lehmann and Joseph P. Romano. *Testing statistical hypotheses*. Springer Texts in Statistics. Springer, third edition, 2005.
- R Duncan Luce. *Individual choice behavior: A theoretical analysis*. New York: Wiley, 1959.
- Horia Mania, Aaditya Ramdas, Martin J Wainwright, Michael I Jordan, and Benjamin Recht. On kernel methods for covariates that are rankings. *Electronic Journal of Statistics*, 12(2): 2537–2577, 2018.

- Lucas Maystre and Matthias Grossglauser. Fast and accurate inference of plackett-luce models. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'15, page 172–180, Cambridge, MA, USA, 2015. MIT Press.
- Sahand Negahban, Sewoong Oh, and Devavrat Shah. Iterative ranking from pair-wise comparisons. In *Advances in neural information processing systems*, pages 2474–2482, 2012.
- Ritesh Noothigattu, Nihar B Shah, and Ariel D Procaccia. Loss functions, axioms, and peer review. *arXiv preprint arXiv:1808.09057*, 2018.
- Robin L. Plackett. The analysis of permutations. *Journal of the Royal Statistical Society*, 24(2):193–202, 1975.
- Arun Rajkumar and Shivani Agarwal. A statistical convergence perspective of algorithms for rank aggregation from pairwise data. In *International Conference on Machine Learning*, pages 118–126, 2014.
- Arun Rajkumar, Suprovat Ghoshal, Lek-Heng Lim, and Shivani Agarwal. Ranking from stochastic pairwise preferences: Recovering condorcet winners and tournament solution sets at the top. In *International Conference on Machine Learning*, pages 665–673, 2015.
- Karthik Raman and Thorsten Joachims. Methods for ordinal peer grading. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1037–1046. ACM, 2014.
- C. Rastogi, S. Balakrishnan, N. Shah, and A. Singh. Two-sample testing on pairwise comparison data and the role of modeling assumptions. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pages 1271–1276, 2020.
- Michel Regenwetter, Jason Dana, and Clinton P Davis-Stober. Transitivity of preferences. *Psychological review*, 118(1):42, 2011.
- Paul R. Rosenbaum. An exact distribution-free test comparing two multivariate distributions based on adjacency. *Journal of the Royal Statistical Society Series B*, 67:515–530, 09 2005.
- Arjun Seshadri and Johan Ugander. Fundamental limits of testing the independence of irrelevant alternatives in discrete choice. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, page 65?66, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450367929.
- Nihar B. Shah and Martin J. Wainwright. Simple, robust and optimal ranking from pairwise comparisons. *Journal of Machine Learning Research*, 18(199):1–38, 2018.
- Nihar B Shah, Joseph K Bradley, Abhay Parekh, Martin Wainwright, and Kannan Ramchandran. A case for ordinal peer-evaluation in moocs. In *NeurIPS Workshop on Data Driven Education*, pages 1–8, 2013.

- Nihar B Shah, Sivaraman Balakrishnan, Joseph Bradley, Abhay Parekh, Kannan Ramchandran, and Martin J Wainwright. Estimation from pairwise comparisons: Sharp minimax bounds with topology dependence. *The Journal of Machine Learning Research*, 17(1):2049–2095, 2016.
- Nihar B Shah, Sivaraman Balakrishnan, Adityanand Guntuboyina, and Martin J Wainwright. Stochastically transitive models for pairwise comparisons: Statistical and computational issues. *IEEE Transactions on Information Theory*, 63(2):934–959, 2017.
- Nihar B Shah, Behzad Tabibian, Krikamol Muandet, Isabelle Guyon, and Ulrike Von Luxburg. Design and analysis of the NIPS 2016 review process. *The Journal of Machine Learning Research*, 19(1):1913–1946, 2018.
- Gabor J. Szekely and Maria L. Rizzo. Testing for equal distributions in high dimensions. *InterStat*, 2004.
- Balázs Szörényi, Róbert Busa-Fekete, Adil Paul, and Eyke Hüllermeier. Online rank elicitation for Plackett-Luce: A dueling bandits approach. In *Advances in Neural Information Processing Systems*, pages 604–612, 2015.
- Louis L Thurstone. A law of comparative judgment. *Psychological Review*, 34(4):273, 1927.
- G. Valiant and P. Valiant. An automatic inequality prover and instance optimal identity testing. *SIAM Journal on Computing*, 46(1):429–455, 2017.
- Paul Valiant. Testing symmetric properties of distributions. *SIAM Journal on Computing*, 40(6):1927–1968, December 2011. ISSN 0097-5397.
- Han LJ Van Der Maas and Eric-Jan Wagenmakers. A psychometric analysis of chess expertise. *American Journal of Psychology*, 118(1):29–60, 2005.
- Jingyan Wang and Nihar B Shah. Your 2 is my 1, your 3 is my 9: Handling arbitrary miscalibrations in ratings. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 864–872. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
- Yichong Xu, Han Zhao, Xiaofei Shi, and Nihar B. Shah. On strategyproof conference peer review. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 616–622. International Joint Conferences on Artificial Intelligence Organization, 7 2019.
- Yichong Xu, Sivaraman Balakrishnan, Aarti Singh, and Artur Dubrawski. Regression with comparisons: Escaping the curse of dimensionality with ordinal information. *Journal of Machine Learning Research*, 21(162):1–54, 2020.