JMLR: Workshop and Conference Proceedings vol 49:1–26, 2016

When Can We Rank Well from Comparisons of $O(n \log n)$ Non-Actively Chosen Pairs?

Arun Rajkumar ARUN_R@CSA.IISC.ERNET.IN
Xerox Research, Bangalore, India

Shivani Agarwal SHIVANI@CSA.IISC.ERNET.IN
Radcliffe Institute for Advanced Study, Harvard University, Cambridge, MA, USA
Indian Institute of Science, Bangalore, India

Abstract

Ranking from pairwise comparisons is a ubiquitous problem and has been studied in disciplines ranging from statistics to operations research and from theoretical computer science to machine learning. Here we consider a general setting where outcomes of pairwise comparisons between items i and j are drawn probabilistically by flipping a coin with unknown bias P_{ij}, and ask under what conditions on these unknown probabilities one can learn a good ranking from comparisons of only $O(n \log n)$ non-actively chosen pairs. Recent work has established this is possible under the Bradley-Terry-Luce (BTL) and noisy permutation (NP) models. Here we introduce a broad family of 'low-rank' conditions on the probabilities P_{ij} under which the resulting preference matrix $\mathbf{P}$ has low rank under some link function, and show these conditions encompass the BTL and Thurstone classes as special cases, but are considerably more general. We then give a new algorithm called low-rank pairwise ranking (LRPR) which provably learns a good ranking from comparisons of only $O(n \log n)$ randomly chosen comparisons under such low-rank models. Our algorithm and analysis make use of tools from the theory of low-rank matrix completion, and provide a new perspective on the problem of ranking from pairwise comparisons in non-active settings.

Keywords: Ranking; pairwise comparisons; low-rank matrix completion.

1. Introduction

Ranking from pairwise comparisons is a ubiquitous problem and has been studied in disciplines ranging from statistics to operations research and from theoretical computer science to machine learning (Thurstone, 1927; Bradley and Terry, 1952; Luce, 1959; Kendall, 1955; Saaty, 1980; David, 1988; Keener, 1993; Ailon et al., 2008; Braverman and Mossel, 2008; Jiang et al., 2011; Lu and Boutilier, 2011; Gleich and Lim, 2011; Ailon, 2012; Jamieson and Nowak, 2011; Yue et al., 2012; Negahban et al., 2012; Wauthier et al., 2013; Busa-Fekete and Hüllermeier, 2014; Rajkumar and Agarwal, 2014; Rajkumar et al., 2015; Shah et al., 2015). Here we are interested in the following general question: There are n items and an unknown preference matrix $\mathbf{P} \in [0,1]^{n \times n}$, with $P_{ij} + P_{ji} = 1 \; \forall i,j$, such that whenever items i and j are compared, item i beats item j with probability P_{ij} and j beats i with probability $P_{ij} = 1 - P_{ij}$. Given the ability to make comparisons among only $O(n \log n)$ pairs, which can be compared more than once but must be chosen *non-actively*, i.e. before observing the outcomes of any comparisons, under what conditions on $\mathbf{P}$ can we learn a good ranking of the n items? This question is important in any pairwise comparison setting where the number of items n is large and decisions on which pairs to compare cannot be made adaptively, as is often the case for example in consumer surveys.

Traditional sorting algorithms in computer science, which apply when $\mathbf{P}$ corresponds to a full deterministic ordering of the n items, i.e. when there exists an ordering or permutation $\sigma \in \mathcal{S}_n$ such that $P_{ij} = 1$ for all i, j with $\sigma(i) < \sigma(j)$, require the $O(n \log n)$ pairs to be chosen in an active manner (some pairs can be chosen only after observing the comparison outcomes of previous pairs). Braverman and Mossel (2008, 2009) showed that if $\mathbf{P}$ follows a *noisy permutation* (NP) model, i.e. if there is a permutation $\sigma \in \mathcal{S}_n$ and a noise parameter $p \in [0, \frac{1}{2})$ such that $P_{ij} = 1 - p$ for all i, j with $\sigma(i) < \sigma(j)$, then one can use *noisy sorting* algorithms to learn a ranking close to σ by observing comparisons of only $O(n \log n)$ pairs (using only one comparison per pair), but again, these algorithms require active selection of pairs. Similarly, the embedding-based algorithms of Jamieson and Nowak (2011), which apply when $\mathbf{P}$ corresponds to either a full deterministic ordering of the n items or a noisy permutation as above, but with the further restriction that the associated permutation can be realized by an embedding of the n items in d dimensions, involve comparisons of only $O(d \log n)$ and $O(d \log^2 n)$ pairs, respectively, but also require these pairs to be chosen actively. Ailon (2012) gives a decomposition-based algorithm that applies when $\mathbf{P}$ corresponds to a deterministic tournament, i.e. when $P_{ij} \in \{0, 1\}\ \forall i \neq j$; the algorithm compares $O(n\, \mathrm{poly}(\log n))$ actively chosen pairs. The dueling bandits literature (e.g. see the recent survey by Busa-Fekete and Hüllermeier (2014)) also involves comparing actively chosen pairs.

On the other hand, statistical approaches to ranking from pairwise comparisons generally start with an observed sample of randomly drawn pairwise comparisons, and infer a ranking from these; no active selection of pairs is involved. For example, two widely-studied statistical models for pairwise comparisons are the *Thurstone* model, under which $\mathbf{P}$ is parametrized by a score vector $\mathbf{s} \in \mathbb{R}^n$ such that $P_{ij} = \Phi(s_i - s_j)\ \forall i, j$ (where $\Phi(\cdot)$ denotes the standard normal CDF), and the *Bradley-Terry-Luce* (BTL) model, under which $\mathbf{P}$ is parametrized by a score vector $\mathbf{w} \in (0, \infty)^n$ such that $P_{ij} = \frac{w_i}{w_i + w_j}\ \forall i, j$; indeed, there has been much work on developing algorithms for maximum likelihood estimation (MLE) of the parameters $\mathbf{s}$ or $\mathbf{w}$ from observed pairwise comparisons drawn according to these models, which can then be used to rank the n items by sorting them in descending order of the estimated scores (Thurstone, 1927; Bradley and Terry, 1952; Luce, 1959; Hunter, 2004). However, most analyses of these algorithms implicitly assume all $\binom{n}{2}$ pairs are compared (e.g. see recent work by Rajkumar and Agarwal (2014) for an analysis of BTL maximum likelihood estimation and other algorithms when all $\binom{n}{2}$ pairs are compared).[1]

Recently, Negahban et al. (2012) proposed a spectral ranking algorithm termed Rank Centrality, and showed that if the underlying preference matrix $\mathbf{P}$ follows a BTL model as above, then comparing only $O(n \log n)$ randomly chosen pairs, each $O(\log n)$ times, is sufficient to ensure that the ranking of the n items learned by the Rank Centrality algorithm is close to their ranking under the BTL score vector. Somewhat along similar lines, Wauthier et al. (2013) analyzed a simple algorithm termed Balanced Rank Estimation (BRE) that ranks items by their observed Borda scores, and showed that if $\mathbf{P}$ follows a noisy permutation (NP) model (see above), then comparing only $O(n \log n)$ randomly chosen pairs (in this case using just one comparison per pair) suffices to ensure that the ranking learned by BRE is close to the underlying NP permutation σ. The BTL and NP classes are both natural but relatively limited classes of pairwise comparison models (see Figure 1); our interest here is in understanding under what other classes of preference matrices $\mathbf{P}$ one can

1. We conjecture that for MLE under BTL (and possibly also under other parametric models for pairwise comparisons), comparisons of $O(n \log n)$ pairs may suffice to obtain a good ranking. However we do not investigate this issue here as it is beyond the scope of our work. Moreover, the algorithms we propose here provide an alternative to MLE for estimating a good ranking from comparisons of $O(n \log n)$ pairs under BTL, Thurstone, and other models.

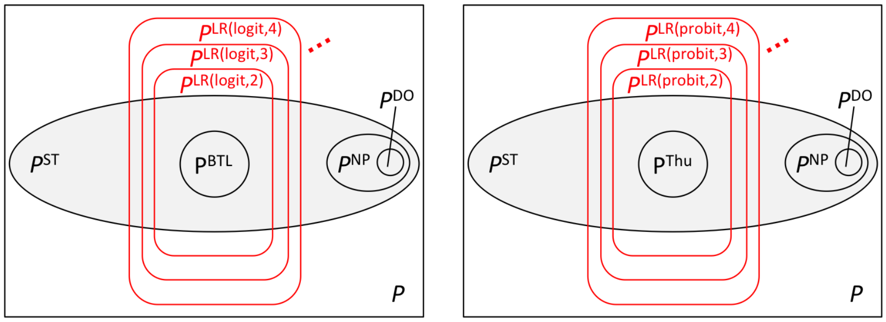

Figure 1: Classes of pairwise preference matrices. See Section 2.2 for detailed definitions.

Table 1: Summary of our results in relation to previous work.

Algorithm	**# Pairs Compared**	**# Comparisons per Pair**	**Class of Models P**	**Selection of Pairs**
Sorting (Computer Science)	$O(n \log n)$	1	$\mathcal{P}^{\text{DO}}$	Active
Noisy Sorting (Braverman and Mossel, 2008)	$O(n \log n)$	1	$\mathcal{P}^{\text{NP}}$	Active
Sorting under Embedding (Jamieson and Nowak, 2011)	$O(d \log n)$	1	$\mathcal{P}^{\text{DO}(\mathbf{X})}$ ($\mathbf{X} \in \mathbb{R}^{d \times n}$)	Active
Noisy Sorting under Embedding (Jamieson and Nowak, 2011)	$O(d \log^2 n)$	1	$\mathcal{P}^{\text{NP}(\mathbf{X})}$ ($\mathbf{X} \in \mathbb{R}^{d \times n}$)	Active
Decomposition-Based Algorithm (Ailon, 2012)	$O(n\, \text{poly}(\log n))$	1	$\mathcal{P}^{\text{DTour}}$	Active
MLE (Statistics)	typically $O(n^2)$	multiple	$\mathcal{P}^{\text{BTL}}$, $\mathcal{P}^{\text{Thu}}$, etc	Non-active
Rank Centrality (Negahban et al., 2012)	$O(n \log n)$	$O(\log n)$	$\mathcal{P}^{\text{BTL}}$	Non-active
BRE (Wauthier et al., 2013)	$O(n \log n)$	1	$\mathcal{P}^{\text{NP}}$	Non-active
This work	$O(nr \log n)$	$O(r \log n)$	$\mathcal{P}^{\text{LR}(\psi,r)} \cap \mathcal{P}^{\text{ST}}$	Non-active

learn a good ranking from comparisons of only $O(n \log n)$ non-actively chosen pairs. The quality of a learned ranking will be measured in terms of the number of pairwise disagreements w.r.t. the underlying preference matrix $\mathbf{P}$; since minimizing the number of pairwise disagreements in general is NP-hard even under knowledge of the exact matrix $\mathbf{P}$ (Alon, 2006; Ailon et al., 2008), we will restrict our attention to pairwise preferences satisfying the *stochastic transitivity* (ST) condition, namely $P_{ij} > \frac{1}{2}, P_{jk} > \frac{1}{2} \implies P_{ik} > \frac{1}{2}$, for which minimizing the number of pairwise disagreements under knowledge of the exact matrix $\mathbf{P}$ can be done efficiently (e.g. by running topological sort on the induced acyclic pairwise preference graph). We note that the previously studied classes of BTL and NP preferences also satisfy the ST condition (see Figure 1).

In this paper, we define broad classes $\mathcal{P}^{\text{LR}(\psi,r)}$ of 'low-rank' preference matrices $\mathbf{P}$ that have rank at most r under a suitable link transform $\psi : [0,1] \to \mathbb{R}$, and design a family of *low-rank pairwise ranking* (LRPR) algorithms which, for stochastically transitive preference matrices $\mathbf{P}$ in $\mathcal{P}^{\text{LR}(\psi,r)}$, require comparisons of only $O(nr \log n)$ randomly chosen pairs to learn a good ranking. We show that the class of BTL preference matrices is a strict subset of the class $(\mathcal{P}^{\text{LR(logit,2)}} \cap \mathcal{P}^{\text{ST}})$ of stochastically transitive matrices that have rank ≤ 2 under the logit link $\psi_{\text{logit}}(p) = \log(\frac{p}{1-p})$, and that the class of Thurstone preferences matrices is a strict subset of the class $(\mathcal{P}^{\text{LR(probit,2)}} \cap \mathcal{P}^{\text{ST}})$ of stochastically transitive matrices that have rank ≤ 2 under the probit link $\psi_{\text{probit}}(p) = \Phi^{-1}(p)$, so that instantiations of our LRPR algorithm apply to both these classes. However our approach is

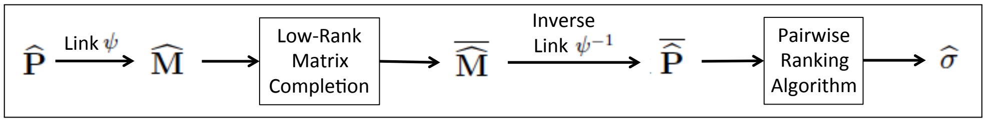

Figure 2: Our proposed *low-rank pairwise ranking* (LRPR) approach first applies a low-rank matrix completion routine to (a link-transformed version of) the incomplete pairwise comparison matrix $\widehat{\mathbf{P}}$ containing $O(nr\log n)$ entries, and then applies a pairwise ranking algorithm to (an inverse link-transformed version of) the completed matrix to obtain a ranking $\widehat{\sigma}$.

considerably more general and yields efficient algorithms for learning rankings from comparisons of $O(n\log n)$ non-actively chosen pairs for other types of preferences as well (see Figure 1 and Table 1). We also give results showing our approach yields good rankings even when the underlying preference matrix $\mathbf{P}$ is only *close to* a low-rank matrix under a link ψ.

Our algorithmic framework is based on tools from the theory of low-rank matrix completion (Candès and Recht, 2009; Candès and Tao, 2010). Specifically, given the outcomes of pairwise comparisons among the (randomly chosen) $O(nr\log n)$ pairs, we construct an *incomplete* pairwise comparison matrix $\widehat{\mathbf{P}}$ that contains $O(nr\log n)$ observed (noisy) entries, apply a low-rank matrix completion algorithm to (a link-transformed version of) $\widehat{\mathbf{P}}$ to obtain a completed comparison matrix $\overline{\widehat{\mathbf{P}}}$, and then apply a pairwise ranking algorithm to $\overline{\widehat{\mathbf{P}}}$ (see Figure 2). One can apply this framework in conjunction with any low-rank matrix completion routine that has exact recovery guarantees under *noisy* observations of matrix entries; we use the *OptSpace* algorithm of Keshavan et al. (2009). Our framework also recovers as a special case the nuclear norm aggregation (NNA) algorithm of Gleich and Lim (2011), which makes use of rank-2 approximations using singular value projection (Jain et al., 2010); however there are no known formal guarantees for NNA in the setting we consider.[2]

We conduct experiments comparing our LRPR algorithms with RC, BRE, and NNA, and find that our LRPR algorithms generally outperform all three.

1.1. Summary of Contributions

In summary, our main contributions in this paper are the following:

- We identify broad classes $\mathcal{P}^{\mathrm{LR}(\psi,r)}$ of 'low-rank' preference matrices $\mathbf{P}$ under which one can rank well from comparisons of $O(n\log n)$ randomly chosen pairs, and establish relationships between the well-studied BTL and Thurstone classes and these new classes that we define;
- We give a family of efficient *low-rank pairwise ranking* (LRPR) algorithms that use tools from low-rank matrix completion to learn a ranking from comparisons of $O(n\log n)$ pairs, and show that these algorithms provably learn a good ranking under the new classes $\mathcal{P}^{\mathrm{LR}(\psi,r)}$;
- We give supporting experimental evidence of the broad applicability of our LRPR algorithms by applying them to pairwise comparisons drawn from various preference structures, where we find our LRPR algorithms generally outperform existing baselines.

1.2. Organization

We start with preliminaries and background in Section 2. Section 3 introduces low-rank preferences. Section 4 outlines our low-rank pairwise ranking (LRPR) algorithmic framework. Section 5 gives formal guarantees on the ability of our LRPR algorithms to learn good rankings from comparisons

2. The recovery guarantee of Gleich and Lim (2011) applies in a 'noiseless' setting where one has access to $O(n\log n)$ exact entries of the underlying matrix $\mathbf{P}$, rather than noisy versions from observed pairwise comparisons.

of $O(n \log n)$ randomly chosen pairs under low-rank preferences. Section 6 gives experimental results. All proofs are deferred to the Appendix.

2. Preliminaries and Background

We describe the problem setup formally in Section 2.1, summarize various classes of preference matrices in Section 2.2, and give some background on low-rank matrix completion in Section 2.3.

2.1. Problem Setup

Let $[n] = \{1, \ldots, n\}$ denote the set of n items to be ranked, and $\binom{[n]}{2} = \{(i, j) : 1 \leq i < j \leq n\}$ denote the set of all $\binom{n}{2}$ item pairs. Let $\mathbf{P} \in [0, 1]^{n \times n}$ (with $P_{ij} + P_{ji} = 1 \;\forall i, j$) denote an unknown preference matrix according to which outcomes of pairwise comparisons are randomly drawn: every time a pair of items (i, j) is compared, i beats j with probability P_{ij} and j beats i with probability $1 - P_{ij}$ (independently of other comparisons). We are interested in this paper in algorithms which, given a pairwise comparison data set of the form $S = \{(i, j, \{y^k_{ij}\}_{k=1}^K)\}_{(i,j) \in E}$, consisting of a (non-actively chosen) set of pairs $E \subseteq \binom{[n]}{2}$ together with K pairwise comparison outcomes $y^k_{ij} \in \{0, 1\}$ ($k \in [K]$) for each pair $(i, j) \in E$ (drawn according to $\mathbf{P}$), where $y^k_{ij} = 1$ denotes that i beats j in the k-th comparison of (i, j) and $y^k_{ij} = 0$ denotes j beats i, learn from S a ranking or permutation of the n items, $\widehat{\sigma} \in \mathcal{S}_n$. We will denote by $m = |E|$ the number of pairs compared; our goal is to understand under what conditions on $\mathbf{P}$ one can learn a 'good' ranking from comparisons of $m = O(n \log n)$ pairs. The quality of the learned ranking $\widehat{\sigma}$ will be measured by the fraction of pairs on which $\widehat{\sigma}$ disagrees with $\mathbf{P}$:

$$\mathrm{dis}(\widehat{\sigma}, \mathbf{P}) \;=\; \frac{1}{\binom{n}{2}} \sum_{i<j} \mathbf{1}\Big(\big(i \succ_{\mathbf{P}} j\big) \wedge \big(\widehat{\sigma}(i) > \widehat{\sigma}(j)\big)\Big) + \mathbf{1}\Big(\big(j \succ_{\mathbf{P}} i\big) \wedge \big(\widehat{\sigma}(i) < \widehat{\sigma}(j)\big)\Big)\,,$$

where we denote

$$i \succ_{\mathbf{P}} j \iff P_{ij} > \tfrac{1}{2}\,.$$

2.2. Classes of Preference Matrices

We will denote by $\mathcal{P}_n$ the set of all pairwise preference matrices over n items:

$$\mathcal{P}_n \;=\; \big\{\mathbf{P} \in [0,1]^{n \times n} \,\big|\, P_{ij} + P_{ji} = 1 \;\forall i, j\big\}\,.$$

An important class of preference matrices is those that satisfy *stochastic transitivity* (ST), also referred to as the directed acyclic graph (DAG) condition by Rajkumar et al. (2015):

$$\mathcal{P}_n^{\mathrm{ST}} \;=\; \big\{\mathbf{P} \in \mathcal{P}_n \,\big|\, i \succ_{\mathbf{P}} j,\, j \succ_{\mathbf{P}} k \implies i \succ_{\mathbf{P}} k\big\}\,.$$

The *deterministic ordering* (DO) model can be described simply as

$$\mathcal{P}_n^{\mathrm{DO}} \;=\; \big\{\mathbf{P} \in \mathcal{P}_n \,\big|\, \exists \sigma \in \mathcal{S}_n : \sigma(i) < \sigma(j) \implies P_{ij} = 1\big\}\,.$$

The *noisy permutation* (NP) model studied by Braverman and Mossel (2009) and Wauthier et al. (2013) can be described as follows:

$$\mathcal{P}_n^{\mathrm{NP}} \;=\; \big\{\mathbf{P} \in \mathcal{P}_n \,\big|\, \exists \sigma \in \mathcal{S}_n, p \in [0, \tfrac{1}{2}) : \sigma(i) < \sigma(j) \implies P_{ij} = 1 - p\big\}\,.$$

It is easy to see that $\mathcal{P}_n^{\mathrm{DO}} \subsetneq \mathcal{P}_n^{\mathrm{NP}} \subsetneq \mathcal{P}_n^{\mathrm{ST}}$. Jamieson and Nowak (2011) consider embedding-based variants of the above classes, where items $i \in [n]$ are assumed to be embedded as points $\mathbf{x}_i \in \mathbb{R}^d$ via an embedding matrix $\mathbf{X} = [\mathbf{x}_1 \ldots \mathbf{x}_n] \in \mathbb{R}^{d \times n}$, and one effectively considers only permutations that can be realized via distances from some point $\mathbf{x} \in \mathbb{R}^d$ to these embedded points:

$$\mathcal{P}_n^{\mathrm{DO}(\mathbf{X})} \;=\; \left\{\mathbf{P} \in \mathcal{P}_n \,\middle|\, \exists \mathbf{x} \in \mathbb{R}^d : \|\mathbf{x} - \mathbf{x}_i\| < \|\mathbf{x} - \mathbf{x}_j\| \implies P_{ij} = 1\right\};$$
$$\mathcal{P}_n^{\mathrm{NP}(\mathbf{X})} \;=\; \left\{\mathbf{P} \in \mathcal{P}_n \,\middle|\, \exists \mathbf{x} \in \mathbb{R}^d, p \in [0, \tfrac{1}{2}) : \|\mathbf{x} - \mathbf{x}_i\| < \|\mathbf{x} - \mathbf{x}_j\| \implies P_{ij} = 1 - p\right\}.$$

The deterministic tournaments studied by Ailon (2012) can be described as

$$\mathcal{P}_n^{\mathrm{Dtour}} \;=\; \left\{\mathbf{P} \in \mathcal{P}_n \,\middle|\, P_{ij} \in \{0, 1\}\ \forall i \neq j\right\}.$$

The Bradley-Terry-Luce (BTL) and Thurstone models studied in statistics can be defined as

$$\mathcal{P}_n^{\mathrm{BTL}} \;=\; \left\{\mathbf{P} \in \mathcal{P}_n \,\middle|\, \exists \mathbf{w} \in \mathbb{R}^n_{++} : P_{ij} = \tfrac{w_i}{w_i + w_j}\ \forall i, j\right\};$$
$$\mathcal{P}_n^{\mathrm{Thu}} \;=\; \left\{\mathbf{P} \in \mathcal{P}_n \,\middle|\, \exists \mathbf{s} \in \mathbb{R}^n : P_{ij} = \Phi(s_i - s_j)\ \forall i, j\right\},$$

where $\Phi(\cdot)$ is the standard normal CDF. Again, it is easy to see that $\mathcal{P}_n^{\mathrm{BTL}} \subsetneq \mathcal{P}_n^{\mathrm{ST}}$ and $\mathcal{P}_n^{\mathrm{Thu}} \subsetneq \mathcal{P}_n^{\mathrm{ST}}$.

When n is clear from context, we will sometimes drop it from the subscript in the above classes, e.g. writing $\mathcal{P}_n$ as $\mathcal{P}$, $\mathcal{P}_n^{\mathrm{ST}}$ as $\mathcal{P}^{\mathrm{ST}}$, and so on.

2.3. Low-Rank Matrix Completion

In low-rank matrix completion, one is given an incomplete matrix $\widehat{\mathbf{M}} \in (\mathbb{R} \cup \{?\})^{n \times n}$, and the goal is to construct from this a complete matrix $\widehat{\widehat{\mathbf{M}}} \in \mathbb{R}^{n \times n}$ that has low rank (usually $\leq$ some target rank r) and that is close to $\widehat{\mathbf{M}}$ on the observed entries $\Omega = \{(i, j) \,|\, \widehat{M}_{ij} \neq ?\}$. One of the remarkable results in applied mathematics in recent years is that if the observed entries in $\widehat{\mathbf{M}}$ come from an underlying low-rank matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$, i.e. if $\widehat{\mathbf{M}}^\Omega = \mathbf{M}^\Omega$ (where for a matrix $\mathbf{A}$ we denote $\mathbf{A}^\Omega = (A_{ij})_{(i,j) \in \Omega}$), and if Ω is sampled uniformly at random from all subsets of $[n] \times [n]$ of size $|\Omega| \approx nr \log n$, then with high probability, one can recover $\mathbf{M}$ from $\widehat{\mathbf{M}}$ *exactly* (Candès and Recht, 2009; Candès and Tao, 2010). Several matrix completion algorithms are known to achieve this, including singular value thresholding (Cai et al., 2010), singular value projection (Jain et al., 2010), and *OptSpace* (Keshavan et al., 2009).

We make use of the *OptSpace* matrix completion algorithm of Keshavan et al. (2009), which comes with (approximate) recovery guarantees even in *noisy* settings, where the observed entries in $\widehat{\mathbf{M}}$ are not directly from a low-rank matrix $\mathbf{M}$ but rather are noisy realizations of such a low-rank matrix. In order to state the result of Keshavan et al. (2009), we will need the following definition:

Definition 1 ((μ_0, μ_1)-Incoherence) *A rank-r matrix $\mathbf{M} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top \in \mathbb{R}^{n \times n}$, where $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{n \times r}$ are orthogonal matrices with $\mathbf{U}^\top\mathbf{U} = \mathbf{V}^\top\mathbf{V} = n\mathbf{I}_r$ and $\mathbf{\Sigma} \in \mathbb{R}^{r \times r}$ is a diagonal matrix, is said to be* (μ_0, μ_1)-incoherent *if the following hold:*

(i) $\forall i \in [n] : \;\; \sum_{k=1}^r U_{ik}^2 \leq \mu_0 r\,, \;\; \sum_{k=1}^r V_{ik}^2 \leq \mu_0 r\,;$

(ii) $\forall i, j \in [n] : \;\; \left|\sum_{k=1}^r U_{ik}\left(\frac{\Sigma_k}{\Sigma_1}\right)V_{jk}\right| \leq \mu_1\sqrt{r}\,.$

Theorem 2 (Keshavan et al. (2009)) *Let $\mathbf{M} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top \in \mathbb{R}^{n \times n}$ be a (μ_0, μ_1)-incoherent matrix of rank r, where $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{n \times r}$ are orthogonal matrices with $\mathbf{U}^\top\mathbf{U} = \mathbf{V}^\top\mathbf{V} = n\mathbf{I}_r$ and $\mathbf{\Sigma} \in \mathbb{R}^{r \times r}$ is a diagonal matrix with $\Sigma_{\min} = \Sigma_1 \leq \ldots \leq \Sigma_r = \Sigma_{\max}$, and let $\kappa = (\Sigma_{\max}/\Sigma_{\min})$. Let $\mathbf{Z} \in \mathbb{R}^{n \times n}$. Let $\Omega \subseteq [n] \times [n]$, and let $\widehat{\mathbf{M}} \in (\mathbb{R} \cup \{?\})^{n \times n}$ be such that $\widehat{\mathbf{M}}^\Omega = (\mathbf{M} + \mathbf{Z})^\Omega$, and $\widehat{M}_{ij} = ?\ \forall (i, j) \notin \Omega$. There exist constants C, C' such that if*

$$|\Omega| \;\geq\; C\kappa^2 n \max\left(\mu_0 r \log(n),\, \mu_0^2 r^2 \kappa^4,\, \mu_1^2 r^2 \kappa^4\right)$$

and if Ω *is drawn uniformly at random from all subsets of* $[n] \times [n]$ *of size* $|\Omega|$*, then with probability at least* $1 - \frac{1}{n^3}$*, the matrix* $\overline{\widehat{\mathbf{M}}}$ *output by running OptSpace on* $\widehat{\mathbf{M}}$ *satisfies*

$$\frac{1}{n}\big\|\overline{\widehat{\mathbf{M}}} - \mathbf{M}\big\|_F \;\leq\; C'\kappa^2 \frac{n\sqrt{r}}{|\Omega|}\|\mathbf{Z}^{\Omega}\|_2\,,$$

provided that the right hand side above is less than $\Sigma_{\min}$.

3. Low-Rank Preferences

We now introduce the notion of 'low-rank' preference matrices.

A pairwise preference matrix $\mathbf{P}$ may not be low-rank itself. Indeed, if we replace the diagonal entries of $\mathbf{P}$ with zeros (which does not affect the entries of interest, i.e. entries P_{ij} for $i \neq j$), then the resulting matrix is always of high rank:

Proposition 3 *Let* $\mathbf{P} \in \mathcal{P}_n$*. Then* $\operatorname{rank}(\mathbf{P} - \frac{1}{2}\mathbf{I}_n) \geq n - 1$.

However, in many settings of interest, a suitably transformed version of $\mathbf{P}$ may have low rank. In particular, we will consider transformation via *link functions*, which are widely used in machine learning and statistics to map probabilities to real numbers and vice-versa:

Definition 4 (Link functions) *A* link function *is any strictly increasing function* $\psi : [0,1] \to \mathbb{R}$.[3]

Two commonly used link functions are the logit and probit links:

$$\psi_{\text{logit}}(p) \;=\; \log\left(\frac{p}{1-p}\right); \qquad \psi_{\text{probit}}(p) \;=\; \Phi^{-1}(p)\,.$$

(Recall that $\Phi(\cdot)$ is the standard normal CDF.) We will be interested in broad classes of preference matrices that have low rank under some link:

Definition 5 (Low-rank preferences under link ψ) *Let* $\psi : [0,1] \to \mathbb{R}$ *be a link function and* $r \in [n]$*. Define the class of* rank-r preference matrices under ψ*, denoted* $\mathcal{P}_n^{\text{LR}(\psi,r)}$*, as*

$$\mathcal{P}_n^{\text{LR}(\psi,r)} \;=\; \big\{\mathbf{P} \in \mathcal{P}_n \;\big|\; \operatorname{rank}(\psi(\mathbf{P})) \leq r\big\}\,.$$

As concrete examples, the following propositions show that all preference matrices $\mathbf{P}$ satisfying the BTL or Thurstone conditions have rank at most 2 under the logit and probit links, respectively; moreover, the corresponding classes of rank-2 preference matrices under these links are strictly more general than the BTL and Thurstone classes:

Proposition 6 $\mathcal{P}_n^{\text{BTL}} \subsetneq (\mathcal{P}_n^{\text{LR(logit,2)}} \cap \mathcal{P}_n^{\text{ST}}) \subsetneq \mathcal{P}_n^{\text{LR(logit,2)}}$.

Proposition 7 $\mathcal{P}_n^{\text{Thu}} \subsetneq (\mathcal{P}_n^{\text{LR(probit,2)}} \cap \mathcal{P}_n^{\text{ST}}) \subsetneq \mathcal{P}_n^{\text{LR(probit,2)}}$.

The following characterization of $(\mathcal{P}_n^{\text{LR(logit,2)}} \cap \mathcal{P}_n^{\text{ST}})$ and $\mathcal{P}_n^{\text{BTL}}$ makes clear the difference between the two classes:

3. Strictly speaking, we allow $\psi : [0,1] \to \overline{\mathbb{R}}$, where $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$; we ignore this issue for simplicity.

Theorem 8 (Characterization of $(\mathcal{P}_n^{\mathbf{LR(logit},2)} \cap \mathcal{P}_n^{\mathbf{ST}})$ and $\mathcal{P}_n^{\mathbf{BTL}}$) *Let* $\mathbf{P} \in \mathcal{P}_n$.
Part 1. $\mathbf{P} \in (\mathcal{P}_n^{\mathrm{LR(logit},2)} \cap \mathcal{P}_n^{\mathrm{ST}})$ *iff* $\exists \mathbf{x} \in \mathbb{R}^n, \mathbf{y} \in \mathbb{R}_+^n$ *with* $\mathbf{x}^\top \mathbf{y} = 0$ *such that*

$$\psi_{\text{logit}}(\mathbf{P}) = \mathbf{x}\mathbf{y}^\top - \mathbf{y}\mathbf{x}^\top .$$

Part 2. $\mathbf{P} \in \mathcal{P}_n^{\mathrm{BTL}}$ *iff* $\exists \mathbf{x} \in \mathbb{R}^n$ *with* $\mathbf{x}^\top \mathbf{e}_n = 0$ *such that*

$$\psi_{\text{logit}}(\mathbf{P}) = \mathbf{x}\mathbf{e}_n^\top - \mathbf{e}_n\mathbf{x}^\top .$$

Thus, a general preference matrix $\mathbf{P}$ in $(\mathcal{P}_n^{\mathrm{LR(logit},2)} \cap \mathcal{P}_n^{\mathrm{ST}})$ is characterized by $2n$ parameters, rather than n parameters as is the case for preference matrices in $\mathcal{P}_n^{\mathrm{BTL}}$. This can be useful for capturing preferences in settings where the probability of an item i beating an item j is determined not just by a single score for each of the two items, but rather by two numbers for each item; for example, in tennis, each player might be characterized by two numbers denoting the forehand and backhand quality, and the probability of one player beating the other might depend on both forehand and backhand quality of both the players. A similar result holds for $(\mathcal{P}_n^{\mathrm{LR(probit},2)} \cap \mathcal{P}_n^{\mathrm{ST}})$ and $\mathcal{P}_n^{\mathrm{Thu}}$:

Theorem 9 (Characterization of $(\mathcal{P}_n^{\mathbf{LR(probit},2)} \cap \mathcal{P}_n^{\mathbf{ST}})$ and $\mathcal{P}_n^{\mathbf{Thu}}$) *Let* $\mathbf{P} \in \mathcal{P}_n$.
Part 1. $\mathbf{P} \in (\mathcal{P}_n^{\mathrm{LR(probit},2)} \cap \mathcal{P}_n^{\mathrm{ST}})$ *iff* $\exists \mathbf{x} \in \mathbb{R}^n, \mathbf{y} \in \mathbb{R}_+^n$ *with* $\mathbf{x}^\top \mathbf{y} = 0$ *such that*

$$\psi_{\text{probit}}(\mathbf{P}) = \mathbf{x}\mathbf{y}^\top - \mathbf{y}\mathbf{x}^\top .$$

Part 2. $\mathbf{P} \in \mathcal{P}_n^{\mathrm{Thu}}$ *iff* $\exists \mathbf{x} \in \mathbb{R}^n$ *with* $\mathbf{x}^\top \mathbf{e}_n = 0$ *such that*

$$\psi_{\text{probit}}(\mathbf{P}) = \mathbf{x}\mathbf{e}_n^\top - \mathbf{e}_n\mathbf{x}^\top .$$

More generally, preference matrices that have a small rank r under some link ψ can be described by a smaller number of parameters than the $\binom{n}{2}$ parameters needed to describe an arbitrary preference matrix in $\mathcal{P}_n$.

Unlike BTL and Thurstone preference matrices which have rank ≤ 2 under the logit and probit links as above, preference matrices in the NP class do not have low rank under any 'skew-symmetric' link (including in particular the logit and probit links, both of which are skew-symmetric):

Proposition 10 *Let* $\mathbf{P} \in \mathcal{P}_n^{\mathrm{NP}}$, *and let* $\psi : [0,1] \to \mathbb{R}$ *be any link function satisfying* $\psi(1-q) = -\psi(q)\ \forall q \in [0,1]$. *Then*

$$\mathrm{rank}(\psi(\mathbf{P})) = \begin{cases} n & \text{if } n \text{ is even} \\ n-1 & \text{if } n \text{ is odd.} \end{cases}$$

We will also consider *approximately* low rank preferences:

Definition 11 (Approximately low-rank preferences under link ψ) *Let* $\psi : [0,1] \to \mathbb{R}$ *be a link function and* $r \in [n]$. *Let* $\beta > 0$. *Define the class of* β-approximately rank-r preference matrices under ψ, *denoted* $\mathcal{P}_n^{\mathrm{LR}(\psi,r,\beta)}$, *as*

$$\mathcal{P}_n^{\mathrm{LR}(\psi,r,\beta)} = \left\{ \mathbf{P} \in \mathcal{P}_n \,\middle|\, \exists \mathbf{M} \in \mathbb{R}^{n \times n} : \mathrm{rank}(\mathbf{M}) \leq r \text{ and } \|\psi(\mathbf{P}) - \mathbf{M}\|_F \leq \beta \right\} .$$

Below we describe our family of low-rank pairwise ranking algorithms that find good rankings under both low-rank and approximately low-rank preferences from comparisons of $O(rn \log n)$ pairs.

4. Low-Rank Pairwise Ranking Algorithm

Algorithm 1 describes our low-rank pairwise ranking (LRPR) algorithm (see also Figure 2). The algorithm is parametrized by a link function ψ and target rank $r \in [n]$.[4] Given pairwise comparisons S, the algorithm first constructs an incomplete empirical comparison matrix $\widehat{\mathbf{P}}$, and applies the link ψ to $\widehat{\mathbf{P}}$ to construct an incomplete matrix $\widehat{\mathbf{M}}$ with entries in $\mathbb{R}$. It then applies a matrix completion subroutine *MC* to $\widehat{\mathbf{M}}$ to obtain a rank-r completed matrix $\overline{\widehat{\mathbf{M}}}$ with entries in $\mathbb{R}$, applies the inverse link ψ^{-1} to this completed matrix to obtain a completed comparison matrix $\overline{\widehat{\mathbf{P}}}$, and then uses a pairwise ranking subroutine *PR* to estimate a ranking $\widehat{\sigma}$ from this completed comparison matrix.

Each setting of ψ, r, *MC* and *PR* in Algorithm 1 yields a specific instantiation of the LRPR algorithmic framework. As a special case, the LRPR framework recovers the nuclear norm aggregation (NNA) algorithm of Gleich and Lim (2011): when used with their 'log-odds' method for obtaining pairwise comparisons, the NNA algorithm can be viewed as applying a logit link ($\psi = \psi_{\text{logit}}$), using a singular value projection (SVP) based matrix completion routine (*MC* $=$ *SVP*) with target rank 2 ($r = 2$), and then constructing a ranking based on Borda scores (*PR* $=$ *Borda*). However there are no known guarantees on the quality of the ranking $\widehat{\sigma}$ returned by NNA when one starts with $\widehat{\mathbf{P}}$ constructed from observed pairwise comparisons S rather than the true preference matrix $\mathbf{P}$.

In our analysis and experiments, we will take *MC* to be the *OptSpace* algorithm of Keshavan et al. (2009) (see Section 2.3). For the pairwise ranking routine *PR*, our theorems will hold for any constant-factor approximate pairwise ranking algorithm (which we define formally in Section 5); in our experiments, we will take *PR* to be the Copeland ranking procedure (Copeland, 1951), which simply ranks items by their Copeland scores (number of wins in input matrix; in our case, the Copeland score of item i is simply $\sum_{j=1}^{n} \mathbf{1}(\overline{\widehat{P}}_{ij} > \frac{1}{2})$), and which has a 5-approximation guarantee. We will consider various choices of link function ψ and target rank r.

5. Analysis

Here we give our main results establishing theoretical guarantees for the LRPR algorithm when instantiated with *MC* $=$ *OptSpace* as the matrix completion routine, and with *PR* taken to be any constant-factor approximate pairwise ranking algorithm, which we define as follows:

Definition 12 (Approximate pairwise ranking algorithm) *Let PR be a pairwise ranking routine that given as input a preference/comparison matrix $\mathbf{Q} \in \mathcal{P}_n$, returns as output a permutation $\widehat{\sigma} \in \mathcal{S}_n$. Let $\gamma > 1$. We will say PR is a* γ-approximate pairwise ranking algorithm *if for all $\mathbf{Q} \in \mathcal{P}_n$, the permutation $\widehat{\sigma}$ returned by PR when given $\mathbf{Q}$ as input satisfies*

$$\text{dis}(\widehat{\sigma}, \mathbf{Q}) \leq \gamma \min_{\sigma \in \mathcal{S}_n} \text{dis}(\sigma, \mathbf{Q}) \,.$$

In particular, any γ-approximation algorithm for the minimum feedback arc set problem in tournaments (MFAST) immediately yields a γ-approximate pairwise ranking algorithm *PR* as follows: given as input a preference/comparison matrix $\mathbf{Q} \in \mathcal{P}_n$, one simply applies the γ-approximation algorithm for MFAST to a 0-1 version of the probabilistic tournament induced by $\mathbf{Q}$, where an edge (i, j) is present with weight 1 if $Q_{ij} > \frac{1}{2}$; if $Q_{ij} = \frac{1}{2}$, one can randomly choose the direction of the associated edge. For example, one could use any of the approximation algorithms of Ailon et al. (2008) or the PTAS of Kenyon-Mathieu and Schudy (2007) in this manner. In our experiments, we

4. Strictly speaking, the target rank r is not necessarily required as one can often estimate it as part of the matrix completion routine, but we include it here for simplicity.

Algorithm 1 Low-Rank Pairwise Ranking (LRPR)

Parameters:
Link function $\psi : [0,1] \to \mathbb{R}$
Target rank $r \in [n]$

Subroutines:
Low-rank matrix completion algorithm *MC*
Pairwise ranking algorithm *PR*

Input: Pairwise comparison data set $S = \{(i, j, \{y_{ij}^k\}_{k=1}^K)\}_{(i,j)\in E}$

• Construct (incomplete) empirical comparison matrix $\widehat{\mathbf{P}} \in ([0,1] \cup \{?\})^{n\times n}$ from S:

$$\widehat{P}_{ij} \quad = \quad \begin{cases} \frac{1}{K}\sum_{k=1}^K y_{ij}^k & \text{if } (i,j) \in E \\ \frac{1}{K}\sum_{k=1}^K (1 - y_{ji}^k) & \text{if } (j,i) \in E \\ \frac{1}{2} & \text{if } i = j \\ ? & \text{otherwise.} \end{cases}$$

• Construct link-transformed matrix: $\widehat{\mathbf{M}} = \psi(\widehat{\mathbf{P}}) \in \big(\mathbb{R} \cup \{?\}\big)^{n\times n}$

• Obtain completed rank-r matrix: $\overline{\widehat{\mathbf{M}}} = MC(\widehat{\mathbf{M}}, r) \in \mathbb{R}^{n\times n}$

• Apply inverse link transform to $\overline{\widehat{\mathbf{M}}}$ to obtain completed comparison matrix $\overline{\widehat{\mathbf{P}}} \in [0,1]^{n\times n}$:

$$\overline{\widehat{P}}_{ij} = \begin{cases} \frac{1}{2} + \min\big(\big|\psi^{-1}(\overline{\widehat{M}}_{ij}) - \frac{1}{2}\big|, \big|\psi^{-1}(\overline{\widehat{M}}_{ji}) - \frac{1}{2}\big|\big) & \text{if } i \neq j \text{ and } \overline{\widehat{M}}_{ij} > \overline{\widehat{M}}_{ji} \\ \frac{1}{2} - \min\big(\big|\psi^{-1}(\overline{\widehat{M}}_{ij}) - \frac{1}{2}\big|, \big|\psi^{-1}(\overline{\widehat{M}}_{ji}) - \frac{1}{2}\big|\big) & \text{if } i \neq j \text{ and } \overline{\widehat{M}}_{ij} < \overline{\widehat{M}}_{ji} \\ \frac{1}{2} & \text{if } i = j\,. \end{cases}$$

• Obtain ranking: $\widehat{\sigma} = PR(\overline{\widehat{\mathbf{P}}}) \in \mathcal{S}_n$

Output: Permutation $\widehat{\sigma} \in \mathcal{S}_n$

will take *PR* to be the Copeland ranking procedure (Copeland, 1951), which simply ranks items i by their Copeland scores (number of wins in input matrix, $\sum_{j=1}^n \mathbf{1}(Q_{ij} > \frac{1}{2})$), and which is known to have a 5-approximation guarantee (Coppersmith et al., 2006).

We will find it useful to define the following quantities associated with a preference matrix $\mathbf{P}$ and link function ψ:

$$P_{\min} \quad = \quad \min_{i\neq j} P_{ij}\,; \qquad \Delta_{\min}^{\mathbf{P},\psi} \quad = \quad \min_{i\neq j} |\psi(P_{ij}) - \psi(\tfrac{1}{2})|\,.$$

5.1. Guarantees for LRPR Algorithm under Low-Rank Preferences

The following result shows that if the underlying preference matrix $\mathbf{P}$ has rank r under some link function, then comparisons of $O(nr\log n)$ randomly chosen pairs, with $O(r\log n)$ comparisons per pair, are sufficient for the LRPR algorithm to return a good ranking:

Theorem 13 (Performance of LRPR algorithm for low-rank preferences) *Let $\psi : [0,1] \to \mathbb{R}$ be a link function and let $r \in [n]$. Let $\mathbf{P} \in (\mathcal{P}_n^{\mathrm{LR}(\psi,r)} \cap \mathcal{P}_n^{\mathrm{ST}})$. Let $\psi(\mathbf{P}) = \mathbf{U\Sigma V}^\top$ be a (μ_0, μ_1)-incoherent matrix, where $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{n\times r}$ are orthogonal matrices with $\mathbf{U}^\top\mathbf{U} = \mathbf{V}^\top\mathbf{V} = n\mathbf{I}_r$ and $\mathbf{\Sigma} \in \mathbb{R}^{r\times r}$ is a diagonal matrix with $\Sigma_{\min} = \Sigma_1 \leq \ldots \leq \Sigma_r = \Sigma_{\max}$, and let $\kappa = (\Sigma_{\max}/\Sigma_{\min})$.*

Let ψ be L-Lipschitz in $[\frac{P_{\min}}{2}, 1 - \frac{P_{\min}}{2}]$. Let MC = OptSpace, and let PR be any γ-approximate pairwise ranking algorithm. Let $0 < \epsilon < \frac{1}{2}$. There exist constants C, C' such that if

$$m \;\geq\; C\kappa^2 n \max\left(\mu_0 r \log(n),\, \mu_0^2 r^2 \kappa^4,\, \mu_1^2 r^2 \kappa^4\right)$$

and

$$K \;\geq\; \log(n) \max\left(\frac{4C' L^2 \kappa^4 r(1+\gamma)}{(\Delta_{\min}^{\mathbf{P},\psi})^2 \epsilon},\, \frac{C' L^2 \kappa^4 r}{\Sigma_{\min}^2},\, \frac{11}{P_{\min}^2}\right),$$

and if $E \subseteq \binom{[n]}{2}$ with $|E| = m$ is chosen uniformly at random from all such subsets of size m and $S = \left\{\left(i, j, \{y_{ij}^k\}_{k=1}^K\right)\right\}_{(i,j)\in E}$ is generated by comparing each pair $(i, j) \in E$ (independently) K times according to $\mathbf{P}$, then with probability at least $1 - \frac{2}{n^3}$, the permutation $\widehat{\sigma}$ produced by running the LRPR(ψ, r) *algorithm on S, with subroutines MC and PR as above, satisfies*

$$\mathrm{dis}(\widehat{\sigma}, \mathbf{P}) \;\leq\; \epsilon\,.$$

The above result requires the link-transformed preference matrix $\psi(\mathbf{P})$ to satisfy incoherence properties. It is not hard to see that rank-2 preference matrices under the logit link (and therefore as a special case, BTL preference matrices) satisfy such incoherence properties:

Lemma 14 (Incoherence of rank-2 preferences under logit link) *Let $\mathbf{P} \in (\mathcal{P}_n^{\text{LR(logit,2)}} \cap \mathcal{P}_n^{\text{ST}})$, with $\mathbf{x} \in \mathbb{R}^n, \mathbf{y} \in \mathbb{R}_+^n$ such that $\mathbf{x}^\top \mathbf{y} = 0$ and $\psi_{\text{logit}}(\mathbf{P}) = \mathbf{x}\mathbf{y}^\top - \mathbf{y}\mathbf{x}^\top$. Let $\mu = \frac{1}{2}\left(\frac{x_{\max}^2}{x_{\min}^2} + \frac{y_{\max}^2}{y_{\min}^2}\right)$, where $x_{\min} = \min_i |x_i|$, $x_{\max} = \max_i |x_i|$, $y_{\min} = \min_i |y_i|$, and $y_{\max} = \max_i |y_i|$. Then $\psi_{\text{logit}}(\mathbf{P})$ is $(\mu, \sqrt{2}\mu)$-incoherent.*

Corollary 15 (Incoherence of BTL preferences) *Let $\mathbf{P} \in \mathcal{P}_n^{\text{BTL}}$, with parameter vector $\mathbf{w} \in \mathbb{R}_{++}^n$. Let $\mu = \frac{1}{2}\left(\frac{(\log w_{\max} - \frac{1}{n}\sum_{j=1}^n \log w_j)^2}{(\log w_{\min} - \frac{1}{n}\sum_{j=1}^n \log w_j)^2} + 1\right)$. Then $\psi_{\text{logit}}(\mathbf{P})$ is $(\mu, \sqrt{2}\mu)$-incoherent.*

We also have the following result on Lipschitz-ness of ψ_{logit}:

Lemma 16 *For any $q \in (0, \frac{1}{2}]$, ψ_{logit} is $(\frac{4}{q})$-Lipschitz in $[\frac{q}{2}, 1 - \frac{q}{2}]$.*

Thus, as a special case of Theorem 13, we have the following results for the performance of the LRPR algorithm for rank-2 preferences under the logit link (Corollary 17), and more specifically, for BTL preferences (Corollary 18):

Corollary 17 (Performance of LRPR algorithm for rank-2 preferences under logit link) *Let $\mathbf{P} \in (\mathcal{P}_n^{\text{LR(logit,2)}} \cap \mathcal{P}_n^{\text{ST}})$, with $\mathbf{x} \in \mathbb{R}^n, \mathbf{y} \in \mathbb{R}_+^n$ such that $\mathbf{x}^\top \mathbf{y} = 0$ and $\psi_{\text{logit}}(\mathbf{P}) = \mathbf{x}\mathbf{y}^\top - \mathbf{y}\mathbf{x}^\top$. Let μ be as in Lemma 14, let $\overline{x} = \sqrt{\frac{1}{n}\sum_{i=1}^n x_i^2}$ and $\overline{y} = \sqrt{\frac{1}{n}\sum_{i=1}^n y_i^2}$, and let $\Delta = \min_{i\neq j} |x_i y_j - y_i x_j|$. Let MC = OptSpace, and let PR be any γ-approximate pairwise ranking algorithm. Let $0 < \epsilon < \frac{1}{2}$. There exist constants C, C' such that if*

$$m \;\geq\; Cn \max\left(2\mu \log(n),\, 8\mu^2\right)$$

and

$$K \;\geq\; \frac{1}{P_{\min}^2} \log(n) \max\left(\left(\frac{512}{9}\right)\frac{C'(1+\gamma)}{\Delta^2 \epsilon},\, \left(\frac{128}{9}\right)\frac{C'}{\overline{x}^2\overline{y}^2},\, 11\right),$$

and if $E \subseteq \binom{[n]}{2}$ with $|E| = m$ is chosen uniformly at random from all such subsets of size m and $S = \left\{\left(i, j, \{y_{ij}^k\}_{k=1}^K\right)\right\}_{(i,j)\in E}$ is generated by comparing each pair $(i, j) \in E$ (independently) K times according to $\mathbf{P}$, then with probability at least $1 - \frac{2}{n^3}$, the permutation $\widehat{\sigma}$ produced by running the LRPR$(\psi_{\text{logit}}, 2)$ *algorithm on S, with subroutines MC and PR as above, satisfies*

$$\mathrm{dis}(\widehat{\sigma}, \mathbf{P}) \;\leq\; \epsilon\,.$$

Corollary 18 (Performance of LRPR algorithm for BTL preferences) *Let* $\mathbf{P} \in \mathcal{P}_n^{\text{BTL}}$*, with parameter vector* $\mathbf{w} \in \mathbb{R}_{++}^n$*. Let* $w_{\max} = \max_i w_i$ *and* $w_{\min} = \min_i w_i$*. Let* $\mu = \frac{1}{2}\Big(\frac{(\log w_{\max} - \frac{1}{n}\sum_{j=1}^n \log w_j)^2}{(\log w_{\min} - \frac{1}{n}\sum_{j=1}^n \log w_j)^2} + 1\Big)$*,* $\overline{x} = \sqrt{\frac{1}{n}\sum_{i=1}^n(\log w_i - \frac{1}{n}\sum_{j=1}^n \log w_j)^2}$*, and* $\Delta = \min_{i\neq j} |\log(\frac{w_i}{w_j})|$*, and let* $b = \frac{w_{\max}}{w_{\min}}$*. Let MC* $=$ *OptSpace, and let PR be any* γ*-approximate pairwise ranking algorithm. Let* $0 < \epsilon < \frac{1}{2}$*. There exist constants* C, C' *such that if*

$$m \;\geq\; Cn \max\big(2\mu\log(n),\, 8\mu^2\big)$$

and

$$K \;\geq\; (b+1)^2 \log(n) \max\bigg(\Big(\frac{512}{9}\Big)\frac{C'(1+\gamma)}{\Delta^2\epsilon},\, \Big(\frac{128}{9}\Big)\frac{C'}{\overline{x}^2},\, 11\bigg),$$

and if $E \subseteq \binom{[n]}{2}$ *with* $|E| = m$ *is chosen uniformly at random from all such subsets of size* m *and* $S = \big\{\big(i, j, \{y_{ij}^k\}_{k=1}^K\big)\big\}_{(i,j)\in E}$ *is generated by comparing each pair* $(i,j) \in E$ *(independently)* K *times according to* $\mathbf{P}$*, then with probability at least* $1 - \frac{2}{n^3}$*, the permutation* $\widehat{\sigma}$ *produced by running the* $\text{LRPR}(\psi_{\text{logit}}, 2)$ *algorithm on* S*, with subroutines MC and PR as above, satisfies*

$$\text{dis}(\widehat{\sigma}, \mathbf{P}) \;\leq\; \epsilon\,.$$

Similar results also hold for rank-2 preferences under the probit link and for Thurstone preferences; we omit details for brevity.

Remark 1 *The above results suggest that the number of comparisons per pair,* K*, increases as* $P_{\min}$ *decreases. This may seem counter-intuitive at first, since in some sense, a smaller value of* $P_{\min}$ *should make the problem of learning a good ranking easier. The reason for* K *increasing as* $P_{\min}$ *decreases is that while learning a good ranking w.r.t.* $\mathbf{P}$ *requires knowledge of only whether the entries* P_{ij} *are larger than or smaller than* $\frac{1}{2}$*, the LRPR algorithm effectively estimates the entries* P_{ij} *themselves (via estimation of the link-transformed values* $\psi(P_{ij})$*). A similar phenomenon is observed for example in the analysis of the Rank Centrality algorithm (Negahban et al., 2012), which effectively estimates the parameters* w_i *of the BTL model assumed (via estimation of the stationary probability vector of an associated Markov chain); in this case too, a smaller value of* $w_{\min}$ *in the BTL model, which amounts to a smaller value of* $P_{\min}$ *in the corresponding preference matrix* $\mathbf{P}$*, leads to a larger number of comparisons* K*. It remains an open question to design algorithms that directly estimate a good ranking from comparisons of* $O(n\log n)$ *pairs under the types of preference classes considered here and to improve the dependence of* K *on* $P_{\min}$*.*

5.2. Guarantees for LRPR Algorithm under Approximately Low-Rank Preferences

The following result shows that the LRPR algorithm finds a good ranking even when the underlying preference matrix $\mathbf{P}$ is only approximately (but reasonably close to) low rank under some link function:

Theorem 19 (Performance of LRPR algorithm for approximately low-rank preferences) *Let* $\psi : [0,1] \to \mathbb{R}$ *be a link function and let* $r \in [n]$ *and* $\beta > 0$*. Let* $\mathbf{P} \in (\mathcal{P}_n^{\text{LR}(\psi,r,\beta)} \cap \mathcal{P}_n^{\text{ST}})$*. Let* $\mathbf{M} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top \in \mathbb{R}^{n\times n}$ *be a* (μ_0, μ_1)*-incoherent matrix of rank* r *such that* $\|\psi(\mathbf{P}) - \mathbf{M}\|_F \leq \beta$*, where* $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{n\times r}$ *are orthogonal matrices with* $\mathbf{U}^\top\mathbf{U} = \mathbf{V}^\top\mathbf{V} = n\mathbf{I}_r$ *and* $\mathbf{\Sigma} \in \mathbb{R}^{r\times r}$ *is a diagonal matrix with* $\Sigma_{\min} = \Sigma_1 \leq \ldots \leq \Sigma_r = \Sigma_{\max}$*, and let* $\kappa = (\Sigma_{\max}/\Sigma_{\min})$*. Let* ψ *be*

L-Lipschitz in $[P_{\min}, 1 - P_{\min}]$. *Let MC = OptSpace, and let PR be any* γ*-approximate pairwise ranking algorithm. Let* $0 < \epsilon < \frac{1}{2}$. *There exist constants* C, C' *such that if* $\beta \leq \alpha \frac{\Sigma_{\min}}{C'\kappa^2\sqrt{r}}$ *for some* $\alpha \in (0, 1)$*, if* m *satisfies the conditions of Theorem 13 and*

$$K \;\geq\; \log(n) \max\left(\frac{4C'L^2\kappa^4 r(1+\gamma)}{(\Delta_{\min}^{\mathbf{P},\psi})^2\epsilon}, \frac{C'L^2\kappa^4 r}{(1-\alpha)^2\Sigma_{\min}^2}, \frac{11}{P_{\min}^2} \right),$$

and if $E \subseteq \binom{[n]}{2}$ *and* $S = \big\{\big(i, j, \{y_{ij}^k\}_{k=1}^K\big)\big\}_{(i,j)\in E}$ *are chosen randomly as described in Theorem 13, then with probability at least* $1 - \frac{2}{n^3}$*, the permutation* $\widehat{\sigma}$ *produced by running the* $\mathrm{LRPR}(\psi, r)$ *algorithm on* S*, with subroutines MC and PR as above, satisfies*

$$\mathrm{dis}(\widehat{\sigma}, \mathbf{P}) \;\leq\; \epsilon \;+\; (1+\gamma)\left(\frac{4\alpha\Sigma_{\min}}{\Delta_{\min}^{\mathbf{P},\psi}}\right)^2 + 8\sqrt{\epsilon(1+\gamma)}\left(\frac{\alpha\Sigma_{\min}}{\Delta_{\min}^{\mathbf{P},\psi}}\right).$$

Clearly, as α (and therefore β) approaches zero, both the second and third terms in the bound on $\mathrm{dis}(\widehat{\sigma}, \mathbf{P})$ above vanish, and we recover the result for low-rank preferences in Theorem 13.

6. Experiments

In this section we describe results of experiments with our LRPR algorithm applied to pairwise comparisons drawn from various low-rank preference matrices. Specifically, we took the number of items to be $n = 500$, and constructed three preference matrices as follows:

(1) $\mathbf{P}^1 \in \mathcal{P}_{500}^{\mathrm{BTL}}$.
This was constructed by generating a score vector $\mathbf{w} \in [0, 1]^{500}$ with entries drawn uniformly at random from $[0, 1]$, and then setting $P_{ij} = \frac{w_i}{w_i + w_j} \forall i, j$.

(2) $\mathbf{P}^2 \in (\mathcal{P}_{500}^{\mathrm{LR(logit,2)}} \cap \mathcal{P}_{500}^{\mathrm{ST}}) \setminus \mathcal{P}_{500}^{\mathrm{BTL}}$.
This was constructed by generating two score vectors $\mathbf{x}, \mathbf{y} \in [0, 1]^{500}$ with entries drawn uniformly at random from $[0, 1]$, and then setting $\mathbf{P} = \psi_{\mathrm{logit}}^{-1}(\mathbf{x}\mathbf{y}^\top - \mathbf{y}\mathbf{x}^\top)$.

(3) $\mathbf{P}^3 \in (\mathcal{P}_{500}^{\mathrm{LR(logit,4)}} \cap \mathcal{P}_{500}^{\mathrm{ST}}) \setminus \mathcal{P}_{500}^{\mathrm{BTL}}$.
This was constructed by generating a skew-symmetric matrix $\mathbf{Y} \in \mathbb{R}^{500\times 500}$ with upper triangular entries drawn randomly from $\mathcal{N}(5, 1)$, finding its rank-4 projection $\overline{\mathbf{Y}}$ (also skew-symmetric; see (Gleich and Lim, 2011, Lemma 2)), and then setting $\mathbf{P} = \psi_{\mathrm{logit}}^{-1}(\overline{\mathbf{Y}})$. [5]

In each case, we generated pairwise comparison data from these preference matrices, in which we varied m, the number of pairs compared, as well as K, the number of times each pair was sampled, and measured $\mathrm{dis}(\widehat{\sigma}, \mathbf{P})$, the fraction of pairs on which the learned ranking $\widehat{\sigma}$ disagreed with the true preference matrix $\mathbf{P}$ where $\mathbf{P} = \mathbf{P}^1$, $\mathbf{P}^2$ or $\mathbf{P}^3$, respectively. For comparison, we also applied the following three algorithms as baselines: the Rank Centrality (RC) algorithm, for which the stationary distribution constructed by the algorithm is known to converge to a stationary distribution associated with the underlying preference matrix $\mathbf{P}$ when $\mathbf{P} \in \mathcal{P}^{\mathrm{BTL}}$ (Negahban et al., 2012); the Balanced Rank Estimation (BRE) algorithm, which has recovery guarantees when $\mathbf{P} \in \mathcal{P}^{\mathrm{NP}}$

5. This procedure can in general lead to a preference matrix $\mathbf{P} \in \mathcal{P}_{500}^{(\mathrm{logit},4)} \setminus \mathcal{P}_{500}^{\mathrm{ST}}$. We checked to make sure that the generated matrix satisfied the stochastic transitivity condition.

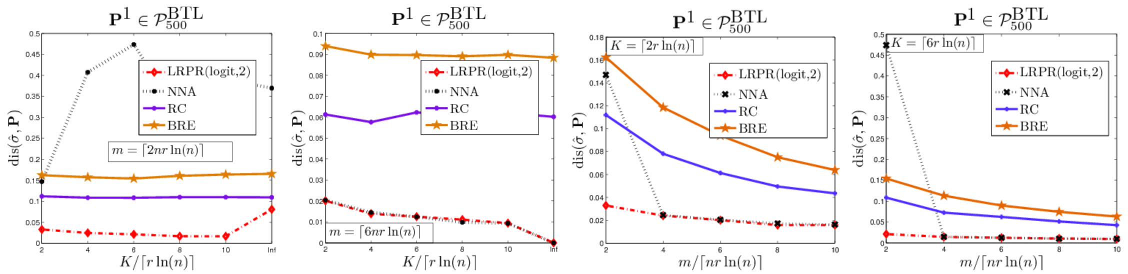

Figure 3: Ranking results when pairwise comparisons come from $\mathbf{P}^1 \in \mathcal{P}_{500}^{\mathrm{BTL}}$.

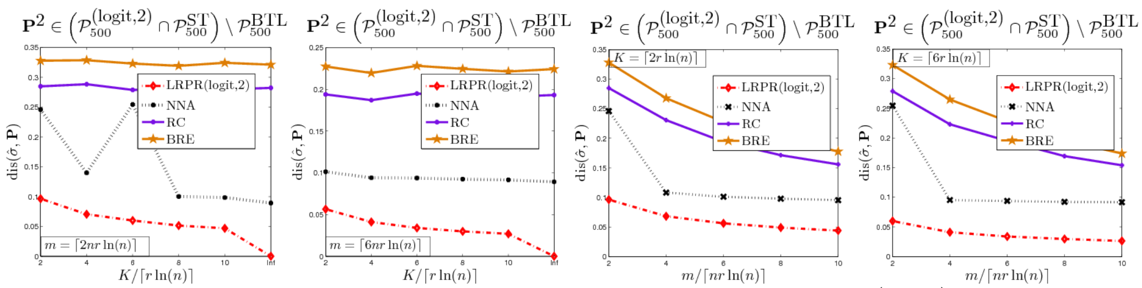

Figure 4: Ranking results when pairwise comparisons come from $\mathbf{P}^2 \in (\mathcal{P}_{500}^{(\mathrm{logit},2)} \cap \mathcal{P}_{500}^{\mathrm{ST}}) \setminus \mathcal{P}_{500}^{\mathrm{BTL}}$.

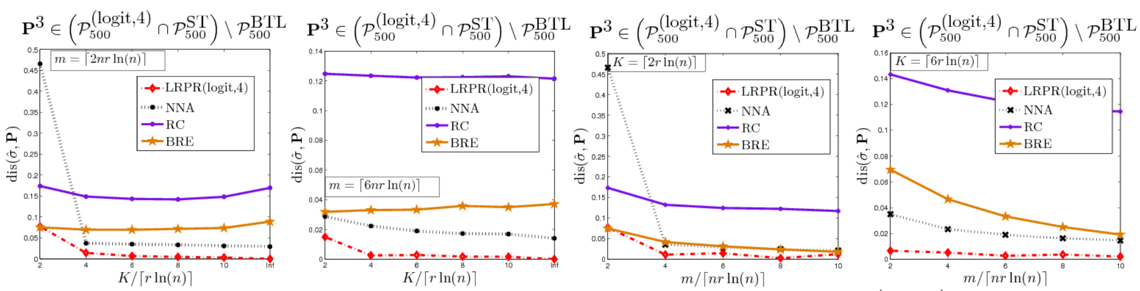

Figure 5: Ranking results when pairwise comparisons come from $\mathbf{P}^3 \in (\mathcal{P}_{500}^{(\mathrm{logit},4)} \cap \mathcal{P}_{500}^{\mathrm{ST}}) \setminus \mathcal{P}_{500}^{\mathrm{BTL}}$.

(Wauthier et al., 2013); and the nuclear norm aggregation (NNA) algorithm (Gleich and Lim, 2011), which also uses a matrix completion approach (and can be viewed as a special case of our LRPR framework) but for which no theoretical guarantees are known in our setting (i.e. in the setting where a ranking is to be learned from noisy comparisons of $O(n \log n)$ pairs).

The results are shown in Figures 3, 4 and 5 (in the plots where m is fixed and K varies, $K = \infty$ corresponds to being given the true preference matrix entries P_{ij} for compared pairs $(i, j) \in E$). As can be seen, for $\mathbf{P}^1$, which satisfies the BTL condition, when the number of pairs compared, m, is very small, our LRPR algorithm outperforms all the others; when m is larger, LRPR and NNA perform similarly, with both outperforming RC and BRE. For $\mathbf{P}^2$ and $\mathbf{P}^3$, LRPR clearly outperforms all three baselines.

7. Conclusion

We have considered the question of when one can learn a good ranking of n items from comparisons of only $O(n \log n)$ non-actively chosen pairs. Previous results have established this possibility only under the Bradley-Terry-Luce (BTL) and noisy permutation (NP) classes of pairwise preferences.

In this paper we have shown this is possible under a broad family of pairwise preference structures that we call *low-rank preferences*, which include for example the BTL and Thurstone classes as special cases but are significantly more general. In particular, our new pairwise ranking algorithm, called *low-rank pairwise ranking* (LRPR), makes use of tools from the theory of low-rank matrix completion to learn provably good rankings from comparisons of $O(n \log n)$ randomly chosen pairs under any such low-rank preference structure, including BTL and Thurstone but also more general classes of preference structures.

Our LRPR algorithmic framework applies a low-rank matrix completion algorithm to a link-transformed version of the empirical pairwise comparison probability matrix. Here we have used the *OptSpace* matrix completion algorithm of Keshavan et al. (2009), which has strong recovery guarantees under noisy observation of matrix entries. It may also be interesting to explore other low-rank matrix completion methods with noisy recovery guarantees, such as those of Lafond (2015).

Acknowledgments

Thanks to Sewoong Oh for discussions related to this work and to the anonymous reviewers for helpful comments. SA thanks DST, the Indo-US Science & Technology Forum, and the Radcliffe Institute for their support.

Appendix A. Proof of Proposition 3

Proof Denote $\mathbf{L} = \mathbf{P} - \frac{1}{2}\mathbf{I}_n$. Denote by $\mathbf{e}_n$ the $n \times 1$ all-ones vector, and by $\mathbf{E}_n$ the $n \times n$ all-ones matrix. Let $\mathbf{B} = [\mathbf{L}^\top \ \mathbf{e}_n] \in \mathbb{R}^{n \times (n+1)}$. Now suppose there exists $\mathbf{x} \in \mathbb{R}^n$ such that $\mathbf{x}^\top \mathbf{B} = 0$. Then we have

$$\mathbf{x}^\top (\mathbf{E}_n - \mathbf{I}_n)\mathbf{x} = \mathbf{x}^\top ((\mathbf{P} + \mathbf{P}^\top) - \mathbf{I}_n)\mathbf{x} = \mathbf{x}^\top \mathbf{L}\mathbf{x} + \mathbf{x}^\top \mathbf{L}^\top \mathbf{x} = 0\,.$$

Since $\mathbf{x}^\top \mathbf{e}_n = 0$, this gives $\mathbf{x}^\top \mathbf{x} = 0$, and therefore $\mathbf{x} = \mathbf{0}$. Thus the columns of $\mathbf{L}^\top$ along with the all-ones vector $\mathbf{e}_n$ span $\mathbb{R}^n$, and therefore rank($\mathbf{L}$) is at least $n-1$. ■

Appendix B. Proof of Proposition 6

Proof It is easy to see that $\mathbf{P}_n^{\text{BTL}} \subseteq \mathcal{P}_n^{\text{ST}}$. To see that $\mathbf{P}_n^{\text{BTL}} \subseteq \mathcal{P}_n^{\text{LR(logit,2)}}$, let $\mathbf{P} \in \mathcal{P}_n^{\text{BTL}}$, with score vector $\mathbf{w} \in \mathbb{R}^n_{++}$ such that $P_{ij} = \frac{w_i}{w_i + w_j} \forall i, j$. Let $\mathbf{Y} = \psi_{\text{logit}}(\mathbf{P})$. Then it is easy to see that $Y_{ij} = \log(w_i) - \log(w_j)\ \forall i, j$, i.e. $\mathbf{Y} = \mathbf{s}\mathbf{e}^\top - \mathbf{e}\mathbf{s}^\top$, where $s_i = \log(w_i)\ \forall i$. Thus $\text{rank}(\mathbf{Y}) = 2$, and therefore $\mathbf{P} \in \mathcal{P}_n^{\text{LR(logit,2)}}$. This gives $\mathbf{P}_n^{\text{BTL}} \subseteq (\mathcal{P}_n^{\text{LR(logit,2)}} \cap \mathcal{P}_n^{\text{ST}})$.

To see that the above containment is strict, consider $\mathbf{Y} = \mathbf{x}\mathbf{y}^\top - \mathbf{y}\mathbf{x}^\top$, where $x_i = i\ \forall i$, and $y_i = 1\ \forall i \in [n-1]$ and $y_n = \frac{1}{2}$. It can be verified that $\mathbf{P} = \psi_{\text{logit}}^{-1}(\mathbf{Y}) \in (\mathcal{P}_n^{\text{LR(logit,2)}} \cap \mathcal{P}_n^{\text{ST}}) \setminus \mathcal{P}_n^{\text{BTL}}$.

To see that $(\mathcal{P}_n^{\text{LR(logit,2)}} \cap \mathcal{P}_n^{\text{ST}}) \subsetneq \mathcal{P}_n^{\text{LR(logit,2)}}$, consider any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ such that

$$\begin{array}{lll} x_1 = -1 & x_2 = 1 & x_3 = 1 \\ y_1 = 1 & y_2 = -2 & y_3 = 1\,, \end{array}$$

and let $\mathbf{P} = \psi_{\text{logit}}^{-1}(\mathbf{x}\mathbf{y}^\top - \mathbf{y}\mathbf{x}^\top)$. Then by construction, $\mathbf{P} \in \mathcal{P}_n^{(\text{logit},2)}$. However it can be verified that

$$P_{12} = \psi_{\text{logit}}^{-1}(1) > \tfrac{1}{2}\,; \quad P_{23} = \psi_{\text{logit}}^{-1}(3) > \tfrac{1}{2}\,; \quad P_{31} = \psi_{\text{logit}}^{-1}(2) > \tfrac{1}{2}\,;$$

i.e. that

$$1 \succ_{\mathbf{P}} 2\,; \quad 2 \succ_{\mathbf{P}} 3\,; \quad 3 \succ_{\mathbf{P}} 1\,,$$

and therefore $\mathbf{P} \notin \mathcal{P}_n^{\text{ST}}$. The claim follows. ■

Appendix C. Proof of Proposition 7

Proof The proof is similar to that of Proposition 6. It is easy to see that $\mathbf{P}_n^{\text{Thu}} \subseteq \mathcal{P}_n^{\text{ST}}$. To see that $\mathbf{P}_n^{\text{Thu}} \subseteq \mathcal{P}_n^{\text{LR(probit,2)}}$, let $\mathbf{P} \in \mathcal{P}_n^{\text{Thu}}$, with score vector $\mathbf{s} \in \mathbb{R}^n$ such that $P_{ij} = \Phi(s_i - s_j)\ \forall i, j$. Let $\mathbf{Y} = \psi_{\text{probit}}(\mathbf{P})$. Then clearly $Y_{ij} = s_i - s_j\ \forall i, j$, i.e. $\mathbf{Y} = \mathbf{s}\mathbf{e}^\top - \mathbf{e}\mathbf{s}^\top$. Thus $\text{rank}(\mathbf{Y}) = 2$, and therefore $\mathbf{P} \in \mathcal{P}_n^{\text{LR(probit,2)}}$. This gives $\mathbf{P}_n^{\text{Thu}} \subsetneq (\mathcal{P}_n^{\text{LR(probit,2)}} \cap \mathcal{P}_n^{\text{ST}})$.

To see that the above containment is strict, consider $\mathbf{Y} = \mathbf{x}\mathbf{y}^\top - \mathbf{y}\mathbf{x}^\top$, where $x_i = i\ \forall i$, and $y_i = 1\ \forall i \in [n-1]$ and $y_n = \frac{1}{2}$. It can be verified that $\mathbf{P} = \psi_{\text{probit}}^{-1}(\mathbf{Y}) \in (\mathcal{P}_n^{\text{LR(probit,2)}} \cap \mathcal{P}_n^{\text{ST}}) \setminus \mathcal{P}_n^{\text{Thu}}$.

To see that $(\mathcal{P}_n^{\text{LR(probit,2)}} \cap \mathcal{P}_n^{\text{ST}}) \subsetneq \mathcal{P}_n^{\text{LR(probit,2)}}$, consider any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ such that

$$\begin{array}{lll} x_1 = -1 & x_2 = 1 & x_3 = 1 \\ y_1 = 1 & y_2 = -2 & y_3 = 1\,, \end{array}$$

and let $\mathbf{P} = \psi_{\text{probit}}^{-1}(\mathbf{x}\mathbf{y}^\top - \mathbf{y}\mathbf{x}^\top)$. Then by construction, $\mathbf{P} \in \mathcal{P}_n^{\text{(probit,2)}}$. However it can be verified that

$$P_{12} = \psi_{\text{probit}}^{-1}(1) > \tfrac{1}{2}\,; \quad P_{23} = \psi_{\text{probit}}^{-1}(3) > \tfrac{1}{2}\,; \quad P_{31} = \psi_{\text{probit}}^{-1}(2) > \tfrac{1}{2}\,;$$

i.e. that

$$1 \succ_{\mathbf{P}} 2\,; \quad 2 \succ_{\mathbf{P}} 3\,; \quad 3 \succ_{\mathbf{P}} 1\,,$$

and therefore $\mathbf{P} \notin \mathcal{P}_n^{\text{ST}}$. The claim follows. ■

Appendix D. Proof of Theorem 8

Proof

Proof of Part 1:

We first prove the 'if' direction. Suppose $\exists \mathbf{x} \in \mathbb{R}^n, \mathbf{y} \in \mathbb{R}_+^n$ with $\mathbf{x}^\top\mathbf{y} = 0$ such that $\psi_{\text{logit}}(\mathbf{P}) = \mathbf{x}\mathbf{y}^\top - \mathbf{y}\mathbf{x}^\top$. Then clearly $\mathbf{P} \in \mathcal{P}_n^{\text{LR(logit,2)}}$. If $n = 2$, then $\mathcal{P}_n = \mathcal{P}_n^{\text{ST}}$ and the claim follows trivially. Therefore assume $n \geq 3$; we will show that $\mathcal{P}_n = \mathcal{P}_n^{\text{ST}}$ in this case as well. Observe that for any $i \neq j$,

$$i \succ_{\mathbf{P}} j \iff P_{ij} > \tfrac{1}{2} \iff \psi_{\text{logit}}(P_{ij}) > 0 \iff x_i y_j > x_j y_i\,. \tag{1}$$

Now consider any three items i, j, k for which $i \succ_{\mathbf{P}} j$ and $j \succ_{\mathbf{P}} k$. We will show that $i \succ_{\mathbf{P}} k$. By Eq. (1), we have

$$x_i y_j \;>\; x_j y_i \tag{2}$$
$$x_j y_k \;>\; x_k y_j\,. \tag{3}$$

We will show that

$$x_i y_k \;>\; x_k y_i\,; \tag{4}$$

the claim that $i \succ_{\mathbf{P}} k$ will then follow from Eq. (1). By Eqs. (2-3), at most one of x_i, x_j, x_k can be zero; if any of them is zero, then Eq. (4) follows trivially. Therefore assume x_i, x_j, x_k are all non-zero. We will consider all 8 possibilities for the signs (positive or negative) of x_i, x_j, x_k, and will see that in each case, either Eq. (4) holds, or the combination of signs is simply not possible as it contradicts Eq. (2) or Eq. (3). The following table explains each of these cases; here $+$ indicates the corresponding element is positive and $-$ indicates that it is negative:

x_i	x_j	x_k	Explanation
$+$	$+$	$+$	By Eqs. (2-3), $(x_i y_j)(x_j y_k) > (x_j y_i)(x_k y_j)$; dividing by $x_j y_j > 0$ gives $x_i y_k > x_k y_i$
$+$	$+$	$-$	Not possible (implies $x_j y_k > 0$ and $x_k y_j < 0$, which contradicts Eq. (3))
$+$	$-$	$+$	Not possible (implies $x_j y_k < 0$ and $x_k y_j > 0$, which contradicts Eq. (3))
$+$	$-$	$-$	In this case $x_i y_k > 0$ and $x_k y_i < 0$; thus $x_i y_k > x_k y_i$
$-$	$+$	$+$	Not possible (implies $x_i y_j < 0$ and $x_j y_i > 0$, which contradicts Eq. (2))
$-$	$+$	$-$	Not possible (implies $x_i y_j > 0$ and $x_j y_i < 0$, which contradicts Eq. (2))
$-$	$-$	$+$	Not possible (implies $x_j y_k < 0$ and $x_k y_j > 0$, which contradicts Eq. (3))
$-$	$-$	$-$	By Eqs. (2-3), $(x_i y_j)(x_j y_k) < (x_j y_i)(x_k y_j)$; dividing by $x_j y_j < 0$ gives $x_i y_k > x_k y_i$

Thus Eq. (4) holds in all realizable settings of x_i, x_j, x_k. This proves the 'if' direction.

Next we prove the 'only if' direction. Suppose $\mathbf{P} \in (\mathcal{P}_n^{\text{LR(logit,2)}} \cap \mathcal{P}_n^{\text{ST}})$. Since $\mathbf{P} \in \mathcal{P}_n^{\text{LR(logit,2)}}$ and ψ_{logit} is 'skew-symmetric' around $\frac{1}{2}$ in that $\psi_{\text{logit}}(p) = -\psi_{\text{logit}}(1-p)\ \forall p$, we have $\exists \mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ such that $\psi_{\text{logit}}(\mathbf{P}) = \mathbf{x}\mathbf{y}^\top - \mathbf{y}\mathbf{x}^\top$. For any $\alpha \in \mathbb{R}$, define

$$\begin{aligned} \mathbf{y}^\alpha &= \mathbf{y} + \alpha \mathbf{x} \\ \mathbf{x}^\alpha &= \mathbf{x} - \frac{((\mathbf{y}^\alpha)^\top \mathbf{x})}{((\mathbf{y}^\alpha)^\top \mathbf{y}^\alpha)} \mathbf{y}^\alpha . \end{aligned}$$

Then it can be verified that for all $\alpha \in \mathbb{R}$, $(\mathbf{x}^\alpha)^\top \mathbf{y}^\alpha = 0$ and $\psi_{\text{logit}}(\mathbf{P}) = \mathbf{x}^\alpha (\mathbf{y}^\alpha)^\top - \mathbf{y}^\alpha (\mathbf{x}^\alpha)^\top$. We will show that there is a choice of α for which $\mathbf{y}^\alpha \in \mathbb{R}_+^n$. Specifically, since $\mathbf{P} \in \mathcal{P}_n^{\text{ST}}$, we have that $\exists i \in [n]$ such that $P_{ij} \geq \frac{1}{2}\ \forall j$, i.e. such that $j \not\succ_{\mathbf{P}} i\ \forall j$. Choose $\alpha = -\frac{y_i}{x_i}$. We claim that $y_k^\alpha \geq 0\ \forall k$. To see this, suppose for the sake of contradiction that $\exists k$ such that $y_k^\alpha < 0$. By definition of $\mathbf{y}^\alpha$, this gives $y_k + \alpha x_k = y_k + (-\frac{y_i}{x_i}) x_k < 0$, which gives $x_i y_k - y_i x_k < 0$. But

$$x_i y_k - y_i x_k < 0 \iff \psi_{\text{logit}}(P_{ik}) < 0 \iff P_{ik} < \tfrac{1}{2} \iff k \succ_{\mathbf{P}} i ,$$

which contradicts the choice of i. Thus for α chosen as above, we have $\mathbf{y}^\alpha \in \mathbb{R}_+^n$. This proves the 'only if' direction.

Proof of Part 2:
For the 'if' direction, suppose $\exists \mathbf{x} \in \mathbb{R}^n$ with $\mathbf{x}^\top \mathbf{e}_n = 0$ such that $\psi_{\text{logit}}(\mathbf{P}) = \mathbf{x}\mathbf{e}_n^\top - \mathbf{e}_n \mathbf{x}^\top$. Define $\mathbf{w} \in \mathbb{R}_+^n$ as $w_i = \exp(x_i)\ \forall i$. Then it can be verified that $P_{ij} = \frac{w_i}{w_i + w_j}\ \forall i, j$, and so $\mathbf{P} \in \mathcal{P}_n^{\text{BTL}}$.

For the 'only if' direction, suppose that $\mathbf{P} \in \mathcal{P}_n^{\text{BTL}}$. Then $\exists \mathbf{w} \in \mathbb{R}_{++}^n$ such that $P_{ij} = \frac{w_i}{w_i + w_j}\ \forall i, j$. Define $\mathbf{x} \in \mathbb{R}^n$ as $x_i = \log(w_i) - \frac{1}{n}\sum_{j=1}^n \log(w_j)$. Then it can be verified that $\mathbf{x}^\top \mathbf{e}_n = 0$ and $\psi_{\text{logit}}(\mathbf{P}) = \mathbf{x}\mathbf{e}_n^\top - \mathbf{e}_n \mathbf{x}^\top$. ■

Appendix E. Proof of Theorem 9

Proof
Proof of Part 1:
This follows the same steps as in the proof of Theorem 8.

Proof of Part 2:
This follows immediately from the definition of $\mathcal{P}_n^{\text{Thu}}$. ■

Appendix F. Proof of Proposition 10

Proof Let $\sigma \in \mathcal{S}_n, p \in [0, \frac{1}{2})$ be such that

$$\sigma(i) < \sigma(j) \implies P_{ij} = 1 - p\,.$$

Without loss of generality, we will assume that σ is the identity permutation, $\sigma(i) = i$ (note that reordering the rows and columns of a matrix according to a permutation preserves the rank). Then

$$\psi(P_{ij}) = \begin{cases} -\psi(p) & \text{if } i < j \\ \psi(p) & \text{if } i > j \\ 0 & \text{if } i = j. \end{cases}$$

In other words, $\psi(\mathbf{P}) = \psi(p) \cdot \mathbf{A}$, where the matrix $\mathbf{A}$ is given by

$$A_{ij} = \begin{cases} -1 & \text{if } i < j \\ 1 & \text{if } i > j \\ 0 & \text{if } i = j. \end{cases}$$

Clearly, $\mathrm{rank}(\psi(\mathbf{P})) = \mathrm{rank}(\mathbf{A})$. It can be seen that if n is even, then the columns of $\mathbf{A}$ are linearly independent, so that $\mathrm{rank}(\mathbf{A}) = n$. Similarly, if n is odd, then the nullspace of $\mathbf{A}$ contains just the span of the vector $(1, -1, 1, -1, \ldots -1, 1)^\top$, so that $\mathrm{rank}(\mathbf{A}) = n-1$. This proves the result. ■

Appendix G. Proof of Theorem 13

For $\mathbf{P}, \mathbf{Q} \in \mathcal{P}_n$, we will denote the fraction of pairs on which $\mathbf{P}$ and $\mathbf{Q}$ disagree as

$$\text{dis}(\mathbf{P}, \mathbf{Q}) \; = \; \frac{1}{\binom{n}{2}} \sum_{i<j} \mathbf{1}\Big(\big(i \succ_{\mathbf{P}} j\big) \wedge \big(j \succ_{\mathbf{Q}} i\big) \Big) + \mathbf{1}\Big(\big(j \succ_{\mathbf{P}} i\big) \wedge \big(i \succ_{\mathbf{Q}} j\big) \Big)\,.$$

We will need the following lemma:

Lemma 20 *Let $\mathbf{P} \in \mathcal{P}_n^{\text{ST}}$ and $\mathbf{Q} \in \mathcal{P}_n$. Let $\gamma > 1$, and let PR be any γ-approximate pairwise ranking algorithm. Let $\widehat{\sigma} = \text{PR}(\mathbf{Q}) \in \mathcal{S}_n$. Then*

$$\text{dis}(\widehat{\sigma}, \mathbf{P}) \; \leq \; (1 + \gamma) \cdot \text{dis}(\mathbf{Q}, \mathbf{P})\,.$$

Proof For any permutation $\sigma \in S_n$, define $\mathbf{B}^\sigma \in \mathcal{P}_n^{\text{DO}}$ as follows:

$$B^\sigma_{ij} = \begin{cases} 1 & \text{if } \sigma(i) < \sigma(j) \\ 0 & \text{if } \sigma(i) > \sigma(j) \\ \frac{1}{2} & \text{if } i = j \,. \end{cases}$$

Similarly, for any $\mathbf{P} \in \mathcal{P}_n$, define $\mathbf{B}^\mathbf{P} \in \mathcal{P}_n^{\text{DTour}}$ as follows:

$$B^\mathbf{P}_{ij} = \begin{cases} 1 & \text{if } i \succ_\mathbf{P} j \\ 0 & \text{if } j \succ_\mathbf{P} i \\ \frac{1}{2} & \text{if } i = j \,. \end{cases}$$

Then for any $\sigma \in \mathcal{S}_n$ and $\mathbf{P}, \mathbf{Q} \in \mathcal{P}_n$, we can write

$$\begin{aligned} \text{dis}(\sigma, \mathbf{P}) &= \frac{1}{2} \|\mathbf{B}^\sigma - \mathbf{B}^\mathbf{P}\|_1 &\quad (5) \\ \text{dis}(\mathbf{P}, \mathbf{Q}) &= \frac{1}{2} \|\mathbf{B}^\mathbf{P} - \mathbf{B}^\mathbf{Q}\|_1 \,. &\quad (6) \end{aligned}$$

Now, for the given $\mathbf{P}$, $\mathbf{Q}$, and $\widehat{\sigma}$, we have

$$\begin{aligned} \text{dis}(\widehat{\sigma}, \mathbf{P}) &= \frac{1}{2} \|\mathbf{B}^\mathbf{P} - \mathbf{B}^{\widehat{\sigma}}\|_1 \quad \text{(from Equation 5)} \\ &\leq \frac{1}{2} \|\mathbf{B}^\mathbf{P} - \mathbf{B}^\mathbf{Q}\|_1 + \frac{1}{2} \|\mathbf{B}^\mathbf{Q} - \mathbf{B}^{\widehat{\sigma}}\|_1 \quad \text{(by triangle inequality)} \\ &\leq \frac{1}{2} \|\mathbf{B}^\mathbf{P} - \mathbf{B}^\mathbf{Q}\|_1 + \frac{\gamma}{2} \cdot \min_\sigma \|\mathbf{B}^\mathbf{Q} - \mathbf{B}^{\widehat{\sigma}}\|_1 \quad (\text{since } \widehat{\sigma} = \mathit{PR}(\mathbf{Q})) \\ &\leq \frac{1}{2} \|\mathbf{B}^\mathbf{P} - \mathbf{B}^\mathbf{Q}\|_1 + \frac{\gamma}{2} \cdot \|\mathbf{B}^\mathbf{Q} - \mathbf{B}^\pi\|_1 \,, \quad \text{where } \pi = \mathit{PR}(\mathbf{P}) \\ &= \frac{1}{2} \|\mathbf{B}^\mathbf{P} - \mathbf{B}^\mathbf{Q}\|_1 + \frac{\gamma}{2} \cdot \|\mathbf{B}^\mathbf{Q} - \mathbf{B}^\mathbf{P}\|_1 \\ &\qquad (\text{since } \mathbf{P} \in \mathcal{P}_n^{\text{ST}}, \text{ we have } \text{dis}(\pi, \mathbf{P}) = 0 \text{ and therefore } \mathbf{B}^\pi = \mathbf{B}^\mathbf{P}) \\ &= \left(\frac{1+\gamma}{2}\right) \cdot \|\mathbf{B}^\mathbf{P} - \mathbf{B}^\mathbf{Q}\|_1 \\ &= (1+\gamma) \cdot \text{dis}(\mathbf{P}, \mathbf{Q}) \,. \end{aligned}$$

This proves the claim. ■

We will also need the following theorem of Keshavan et al. (2009), which bounds the spectral norm of any (incomplete) matrix in terms of the maximum size of its entries:

Theorem 21 (Keshavan et al. (2009)) *For any matrix $\mathbf{Z} \in \mathbb{R}^{n \times n}$ and any set $\Omega \subseteq [n] \times [n]$,*

$$\|\mathbf{Z}^\Omega\|_2 \;\leq\; \frac{2|\Omega|}{n} \max_{(i,j) \in \Omega} |Z_{ij}| \,.$$

Proof [Proof of Theorem 13] Let C, C' be the constants given by Theorem 2. Let m and K satisfy the given conditions, and let $\widehat{\mathbf{P}}$ denote the (incomplete) empirical comparison matrix constructed from S in the LRPR algorithm. We can write

$$\psi(\widehat{\mathbf{P}})^E \;=\; \psi(\mathbf{P})^E + \underbrace{(\psi(\widehat{\mathbf{P}}) - \psi(\mathbf{P}))^E}_{\mathbf{Z}^E}.$$

We will show that with high probability, $\|\mathbf{Z}^E\|_2 = \|(\psi(\widehat{\mathbf{P}}) - \psi(\mathbf{P}))^E\|_2$ is small; by Theorem 2, it will then follow that the completed matrix $\widehat{\widehat{\mathbf{M}}}$ obtained by running OptSpace on $\widehat{\mathbf{M}} = \psi(\widehat{\mathbf{P}})^E$ in the LRPR algorithm is close to $\psi(\mathbf{P})$, which in turn will imply that the number of disagreements between the completed comparison matrix $\widehat{\widehat{\mathbf{P}}}$ and the true preference matrix $\mathbf{P}$ is small. This will then allow us to show (by virtue of Lemma 20) that the ranking $\widehat{\sigma}$ produced by running *PR* on $\widehat{\widehat{\mathbf{P}}}$ is also good w.r.t. $\mathbf{P}$.

For any fixed $E \subseteq [n] \times [n]$ of size $|E| = m$, consider the following event:

$$\begin{aligned} A_E \;&=\; \Big\{ |\widehat{P}_{ij} - P_{ij}| < \frac{P_{\min}}{2} \quad \forall (i,j) \in E \Big\} \\ &\subseteq\; \Big\{ \widehat{P}_{ij} \in \Big(\frac{P_{\min}}{2}, 1 - \frac{P_{\min}}{2}\Big) \quad \forall (i,j) \in E \Big\}. \end{aligned}$$

Since $K \geq \frac{11\log(n)}{P_{\min}^2}$, it is easy to see (using Hoeffding's inequality) that

$$P(A_E) \;\geq\; 1 - \frac{1}{2n^3}. \tag{7}$$

Now, let

$$\tau \;=\; \frac{m}{C'\kappa^2 n\sqrt{r}} \min\Big(\sqrt{\frac{\epsilon}{1+\gamma}} \frac{\Delta_{\min}^{\mathbf{P},\psi}}{2}, \Sigma_{\min} \Big).$$

Then (for fixed E), we have,

$$\begin{aligned}
\mathbf{P}\big(\|\mathbf{Z}^E\|_2 > \tau\big) \;&\leq\; \mathbf{P}\Big(\max_{(i,j)\in E} |Z_{ij}^E| > \frac{n}{m}\tau \Big) \quad \text{(by Theorem 21)} \\
&=\; \mathbf{P}\Big(\max_{(i,j)\in E} |\psi(\widehat{P}_{ij}) - \psi(P_{ij})| > \frac{n}{m}\tau \Big) \\
&=\; \mathbf{P}\Big(\exists (i,j) \in E : |\psi(\widehat{P}_{ij}) - \psi(P_{ij})| > \frac{n}{m}\tau \Big) \\
&\leq\; \mathbf{P}\Big(\exists (i,j) \in E : |\psi(\widehat{P}_{ij}) - \psi(P_{ij})| > \frac{n}{m}\tau \,\Big|\, A_E \Big) \mathbf{P}\big(A_E\big) + \mathbf{P}\big(A_E^c\big) \\
&\leq\; \Big(\sum_{(i,j)\in E} \mathbf{P}\Big(|\psi(\widehat{P}_{ij}) - \psi(P_{ij})| > \frac{n}{m}\tau \,\Big|\, A_E \Big) \mathbf{P}\big(A_E\big) \Big) + \mathbf{P}\big(A_E^c\big) \quad \text{(by union bound)} \\
&\leq\; \Big(\sum_{(i,j)\in E} \mathbf{P}\Big(|\psi(\widehat{P}_{ij}) - \psi(P_{ij})| > \frac{n}{m}\tau \,\Big|\, A_E \Big) \mathbf{P}\big(A_E\big) \Big) + \frac{1}{2n^3} \quad \text{(by Eq. (7))} \\
&\leq\; \Big(\sum_{(i,j)\in E} \mathbf{P}\Big(|\widehat{P}_{ij} - P_{ij}| > \frac{n}{mL}\tau \,\Big|\, A_E \Big) \mathbf{P}\big(A_E\big) \Big) + \frac{1}{2n^3} \\
&\qquad\qquad \text{(by definition of } A_E \text{ and since } \psi \text{ is } L\text{-Lipschitz in } [\tfrac{P_{\min}}{2}, 1 - \tfrac{P_{\min}}{2}]) \\
&\leq\; \Big(\sum_{(i,j)\in E} \mathbf{P}\Big(|\widehat{P}_{ij} - P_{ij}| > \frac{n}{mL}\tau \Big) \Big) + \frac{1}{2n^3} \\
&\leq\; \frac{1}{2n^3} + \frac{1}{2n^3} = \frac{1}{n^3} \quad \text{(by Hoeffding's inequality, since } K \geq \tfrac{m^2L^2\log(n)}{n^2\tau^2}).
\end{aligned}$$

Since the above holds for all fixed E, we have that it holds under the random choice of E as well. Therefore, under the random choice of both E and S, we have that with probability at least $1 - \frac{1}{n^3}$,

$$\|\mathbf{Z}^E\|_2 \;\leq\; \frac{m}{C'\kappa^2 n\sqrt{r}} \min\left(\sqrt{\frac{\epsilon}{1+\gamma}} \frac{\Delta^{\mathbf{P},\psi}_{\text{min}}}{2}\,, \Sigma_{\text{min}}\right).$$

Thus, by Theorem 2, we have that with probability at least $1 - \frac{2}{n^3}$,

$$\|\overline{\widehat{\mathbf{M}}} - \psi(\mathbf{P})\|_F \;\leq\; n\sqrt{\frac{\epsilon}{1+\gamma}} \frac{\Delta^{\mathbf{P},\psi}_{\text{min}}}{2}\,. \tag{8}$$

The above inequality implies that there can be at most $\frac{\epsilon}{1+\gamma}n^2$ pairs for which $|\overline{\widehat{M}}_{ij} - \psi(P_{ij})|$ is greater than or equal to $\frac{\Delta^{\mathbf{P},\psi}_{\text{min}}}{2}$. Thus, for at least $\frac{\epsilon}{1+\gamma}$ fraction of pairs $i < j$, we must have

$$|\overline{\widehat{M}}_{ij} - \psi(P_{ij})| < \frac{\Delta^{\mathbf{P},\psi}_{\text{min}}}{2} \quad \text{and} \quad |\overline{\widehat{M}}_{ji} - \psi(P_{ji})| < \frac{\Delta^{\mathbf{P},\psi}_{\text{min}}}{2}\,, \tag{9}$$

which gives for these pairs

$$\begin{aligned}
i \succ_{\mathbf{P}} j \quad &\iff \quad P_{ij} > \tfrac{1}{2} \\
&\iff \quad \psi(P_{ij}) > \psi(\tfrac{1}{2}) \quad \text{(since } \psi \text{ is strictly increasing)} \\
&\iff \quad \overline{\widehat{M}}_{ij} > \psi(\tfrac{1}{2}) \text{ and } \overline{\widehat{M}}_{ji} < \psi(\tfrac{1}{2}) \quad \text{(from Eq. (9) and definition of } \Delta^{\mathbf{P},\psi}_{\text{min}}) \\
&\iff \quad \overline{\widehat{P}}_{ij} > \tfrac{1}{2} \quad \text{(by construction of } \overline{\widehat{\mathbf{P}}}\text{, and since } \psi \text{ is strictly increasing)} \\
&\iff \quad i \succ_{\overline{\widehat{\mathbf{P}}}} j\,,
\end{aligned}$$

and similarly $j \succ_{\mathbf{P}} i \iff j \succ_{\overline{\widehat{\mathbf{P}}}} i$. Thus, we have that with probability at least $1 - \frac{2}{n^3}$, the fraction of pairs on which $\overline{\widehat{\mathbf{P}}}$ and $\mathbf{P}$ and disagree is at most $\frac{\epsilon}{1+\gamma}$:

$$\text{dis}(\overline{\widehat{\mathbf{P}}}, \mathbf{P}) \;\leq\; \frac{\epsilon}{1+\gamma}\,.$$

By Lemma 20, we thus have with probability at least $1 - \frac{2}{n^3}$, $\text{dis}(\widehat{\sigma}, \mathbf{P}) \;\leq\; \epsilon\,.$ ■

Appendix H. Proof of Lemma 14

Proof We have $\text{rank}(\psi_{\text{logit}}(\mathbf{P})) \leq 2 = r$. Let

$$\mathbf{u}_1 = \frac{\sqrt{n}\mathbf{x}}{\|\mathbf{x}\|_2},\, \mathbf{u}_2 = \frac{\sqrt{n}\mathbf{y}}{\|\mathbf{y}\|_2},\, \mathbf{v}_1 = \mathbf{u}_2,\, \mathbf{v}_2 = -\mathbf{u}_1,\, \Sigma_1 = \Sigma_2 = \frac{\|\mathbf{x}\|_2\|\mathbf{y}\|_2}{n}\,.$$

Then we have

$$\begin{aligned}
\psi_{\text{logit}}(\mathbf{P}) \quad &= \quad \mathbf{x}\mathbf{y}^\top - \mathbf{y}\mathbf{x}^\top \\
&= \quad \Sigma_1\mathbf{u}_1\mathbf{v}_1^\top + \Sigma_2\mathbf{u}_2\mathbf{v}_2^\top \\
&= \quad \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top\,,
\end{aligned}$$

where $\mathbf{U} = [\mathbf{u}_1\, \mathbf{u}_2]$ with $\mathbf{U}^\top\mathbf{U} = n\mathbf{I}_2$, $\mathbf{V} = [\mathbf{v}_1\, \mathbf{v}_2]$ with $\mathbf{V}^\top\mathbf{V} = n\mathbf{I}_2$, and $\mathbf{\Sigma} = \mathrm{diag}(\Sigma_1, \Sigma_2)$. Now, for all $i \in [n]$, we have

$$\begin{aligned}\sum_{k=1}^{2} U_{ik}^2 &= \left(\frac{\sqrt{n}x_i}{\|\mathbf{x}\|_2}\right)^2 + \left(\frac{\sqrt{n}y_i}{\|\mathbf{y}\|_2}\right)^2 \\ &\le \left(\frac{x_{\max}}{x_{\min}}\right)^2 + \left(\frac{y_{\max}}{y_{\min}}\right)^2 \;=\; 2\mu\,.\end{aligned}$$

Moreover, for all $i, j \in [n]$, we have

$$\begin{aligned}\Big|\sum_{k=1}^{2} U_{ik}\Big(\frac{\Sigma_k}{\Sigma_1}\Big)V_{jk}\Big| &= \Big|\sum_{k=1}^{2} U_{ik}V_{jk}\Big| \\ &= \left|\frac{nx_iy_j}{\|\mathbf{x}\|_2\|\mathbf{y}\|_2} - \frac{nx_jy_i}{\|\mathbf{x}\|_2\|\mathbf{y}\|_2}\right| \\ &\le \left|\frac{nx_iy_j}{\|\mathbf{x}\|_2\|\mathbf{y}\|_2}\right| + \left|\frac{nx_jy_i}{\|\mathbf{x}\|_2\|\mathbf{y}\|_2}\right| \\ &\le 2\Big(\frac{x_{\max}}{x_{\min}}\Big)\Big(\frac{y_{\max}}{y_{\min}}\Big) \\ &\le \Big(\frac{x_{\max}}{x_{\min}}\Big)^2 + \Big(\frac{y_{\max}}{y_{\min}}\Big)^2 \quad \text{(since } 2ab \le a^2 + b^2\text{)} \\ &= \sqrt{2}\big(\sqrt{2}\mu\big)\,.\end{aligned}$$

Thus $\psi_{\text{logit}}(\mathbf{P})$ is $(\mu, \sqrt{2}\mu)$-incoherent. ■

Appendix I. Proof of Corollary 15

Proof This follows by setting $\mathbf{x} \in \mathbb{R}^n$ as $x_i = \log(w_i) - \frac{1}{n}\sum_{j=1}^{n}\log(w_j)$ and $\mathbf{y} \in \mathbb{R}_+^n$ as $\mathbf{y} = \mathbf{e}_n$, and then proceeding as in Lemma 14. ■

Appendix J. Proof of Lemma 16

Proof Let $q \in (0, \frac{1}{2}]$. We know that ψ_{logit} is L'-Lipschitz in $[\frac{q}{2}, 1 - \frac{q}{2}]$ for any $L' \ge L$, where $L = \sup_{p\in[\frac{q}{2},1-\frac{q}{2}]} |\psi'_{\text{logit}}(p)|$. We have $\psi'_{\text{logit}}(p) = \frac{1}{p(1-p)}$, and therefore

$$\begin{aligned}L &= \sup_{p\in[\frac{q}{2},1-\frac{q}{2}]} \frac{1}{p(1-p)} \\ &\le \frac{1}{\frac{q}{2}(1-\frac{q}{2})} \quad \text{(since } \tfrac{1}{p(1-p)} \text{ in the above interval is maximized at } p = \tfrac{q}{2} \text{ and } p = 1 - \tfrac{q}{2}\text{)} \\ &= \frac{4}{q(2-q)} \\ &\le \frac{8}{3q} \quad \text{(since } q \le \tfrac{1}{2} \text{ and therefore } 2 - q \ge \tfrac{3}{2}\text{)}.\end{aligned}$$

The result follows. ■

Appendix K. Proof of Corollary 17

Proof As in Lemma 14, we can define matrices $\mathbf{U}, \mathbf{V}, \mathbf{\Sigma}$ such that $\psi_{\text{logit}}(\mathbf{P}) = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ is a rank-2, $(\mu, \sqrt{2}\mu)$-incoherent matrix with $\Sigma_{\min} = \Sigma_{\max} = \overline{x}\,\overline{y}$, and therefore $\kappa = \frac{\Sigma_{\max}}{\Sigma_{\min}} = 1$. Moreover, by Lemma 16, we have that ψ_{logit} is $(\frac{8}{3P_{\min}})$-Lipschitz in $[\frac{P_{\min}}{2}, 1 - \frac{P_{\min}}{2}]$, and by definition, we have $\Delta_{\min}^{\mathbf{P},\psi_{\text{logit}}} = \Delta$. The result then follows by substituting the above quantities in Theorem 13. ■

Appendix L. Proof of Corollary 18

Proof The proof follows from Corollary 17, by observing that here $\psi_{\text{logit}}(\mathbf{P}) = \mathbf{x}\mathbf{y}^\top - \mathbf{y}\mathbf{x}^\top$ and $\mathbf{x}^\top\mathbf{y} = 0$ with $x_i = \log w_i - \frac{1}{n}\sum_{j=1}^n \log(w_j)\ \forall i$ and $y_i = 1\ \forall i$, and $P_{\min} = \frac{1}{b+1}$. ■

Appendix M. Proof of Theorem 19

Proof The proof is broadly similar to that of Theorem 13. Let C, C' be the constants given by Theorem 2. Let $\beta \leq \alpha \frac{\Sigma_{\min}}{C'\kappa^2\sqrt{r}}$ for some $\alpha \in (0,1)$. Let m and K satisfy the given conditions, and let $\widehat{\mathbf{P}}$ denote the (incomplete) empirical comparison matrix constructed from S in the LRPR algorithm. Here we can write

$$\begin{aligned}
\psi(\widehat{\mathbf{P}})^E &= \mathbf{M}^E + \underbrace{(\psi(\widehat{\mathbf{P}}) - \mathbf{M})^E}_{\mathbf{Z}^E} \\
&= \mathbf{M}^E + \underbrace{(\psi(\widehat{\mathbf{P}}) - \psi(\mathbf{P}))^E}_{\mathbf{X}^E} + \underbrace{(\psi(\mathbf{P}) - \mathbf{M})^E}_{\mathbf{Y}^E}.
\end{aligned}$$

By assumption, we have

$$\max_{(i,j)\in E} |Y_{ij}^E| \;\leq\; \|\mathbf{Y}_E\|_F \;\leq\; \beta \;\leq\; \alpha \frac{\Sigma_{\min}}{C'\kappa^2\sqrt{r}}.$$

As in the proof of Theorem 13, we can show that under the random choice of both E and S, we have that with probability at least $1 - \frac{1}{n^3}$,

$$\max_{(i,j)\in E} |X_{ij}^E| \;\leq\; \frac{1}{C'\kappa^2\sqrt{r}} \min\left(\sqrt{\frac{\epsilon}{1+\gamma}}\frac{\Delta_{\min}^{\mathbf{P},\psi}}{2}, (1-\alpha)\Sigma_{\min}\right).$$

Thus, by Theorem 21, we have with probability at least $1 - \frac{1}{n^3}$,

$$
\begin{aligned}
\|\mathbf{Z}^E\|_2 &\le \frac{2m}{n} \max_{(i,j)\in E} |Z_{ij}^E| \\
&\le \frac{2m}{n}\Big(\max_{(i,j)\in E} |X_{ij}^E| + \max_{(i,j)\in E} |Y_{ij}^E| \Big) \\
&\le \frac{2m}{C'\kappa^2 n\sqrt{r}} \min\left(\sqrt{\frac{\epsilon}{1+\gamma}} \frac{\Delta_{\min}^{\mathbf{P},\psi}}{2} + \alpha\Sigma_{\min}\,, \Sigma_{\min} \right).
\end{aligned}
$$

Now, by Theorem 2, we have that with probability at least $1 - \frac{1}{n^3}$,

$$
\|\widehat{\overline{\mathbf{M}}} - \mathbf{M}\|_F \le C'\kappa^2 \frac{n^2\sqrt{r}}{2m} \|\mathbf{Z}^E\|_2\,,
$$

provided that $\|\mathbf{Z}^E\|_2 \le \frac{2m}{C'\kappa^2 n\sqrt{r}}\Sigma_{\min}$. Thus, combining the above two statements, we have that with probability at least $1 - \frac{2}{n^3}$,

$$
\begin{aligned}
\|\widehat{\overline{\mathbf{M}}} - \mathbf{M}\|_F &\le C'\kappa^2 \frac{n^2\sqrt{r}}{2m} \left(\frac{2m}{C'\kappa^2 n\sqrt{r}} \left(\sqrt{\frac{\epsilon}{1+\gamma}} \frac{\Delta_{\min}^{\mathbf{P},\psi}}{2} + \alpha\Sigma_{\min} \right) \right) \\
&= n \left(\sqrt{\frac{\epsilon}{1+\gamma}} \frac{\Delta_{\min}^{\mathbf{P},\psi}}{2} + \alpha\Sigma_{\min} \right).
\end{aligned}
$$

This gives that with probability at least $1 - \frac{2}{n^3}$,

$$
\begin{aligned}
\|\widehat{\overline{\mathbf{M}}} - \psi(\mathbf{P})\|_F &\le \|\widehat{\overline{\mathbf{M}}} - \mathbf{M}\|_F + \|\mathbf{M} - \psi(\mathbf{P})\|_F \\
&\le n \left(\sqrt{\frac{\epsilon}{1+\gamma}} \frac{\Delta_{\min}^{\mathbf{P},\psi}}{2} + \alpha\Sigma_{\min} \right) + \alpha \frac{\Sigma_{\min}}{C'\kappa^2\sqrt{r}} \\
&= n \left(\sqrt{\frac{\epsilon}{1+\gamma}} + \frac{2\alpha\Sigma_{\min}}{\Delta_{\min}^{\mathbf{P},\psi}} \left(1 + \frac{1}{C'\kappa^2 n\sqrt{r}} \right) \right) \frac{\Delta_{\min}^{\mathbf{P},\psi}}{2} \\
&\le n \left(\sqrt{\frac{\epsilon}{1+\gamma}} + \frac{4\alpha\Sigma_{\min}}{\Delta_{\min}^{\mathbf{P},\psi}} \right) \frac{\Delta_{\min}^{\mathbf{P},\psi}}{2} \quad \text{(since } \frac{1}{C'\kappa^2 n\sqrt{r}} \le 1\text{)}.
\end{aligned}
$$

Similar to the argument in Theorem 13, this gives that with probability at least $1 - \frac{2}{n^3}$,

$$
\text{dis}(\widehat{\overline{\mathbf{P}}}, \mathbf{P}) \le \frac{\epsilon}{1+\gamma} + \left(\frac{4\alpha\Sigma_{\min}}{\Delta_{\min}^{\mathbf{P},\psi}} \right)^2 + 8\sqrt{\frac{\epsilon}{1+\gamma}} \left(\frac{\alpha\Sigma_{\min}}{\Delta_{\min}^{\mathbf{P},\psi}} \right).
$$

Again, by Lemma 20, this gives with probability at least $1 - \frac{2}{n^3}$,

$$
\text{dis}(\widehat{\sigma}, \mathbf{P}) \le \epsilon + (1+\gamma) \left(\frac{4\alpha\Sigma_{\min}}{\Delta_{\min}^{\mathbf{P},\psi}} \right)^2 + 8\sqrt{\epsilon(1+\gamma)} \left(\frac{\alpha\Sigma_{\min}}{\Delta_{\min}^{\mathbf{P},\psi}} \right).
$$

■

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.

Nir Ailon, Moses Charikar, and Alantha Newman. Aggregating inconsistent information: Ranking and clustering. *Journal of the ACM*, 55(5):23, 2008.

Noga Alon. Ranking tournaments. *SIAM Journal of Discrete Mathematics*, 20(1):137–142, 2006.

Ralph A. 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 19th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2008.

Mark Braverman and Elchanan Mossel. Sorting from noisy information. *arXiv preprint arXiv:0910.1191*, 2009.

Róbert Busa-Fekete and Eyke Hüllermeier. A survey of preference-based online learning with bandit algorithms. In *Proceedings of the 25th International Conference on Algorithmic Learning Theory*, 2014.

Jian-Feng Cai, Emmanuel J. Candès, and Zuowei Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2010.

Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.

Emmanuel J. Candès and Terence Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.

A. H. Copeland. A 'reasonable' social welfare function. In *Seminar on Mathematics in Social Sciences, University of Michigan*, 1951.

Don Coppersmith, Lisa Fleischer, and Atri Rudra. Ordering by weighted number of wins gives a good ranking for weighted tournaments. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2006.

H. A. David. *The Method of Paired Comparisons*. Oxford University Press, 1988.

David F. Gleich and Lek-Heng Lim. Rank aggregation via nuclear norm minimization. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2011.

David R. Hunter. MM algorithms for generalized Bradley-Terry models. *The Annals of Statistics*, 32(1):384–406, 2004.

Prateek Jain, Raghu Meka, and Inderjit S. Dhillon. Guaranteed rank minimization via singular value projection. In *Advances in Neural Information Processing Systems*, 2010.

Kevin G. Jamieson and Rob Nowak. Active ranking using pairwise comparisons. In *Advances in Neural Information Processing Systems*, 2011.

Xiaoye Jiang, Lek-Heng Lim, Yuan Yao, and Yinyu Ye. Statistical ranking and combinatorial Hodge theory. *Mathematical Programming*, 127(1):203–244, 2011.

James P. Keener. The Perron-Frobenius theorem and the ranking of football teams. *SIAM Review*, 35(1):80–93, 1993.

Maurice G. Kendall. Further contributions to the theory of paired comparisons. *Biometrics*, 11(1): 43–62, 1955.

Claire Kenyon-Mathieu and Warren Schudy. How to rank with few errors. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, 2007.

Raghunandan Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from noisy entries. In *Advances in Neural Information Processing Systems*, 2009.

Jean Lafond. Low rank matrix completion with exponential family noise. In *Proceedings of the 28th Conference on Learning Theory*, 2015.

Tyler Lu and Craig Boutilier. Learning Mallows models with pairwise preferences. In *Proceedings of the 28th International Conference on Machine Learning*, 2011.

R. D. Luce. *Individual Choice Behavior: A Theoretical Analysis*. Wiley, 1959.

Sahand Negahban, Sewoong Oh, and Devavrat Shah. Rank Centrality: Ranking from pair-wise comparisons. *arXiv preprint arXiv:1209.1688*, 2012.

Arun Rajkumar and Shivani Agarwal. A statistical convergence perspective of algorithms for rank aggregation from pairwise data. In *Proceedings of the 31st International Conference on Machine Learning*, 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 *Proceedings of the 32nd International Conference on Machine Learning*, 2015.

Thomas L. Saaty. *The Analytic Hierarchy Process*. McGraw Hill, 1980.

Nihar B. Shah, Sivaraman Balakrishnan, Adityanand Guntuboyina, and Martin J. Wainwright. Stochastically transitive models for pairwise comparisons: Statistical and computational issues. *arXiv preprint arXiv:1510.05610.v2*, 2015.

Louis L. Thurstone. A law of comparative judgment. *Psychological Review*, 34(4):273, 1927.

Fabian L. Wauthier, Michael I. Jordan, and Nebojsa Jojic. Efficient ranking from pairwise comparisons. In *Proceedings of the 30th International Conference on Machine Learning*, 2013.

Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. The K-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012.