

Estimation from Pairwise Comparisons: Sharp Minimax Bounds with Topology Dependence

Nihar B. Shah[†]

NIHAR@EECS.BERKELEY.EDU

Sivaraman Balakrishnan[#]

SIVA@STAT.CMU.EDU

Joseph Bradley[†]

JOSEPH.KURATA.BRADLEY@GMAIL.COM

Abhay Parekh[†]

PAREKH@BERKELEY.EDU

Kannan Ramchandran[†]

KANNANR@EECS.BERKELEY.EDU

Martin J. Wainwright[†] *

WAINWRIG@BERKELEY.EDU

[†] *Department of Electrical Engineering and Computer Sciences*

* *Department of Statistics,*

University of California, Berkeley

Berkeley, CA-94720, USA

[#] *Department of Statistics,*

Carnegie Mellon University

Pittsburgh, PA-15213, USA

Editor: Moritz Hardt

Keywords: Pairwise comparisons, graph, topology, ranking, crowdsourcing

Abstract

Data in the form of pairwise comparisons arises in many domains, including preference elicitation, sporting competitions, and peer grading among others. We consider parametric ordinal models for such pairwise comparison data involving a latent vector $w^* \in \mathbb{R}^d$ that represents the “qualities” of the d items being compared; this class of models includes the two most widely used parametric models—the Bradley-Terry-Luce (BTL) and the Thurstone models. Working within a standard minimax framework, we provide tight upper and lower bounds on the optimal error in estimating the quality score vector w^* under this class of models. The bounds depend on the topology of the comparison graph induced by the subset of pairs being compared, via the spectrum of the Laplacian of the comparison graph. Thus, in settings where the subset of pairs may be chosen, our results provide principled guidelines for making this choice. Finally, we compare these error rates to those under cardinal measurement models and show that the error rates in the ordinal and cardinal settings have identical scalings apart from constant pre-factors.

1. Introduction

In an increasing range of applications, it is of interest to elicit judgments from non-expert humans. For instance, in marketing, elicitation of preferences of consumers about products, either directly or indirectly, is a common practice (Green et al., 1981). The gathering of this and related data types has been greatly facilitated by the emergence of “crowdsourcing”

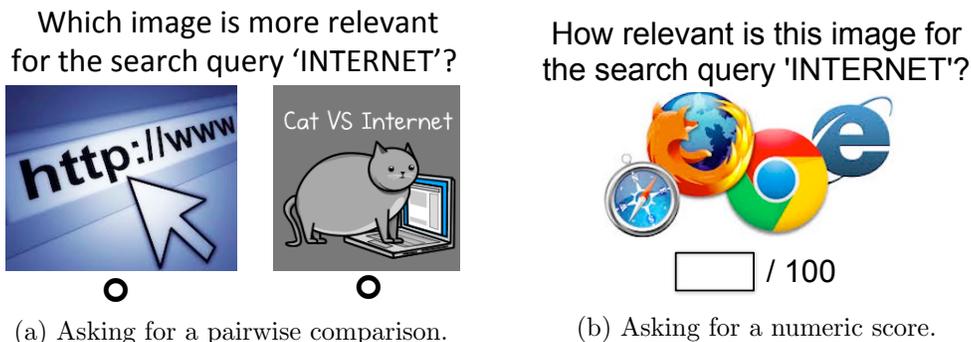


Figure 1: An example of eliciting judgments from people: rating the relevance of the result of a search query.

platforms such as Amazon Mechanical Turk: they have become powerful, low-cost tools for collecting human judgments (Khatib et al., 2011; Lang and Rio-Ross, 2011; von Ahn et al., 2008). Crowdsourcing is employed not only for collection of consumer preferences, but also for other types of data, including counting the number of malaria parasites in an image of a blood smear (Luengo-Oroz et al., 2012); rating responses of an online search engine to search queries (Kazai, 2011); or for labeling data for training machine learning algorithms (Hinton et al., 2012; Raykar et al., 2010; Deng et al., 2009). In a different domain, competitive sports can be understood as a mechanism for sequentially performing comparisons between individuals or teams (Ross, 2007; Herbrich et al., 2007). Finally, peer-grading in massive open online courses (MOOCs) (Piech et al., 2013) can be viewed as another form of elicitation.

A common method of elicitation is through pairwise comparisons. For instance, the decision of a consumer to choose one product over another constitutes a pairwise comparison between the two products. Workers in a crowdsourcing setup are often asked to compare pairs of items: for instance, they might be asked to identify the better of two possible results of a search engine, as shown in Figure 1a. Competitive sports such as chess or basketball also involve sequences of pairwise comparisons. From a modeling point of view, we can think of pairwise comparisons as a means of estimating the underlying “qualities” or “weights” of the items being compared (e.g., skill levels of chess players, relevance of search engine results, etc.). Each pairwise comparison can be viewed as a noisy sample of some function of the underlying pair of (real-valued) weights. Noise can arise from a variety of sources. When objective questions are posed to human subjects, noise can arise from their differing levels of expertise. In a sports competition, many sources of randomness can influence the outcome of any particular match between a pair of competitors. Thus, one important goal is to estimate the latent qualities based on noisy data in the form of pairwise comparisons. A related problem is that of experimental design: assuming that we can choose the subset of pairs to be compared (e.g., in designing a chess tournament), what choice leads to the most accurate estimation? Characterizing the fundamental difficulty of estimating the weights allow us to make this choice judiciously. These tasks are the primary focus of this paper.

In more detail, the focus of this paper is the aggregation from pairwise comparisons in a fairly broad class of parametric models. This class includes as special cases the two most

popular models for pairwise comparisons—namely, the Thurstone (Case V) (Thurstone, 1927) and the Bradley-Terry-Luce (BTL) (Bradley and Terry, 1952; Luce, 1959) models. The Thurstone (Case V) model has been used in a variety of both applied (Swets, 1973; Ross, 2007; Herbrich et al., 2007) and theoretical papers (Bramley, 2005; Krabbe, 2008; Nosofsky, 1985). Similarly, the BTL model has been popular in both theory and practice (Nosofsky, 1985; Atkinson et al., 1998; Koehler and Ridpath, 1982; Heldsinger and Humphry, 2010; Loewen et al., 2012; Green et al., 1981; Khairullah and Zions, 1987).

1.1 Related work

There is a vast literature on the Thurstone and BTL models, and we focus on those most closely related to our own work. Negahban et al. (2012) provide minimax bounds for the BTL model in the special case where the comparisons are evenly distributed. They focus on this case in order to complement their analysis of an algorithm based on a random walk. In their analysis, there is a gap between the achievable rate of the MLE and the lower bound. In contrast, our analysis eliminates this discrepancy and shows that MLE is an optimal estimator (up to constant factors) and achieves the minimax rate. In a work concurrent with our initial submission to arXiv (Shah et al., 2014), Hajek et al. (2014) consider the problem of estimation in the Plackett-Luce model, which extends the BTL model to comparisons of two or more items. They derive bounds on the minimax error rates under this model which, under certain conditions on the comparison graphs, are tight up to logarithmic factors. In general, their topology-dependent bounds rely on the degrees of the vertices in the comparison graph which makes the bounds quite loose. In contrast, our results are tight up to constants and, as we detail in the following sections, provide deeper insights into the role of the topology of the comparison graph. We elaborate on these differences in the appropriate sections in the sequel. We also note that the models studied in the present paper as well as in the aforementioned works fall under the broader paradigm of random utility models (Thurstone, 1927; Train, 2009; Azari Soufiani et al., 2013).

In other related works, Jagabathula and Shah (2008) design an algorithm for aggregating ordinal data when the underlying distribution over the permutations is assumed to be sparse. Ammar and Shah (2011) employ a different, maximum entropy approach towards parameterization and inference from partially ranked data. Rajkumar and Agarwal (2014) study the statistical convergence properties of several rank aggregation algorithms.

Our work assumes a fixed design setup. In this setup, the choice of which pairs to compare and the number of times to compare them is chosen ahead of time in a non-adaptive fashion. There is a parallel line of literature on “sorting” or “active ranking” from pairwise comparisons. For instance, Braverman and Mossel (2008) assume a noise model where the outcome of a pairwise comparison depends only on the relative ranks of the items being compared, and not on their actual ranks or values. On the other hand, Jamieson and Nowak (2011) consider the problem of ranking a set of items assuming that items can be embedded into a smaller-dimensional Euclidean space, and that the outcomes of the pairwise comparisons are based on the relative distances of these items from a fixed reference point in the Euclidean space.

A recent line of work considers a variant of the BTL and the Thurstone models where the comparisons may depend on some auxiliary unknown variable in addition to the items being compared; for instance, the accuracy of the individual making the comparison in an objective task. Chen et al. (2013) consider a crowdsourcing setup where the outcome depends on the worker’s expertise. They present algorithms for inference under such a model and present empirical evaluations. Yi et al. (2013) consider a problem in the spirit of collaborative filtering where certain unknown preferences of a certain user must be predicted based on the preferences of other users as well as of that user over other items. Lee et al. (2011) consider the inverse problem of measuring the expertise of individuals based on the rankings submitted by them, and the proposed algorithms assume an underlying Thurstone model. Shah et al. (2013) make a case for ordinal evaluations in certain types of MOOC homeworks/exams and study a variant of the BTL model.

1.2 Our contributions

Both the Thurstone (Case V) and BTL models involve an unknown vector $w^* \in \mathbb{R}^d$ corresponding to the underlying qualities of d items, and in a pairwise comparison between items j and k , the probability of j being ranked above k is some function F of the difference $w_j^* - w_k^*$. The Thurstone (Case V) and BTL are based on different choices of F , and both belong to the broader class of models analyzed in this paper, in which F is required only to be strongly log-concave.

With this context, the main contributions of this paper are to provide some answers to the following questions:

- How does the minimax error for estimating the weight vector w^* in various norms scale with the problem dimension (the number of items) and the number of observations?
 - We derive upper and lower bounds on the minimax estimation rates under the model described above. Our upper/lower bounds on the estimation error agree up to constant factors: to the best of our knowledge, despite the voluminous literature on these two models, this provides the first sharp characterization of the associated minimax rates. Moreover, our error guarantees provide guidance to the practitioner in assessing the number of pairwise comparisons to be made in order to guarantee a pre-specified accuracy.
- Given a budget of n comparisons, which pairs of items should be compared?
 - The bounds that we derive depend on the comparison graph induced by the subset of pairs that are compared. Our theoretical analysis reveals that the spectral gap of a certain scaled version of the graph Laplacian plays a fundamental role, and provides guidelines for the practitioner on how to choose the subset of comparisons to be made.
- When is it better to elicit pairwise comparisons versus numeric scores?
 - When eliciting data, one often has the liberty to ask for either cardinal values (Figure 1b) or for pairwise comparisons (Figure 1a) from the human subjects. One would like to adopt the approach that would lead to a better estimate. One may be tempted to think that cardinal elicitation methods are superior, since each cardinal measurement

gives a real-valued number whereas an ordinal measurement provides at most one bit of information. Our bounds show, however, that the scaling of the error in the cardinal and ordinal settings is identical up to constant pre-factors. As we demonstrate, this result allows for a comparison of cardinal and ordinal data elicitation methods in terms of the per-measurement noise alone, *independent* of the number of measurements and the number of items. A priori, there is no obvious reason for the relative performance to be independent of the number of measurements and items.

Notation: For any symmetric matrix M of size $(m \times m)$, we let $\lambda_1(M) \leq \lambda_2(M) \leq \dots \leq \lambda_m(M)$ denote its ordered eigenvalues. We use the notation $D_{\text{KL}}(\mathbb{P}_1 \parallel \mathbb{P}_2)$ to denote the Kullback-Leibler divergence between the two distributions \mathbb{P}_1 and \mathbb{P}_2 . For any integer m , we let $[m]$ denote the set $\{1, \dots, m\}$. For any pair of vectors u and v of the same length, we let $\langle u, v \rangle$ denote their inner product.

2. Problem formulation

We begin with some background followed by a precise formulation of the problem.

2.1 Generative models for ranking

Given a collection of d items to be evaluated, we suppose that each item has a certain numeric *quality score*, and a comparison of any pair of items is generated via a comparison of the two quality scores in the presence of noise. We represent the quality scores as a vector $w^* \in \mathbb{R}^d$, so item $j \in [d]$ has quality score w_j^* . Now suppose that we make n pairwise comparisons: if comparison $i \in [n]$ pertains to comparing item a_i with item b_i , then it can be described by a differencing vector $x_i \in \mathbb{R}^d$, with entry a_i equal to one, entry b_i equal to -1 , and the remaining entries set to 0.

With this notation, we study the problem of estimating the weight vector w^* based on observing a collection of n independent samples $y_i \in \{-1, 1\}$ drawn from the distribution

$$\mathbb{P}[y_i = 1 | x_i, w^*] = F\left(\frac{\langle x_i, w^* \rangle}{\sigma}\right) \quad \text{for } i \in [n], \quad (\text{ORDINAL})$$

where F is a known function taking values in $[0, 1]$. Since the probability of item a_i dominating b_i should be independent of the order of the two items being compared, we require throughout that $F(x) = 1 - F(-x)$.

In any model of the general form (ORDINAL), the parameter $\sigma > 0$, assumed to be known, plays the role of a noise parameter, with a higher value of σ leading to more uncertainty in the comparisons. Moreover, we assume that F is *strongly log-concave* in a neighborhood of the origin, meaning that there is some curvature parameter $\gamma > 0$ such that

$$\frac{d^2}{dt^2}(-\log F(t)) \geq \gamma \quad \text{for all } t \in [-2B/\sigma, 2B/\sigma]. \quad (1)$$

Here the known parameter B denotes a bound on the ℓ_∞ -norm of the weight vector, namely

$$\|w^*\|_\infty \leq B.$$

As our analysis shows, a bound of this form is fundamental: the minimax error for estimating w^* diverges to infinity if we are allowed to consider models in which B is arbitrarily large (see Proposition 17 in Appendix G). Informally, this behavior is related to the difficulty of estimating very small (or very large) probabilities that can arise in the two models for large $\|w^*\|_\infty$. Note that any model of the form (ORDINAL) is invariant to shifts in w^* , that is, it does not differentiate between the vector w^* and the shifted vector $w^* + 1$, where 1 denotes the vector of all ones. Therefore, in order to ensure identifiability of w^* , we assume throughout that $\langle 1, w^* \rangle = 0$. We use the notation \mathcal{W}_B to denote the set of permissible quality score vectors

$$\mathcal{W}_B := \{w \in \mathbb{R}^d \mid \|w\|_\infty \leq B, \text{ and } \langle 1, w \rangle = 0\}. \quad (2)$$

Both the Thurstone (Case V) model with Gaussian noise (Thurstone, 1927) and Bradley-Terry-Luce (BTL) models (Bradley and Terry, 1952; Luce, 1959) are special cases of this general set-up, as we now describe.

Thurstone (Case V): This model is a special case of the family (ORDINAL), obtained by setting

$$F(t) = \int_{-\infty}^t \frac{1}{\sqrt{2\pi}} e^{-u^2/2} du, \quad (3)$$

corresponding to the CDF of the standard normal distribution. Consequently, the Thurstone model can alternatively be written as making n i.i.d. observations of the form

$$y_i = \text{sign} \left\{ \langle x_i, w^* \rangle + \epsilon_i \right\}, \quad \text{for } i \in [n], \quad (\text{THURSTONE})$$

where $\epsilon_i \sim N(0, \sigma^2)$ is observation noise. It can be verified that the THURSTONE model is strongly log-concave over the set \mathcal{W}_B (e.g., see Tsukida and Gupta (2011)).

Bradley-Terry-Luce: The Bradley-Terry-Luce (BTL) model (Bradley and Terry, 1952; Luce, 1959) is another special case in which

$$F(t) = \frac{1}{1 + e^{-t}},$$

and hence

$$\mathbb{P}[y_i = 1 \mid x_i, w^*] = \frac{1}{1 + \exp\left(-\frac{\langle x_i, w^* \rangle}{\sigma}\right)} \quad \text{for } i \in [n]. \quad (\text{BTL})$$

It can also be verified that the BTL model is strongly log-concave over the set \mathcal{W}_B .

Cardinal observation models: While our primary focus is on the pairwise-comparison setting, for comparison purposes we also analyze analogous cardinal settings where each observation is real valued. In particular, we consider the following two cardinal analogues of the THURSTONE model. In the CARDINAL model we consider, each observation $i \in [n]$ consists of a numeric evaluation $y_i \in \mathbb{R}$ of a single item,

$$y_i = \langle u_i, w^* \rangle + \epsilon_i \quad \text{for } i \in [n], \quad (\text{CARDINAL})$$

where u_i in this case is a coordinate vector with one of its entries equal to 1 and remaining entries equal to 0, and ϵ_i is independent Gaussian noise $N(0, \sigma^2)$. One may alternatively elicit cardinal values of the differences between pairs of items

$$y_i = \langle x_i, w^* \rangle + \epsilon_i \quad \text{for } i \in [n], \quad (\text{PAIRED CARDINAL})$$

where ϵ_i are i.i.d. $N(0, \sigma^2)$. We term this model the PAIRED CARDINAL model.

2.2 Fixed design and the graph Laplacian

We analyze the estimation error when a fixed subset of pairs is chosen for comparison. Of interest to us is the *comparison graph* defined by these chosen pairs, with each pair inducing an edge in the graph. Edge weights are determined by the fraction of times a given pair is compared. The analysis in the sequel reveals the central role played by the Laplacian of this weighted graph. Note that we are operating in a fixed-design setup where the graph is constructed offline and does not depend on the observations.

In the ordinal models, the i^{th} measurement is related to the difference between the two items being compared, as defined by the measurement vector $x_i \in \mathbb{R}^d$. We let $X \in \mathbb{R}^{n \times d}$ denote the measurement matrix with the vector x_i^T as its i^{th} row. The Laplacian matrix L associated with this differencing matrix is given by

$$L := \frac{1}{n} X^T X = \frac{1}{n} \sum_{i=1}^n x_i x_i^T. \quad (4)$$

By construction, for any vector $v \in \mathbb{R}^d$, we have $v^T L v = \sum_{j \neq k} L_{jk} (v_j - v_k)^2$, where L_{jk} is the fraction of the measurement vectors $\{x_i\}_{i=1}^n$ in which items (j, k) are compared.

The Laplacian matrix is positive semidefinite, and has at least one zero-eigenvalue, corresponding to the all-ones eigenvector. The Laplacian matrix induces a graph on the vertex set $\{1, \dots, d\}$, in which a given pair (j, k) is included as an edge if and only if $L_{jk} \neq 0$, and the weight on an edge (j, k) equals L_{jk} . We emphasize that throughout our analysis, we assume that the comparison graph is *connected*, since otherwise, the quality score vector w^* is not identifiable. Note that the Laplacian matrix L induces a semi-norm¹ on \mathbb{R}^d , given by

$$\|u - v\|_L := \sqrt{(u - v)^T L (u - v)}. \quad (5)$$

We study optimal rates of estimation in this semi-norm, as well as the usual ℓ_2 -norm. As will be clearer in the sequel the L semi-norm is a natural metric in our setup, and estimation in this induced metric can be done at a topology independent rate. The estimation error in the L semi-norm is closely related to the prediction risk in generalized linear models. In particular, for an estimate \hat{w} of w^* from n_{ij} comparisons between each pair of items (i, j) , we have $\|\hat{w} - w^*\|_L^2 = \sum_{i < j} n_{ij} ((\hat{w}_i - \hat{w}_j) - (w_i^* - w_j^*))^2$. It arises naturally when one is interested in predicting the probability of a certain outcome for a new comparison.

1. A semi-norm differs from a norm in that the semi-norm of a non-zero element is allowed to be zero.

3. Bounds on the minimax risk

In this section, we state the main results of the paper, and discuss some of their consequences.

3.1 Minimax rates in the squared L semi-norm

Our first main result provides bounds on the minimax risk under the squared L semi-norm (5) in the pairwise comparison models introduced earlier. In all of the statements, we use c_1, c_2 , etc. to denote positive numerical constants, independent of the sample size n , number of items d and other problem-dependent parameters.

Apart from the curvature parameter γ defined earlier in (1), the bounds presented subsequently depend on F through a second parameter ζ , defined as

$$\zeta := \frac{\max_{x \in [0, 2B/\sigma]} F'(x)}{F(2B/\sigma)(1 - F(2B/\sigma))}. \quad (6)$$

In the BTL and the THURSTONE models, we have $\zeta := \frac{F'(0)}{F(2B/\sigma)(1 - F(2B/\sigma))}$. For instance, when $B = \sigma = 1$, then under the BTL model we have $\gamma = 0.25$ and $\zeta = 1.43$. As observed in Negahban et al. (2015), the reader may consider the parameters B, σ, γ and ζ as having $\mathcal{O}(1)$ values: a fixed value of B is the hardest regime, and furthermore, in situations of practical interest, the problem dependent parameters σ, γ and ζ , are typically independent of d and n . Consequently, in the sequel, we treat these parameters as fixed.

Theorem 1 (Bounds on minimax rates in L semi-norm)

- (a) For a sample size $n \geq \frac{c_1 \sigma^2 \text{tr}(L^\dagger)}{\zeta B^2}$, any estimator \tilde{w} based on n samples from the ORDINAL model has Laplacian squared error lower bounded as

$$\sup_{w^* \in \mathcal{W}_B} \mathbb{E} \left[\|\tilde{w} - w^*\|_L^2 \right] \geq \frac{c_{1\ell}}{\zeta} \sigma^2 \frac{d}{n}. \quad (7a)$$

- (b) For any instance of the ORDINAL model with γ -strong log-concavity and any $w^* \in \mathcal{W}_B$, the maximum likelihood estimator satisfies the bound

$$\mathbb{P} \left[\|\hat{w}_{ML} - w^*\|_L^2 > t \frac{c\zeta^2 \sigma^2 d}{\gamma^2 n} \right] \leq e^{-t} \quad \text{for all } t \geq 1,$$

and consequently

$$\sup_{w^* \in \mathcal{W}_B} \mathbb{E} \left[\|\hat{w}_{ML} - w^*\|_L^2 \right] \leq \frac{c_{1u} \zeta^2}{\gamma^2} \sigma^2 \frac{d}{n}. \quad (7b)$$

The results of Theorem 1 characterize the minimax risk in the squared L semi-norm up to constant factors. The upper bounds follow from an analysis of the maximum likelihood estimator, which turns out to be a convex optimization problem. On the other hand, the lower bounds are based on a combination of information-theoretic techniques and carefully constructed packings of the parameter set \mathcal{W}_B . The main technical difficulty is in constructing a packing in the semi-norm induced by the Laplacian L . See Appendix A for the full proof.

3.2 Minimax rates in the squared ℓ_2 -norm

Let us now turn to optimizing the minimax risk under the squared Euclidean norm. Theorem 2 below presents upper and lower bounds on this quantity.

Theorem 2 (Bounds on minimax rates in ℓ_2 -norm)

- (a) For a sample size $n \geq \frac{c_2 \sigma^2 \text{tr}(L^\dagger)}{\zeta B^2}$, any estimator \tilde{w} based on n samples from the ORDINAL model has squared Euclidean error lower bounded as

$$\sup_{w^* \in \mathcal{W}_B} \mathbb{E} \left[\|\tilde{w} - w^*\|_2^2 \right] \geq c_{2\ell} \frac{\sigma^2}{n} \max \left\{ d^2, \max_{d' \in \{2, \dots, d\}} \sum_{i=\lceil 0.99d' \rceil}^{d'} \frac{1}{\lambda_i(L)} \right\}. \quad (8a)$$

- (b) For any instance of the ORDINAL model with γ -strong log-concavity and any $w^* \in \mathcal{W}_B$, the maximum likelihood estimator satisfies the bound

$$\mathbb{P} \left[\|\hat{w}_{ML} - w^*\|_2^2 > t \frac{c\zeta^2 \sigma^2}{\gamma^2} \frac{d}{\lambda_2(L)n} \right] \leq e^{-t} \quad \text{for all } t \geq 1, \quad (8b)$$

and consequently

$$\sup_{w^* \in \mathcal{W}_B} \mathbb{E} \left[\|\hat{w}_{ML} - w^*\|_2^2 \right] \leq \frac{c_{2u} \zeta^2}{\gamma^2} \sigma^2 \frac{d}{\lambda_2(L)n}. \quad (8c)$$

See Appendix B for the proof of this theorem. As we describe in the next section, the upper and lower bounds on minimax risk from Theorem 2 to identify the comparison graph(s) that lead to the best possible minimax risk over all possible graph topologies.

Figure 2 depicts results from simulations under the THURSTONE model, depicting the squared ℓ_2 error for the maximum likelihood estimator for various values of n and d . In the simulations, the true vector w^* is generated by first drawing a d -length vector uniformly at random from $[-1, 1]^d$, followed by a scale and shift to ensure $w^* \in \mathcal{W}_B$. The n pairs are chosen uniformly (with replacement) at random from the set of $\binom{d}{2}$ possible pairs of items. The value of σ and B are both fixed to be 1. Given the n samples, inference is performed via the maximum likelihood estimator for the THURSTONE model. Each point in the plots is an average of 20 such trials.

The error in Figure 2 reduces linearly with n , exactly as predicted by our Theorem 2. For the complete graph, $\frac{1}{\lambda_2(L)} = \frac{d-1}{2}$. Theorem 2 thus predicts a quadratic increase in the error with d . As predicted, the error when normalized by $\frac{1}{d^2}$ in Figure 2 converges to the same curve for all values of d .

Detailed comparison with other work: Having stated our main theoretical results we are now in a position to revisit the results of the earlier works of Negahban et al. (2012, 2015); Hajek et al. (2014). The papers by Negahban et al. (2012, 2015) consider the BTL model under a restricted sampling setting in which every pair considered is compared the same number of times. For the problem of recovering the vector w^* (as considered in the present paper), they derive upper bounds when the pairs considered for comparisons arise

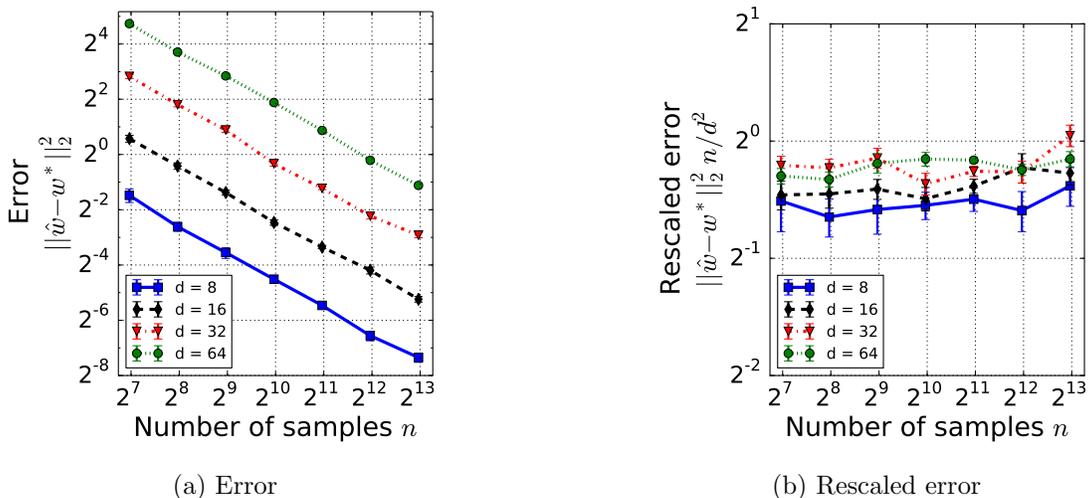


Figure 2: Simulation results under the THURSTONE model. The comparison topology chosen here is the complete graph.

from an Erdős-Rényi comparison graph; setting $t = \log d$ in equation (8b) in Theorem 2 above recovers their results. The two papers also provide results pertaining to inference of a set of parameters that are an exponentiated form of w^* . Their upper bounds for these parameters are derived in terms of the random-walk normalized Laplacian matrix of the comparison graph, while our bounds for w^* are derived in terms of the (combinatorial) Laplacian matrix. Their associated information-theoretic lower bounds also assume an Erdős-Rényi comparison graph whereas our lower bounds apply to general comparison graphs.

A concurrent paper by Hajek et al. (2014) also considers the specific BTL model, but a general comparison topology. Their high-probability upper bounds can be recovered by setting $t = \log d$ in equation (8b) of Theorem 2, while their upper bounds on the expected error are loose by a logarithmic factor. On the other hand, the lower bounds of Hajek et al. (2014), although dependent on topology, are quite loose due to their reliance on the degrees of the vertices in the comparison graph. On the other hand, our results are derived in terms of the graph Laplacian which better captures the critical aspects of the problem. For instance, considering a disconnected graph where every node has at least one neighbor as a sanity check, the degree-based lower bounds of Hajek et al. (2014) do not reflect the non-identifiability of w^* due to their reliance on only the degrees of the vertices, whereas the presence of the spectral gap in our bound (8a) indicates non-identifiability. As another example, the bounds of Hajek et al. (2014) cannot distinguish between a star graph and a path graph, whereas our results establish the star graph as an optimal comparison graph and the path graph as strictly suboptimal. In Section 4 below, we further describe deeper insights on the graph topology derived from our analytical results.

The PAIRED CARDINAL model: Before concluding this section, we also look at the PAIRED CARDINAL model (Section 2.1), the cardinal analogue of the THURSTONE model.

Theorem 3 (Bounds on minimax rates in ℓ_2 -norm) *For the PAIRED CARDINAL model, the minimax risk is sandwiched as*

$$c_{3\ell} \sigma^2 \frac{\text{tr}(L^\dagger)}{n} \leq \inf_{\hat{w}} \sup_{w^* \in \mathcal{W}_\infty} \mathbb{E} \left[\|\hat{w} - w^*\|_2^2 \right] \leq c_{3u} \sigma^2 \frac{\text{tr}(L^\dagger)}{n}. \quad (9)$$

The proof of Theorem 3 is available in Appendix C.

We conjecture that the dependence of the squared ℓ_2 minimax risk under the ORDINAL models on the problem parameters n , d and the graph topology is identical to that derived in Theorem 3 for the PAIRED CARDINAL model, i.e., is proportional to $\frac{\text{tr}(L^\dagger)}{n}$. Note that the condition $\text{tr}(L) = 2$ implies that $\frac{d^2}{9} \leq \text{tr}(L^\dagger) \leq \frac{d}{\lambda_2(L)}$.

3.3 Extension to m -ary comparisons

Suppose instead of eliciting pairwise comparisons, one can instead ask the workers to make comparisons between more than two options. In particular, we assume that each sample is a selection of the item with the largest perceived quality among some m presented items. The setting of pairwise comparisons is a special case with $m = 2$. Recall from Theorem 2 that the minimum squared ℓ_2 minimax risk in the pairwise comparison setting is of the order $\frac{d^2}{n}$. Our goal in this section is to bring the concept of multiple-item comparison under the same framework as the pairwise case, and via a generalization of our earlier theoretical analysis, understand how the error exponent depends on m .

Consider d items, where every item $j \in [d]$ has a certain underlying quality score $w_j^* \in [-B, B]$. Suppose you have access to n samples, with each sample being a selection of the item with the largest perceived value among some m presented items.

Consider $(d \times m)$ matrices E_1, \dots, E_n such that for each $i \in [n]$, the m columns of E_i are distinct unit vectors. The positions of the non-zero elements in the m columns of E_i represent the identities of the m items compared in the i^{th} sample. One can visualize the choices of the items compared as a hyper-graph, with d vertices representing the d items and hyper-edge $i \in [n]$ containing the m items compared in observation i .

Let R_1, \dots, R_m be $(m \times m)$ permutation matrices representing m cyclic shifts in an arbitrary (but fixed) direction. Consider the observation model

$$\mathbb{P}(y_i = j | w^*, E_i) = F((w^*)^T E_i R_j)$$

for all $j \in [m]$, where $F : [-B, B]^m \rightarrow [0, 1]$ represents the probability of choosing the *first* among the m items presented. We also assume that F does not depend on the *order* of the last $(m - 1)$ coordinates in its input, meaning that the likelihood of choosing an item is independent of the ordering of the m items in the argument to F . For every $x \in [-B, B]^m$, $F(x)$ is assumed to satisfy:

- Shift-invariance: the probabilities depend only on the *differences* in the weights of the items presented, i.e, $F(x)$ depends only on $\{x_i - x_j\}_{i,j \in [m]}$.
- Strong log-concavity: $\nabla^2(-\log F(x)) \succeq H$ for some $(m \times m)$ symmetric matrix H with $\lambda_2(H) > 0$.

Note that the shift-invariance assumption implies $1 \in \text{nullspace}(\nabla^2(-\log F(x)))$, thereby necessitating $\text{nullspace}(H) = \text{span}(1)$ and $\lambda_1(H) = 0$. One can also verify that the model proposed here reduces to the ORDINAL model of Section 2.1 when $m = 2$.

For any hope of inferring the true weights w^* , we must ensure that the comparison hyper-graph is “connected”, i.e., for every pair of items $i, j \in [d]$, there must exist a path connecting item i and item j in the comparison hyper-graph. We assume this condition is satisfied. We also continue to assume that $w^* \in \mathcal{W}_B := \{w \in \mathbb{R}^d \mid \|w\|_\infty \leq B, \langle w, 1 \rangle = 0\}$.

The popular Plackett-Luce model falls in this class, as illustrated below.

Example 1 (Plackett-Luce model (Plackett, 1975; Luce, 1959)) *The Plackett-Luce model concerns the process of choosing an item from a given set. Specifically, given m items with quality scores w_1^*, \dots, w_m^* respectively, the likelihood of choosing item $\ell \in [m]$ under this model is given by*

$$\frac{e^{w_\ell^*}}{\sum_{j=1}^m e^{w_j^*}} =: F([w_\ell^*, w_1^*, \dots, w_{\ell-1}^*, w_{\ell+1}^*, \dots, w_m^*]).$$

Every choice is made independent of all other choices.

It is easy to verify that the Plackett-Luce model satisfies shift invariance. Furthermore, the function F does not depend on the ordering of the last $(m - 1)$ coordinates in its argument. We now show that it also satisfies strong log-concavity. A little algebra gives

$$\nabla^2(-\log F(x)) = \frac{\langle e^x, 1 \rangle \text{diag}(e^x) - e^x (e^x)^T}{(\langle e^x, 1 \rangle)^2},$$

where $e^x := [e^{x_1} \dots e^{x_m}]^T$. We now derive a lower bound for the expression above. An application of the Cauchy-Schwarz inequality yields that for any vector $v \in \mathbb{R}^m$,

$$v^T (e^x (e^x)^T) v \leq v^T \text{diag}(e^x) \langle e^x, 1 \rangle v,$$

with equality if and only if $v \in \text{span}(1)$. It follows that $\lambda_2(\nabla^2(-\log F(x))) > 0$ for all $x \in [-B, B]^m$. Defining the scalar $\beta := \min_{x \in [-B, B]^m} \lambda_2(\frac{\langle e^x, 1 \rangle \text{diag}(e^x) - e^x (e^x)^T}{(\langle e^x, 1 \rangle)^2})$, one can see that setting $H = \beta(I - 11^T)$ satisfies the strong log-concavity conditions.

Our goal is to capture the scaling of the minimax error with respect to the number of observations n , the dimension d of the problem, and the choice of the subsets compared $\{E_i\}_{i \in [n]}$. It is well understood (Miller, 1956; Kiger, 1984; Shiffrin and Nosofsky, 1994; Saaty and Ozdemir, 2003) that humans have a limited information storage and processing capacity, which makes it difficult to compare more than a small number of items. For instance, Saaty and Ozdemir (2003) recommend eliciting preferences over *no more than seven options*. Thus in this work we restrict our attention to $m = \mathcal{O}(1)$. Moreover, the amount of noise in the selection process also depends on the number of items m presented at a time: the higher the number, the greater the noise. We thus do not use a “noise parameter σ ” in this setting, and assume the noise to be incorporated in the function F , which itself is a function of m .

Our results involve the Laplacian of the comparison graph, defined for the m -wise comparison setting as follows. Let L be an $(d \times d)$ matrix that depends on the choice of the comparison topology as

$$L := \frac{1}{n} \sum_{i=1}^n E_i (mI - 11^T) E_i^T. \quad (10)$$

We call L the Laplacian of the comparison hyper-graph. One can verify that when applied to the special case of $m = 2$, the matrix L defined in (10) reduces to the Laplacian of the pairwise-comparison graph defined earlier in (4).

The following theorem presents our main results for the m -wise comparison setting.

Theorem 4 *For the m -WISE model, the minimax risk is sandwiched as*

$$c_{3\ell} \frac{\inf_z F(z)}{m^2 \lambda_m(H) \sup_z \|\nabla F(z)\|_{H^\dagger}^2} \frac{d}{n} \leq \inf_{\hat{w}} \sup_{w^* \in \mathcal{W}_B} \mathbb{E} \left[\|\hat{w} - w^*\|_L^2 \right] \leq c_{3u} \frac{m^2 \sup_z \|\nabla \log F(z)\|_2^2}{\lambda_2(H)^2} \frac{d}{n},$$

in the squared L semi-norm and as

$$c_{4\ell} \frac{\inf_z F(z)}{m^2 \lambda_m(H) \sup_z \|\nabla F(z)\|_{H^\dagger}^2} \frac{d^2}{n} \leq \inf_{\hat{w}} \sup_{w^* \in \mathcal{W}_B} \mathbb{E} \left[\|\hat{w} - w^*\|_2^2 \right] \leq c_{4u} \frac{m^2 \sup_z \|\nabla \log F(z)\|_2^2}{\lambda_2(H)^2} \frac{d}{\lambda_2(L)n},$$

in the squared ℓ_2 norm. Here we assume $n \geq c_5 \frac{\text{tr}(L^\dagger) \inf_z F(z)}{B^2 \lambda_m(H) \sup_z \|\nabla F(z)\|_{H^\dagger}^2}$ for both the lower bounds, and where the suprema and infima with respect to the parameter z are taken over the set $[-B, B]^m$.

The proof of Theorem 4 is provided in Appendix D. Our results establish that the dependence of the squared L semi-norm and squared Euclidean minimax error on m occurs only as multiplicative pre-factors, and the error exponent is independent of m . Thus, if one follows the standard recommendation in the psychology literature Miller (1956); Kiger (1984); Shiffrin and Nosofsky (1994); Saaty and Ozdemir (2003)—namely to choose $m = \mathcal{O}(1)$ —then the best possible scaling of the squared L semi-norm minimax risk with respect to d and n is always $\frac{d}{n}$, that of the squared Euclidean minimax risk is always $\frac{d^2}{n}$, and evenly spreading the samples across all possible choices of m items is optimal. Nevertheless, a more refined modeling and analysis is required to understand the precise tradeoffs governing the choice of the number m of items presented to the user.

Finally, when specialized to the case of $m = 2$, the upper bounds of Theorem 4 are identical (up to constant factors) to those of Theorem 1 and Theorem 2. The lower bound for the L semi-norm is identical to that of Theorem 1. The additional generality results in a lower bound for the ℓ_2 norm that is weaker in general as compared to Theorem 2, but is tight when the underlying hypergraph forms a complete graph.

4. Role of graph topology

We now return to the setting of pairwise comparisons. In certain applications, one may have the liberty to decide which pairs are compared. The results of the previous section

demonstrated the role played by the Laplacian of the comparison graph in the estimation error. We now employ these results to derive guidelines towards designing the comparison graph. Let us focus on the estimation error in the squared ℓ_2 norm in the ordinal setting. As discussed earlier, we assume that the graph induced by the comparisons is connected. An application of Theorem 2 lets us identify good topologies for pairwise comparisons in the fixed-design setup.

A popular class of comparison topologies is that of evenly distributed samples on an unweighted graph (e.g., Negahban et al. (2012)). Consider any fixed, unweighted graph $G = (V, E)$. We assume that the samples are distributed evenly along the edges E of G , and that the sample size n is sufficiently large. Using standard matrix concentration inequalities, it is straightforward to extend our analysis to the setting of random chosen comparisons from a fixed graph (see, for instance, Oliveira (2009)). Let L' denote the Laplacian of G . We define the *scaled Laplacian* of G as

$$L := \frac{1}{|E|} L'.$$

One can verify that the matrix L defined here is identical to what was defined in (4) in a more general context. In order to differentiate from L , we refer to L' as the *regular Laplacian* of the graph G . For a given budget n on the number of samples, we say that a comparison graph is *optimal* if the error under this graph is the smallest (up to constants) among all graphs.

4.1 Analytical results

Consider the ORDINAL model and the squared ℓ_2 -norm as the metric of interest. We claim that in order to determine whether a given comparison graph achieves minimax risk (up to a constant pre-factor), it suffices to examine the eigen-spectrum of the scaled Laplacian matrix. In particular, we claim that:

- If the scaled Laplacian has a second smallest eigenvalue that scales as $\frac{1}{\lambda_2(L)} = \Theta(d)$, then the comparison graph is optimal, and leads to the smallest possible minimax risk, in particular one that scales as $\frac{d^2}{n}$.
- Conversely, if the scaled Laplacian matrix has an eigen-spectrum satisfying

$$d^2 = o\left(\max_{d' \in \{2, \dots, d\}} \sum_{i=\lfloor 0.99d' \rfloor}^{d'} \frac{1}{\lambda_i(L)}\right), \quad (11)$$

then the associated estimation error is strictly larger than the minimax risk. In particular, this sub-optimality holds whenever $d^2 = o(\frac{1}{\lambda_2(L)})$.

In order to verify these claims, we note that by definition (4) of the Laplacian matrix, we have

$$\text{tr}(L) = \frac{1}{n} \sum_{i=1}^n \text{tr}(x_i x_i^T) = 2.$$

It follows that $\lambda_2(L) \leq \frac{2}{d-1}$, i.e., that $\frac{1}{\lambda_2(L)} = \Omega(d)$. As we will see shortly, several classes of graphs satisfy $\frac{1}{\lambda_2(L)} = \Theta(d)$. Comparing the lower bound of $\Omega(\frac{d^2}{n})$ on the minimax risk (8a) with the upper bound (8c) gives the sufficient condition of $\frac{1}{\lambda_2(L)} = \Theta(d)$ for optimality, and the smallest minimax risk as $\Theta(\frac{d^2}{n})$. The lower bound (8a) now also gives the claimed condition for strict sub-optimality.

In order to illustrate these claims, let us consider a few canonical classes of graphs, and study how the estimation error under the squared Euclidean norm scales in the ORDINAL model. The spectra of the regular Laplacian matrices of these graphs can be found in various standard texts on spectral graph theory (e.g., Brouwer and Haemers (2011)).

- **Complete graph.** A complete graph has one edge between every pair of nodes. The spectrum of the regular Laplacian of the complete graph is $0, d, \dots, d$, and hence the spectrum of the scaled Laplacian L is $0, \frac{2}{d-1}, \dots, \frac{2}{d-1}$. Substituting $\lambda_2(L) = \frac{2}{d-1}$ in Theorem 2b gives an upper bound of $\Theta(\frac{d^2}{n})$ on the minimax risk, and Theorem 2 gives a matching lower bound. The sufficiency condition discussed above proves optimality.
- **Constant-degree expander.** The spectrum of the regular Laplacian of a constant-degree expander graph is $0, \Theta(d), \Omega(d), \dots, \Omega(d)$. Since the number of edges is $\Theta(d)$, the spectrum of the scaled Laplacian equals $0, \Theta(\frac{1}{d}), \Omega(\frac{1}{d}), \dots, \Omega(\frac{1}{d})$. The evaluation of this class of graphs with respect to the minimax risk is identical to that of complete graphs, giving a lower and upper bound of $\Theta(\frac{d^2}{n})$ on the minimax risk, and guaranteeing optimality.
- **Complete bipartite.** The d nodes are partitioned into two sets comprising, say, m_1 and m_2 nodes. There is an edge between every pair of nodes in different sets, and there are no edges between any two nodes in the same set. The eigenvalues of the regular Laplacian of this graph are $0, \underbrace{m_2, \dots, m_2}_{m_1-1}, \underbrace{m_1, \dots, m_1}_{m_2-1}, m_1 + m_2$. Since the total number of edges is $m_1 m_2$, the scaled Laplacian L has a spectrum $0, \underbrace{\frac{1}{m_1}, \dots, \frac{1}{m_1}}_{m_1-1}, \underbrace{\frac{1}{m_2}, \dots, \frac{1}{m_2}}_{m_2-1}, \frac{1}{m_1} + \frac{1}{m_2}$. Suppose without loss of generality that $m_1 \geq m_2$. Also suppose that $m_2 > 1$ (the case of $m_2 = 1$ is the star graph discussed below). Then we have $\frac{1}{m_1} \leq \frac{1}{m_2} \leq \frac{1}{m_1} + \frac{1}{m_2}$ and that $d > m_1 \geq \frac{d}{2}$. Furthermore since $m_2 > 1$, the multiplicity of $\frac{1}{m_1}$ in the spectrum of the scaled Laplacian is at least 1. Thus we have $\lambda_2(L) = \Theta(\frac{1}{d})$. Theorem 2 then gives lower and upper bounds on the minimax risk as $\Theta(\frac{d^2}{n})$ and the sufficiency condition discussed above guarantees its optimality.
- **Star.** A star graph has one central node with edges to every other node. It is a special case of the complete bipartite graph with $m_1 = d - 1$ and $m_2 = 1$. The spectrum of the regular Laplacian is $0, 1, \dots, 1, d$. Since there are $(d - 1)$ edges, the spectrum of the scaled Laplacian is $0, \frac{1}{d-1}, \dots, \frac{1}{d-1}, \frac{d}{d-1}$. Theorem 2 and the sufficiency condition discussed above imply that this class of graphs is optimal and is associated to a minimax risk of $\Theta(\frac{d^2}{n})$.
- **Path.** A path graph is associated to an arbitrary ordering of the d nodes with edges between pairs j and $(j + 1)$ for every $j \in \{1, \dots, d - 1\}$. The spectrum of the regular

Laplacian is given by $2(1 - \cos(\frac{\pi i}{d}))$, $i \in \{0, \dots, d-1\}$, and that of the scaled Laplacian is thus $\frac{2}{d-1}(1 - \cos(\frac{\pi i}{d}))$, $i \in \{0, \dots, d-1\}$. The relation $(1 - \cos x) = \sin^2 \frac{x}{2}$ and the approximation $\sin x \approx x$ for values of x close to zero gives $\lambda_2(L) = \Theta(\frac{1}{d^3})$. The minimax risk is thus upper bounded as $\mathcal{O}(\frac{d^4}{n})$ and lower bounded as $\Omega(\frac{d^3}{n})$. This class of graphs is thus strictly suboptimal.

- **Cycle.** A cycle is identical to a path except for an additional edge between node d and node 1. The spectrum of the regular Laplacian is given by $2(1 - \cos(\frac{2\pi i}{d}))$, $i \in \{0, \dots, d-1\}$, and that of the scaled Laplacian is thus $\frac{2}{d}(1 - \cos(\frac{2\pi i}{d}))$, $i \in \{0, \dots, d-1\}$. The relation $(1 - \cos x) = \sin^2 \frac{x}{2}$ and the approximation $\sin x \approx x$ for values of x close to zero gives $\lambda_2(L) = \Theta(\frac{1}{d^3})$. The minimax risk is thus upper bounded as $\mathcal{O}(\frac{d^4}{n})$ and lower bounded as $\Omega(\frac{d^3}{n})$. This class of graphs is thus strictly suboptimal.
- **Barbell.** The nodes are partitioned into two sets of $\frac{d}{2}$ nodes each, and there is an edge between every pair of nodes within each set. In addition, there is exactly one edge across the sets. The spectrum of the regular Laplacian can be computed as $0, \Theta(\frac{1}{d}), \Theta(d), \dots, \Theta(d)$. Since there are $\Theta(d^2)$ edges, the spectrum of the scaled Laplacian turns out to become $0, \Theta(\frac{1}{d^3}), \Theta(\frac{1}{d}), \dots, \Theta(\frac{1}{d}), \Omega(\frac{1}{d})$. Applying the results derived earlier in the paper, we get that a lower bound of $\Omega(\frac{d^3}{n})$ and an upper bound of $\mathcal{O}(\frac{d^4}{n})$ on the minimax risk, thereby also establishing the sub-optimality of this class of graphs.
- **2D Lattice.** An $(m_1 \times m_2)$ lattice has $d = m_1 m_2$ vertices arranged as a $(m_1 \times m_2)$ grid. Assume $m_1 = \Theta(d)$ and $m_2 = \Theta(d)$. This class of graphs can be written as a Cartesian product of a path graph of length m_1 and a second path graph of length m_2 . As a result, the spectrum of the scaled Laplacian is $\frac{2}{d}(2 - \cos(\frac{\pi i}{m_1}) - \cos(\frac{\pi j}{m_2}))$, $i \in \{0, \dots, m_1-1\}, j \in \{0, \dots, m_2-1\}$. Again, using the small angle approximation of the sinusoid, one can compute an upper bound on the minimax risk as $\mathcal{O}(\frac{d^3}{n})$ and a lower bound of $\Omega(\frac{d^2}{n})$. We do not know at this point whether the 2D lattice minimizes the minimax risk.
- **Hypercube.** Assume $d = 2^m$ for some integer m . Representing each node as a distinct m -length binary vector, an edge exists between the nodes corresponding to any pair of vectors within a Hamming distance of one. The hypercube is an m -fold Cartesian product of a path with two nodes, and hence the regular Laplacian has an eigenvalue of $2i$ with multiplicity $\binom{m}{i}$, for $i \in \{0, \dots, m\}$. The scaled Laplacian has an eigenvalue of $\frac{2i}{d \log d}$ with multiplicity $\binom{m}{i}$, for $i \in \{0, \dots, m\}$. A lower bound on the minimax risk is $\Omega(\frac{d^2}{n})$ and an upper bound is $\mathcal{O}(\frac{d^2 \log d}{n})$. We do not know if the hypercube is optimal, our bounds do tell us that any sub-optimality is bounded by at most a logarithmic factor.

Observe that the degree- k expander requires $n \geq kd$ samples while the complete graph requires $n \geq \binom{d}{2}$ samples, so in practical applications at least for small sample sizes we should prefer a low-degree expander.

Finally, if the conjecture discussed at the end of Section 3.2 were true, namely that the ℓ_2 minimax risk scales as $\sigma^2 \text{tr}(L^\dagger)/n$, then the condition $\text{tr}(L^\dagger) = \Theta(d^2)$ would be necessary and sufficient for optimality of a comparison graph with the scaled Laplacian L . Observe that the graphs designated as ‘optimal’ in the discussion above indeed satisfy this condition. On the other hand, the graphs established as strictly suboptimal have $\text{tr}(L^\dagger) = \Omega(d^3)$.

4.2 Experiments and simulations

This section evaluates the dependence of the squared ℓ_2 -error on the topology of the comparison graph. We consider the following five topologies: path, barbell, complete, expander and 2D-lattice. In order to form an expander graph, we used the construction due to Gabber and Galil (1981). For any chosen graph topology, the n difference vectors are selected as one edge each chosen uniformly at random (with replacement) from the comparison graph. Recall that our theory predicts that the complete and expander graphs yield the best performance, and that the line and dumbbell graphs fare the worst. Also recall that our theory predicts the error scales as $\|w^* - \hat{w}\|_2^2$ scales with n as $1/n$ in the complete and expander topologies.

4.2.1 EXPERIMENTS ON SYNTHETIC DATA

This section describes simulations using data generated synthetically from the THURSTONE model. In the simulations, we first generate a quality score vector $w^* \in \mathcal{W}_B$ using one of the procedures described below. Once w^* is chosen, the n pairwise comparisons for any given topology are generated as follows. An edge is selected uniformly (with replacement) at random from the underlying graph, and the chosen edge determines the pair of items compared. The outcome of the comparison is generated as per the THURSTONE model with the chosen w^* as the underlying quality score. Finally, the maximum likelihood estimator for the THURSTONE model is employed to estimate w^* . Every point in the plots is an average across 40 trials.

The following six procedures are employed to generate the true quality score vector w^* in the six respective subfigures of Figure 3.

- (a) Gaussian: w^* is drawn from the standard normal distribution $\mathcal{N}(0, I)$.
- (b) Uniform: w^* is drawn uniformly at random from the set $[-1, 1]^d$.
- (c) Packing set for the path graph: We first choose a vector z as by setting a value of 0 in the first coordinate, a value -1 in $\frac{d}{2}$ of the other coordinates chosen uniformly at random, and a value 1 in the remaining coordinates. Letting $L = U^T \Lambda U$ denote the eigen-decomposition of the Laplacian matrix of the path graph, w^* is set as $U^T \Lambda^\dagger z$, where Λ^\dagger is the Moore-Penrose pseudoinverse of Λ . This generation process mimics a construction used to prove the lower bound in Theorem 2, and tailors the construction for the path graph.
- (d) Packing set for the barbell graph: The procedure is identical to that in (c), except that the Laplacian matrix used is that of the barbell graph.
- (e) Packing set for the complete graph: The procedure is identical to that in (c), except that the Laplacian matrix used is that of the complete graph.
- (f) Packing set for the star graph: The procedure is identical to that in (c), except that the Laplacian matrix used is that of the star graph.

The vector w^* generated in this procedure is then scaled and shifted to ensure $w^* \in \mathcal{W}_B$. The values of B and σ are set as 1.

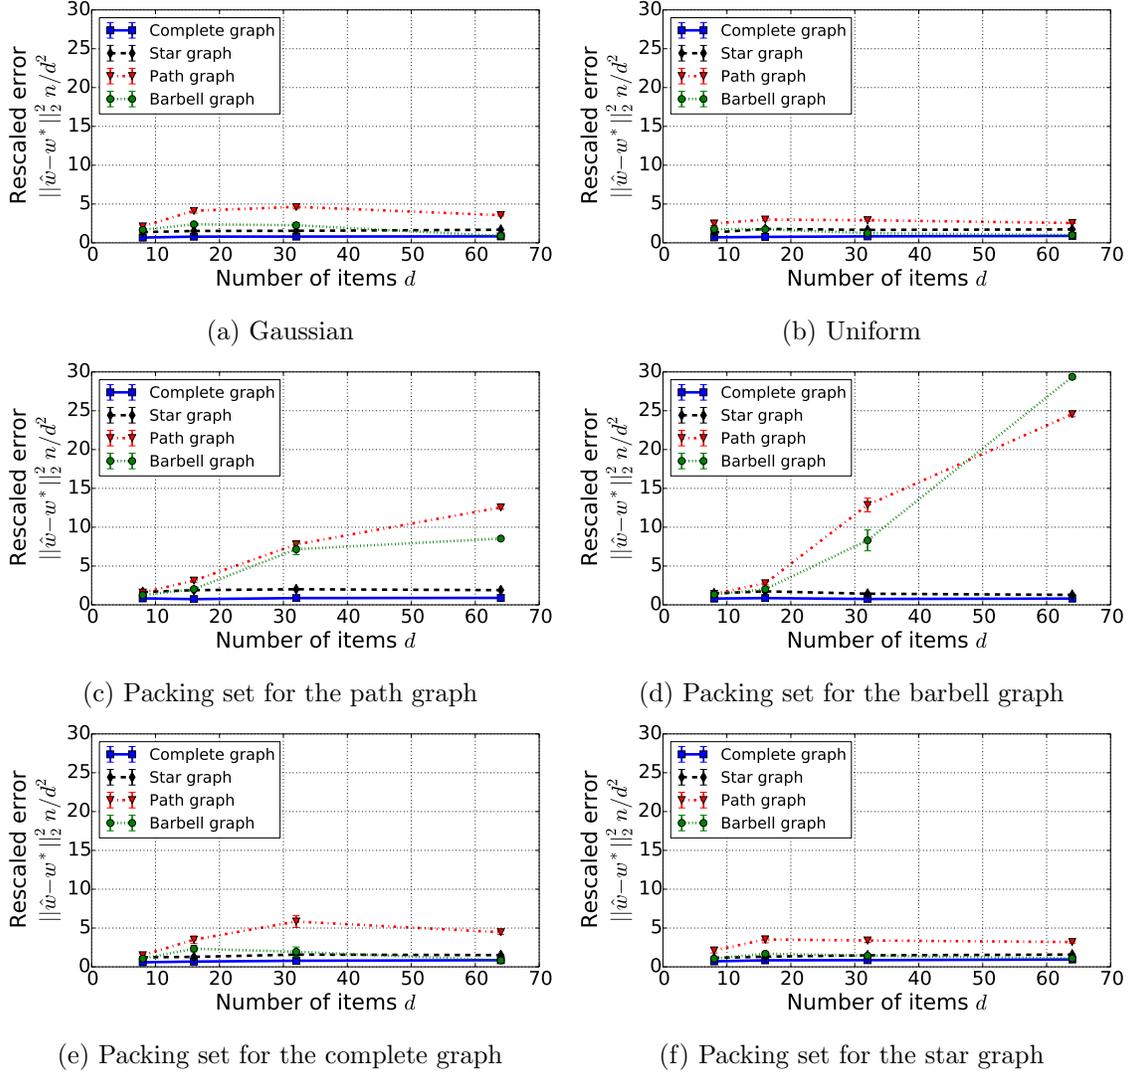


Figure 3: Estimation error under different topologies for different generative processes in the synthetic simulations.

Figure 3 plots the estimation error under various topologies of the comparison graph. Observe in the figure that the error is the lowest under the complete and the star graphs, and the highest under the barbell and the path graphs. In particular, the error consistently varies as $\Theta(d^2/n)$ for the complete and star graphs – this phenomenon holds even in plots (e) and (f) where the procedure to choose w^* forms the worst case for the complete and star graphs respectively according to the proof of Theorem 2. On the other hand, the minimax error varies as $\Omega(d^3/n)$ in the worst case for the path and the barbell graphs. Finally, observe that in the simulations, the (constant) multiplicative factors to the term $\frac{d^2}{n}$ in the error turn out to be rather small, in the range of 0 to 9.

Although not the primary focus of this paper, we note that our implementation for computing the maximum likelihood estimate under the Thurstone model requires several tens of minutes for modest problem dimensions. Computing the MLE is a convex optimization problem. Our implementation, which is not optimized for speed, employs the sequential least squares programming subroutine of the Scipy package in the Python programming language.

4.2.2 EXPERIMENTS ON MTURK

In this section, we describe the results of experiments conducted on the popular Amazon Mechanical Turk (<https://www.mturk.com/>; henceforth referred to as “MTurk”) commercial crowdsourcing platform, evaluating the effects of the choice of the topology. MTurk is an online platform where individuals or businesses can put up a task, and any individual can log in and complete the tasks in exchange for a payment that is specified along with the task. In our experiments, each worker was offered 20 cents per completed task. A worker was allowed to do no more than one task in an experiment. Workers were required to answer all the questions in a task. Only those workers who had 100 or more prior approved works and an approval rate of 95% or higher were allowed. Workers from any country were allowed to participate, except for the task of estimating distances between cities (for which only USA-based workers were permitted since all questions involved American cities).

We conducted three experiments that required the workers to make ordinal choices.

- (a) *Estimating areas of circles*: In each question, the worker was shown a circle in a bounding box (Figure 5a), and the worker was required to identify the fraction of the box’s area that the circle occupied.
- (b) *Estimating age of people from photographs*: The worker was shown photographs of people (Figure 5b) and was asked to estimate their ages.
- (c) *Estimating distances between pairs of cities*: Pairs of cities were listed (Figure 5c) and for each pair, the worker had to estimate the distance between them.

For each experiment, we recruited 140 workers on MTurk, and assigned them to one of the five topologies uniformly at random. In this experiment and others involving aggregation of ordinal data from MTurk, the aggregation procedure follows maximum likelihood estimation under the THURSTONE model, and the estimator is supplied the best-fitting value of σ obtained via 3-fold cross-validation. Each run of the estimation procedure employs the data provided by five randomly chosen workers from the pool of workers who performed

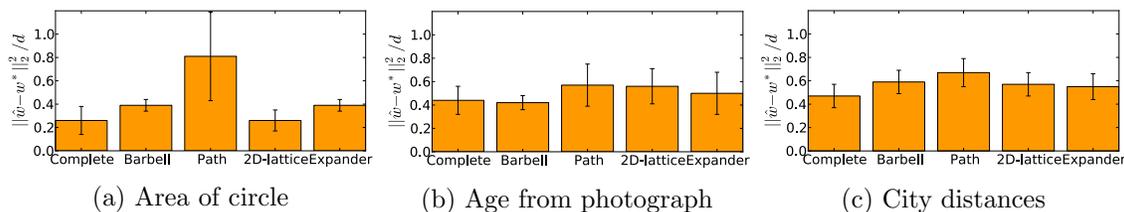


Figure 4: Estimation error under different topologies in the experiments conducted on MTurk. Also shown is the standard deviation across the estimates (The error bars are smaller by a factor of $\sqrt{\text{number of samples}}$).

that task. (The number five is inspired by practical systems as in Wang et al. (2011); Piech et al. (2013).) The complete dataset pertaining to these experiments is available on the first author’s website.

Figure 4 plots the squared ℓ_2 estimation error for the three experiments under the five topologies considered, and the associated standard deviation. We see that the relative errors are generally consistent with our theory, with the complete graph exhibiting the best performance and the path graph faring the worst. On real datasets, model misspecification can in some cases cause the outcomes to differ from our theoretical predictions. Understanding the effect of model misspecification, especially on topology considerations, is an important question we hope to address in future work.

5. Cardinal versus ordinal measurements

In this section, we compare two approaches towards eliciting data: a score-based “cardinal” approach and a comparison-based “ordinal” approach. In a cardinal approach, evaluators directly enter numeric scores as their answers (Figure 1b), while an ordinal approach involves comparing (pairs of) items (Figure 1a).

There are obvious advantages and disadvantages associated with either approach. On one hand, the cardinal approach allows for very fine measurements. For instance, the cardinal measurements in Figure 1 can take any value between 0 and 100, whereas an ordinal measurement is binary. One might be tempted to go even further and argue that ordinal measurements necessarily give less information, for one can always convert a set of cardinal measurements into ordinal, simply by ordering the measurements by value. If this conversion were valid, the data processing inequality (Cover and Thomas, 2012), would then guarantee that estimators based on ordinal data can never outperform estimators based on cardinal data. However, this conversion assumes that cardinal and ordinal measurements suffer from the same type of statistical fluctuation. The following set of experiments show this assumption is false.

5.1 Raw data from MTurk

We conducted seven different experiments on MTurk to investigate the possibility of a “data-processing inequality” between the elicited cardinal and ordinal responses: Are responses elicited in ordinal form equivalent to data obtained by first eliciting cardinal responses and then subtracting pairs of items? Our experiments lead us to conclude that this is generally

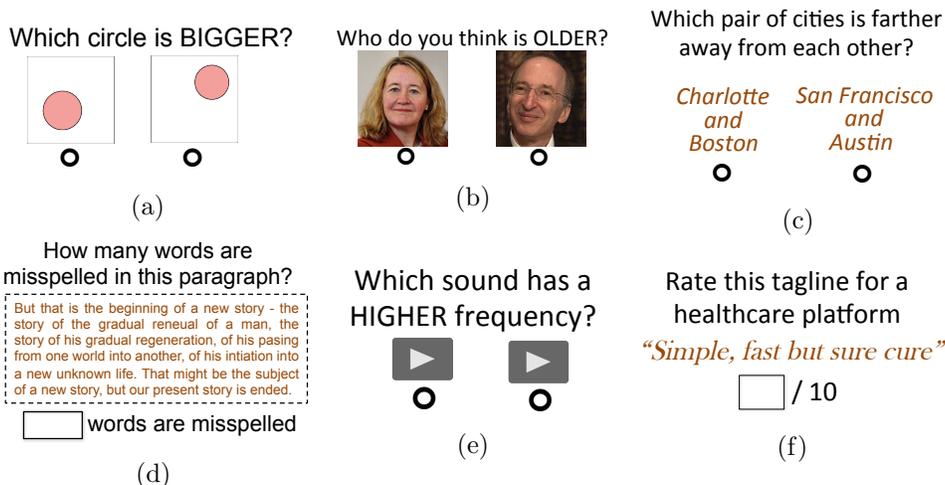


Figure 5: Screenshots of the tasks presented to the subjects. For each task, only one version (cardinal or ordinal) is shown here.

not the case: converting cardinally collected data into ordinal (by subtracting pairs of responses) often leads to a higher amount of noise as compared to that in data that is elicited directly in ordinal form.

The tasks were selected to have a broad coverage of several important subjective judgment paradigms such as preference elicitation, knowledge elicitation, audio and visual perception and skill utilization.

In addition to the three experiments described in Section 4.2.2, we conducted the following four experiments.

- (d) *Finding spelling mistakes in text:* The worker had to identify the number of words that were misspelled in each paragraph shown (Figure 5d).
- (e) *Identifying sounds:* The worker was presented with audio clips, each of which was the sound of a single key on a piano (which corresponds to a single frequency). The worker had to estimate the frequency of the sound in each audio clip (Figure 5e).
- (f) *Rating tag-lines for a product:* A product was described and tag-lines for this product were shown (Figure 5f). The worker had to rate each of these tag-lines in terms of its originality, clarity and relevance to this product.
- (g) *Rating relevance of the results of a search query:* Results for the query ‘Internet’ for an image search were shown (Figure 1) and the worker had to rate the relevance of these results with respect to the given query.

Note that the data collected for experiments (a)–(c) here was different and independent of the data collected for these tasks in Section 4.2.2.

The number of items d in the experiments ranged from 10 to 25. For each of the seven experiments, we recruited 100 workers, and assigned each worker to either the ordinal or the cardinal version of the task at random. Upon obtaining the data, we first reduced the

Task	Circle	Age	Distance	Spelling	Audio	Tagline	Relevance
Error in Ordinal	6%	13%	17%	40%	20%	44%	31%
Std. dev.	.23	.33	.38	.49	.40	.47	.44
Error in Cardinal	17%	17%	20%	42%	29%	42%	35%
Std. dev.	.31	.38	.38	.46	.43	.46	.44
Time in Ordinal	98s	31s	84s	316s	66s	251s	105s
Std. dev.	21.1	14.3	62.1	33.2	11.1	28.1	13.1
Time in Cardinal	181s	70s	144s	525s	134s	342s	185s
Std. dev.	39.9	33.1	56.2	46.0	12.4	44.6	28.2

Table 1: Comparison of the average amount of error when ordinal data is collected directly versus when cardinal data is collected and converted to ordinal. Also tabulated is the median time (in seconds) taken to complete a task by a subject in either type of task.

cardinal data obtained from the experiments into ordinal form by comparing answers given by the subjects to consecutive questions. For five of the experiments ((a) through (e)), we had access to the “ground truth” solutions, using which we computed the fraction of answers that were incorrect in the ordinal and the cardinal-converted-to-ordinal data (any tie in the latter case was counted as half an error). For the two remaining experiments ((f) and (g)) for which there is no ground truth, we computed the ‘error’ as the fraction of (ordinal or cardinal-converted-to-ordinal) answers provided by the subjects that disagreed with each other. It is important to note that in the experiments in this section, we did *not* run any estimation procedure on the data: we only measured the noise in the *raw responses*. The entire data pertaining to these experiments, including the interface seen by the workers and the data obtained from their work, is available on the first author’s website.

The results are summarized in Table 1. If the cardinal measurements could always be converted to ordinal ones with the same noise level as directly eliciting ordinal responses, then it would be unlikely for the amount of error in the ordinal setting to be smaller than that in the cardinal setting. Table 1 shows that converting cardinal data to an ordinal form very often results in a higher (and sometimes significantly higher) per-sample error in the (raw) responses than direct elicitation of ordinal evaluations. Such an outcome may be explained by the argument that the inherent evaluation process in humans is not the same in the cardinal and ordinal cases: humans do *not* perform an ordinal evaluation by first performing cardinal evaluations and then comparing them (Barnett, 2003; Stewart et al., 2005). One can also see from Table 1 that the amount of time required for cardinal evaluations was typically (much) higher than for ordinal evaluations. One can thus assume that we typically have the per-observation error in the ordinal case lower than that in the cardinal case. In particular, considering the THURSTONE and the CARDINAL models (introduced in Section 2.1) with σ and σ_c denoting the standard deviations of the noise in the THURSTONE and the CARDINAL models respectively, the above empirical results imply that $\sigma < \sqrt{2}\sigma_c$.

5.2 Analytical comparison of Cardinal versus Ordinal

As discussed earlier, while cardinal measurements allow more flexibility in the range of responses, ordinal measurements contain a lower per-sample error. Ordinal measurements have additional benefits in that they avoid calibration issues that are frequently encountered in cardinal measurements (Tsukida and Gupta, 2011), such as the evaluators' inherent (and possibly time-varying) biases, or tendencies to give inflated or conservative evaluations. Ordinal measurements are also recognized to be easier or faster for humans to make (Barnett, 2003; Stewart et al., 2005), allowing for more evaluations with the same amount of time, effort and cost.

The lack of clarity regarding when to use a cardinal versus an ordinal approach forms the motivation of this section. Can we make as reliable estimates from paired comparisons as from numeric scores? How much lower does the noise have to be for comparative measurements to be preferred over cardinal measurements? The answers to these questions help to determine how responses should be elicited.

In order to compare the cardinal and ordinal methods of data elicitation, we focus on a setting with evenly budgeted measurements. In accordance with the fixed-design setup assumed throughout the paper, we choose the vectors x_i a priori. Suppose that n is large enough, and that in the ordinal case we compare each pair $n/\binom{d}{2}$ times. In the cardinal case suppose that we evaluate the quality of each item n/d times. We consider the Gaussian-noise models THURSTONE and CARDINAL introduced earlier in Section 2.1. In order to capture the fact that the amount of noise is different in the cardinal and ordinal settings, we denote the standard deviation of the noise in the cardinal setting as σ_c , and retain our notation of σ for the noise in the ordinal setting. In order to bring the two models on the same footing, we measure the error in terms of the squared ℓ_2 -norm.

Let γ_G and ζ_G denote the parameters γ and ζ (defined in (1) and (6) respectively) specialized to the Gaussian distribution. Define $b_\ell(\sigma, B) := \frac{c_2\ell}{\zeta_G(B,\sigma)}$, $b_u(\sigma, B) := \frac{c_2u\zeta_G(B,\sigma)}{\gamma_G(B,\sigma)}$ and $b(\sigma, B) := \left\lceil \frac{c_2\sigma^2}{\zeta_G B^2} \right\rceil$. Observe that b_ℓ , b_u and b are independent of the parameters n and d .

With these preliminaries in place, we now compare the minimax error in the estimation under the cardinal and ordinal settings.

Proposition 5 *Given a sample size n that is a multiple of $d(d-1)b(\sigma, B)$, suppose that we observe each coordinate n/d times under the CARDINAL model. Then the minimax risk is given by*

$$\inf_{\hat{w}} \sup_{w^* \in \mathcal{W}_B} \mathbb{E} \left[\|\hat{w} - w^*\|_2^2 \right] = \sigma_c^2 \frac{d}{n}. \quad (12a)$$

Similarly, if we observe each pair $n/\binom{d}{2}$ times in the THURSTONE model, then the minimax risk is sandwiched as

$$\sigma^2 b_\ell(\sigma, B) \frac{d}{n} \leq \inf_{\hat{w}} \sup_{w^* \in \mathcal{W}_B} \mathbb{E} \left[\|\hat{w} - w^*\|_2^2 \right] \leq \sigma^2 b_u(\sigma, B) \frac{d}{n}. \quad (12b)$$

In the cardinal case, when each coordinate is measured the same number of times, the CARDINAL model reduces to the well-studied normal location model, for which the MLE

is known to be the minimax estimator and its risk is straightforward to characterize (see Lehmann and Casella (1998) for instance). In the ordinal case, the result follows from the general treatment in Section 3.

Let us now return to the question deciding between the cardinal and the ordinal methods of data elicitation. Suppose that we believe the Gaussian-noise models to be reasonably correct, and the per-observation errors σ and σ_c under the two settings are known or can be separately measured. Proposition 5 shows that the scaling of the minimax error in the cardinal and ordinal settings is identical in terms of the problem parameters n and d . As an important consequence, our result thus allows for the choice to be made based *only* on the parameters (σ, σ_c, B) , and independent of n and d : the ordinal approach incurs a lower minimax error when $b_u(\sigma, B)\sigma^2 < \sigma_c^2$ while the cardinal approach is better off in terms of minimax error whenever $b_\ell(\sigma, B)\sigma^2 > \sigma_c^2$. Establishing the exact decision boundary would require tightening the constants in the bounds, a task we leave for future work.

5.3 Aggregate estimation error in experiments on MTurk

For the sake of completeness, we also computed the estimation error in the cardinal and ordinal settings. We consider data from the three experiments (c), (d) and (e).² We normalize the true vector to have $\|w^*\|_\infty = 1$ and set $B = 1$. For each of the three experiments, we execute 100 iterations of the following procedure. Select five workers from the cardinal and five from the ordinal pool of workers uniformly at random. (The number five is inspired by practical systems as in Wang et al. (2011); Piech et al. (2013).) We run the maximum-likelihood estimator of the CARDINAL model on the data from the five workers selected from the cardinal pool, and the maximum-likelihood estimator of the THURSTONE model on the data from the five workers of the ordinal pool. Note that unlike Section 5.1, the cardinal data here is *not* converted to ordinal.

Task	Spelling	Distance	Audio
$\frac{\ w^* - \hat{w}\ _2^2}{d}$ in Ordinal	0.358 \pm 0.035	0.168 \pm 0.026	0.444 \pm 0.055
$\frac{\ w^* - \hat{w}\ _2^2}{d}$ in Cardinal	0.350 \pm 0.045	0.330 \pm 0.028	0.508 \pm 0.053
Kendall’s tau coefficient in Ordinal	0.277 \pm 0.049	0.547 \pm 0.034	0.513 \pm 0.047
Kendall’s tau coefficient in Cardinal	0.129 \pm 0.046	0.085 \pm 0.038	0.304 \pm 0.049

Table 2: Evaluation of the inferred solution from the data received from multiple workers (mean and standard deviation shown).

The results are tabulated in Table 2. To put the results in perspective of the rest of the paper, let us also recall the per-sample errors in these experiments from Table 1. Observe that among these three experiments, the per-sample noise in the cardinal data was closest to that in the ordinal data in the experiment on identifying the number of spelling mistakes. The gap was larger in the two remaining experiments. This fact is reflected in

2. We restrict attention to these three experiments for the following reasons. There is no ground truth for experiments (f) and (g). In experiment (a), the size of each circle in each question is chosen independently from a continuous distribution, making all questions different and preventing aggregation. Experiment (b) employs a disconnected topology.

the results of Table 2 where the estimator on the cardinal data incurs a lower ℓ_2 -error than the estimator on the ordinal data in the experiment on identifying the number of spelling mistakes, whereas the outcome goes the other way in the two remaining experiments. Our theory needs to tighten the constants in order to address this regime. With respect to the Kendall’s tau correlation coefficient—a particular type of ordinal metric—the estimator on the ordinal data consistently gives a higher accuracy as compared to the cardinal case.

6. Conclusions

In this paper, we presented topology-aware minimax error bounds under a broad class of preference-elicitation models. We demonstrated the utility of these results in guiding the selection of comparisons and in guiding the choice of the elicitation paradigm (cardinal versus ordinal) when these options are available. A direction for future work would be to characterize the precise thresholds for making the choice between the cardinal and ordinal approaches. Secondly, the Thurstone and BTL models are parametric idealizations that have proved useful in a wide variety of applications. In future work we would like to investigate more flexible semi-parametric and non-parametric pairwise comparison models (see, for instance, Chatterjee (2014); Braverman and Mossel (2008)).

Acknowledgments

The authors are grateful to the Action Editor Moritz Hardt and anonymous referees for their suggestions that helped improve the manuscript.

This work was partially supported by Office of Naval Research MURI grant N00014-11-1-0688, Air Force Office of Scientific Research Grant AFOSR-FA9550-14-1-0016, and National Science Foundation Grants CIF-31712-23800, DMS-1107000 and CIF-81652-23800. The work of NBS was also partially supported by a Microsoft Research PhD fellowship.

Appendix A. Proof of Theorem 1

The following two sections prove the lower and upper bounds (respectively) on the minimax risk of ORDINAL model under the squared L semi-norm.

A.1 Lower bound

Our lower bounds are based on the Fano argument, which is a standard method in minimax analysis (see for instance Tsybakov (2008)). Suppose that our goal is to bound the minimax risk of estimating a parameter w over an indexed class of distributions $\mathcal{P} = \{\mathbb{P}_w \mid w \in \mathcal{W}\}$ in the square of a pseudo-metric ρ . Consider a collection of vectors $\{w^1, \dots, w^M\}$ contained within \mathcal{W} such that

$$\min_{\substack{j,k \in [M] \\ j \neq k}} \rho(w^j, w^k) \geq \delta \quad \text{and} \quad \frac{1}{\binom{M}{2}} \sum_{\substack{j,k \in [M] \\ j \neq k}} D_{\text{KL}}(\mathbb{P}_{w^j} \parallel \mathbb{P}_{w^k}) \leq \beta.$$

We refer to any such subset as an (δ, β) -packing set.

Lemma 6 (Pairwise Fano minimax lower bound) *Suppose that we can construct a (δ, β) -packing with cardinality M . Then the minimax risk is lower bounded as*

$$\inf_{\hat{w}} \sup_{w^* \in \mathcal{W}} \mathbb{E} \left[\rho(\hat{w}, w^*)^2 \right] \geq \frac{\delta^2}{2} \left(1 - \frac{\beta + \log 2}{\log M} \right). \quad (13)$$

The main difficulty in deriving the lower bound is the construction of a suitable packing set for application in Lemma 6. To this end, given a scalar $\alpha \in (0, \frac{1}{4})$ whose value will be specified later, define the integer

$$M(\alpha) := \left\lceil \exp \left\{ \frac{d}{2} (\log 2 + 2\alpha \log 2\alpha + (1 - 2\alpha) \log(1 - 2\alpha)) \right\} \right\rceil. \quad (14)$$

The following two lemmas aid in our construction of a packing set. The first lemma is a straightforward consequence of the Gilbert-Varshamov bound (Gilbert, 1952; Varshamov, 1957).

Lemma 7 *For any $\alpha \in (0, \frac{1}{4})$, there exists a set of $M(\alpha)$ binary vectors $\{z^1, \dots, z^{M(\alpha)}\} \subset \{0, 1\}^d$ such that*

$$\alpha d \leq \|z^j - z^k\|_2^2 \leq d \quad \text{for all } j \neq k \in [M(\alpha)], \text{ and} \quad (15a)$$

$$\langle e_1, z^j \rangle = 0 \quad \text{for all } j \in [M(\alpha)], \quad (15b)$$

where e_1 denotes the first canonical basis vector.

The next lemma derives an upper bound on the Kullback-Leibler divergence between the probability distributions induced by any pair of quality score vectors.

Lemma 8 *For any pair of quality score vectors w^j and w^k , and for*

$$\zeta := \frac{\max_{x \in [0, 2B/\sigma]} F'(x)}{F(2B/\sigma)(1 - F(2B/\sigma))},$$

we have

$$D_{\text{KL}}(\mathbb{P}_{w^j} \parallel \mathbb{P}_{w^k}) \leq \frac{n\zeta}{\sigma^2} (w^j - w^k)^T L (w^j - w^k). \quad (16)$$

We prove these two lemmas at the end of this section.

Taking these two lemmas as given for the moment, consider the set $\{z^1, \dots, z^{M(\alpha)}\}$ of d -dimensional binary vectors given by Lemma 7. The Laplacian L of the comparison graph is symmetric and positive-semidefinite, and so has a diagonalization of the form $L = U^T \Lambda U$ where $U \in \mathbb{R}^{d \times d}$ is an orthonormal matrix, and Λ is a diagonal matrix of nonnegative eigenvalues.

Letting matrix Λ^\dagger denote the Moore-Penrose pseudo-inverse of Λ , consider the collection $\{w^1, \dots, w^{M(\alpha)}\}$ of vectors given by $w^j := \frac{\delta}{\sqrt{d}} U^T \sqrt{\Lambda^\dagger} z^j$ for each $j \in [M(\alpha)]$. Since $1 \in$

nullspace(L), we are guaranteed that $\langle 1, w^j \rangle = \frac{\delta}{\sqrt{d}} \mathbf{1}^T U^T \sqrt{\Lambda^\dagger} z^j = 0$. On the other hand,

$$\begin{aligned} (w^j - w^k)^T L (w^j - w^k) &\leq \frac{\delta^2}{d} (z^j - z^k)^T \sqrt{\Lambda^\dagger} U L U^T \sqrt{\Lambda^\dagger} (z^j - z^k) \\ &= \frac{\delta^2}{d} (z^j - z^k)^T \sqrt{\Lambda^\dagger} \Lambda \sqrt{\Lambda^\dagger} (z^j - z^k) \\ &= \frac{\delta^2}{d} \|z^j - z^k\|_2^2, \end{aligned}$$

Here the last step makes use of the fact that the first coordinate of each vector z^j and z^k is zero. It follows that $\alpha\delta^2 \leq \|w^j - w^k\|_L^2 \leq \delta^2$.

Setting $\delta^2 := 0.01 \frac{\sigma^2 d}{n\zeta}$, we find that

$$\|w^j\|_\infty \leq \frac{\delta}{\sqrt{d}} \|\sqrt{\Lambda^\dagger} z^j\|_2 \stackrel{(i)}{\leq} \frac{\delta}{\sqrt{d}} \sqrt{\text{tr}(\Lambda^\dagger)} \stackrel{(ii)}{=} \frac{\delta}{\sqrt{d}} \sqrt{\text{tr}(L^\dagger)} \stackrel{(iii)}{\leq} B,$$

where inequality (i) follows from the fact that z^j has entries in $\{0, 1\}$; equation (ii) follows since $L^\dagger = U^T \Lambda^\dagger U$ by definition; and inequality (iii) follows from our choice of δ and our assumption $n \geq \frac{c\sigma^2 \text{tr}(L^\dagger)}{\zeta B^2}$ on the sample size with $c = 0.01$. We have thus verified that each vector w^j also satisfies the boundedness constraint $\|w^j\|_\infty \leq B$ required for membership in \mathcal{W}_B . Finally, observe that

$$\max_{j \neq k} D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) \leq \frac{n\zeta\delta^2}{\sigma^2}, \quad \text{and} \quad \min_{j \neq k} \|w^j - w^k\|_L^2 \geq \alpha\delta^2.$$

We have thus constructed a suitable packing set for applying Lemma 6, which yields the lower bound

$$\mathbb{E}[\|\hat{w} - w^*\|_L^2] \geq \frac{\alpha}{2} \delta^2 \left\{ 1 - \frac{\delta^2 \zeta n + \log 2}{\log M(\alpha)} \right\}.$$

Substituting our choice of δ and setting $\alpha = 0.01$ proves the claim for $d > 9$.

In order to handle the case $d \leq 9$, we consider the set of the three d -length vectors given by $z^1 = [0 \ \cdots \ 0 \ -1]$, $z^2 = [0 \ \cdots \ 0 \ 1]$ and $z^3 = [0 \ \cdots \ 0 \ 0]$. Construct the packing set $\{w^1, w^2, w^3\}$ from these three vectors $\{z^1, z^2, z^3\}$ as done above for the case of $d > 9$. From the calculations made for the general case above, we have for all pairs $\min_{j \neq k} \|w^j - w^k\|_L^2 \geq \frac{\delta^2}{9}$ and $\max_{j,k} \|w^j - w^k\|_L^2 \leq 4\delta^2$, and as a result $\max_{j,k} D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) \leq \frac{4n\zeta\delta^2}{\sigma^2}$. Choosing $\delta^2 = \frac{\sigma^2 \log 2}{8n\zeta}$ and applying Lemma 6 proves the theorem.

The only remaining detail is to prove and Lemma 7 and Lemma 8.

Proof of Lemma 7: The Gilbert-Varshamov bound Gilbert (1952); Varshamov (1957) guarantees the existence of a binary code $\{\tilde{z}^1, \dots, \tilde{z}^N\}$ in dimension $(d-1)$, minimum Hamming distance $\lceil \alpha d \rceil$, and the number of code words N at least

$$N \geq \frac{2^{d-1}}{\sum_{\ell=0}^{\lceil \alpha d \rceil - 1} \binom{d-1}{\ell}}.$$

Since $d \geq 2$ and $\alpha \in (0, \frac{1}{4})$, we have

$$\frac{\lceil \alpha d \rceil - 1}{d - 1} \leq 2\alpha \leq \frac{1}{2}.$$

Applying standard bounds on the tail of the binomial distribution gives

$$\begin{aligned} \frac{1}{2^{d-1}} \sum_{\ell=0}^{\lceil \alpha d \rceil - 1} \binom{d-1}{\ell} &\leq \exp\left(- (d-1) D_{\text{KL}}\left(\frac{\lceil \alpha d \rceil - 1}{d-1} \parallel \frac{1}{2}\right)\right) \\ &\leq \exp\left(- (d-1) D_{\text{KL}}\left(2\alpha \parallel \frac{1}{2}\right)\right), \end{aligned}$$

and hence $N \geq M(\alpha)$.

We now define the set of vectors $\{z^1, \dots, z^{M(\alpha)}\}$ as $(z^i)^T = [0 \ (\tilde{z}^i)^T]$ for every $i \in [M(\alpha)]$. Given this condition, it is easy to see that $\langle e_1, z^j \rangle = 0$ for every vector z^j in this set. Finally, the minimum distance condition gives the desired constraints on the difference between any pair of vectors in this set under the squared ℓ_2 metric.

Proof of Lemma 8: For any pair of quality score vectors w^j and w^k , the KL divergence between the distributions \mathbb{P}_{w^j} and \mathbb{P}_{w^k} is given by

$$D_{\text{KL}}(\mathbb{P}_{w^j} \parallel \mathbb{P}_{w^k}) = \sum_{i=1}^n F(\langle w^j, x_i \rangle / \sigma) \log \frac{F(\langle w^j, x_i \rangle / \sigma)}{F(\langle w^k, x_i \rangle / \sigma)} + (1 - F(\langle w^j, x_i \rangle / \sigma)) \log \frac{1 - F(\langle w^j, x_i \rangle / \sigma)}{1 - F(\langle w^k, x_i \rangle / \sigma)}.$$

For any $a, b \in (0, 1)$, we have the elementary inequality $a \log \frac{a}{b} \leq (a - b) \frac{a}{b}$. Applying this inequality to our expression above gives

$$\begin{aligned} D_{\text{KL}}(\mathbb{P}_{w^j} \parallel \mathbb{P}_{w^k}) &\leq \sum_{i=1}^n (F(\langle w^j, x_i \rangle / \sigma) - F(\langle w^k, x_i \rangle / \sigma)) \frac{F(\langle w^j, x_i \rangle / \sigma)}{F(\langle w^k, x_i \rangle / \sigma)} \\ &\quad - \left\{ F(\langle w^j, x_i \rangle / \sigma) - F(\langle w^k, x_i \rangle / \sigma) \right\} \frac{1 - F(\langle w^j, x_i \rangle / \sigma)}{1 - F(\langle w^k, x_i \rangle / \sigma)} \\ &\leq \sum_{i=1}^n \frac{(F(\langle w^j, x_i \rangle / \sigma) - F(\langle w^k, x_i \rangle / \sigma))^2}{F(\langle w^k, x_i \rangle / \sigma)(1 - F(\langle w^k, x_i \rangle / \sigma))}. \end{aligned}$$

Since $\max\{\|w^j\|_\infty, \|w^k\|_\infty\} \leq B$, and since F is a non-decreasing function, we have

$$D_{\text{KL}}(\mathbb{P}_{w^j} \parallel \mathbb{P}_{w^k}) \leq \sum_{i=1}^n \frac{(F(\langle w^j, x_i \rangle / \sigma) - F(\langle w^k, x_i \rangle / \sigma))^2}{F(2B/\sigma)(1 - F(2B/\sigma))}.$$

Finally, applying the mean value theorem and recalling the definition of ζ (from (6)) yields

$$D_{\text{KL}}(\mathbb{P}_{w^j} \parallel \mathbb{P}_{w^k}) \leq \sum_{i=1}^n \zeta(\langle w^j, x_i \rangle / \sigma - \langle w^k, x_i \rangle / \sigma)^2 = \frac{n\zeta}{\sigma^2} (w^j - w^k)^T L (w^j - w^k),$$

as claimed.

A.2 Upper bound

For the ORDINAL model, the MLE is given by $\hat{w} \in \arg \min_{w \in \mathcal{W}_B} \ell(w)$, where

$$\ell(w) = -\frac{1}{n} \sum_{i=1}^n \left\{ \mathbf{1}[y_i = 1] \log F\left(\frac{\langle x_i, w \rangle}{\sigma}\right) + \mathbf{1}[y_i = -1] \log \left(1 - F\left(\frac{\langle x_i, w \rangle}{\sigma}\right)\right) \right\}, \quad \text{and} \quad (17a)$$

$$\mathcal{W}_B := \{w \in \mathbb{R}^d \mid \langle 1, w \rangle = 0, \quad \text{and} \quad \|w\|_\infty \leq B\}. \quad (17b)$$

Our goal is to bound the estimation error of the MLE in the squared semi-norm $\|v\|_L^2 = v^T L v$.

For the purposes of this proof (as well as subsequent ones), let us state and prove an auxiliary lemma that applies more generally to M -estimators that are based on minimizing an arbitrary convex and differentiable function over some subset \mathcal{W} of the set $\mathcal{W}_\infty := \{w \in \mathbb{R}^d \mid \langle 1, w \rangle = 0\}$. The MLE under consideration here is a special case. This lemma requires that ℓ is differentiable and strongly convex at w^* with respect to the semi-norm $\|\cdot\|_L$, meaning that there is some constant $\kappa > 0$ such that

$$\ell(w^* + \Delta) - \ell(w^*) - \langle \nabla \ell(w^*), \Delta \rangle \geq \kappa \|\Delta\|_L^2 \quad (18)$$

for all perturbations $\Delta \in \mathbb{R}^d$ such that $(w^* + \Delta) \in \mathcal{W}$. Finally, it is also convenient to introduce the semi-norm $\|u\|_{L^\dagger} = \sqrt{u^T L^\dagger u}$, where L^\dagger is the Moore-Penrose pseudo-inverse of L .

Lemma 9 (Upper bound for M -estimators) *Consider the M -estimator*

$$\hat{w} \in \arg \min_{w \in \mathcal{W}} \ell(w), \quad \text{where } \mathcal{W} \text{ is any subset of } \mathcal{W}_\infty, \quad (19)$$

and ℓ is a differentiable cost function satisfying the κ -strong convexity condition (18) at some $w^* \in \mathcal{W}$. Then

$$\|\hat{w} - w^*\|_L \leq \frac{1}{\kappa} \|\nabla \ell(w^*)\|_{L^\dagger}. \quad (20)$$

Proof Since \hat{w} and w^* are optimal and feasible, respectively, for the original optimization problem, we have $\ell(\hat{w}) \leq \ell(w^*)$. Defining the error vector $\Delta = \hat{w} - w^*$, adding and subtracting the quantity $\langle \nabla \ell(w^*), \Delta \rangle$ yields the bound

$$\ell(w^* + \Delta) - \ell(w^*) - \langle \nabla \ell(w^*), \Delta \rangle \leq -\langle \nabla \ell(w^*), \Delta \rangle.$$

By the κ -convexity condition, the left-hand side is lower bounded by $\kappa \|\Delta\|_L^2$. As for the right-hand side, note that Δ satisfies the constraint $\langle 1, \Delta \rangle = 0$, and thus is orthogonal to the nullspace of the Laplacian matrix L . Therefore, by Lemma 16 (in Appendix F), we have $|\langle \nabla \ell(w^*), \Delta \rangle| \leq \|\nabla \ell(w^*)\|_{L^\dagger} \|\Delta\|_L$. Combining the pieces yields the claimed inequality (20). \blacksquare

In order to apply Lemma 9 to the MLE for the ORDINAL model, we need to verify that the negative log likelihood (17a) satisfies the strong convexity condition, and we need to bound the random variable $\|\nabla \ell(w^*)\|_{L^\dagger}$ defined in the dual norm $\|\cdot\|_{L^\dagger}$.

Verifying strong convexity: By chain rule, the Hessian of ℓ is given by

$$\nabla^2 \ell(w) = \frac{1}{n\sigma^2} \sum_{i=1}^n \left\{ \mathbf{1}[y_i = 1]T_{i1} + \mathbf{1}[y_i = -1]T_{i2} \right\} x_i x_i^T,$$

where

$$T_{i1} := \frac{F'(\frac{\langle w, x_i \rangle}{\sigma})^2 - F(\frac{\langle w, x_i \rangle}{\sigma})F''(\frac{\langle w, x_i \rangle}{\sigma})}{F(\frac{\langle w, x_i \rangle}{\sigma})^2}, \quad \text{and} \quad T_{i2} := \frac{F'(\frac{\langle w, x_i \rangle}{\sigma})^2 + (1 - F(\frac{\langle w, x_i \rangle}{\sigma}))F''(\frac{\langle w, x_i \rangle}{\sigma})}{(1 - F(\frac{\langle w, x_i \rangle}{\sigma}))^2}.$$

Observe that the term T_{i1} is simply the second derivative of $\log F$ evaluated at $\frac{\langle w, x_i \rangle}{\sigma}$, and hence the strong log-concavity of F implies $T_{i1} \geq \gamma$. On the other hand, the term T_{i2} is the second derivative of $\log(1 - F)$. Since $F(-x) = 1 - F(x)$ for all x , it follows that the function $x \mapsto 1 - F(x)$ is also strongly log-concave with parameter γ and hence $T_{i2} \geq \gamma$. Putting together the pieces, we conclude that

$$v^T \nabla^2 \ell(w) v \geq \frac{\gamma}{n\sigma^2} \|Xv\|_2^2 \quad \text{for all } v, w \in \mathcal{W}_B,$$

where $X \in \mathbb{R}^{n \times d}$ has the differencing vector $x_i \in \mathbb{R}^d$ as its i^{th} row.

Thus, if we introduce the error vector $\Delta := \hat{w} - w^*$, then we may conclude that

$$\ell(w^* + \Delta) - \ell(w^*) - \langle \nabla \ell(w^*), \Delta \rangle \geq \frac{\gamma}{n\sigma^2} \|X\Delta\|_2^2 = \frac{\gamma}{\sigma^2} \|\Delta\|_L^2,$$

showing that ℓ is strongly convex around w^* with parameter $\kappa = \frac{\gamma}{\sigma^2}$. An application of Lemma 9 then gives $\|\Delta\|_L^2 \leq \frac{\sigma^4}{\gamma^2} \|\nabla \ell(w^*)\|_{L^\dagger}^2$.

Bounding the dual norm: In order to obtain a concrete bound, it remains to control the quantity $\nabla \ell(w^*)^T L^\dagger \nabla \ell(w^*)$. Observe that the gradient takes the form

$$\nabla \ell(w^*) = \frac{-1}{n\sigma} \sum_{i=1}^n \left[\mathbf{1}[y_i = 1] \frac{F'(\langle w^*, x_i \rangle / \sigma)}{F(\langle w^*, x_i \rangle / \sigma)} - \mathbf{1}[y_i = -1] \frac{F'(\langle w^*, x_i \rangle / \sigma)}{1 - F(\langle w^*, x_i \rangle / \sigma)} \right] x_i.$$

Define a random vector $V \in \mathbb{R}^n$ with independent components as

$$V_i = \begin{cases} \frac{F'(\langle w^*, x_i \rangle / \sigma)}{F(\langle w^*, x_i \rangle / \sigma)} & \text{w.p. } F(\langle w^*, x_i \rangle / \sigma) \\ \frac{-F'(\langle w^*, x_i \rangle / \sigma)}{1 - F(\langle w^*, x_i \rangle / \sigma)} & \text{w.p. } 1 - F(\langle w^*, x_i \rangle / \sigma). \end{cases}$$

With this notation, we have $\nabla \ell(w^*) = -\frac{1}{n\sigma} X^T V$. One can verify that $\mathbb{E}[V] = 0$ and

$$|V_i| \leq \sup_{z \in [-2B/\sigma, 2B/\sigma]} \max \left\{ \frac{F'(z)}{F(z)}, \frac{F'(z)}{1 - F(z)} \right\} \leq \sup_{z \in [-2B/\sigma, 2B/\sigma]} \frac{F'(z)}{F(z)(1 - F(z))} \leq \zeta, \quad (21)$$

where ζ is as defined in (6). Defining the n -dimensional square matrix $M := \frac{\sigma^2}{\gamma^2 n^2} X L^\dagger X^T$, our definitions and previous bounds imply that $\|\Delta\|_L^2 \leq V^T M V$.

Consequently, our problem has been reduced to controlling the fluctuations of the quadratic form $V^T M V$; in order to do so, we apply the Hanson-Wright inequality (see Lemma 13 in Appendix E). A straightforward calculation yields

$$\|M\|_{\text{fro}}^2 = (d-1) \frac{\sigma^4}{\gamma^4 n^2} \quad \text{and} \quad \|M\|_{\text{op}} = \frac{\sigma^2}{\gamma^2 n},$$

where we have used the fact that $L = \frac{1}{n} X^T X$. Moreover, since the components of V are independent and of zero mean, a straightforward calculation yields that $\mathbb{E}[V^T M V] \leq \mathbb{E}[\|V\|_\infty^2 \text{tr}(M)] \leq \frac{\zeta^2 \sigma^2 d}{\gamma^2 n}$.

Since $|V_i| \leq \zeta$, the variables are ζ -sub-Gaussian, and hence the Hanson-Wright inequality implies that

$$\mathbb{P}\left[V^T M V - \frac{\zeta^2 \sigma^2 d}{\gamma^2 n} > t\right] \leq 2 \exp\left(-c \min\left\{\frac{t^2 \gamma^4 n^2}{\zeta^4 (d-1) \sigma^4}, \frac{t \gamma^2 n}{\zeta^2 \sigma^2}\right\}\right) \quad \text{for all } t > 0.$$

Consequently, after some simple algebra, we conclude that

$$\mathbb{P}\left(\|\Delta\|_L^2 > t \frac{c \zeta^2 \sigma^2 d}{\gamma^2 n}\right) \leq e^{-t} \quad \text{for all } t \geq 1,$$

for some universal constant c . Integrating this tail bound yields the bound on the expectation.

Appendix B. Proof of Theorem 2

The following two sections prove the upper and lower bounds (respectively) on the minimax risk in the squared Euclidean norm for ORDINAL model. We prove the lower bound in two parts corresponding to the two components of the ‘‘max’’ in the statement of the theorem.

B.1 Upper bound

The proof of the upper bound under the Euclidean norm follows directly from the upper bound under the L semi-norm proved in Theorem 1. From the setting described in Section 2, we have that the nullspace of the matrix L is given by the span of the all ones vector. Furthermore, we have $\langle w^* - \hat{w}, 1 \rangle = 0$, and $\|w^* - \hat{w}\|_L^2 \geq \lambda_2(L) \|w^* - \hat{w}\|_2^2$. Substituting this inequality into the upper bound (7b) gives the desired result.

B.2 Lower bound: Part I

Since the Laplacian L of the comparison graph is symmetric and positive-semidefinite. By diagonalization, we can write $L = U^T \Lambda U$ where $U \in \mathbb{R}^{d \times d}$ is an orthonormal matrix, and Λ is a diagonal matrix of nonnegative eigenvalues with $\Lambda_{jj} = \lambda_j(L)$.

We first use the Fano method (Lemma 6) to prove that the minimax risk is lower bounded as $c \sigma^2 \frac{d^2}{n}$. For scalars $\alpha \in (0, \frac{1}{4})$ and $\delta > 0$ whose values will be specified later, recall the set $\{z^1, \dots, z^{M(\alpha)}\}$ of vectors in the Boolean hypercube $\{0, 1\}^d$ given by Lemma 7. We then define a second set $\{w^j, j \in [M(\alpha)]\}$ via $w^j := \frac{\delta}{\sqrt{d}} U^T P z^j$, where P is a permutation matrix to be specified momentarily. At this point, the only constraint imposed on P is

that it keeps the first coordinate constant. By construction, for each $j \neq k$, we have $\|w^j - w^k\|_2^2 = \frac{\delta^2}{d} \|z^j - z^k\|_2^2 \geq \alpha\delta^2$, where the final inequality follows from the fact that the set $\{z^1, \dots, z^{M(\alpha)}\}$ comprises binary vectors with a minimum Hamming distance at least αd .

Consider any distinct $j, k \in [M(\alpha)]$. Then, for some $\{i_1, \dots, i_r\} \subseteq \{2, \dots, d\}$ with $\alpha d \leq r \leq d$, it must be that

$$\|w^j - w^k\|_L^2 = \frac{\delta^2}{d} \|U^T P z^j - U^T P z^k\|_L^2 = \frac{\delta^2}{d} \|z^j - z^k\|_\Lambda^2 = \frac{\delta^2}{d} \sum_{m=1}^r \lambda_{i_m}(L).$$

It follows that for some non-negative numbers a_2, \dots, a_d such that $\alpha d \leq \sum_{i=2}^d a_i \leq d$,

$$\frac{1}{\binom{M(\alpha)}{2}} \sum_{j \neq k} \|w^j - w^k\|_L^2 = \frac{\delta^2}{d} \sum_{i=2}^d a_i \lambda_i(L).$$

We choose the permutation matrix P such that the last $(d-1)$ coordinates are permuted to have $a_2 \geq \dots \geq a_d$ and the d^{th} coordinate remains fixed. With this choice, we get

$$\frac{1}{\binom{M(\alpha)}{2}} \sum_{j \neq k} \|w^j - w^k\|_L^2 \leq \frac{\delta^2}{d} \frac{d}{d-1} \text{tr}(L) \leq \frac{2\delta^2}{d} \text{tr}(L).$$

Lemma (14) (Appendix F) gives the trace constraint $\text{tr}(L) = 2$, which in turn guarantees that $\frac{1}{\binom{M(\alpha)}{2}} \sum_{j \neq k} \|w^j - w^k\|_L^2 \leq \frac{4\delta^2}{d}$. For the choice of P specified above, we have for every $j \in [M(\alpha)]$,

$$\langle 1, w^j \rangle = \frac{\delta}{\sqrt{d}} e_1^T P z^j = e_1^T z^j = 0,$$

where the final equation employed the property (15b).

Setting $\delta^2 = 0.01 \frac{\sigma^2 d^2}{4n\zeta}$, we have $\|w^j\|_\infty \leq \frac{\delta}{\sqrt{d}} \|z^j\|_2 \stackrel{(i)}{\leq} \delta \stackrel{(ii)}{\leq} B$, where inequality (i) follows from the fact that z^j has entries in $\{0, 1\}$; inequality (ii) follows from our choice of δ and our assumption $n \geq \frac{c\sigma^2 \text{tr}(L^1)}{\zeta B^2}$ on the sample size with $c = 0.002$, where Lemma 14 guarantees $n \geq \frac{c\sigma^2 d^2}{4\zeta B^2}$. We have thus verified that each vector w^j also satisfies the boundedness constraint $\|w^j\|_\infty \leq B$ required for membership in \mathcal{W}_B .

From the proof of Theorem 1, we have that for any distinct $D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) \leq \frac{n\zeta}{\sigma^2} \|w^j - w^k\|_L^2$, and hence

$$\frac{1}{\binom{M(\alpha)}{2}} \sum_{j \neq k} D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) \leq \frac{n\zeta}{\sigma^2} \frac{4\delta^2}{d} = 0.01 d,$$

where we have substituted our previous choice of δ .

Applying Lemma 6 with the packing set $\{w^1, \dots, w^{M(\alpha)}\}$ gives that any estimator \tilde{w} must incur an error lower bounded as

$$\sup_{w^* \in \mathcal{W}_B} \mathbb{E}[\|\tilde{w} - w^*\|_2^2] \geq \frac{\alpha\delta^2}{2} \left(1 - \frac{0.01d + \log 2}{\log M(\alpha)}\right).$$

Substituting our choice of δ and setting $\alpha = 0.01$ proves the claim for $d > 9$.

For the case of $d \leq 9$, consider the set of the three d -length vectors $z^1 = [0 \ \cdots \ 0 \ -1]$, $z^2 = [0 \ \cdots \ 0 \ 1]$ and $z^3 = [0 \ \cdots \ 0 \ 0]$. Construct the packing set $\{w^1, w^2, w^3\}$ from these three vectors $\{z^1, z^2, z^3\}$ as done above for the case of $d > 9$. From the calculations made for the general case above, we have for all pairs $\min_{j \neq k} \|w^j - w^k\|_2^2 \geq \frac{\delta^2}{9}$ and $\max_{j,k} \|w^j - w^k\|_L^2 \leq 4\delta^2$, and as a result $\max_{j,k} D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) \leq \frac{4n\zeta\delta^2}{\sigma^2}$. Choosing $\delta^2 = \frac{\sigma^2 \log 2}{8n\zeta}$ and applying Lemma 6 yields the claim.

B.3 Lower bound: Part II

Given an integer $d' \in \{2, \dots, d\}$, and scalars $\alpha \in (0, \frac{1}{4})$ and $\delta > 0$, define the integer

$$M'(\alpha) := \left\lceil \exp \left\{ \frac{d'}{2} (\log 2 + 2\alpha \log 2\alpha + (1 - 2\alpha) \log(1 - 2\alpha)) \right\} \right\rceil. \quad (22)$$

Applying Lemma 7 with d' as the dimension yields a subset $\{z^1, \dots, z^{M'(\alpha)}\}$ of the Boolean hypercube $\{0, 1\}^{d'}$ with the stated properties. We then define a set of d -length vectors $\{\tilde{w}^1, \dots, \tilde{w}^{M'(\alpha)}\}$ via

$$\tilde{w}^j = [0 \ (z^j)^T \ 0 \ \cdots \ 0]^T \quad \text{for each } j \in [M(\alpha)].$$

For each $j \in [M(\alpha)]$, let us define $w^j := \frac{\delta}{\sqrt{d'}} U^T \sqrt{\Lambda^\dagger} \tilde{w}^j$. Now, letting $e_1 \in \mathbb{R}^d$ denote the first standard basis vector, we have $\langle 1, w^j \rangle = \frac{\delta}{\sqrt{d'}} 1^T U^T \sqrt{\Lambda^\dagger} \tilde{w}^j = 0$. where we have used the fact that $1 \in \text{nullspace}(L)$. Furthermore, for any $j \neq k$, we have

$$\|w^j - w^k\|_2^2 = \frac{\delta^2}{d'} (\tilde{w}^j - \tilde{w}^k)^T \Lambda^\dagger (\tilde{w}^j - \tilde{w}^k) \geq \frac{\delta^2}{d'} \sum_{i=\lceil(1-\alpha)d'\rceil}^{d'} \frac{1}{\lambda_i}.$$

Thus, setting $\delta^2 = 0.01 \frac{\sigma^2 d'}{n\zeta}$ yields

$$\|w^j\|_\infty \leq \frac{\delta}{\sqrt{d'}} \|\sqrt{\Lambda^\dagger} \tilde{w}^j\|_2 \stackrel{(i)}{\leq} \frac{\delta}{\sqrt{d'}} \sqrt{\text{tr}(\Lambda^\dagger)} \stackrel{(ii)}{=} \frac{\delta}{\sqrt{d'}} \sqrt{\text{tr}(L^\dagger)} \stackrel{(iii)}{\leq} B,$$

where inequality (i) follows from the fact that z^j has entries in $\{0, 1\}$; step (ii) follows because the matrices $\sqrt{\Lambda^\dagger}$ and $\sqrt{L^\dagger}$ have the same eigenvalues; and inequality (iii) follows from our choice of δ and our assumption $n \geq \frac{c\sigma^2 \text{tr}(L^\dagger)}{\zeta B^2}$ on the sample size with $c = 0.01$. We have thus verified that each vector w^j also satisfies the boundedness constraint $\|w^j\|_\infty \leq B$ required for membership in \mathcal{W}_B . Furthermore, for any pair of distinct vectors in this set, we have

$$\|w^j - w^k\|_L^2 = \frac{\delta^2}{d'} \|z^j - z^k\|_2^2 \leq \delta^2.$$

From the proof of Theorem 1, we $D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) \leq \frac{n\zeta}{\sigma^2} \|w^j - w^k\|_L^2 \leq 0.01d'$. Applying Lemma 6 with the packing set $\{w^1, \dots, w^{M'(\alpha)}\}$ gives that any estimator \tilde{w} must incur an error lower bounded as

$$\sup_{w^* \in \mathcal{W}_B} \mathbb{E}[\|\tilde{w} - w^*\|_2^2] \geq \frac{\delta^2 \sum_{i=\lceil(1-\alpha)d'\rceil}^{d'} \frac{1}{\lambda_i}}{2} \left(1 - \frac{0.01d' + \log 2}{\log M'(\alpha)}\right).$$

Substituting our choice of δ and setting $\alpha = 0.01$ proves the claim for $d' > 9$.

For the case of $d' \leq 9$, we now prove a lower bound of $\frac{c\sigma^2}{n} \frac{9}{\lambda_2(L)}$ for a universal constant $c > 0$. This quantity is at least as large as the claimed lower bound. Consider the packing set of three d -length vectors $w^1 = \delta U \sqrt{\Lambda^\dagger} [0 \ 1 \ 0 \ \dots \ 0]^T$, $w^2 = -w^1$ and $w^3 = [0 \ \dots \ 0]^T$ for some $\delta > 0$. Then for every $j \neq k$, one can verify that $\|w^j - w^k\|_L^2 \leq 4\delta^2$, $\|w^j - w^k\|_2^2 \geq \frac{\delta^2}{\lambda_2(L)}$. Choosing $\delta^2 = \frac{\sigma^2 \log 2}{8n\zeta}$ and applying Lemma 6 proves the claim for $d' \leq 9$.

Finally, taking the maximum over all values of $d' \in \{2, \dots, d\}$ gives the claimed lower bound.

Appendix C. Proof of Theorem 3

We now turn to the proof of Theorem 3 on the minimax rate for the PAIRED CARDINAL model. Recall that this observation model takes the standard linear model, $y = Xw^* + \epsilon$, where $y \in \mathbb{R}^n$, $w \in \mathbb{R}^d$ and $\epsilon \sim N(0, \sigma^2 I)$.

C.1 Upper bound under the squared L semi-norm

The maximum likelihood estimate in the PAIRED CARDINAL model is a special case of the general M -estimator (19) with $\ell(w) := \frac{1}{2n} \sum_{i=1}^n (y_i - \langle x_i, w \rangle)^2$. For this quadratic objective function, it is easy to verify that the γ -convexity condition holds with $\gamma = 1$. (In particular, note that the Hessian of ℓ is given by $L = X^T X/n$.)

Given the result of Lemma 9, it remains to upper bound $\|\nabla \ell(w^*)\|_{L^\dagger}$. A straightforward computation yields $\|\nabla \ell(w^*)\|_{L^\dagger}^2 = \frac{\epsilon^T Q \epsilon}{\sigma^2}$ where $Q := \frac{\sigma^2}{n^2} X L^\dagger X^T$. Consequently, the random variable $\|\nabla \ell(w^*)\|_{L^\dagger}^2$ is quadratic form in the standard Gaussian random vector $\frac{\epsilon}{\sigma}$. An application of Lemma 15 (Appendix F) gives $\text{tr}(Q) = \frac{\sigma^2}{n} (d-1)$ and $\|Q\|_{\text{op}} = \frac{\sigma^2}{n}$, and then applying a known tail bound on Gaussian quadratic forms (see Lemma 12 in Appendix E) yields

$$\mathbb{P} \left[\frac{\|\nabla \ell(w^*)\|_{L^\dagger}^2}{\sigma^2} \geq \left(\sqrt{\frac{d}{n}} + \frac{\delta}{\sqrt{n}} \right)^2 \right] \leq e^{-\frac{\delta^2}{2}} \quad \text{for all } \delta \geq 0.$$

Since $d \geq 2$, we have $(\sigma \sqrt{\frac{d}{n}} + \frac{\sigma}{\sqrt{n}} \delta)^2 \leq \frac{2\sigma^2 d \delta^2}{n}$ for all $\delta \geq 4$, which yields

$$\mathbb{P} \left[\|\nabla \ell(w^*)\|_{L^\dagger}^2 \geq t \frac{4\sigma^2 d}{n} \right] \leq e^{-t} \quad \text{for all } t \geq 8.$$

Integrating this tail bound yields that $\mathbb{E} \left[\|\nabla \ell(w^*)\|_{L^\dagger}^2 \right] \leq c\sigma^2 \frac{d}{n}$, from which the claim follows.

C.2 Lower bound under the squared L semi-norm

Based on the pairwise Fano lower bound previously stated in Lemma 6, we need to construct a suitable (δ, β) -packing, where the semi-norm $\rho(w^j, w^k) = \|w^j - w^k\|_L$ is defined by the Laplacian. Given the additive Gaussian noise observation model, we also have

$$D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) = \frac{n}{2\sigma^2} \|w^j - w^k\|_L^2, \quad (23)$$

The construction of the packing and the remainder of the proof proceeds in a manner identical to the proof of the lower bound in Theorem 1, except for the absence of the requirement of $\|w^j\|_\infty \leq B$ on the elements $\{w^j\}$ of the packing set.

C.3 Upper bound under the squared Euclidean norm

The upper bound follows by direct analysis of the (unconstrained) least-squares estimate, which has the explicit form $\hat{w} = \frac{1}{n}L^\dagger X^T y$, and thus

$$\mathbb{E}\|\hat{w} - w^*\|_2^2 = \mathbb{E}\left\|\frac{1}{n}L^\dagger X^T \epsilon\right\|_2^2 = \sigma^2 \text{tr}\left(\frac{1}{n^2}L^\dagger X^T X L^\dagger\right)$$

where we have used the fact that $\epsilon \sim N(0, \sigma^2 I_n)$. Since $L = X^T X/n$ by definition, we conclude that $\mathbb{E}\|\hat{w} - w^*\|_2^2 = \frac{\sigma^2 \text{tr}(L^\dagger)}{n}$ as claimed.

C.4 Lower bound under the squared Euclidean norm

We obtain the lower bound by computing the Bayes risk with respect to a suitably defined (proper) prior distribution over the weight vector w^* . In particular, if we impose the prior $w^* \sim N(0, \frac{\sigma^2}{n}L^\dagger)$, Bayes' rule then leads to the posterior distribution

$$\mathbb{P}(w \mid y; X) \propto \exp\left(\frac{-1}{2\sigma^2}\|y - Xw\|_2^2\right) \exp\left(\frac{-n}{2\sigma^2}w^T L w\right) \mathbf{1}\{\langle w, 1 \rangle = 0\}.$$

Thus conditioned on y , w is distributed as $N\left((X^T X + nL)^{-1}X^T y, \frac{\sigma^2}{2}L^\dagger\right)$. By applying iterated expectations, the Bayes risk is given by $\mathbb{E}\|w - \frac{1}{2}L^\dagger X^T y\|_2^2 = \frac{\sigma^2}{2}\text{tr}(L^\dagger)$, which completes the proof.

Appendix D. Proof of Theorem 4

This section presents the proof of Theorem 4 for the setting of m -wise comparisons. We first state some simple properties of the model introduced in Section 3.3, which we use subsequently in the proofs of the results.

Lemma 10 *The Laplacian of the underlying pairwise-comparison graph satisfies the trace constraints $\text{nullspace}(L) = 1$, $\lambda_2(L) > 0$ and $\text{tr}(L) = m(m-1)$.*

Lemma 11 *For any $j \in [m]$, $i \in [n]$ and any vector $v \in \mathbb{R}^m$, we have*

$$\frac{\lambda_2(H)}{m}v^T(mI - 11^T)v \leq v^T R_j H R_j^T v \leq \frac{\lambda_{\max}(H)}{m}v^T(mI - 11^T)v.$$

See Section D.2 for the proof of these auxiliary lemmas.

D.1 Upper bound under the squared L semi-norm

We prove this upper bound by applying Lemma 9. In this case, the rescaled negative log likelihood takes the form

$$\ell(w) = -\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^m \mathbf{1}[y_i = j] \log F(w^T E_i R_j),$$

and the MLE is obtained by constrained minimization over the set $\mathcal{W}_B := \{w \in \mathbb{R}^d \mid \langle \mathbf{1}, w \rangle = 0, \text{ and } \|w\|_\infty \leq B\}$. As in our proof of the upper bound in Theorem 1, we need to verify the κ -strong convexity condition, and to control the dual norm $\|\nabla \ell(w^*)\|_{L^\dagger}$.

Verifying strong convexity: The gradient of the negative log likelihood is

$$\nabla \ell(w) = -\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^m \mathbf{1}[y_i = j] E_i R_j \nabla \log F(v) \Big|_{v=w^T E_i R_j}.$$

The Hessian of the negative log likelihood can be written as

$$\nabla^2 \ell(w) = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^m \mathbf{1}[y_i = j] E_i R_j \nabla^2 \log F(v) \Big|_{v=w^T E_i R_j} R_j^T E_i^T.$$

Using our strongly log-concave assumption on F , we have that for any vector $z \in \mathbb{R}^d$,

$$\begin{aligned} z^T \nabla^2 \ell(w) z &= -\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^m \mathbf{1}[y_i = j] z^T E_i R_j \nabla^2 \log F(v) \Big|_{v=w^T E_i R_j} R_j^T E_i^T z \\ &\geq \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^m \mathbf{1}[y_i = j] z^T E_i R_j H R_j^T E_i^T z \\ &\geq \frac{\lambda_2(H)}{m} \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^m \mathbf{1}[y_i = j] z^T E_i (mI - \mathbf{1}\mathbf{1}^T) E_i^T z, \end{aligned}$$

where the last step follows from Lemma 11. The definition (10) of L implies that

$$z^T \nabla^2 \ell(w) z \geq \frac{\lambda_2(H)}{m} z^T L z = \frac{\lambda_2(H)}{m} \|z\|_L^2.$$

Consequently, the κ -convexity condition holds around w^* with $\kappa = \frac{\lambda_2(H)}{m}$. An application of Lemma 9 then yields

$$\|\widehat{w}_{\text{ML}} - w^*\|_L^2 \leq \frac{m^2}{\lambda_2(H)^2} \|\nabla \ell(w^*)\|_{L^\dagger}^2 = \frac{m^2}{\lambda_2(H)^2} \nabla \ell(w^*)^T L^\dagger \nabla \ell(w^*). \quad (24)$$

Controlling the dual norm: The gradient of the negative log likelihood can then be rewritten as $\nabla \ell(w^*) = -\frac{1}{n} \sum_{i=1}^n E_i V_i$, where each index $i \in [n]$, the random vector $V_i \in \mathbb{R}^m$ is given by $V_i := \sum_{j=1}^m \mathbf{1}[y_i = j] R_j \nabla \log F(\langle w^*, E_i \rangle R_j)$. Now observe that the matrix $M := I - \frac{1}{m} \mathbf{1}\mathbf{1}^T$ is symmetric and positive semi-definite with rank $(m-1)$, eigenvalues $\{1, \dots, 1, 0\}$, its nullspace equals the span of the all-ones vector, and that $M^\dagger = M$. Using this matrix, we define the transformed vector $\widetilde{V}_i := (M^\dagger)^{\frac{1}{2}} V_i$ for each $i \in [n]$.

Consider a vector x and its shifted version $x+t\mathbf{1}$, where $t \in \mathbb{R}$ and $\mathbf{1}$ denotes the vector of all ones. By the shift invariance property, the function $g(t) = F(x+t\mathbf{1}) - F(x)$ is constant, and hence

$$g'(0) = \langle \nabla F(x), \mathbf{1} \rangle = 0, \quad \text{and} \quad g''(0) = \langle \mathbf{1}, (\nabla^2 F(x)) \mathbf{1} \rangle = 0, \quad (25)$$

which implies that $1 \in \text{nullspace}(\nabla^2 F(x))$. Continuing on, we also have that $\langle \nabla \log F(x), 1 \rangle = \frac{1}{F(x)} \langle \nabla F(x), 1 \rangle = 0$. Consequently, $\langle V_i, 1 \rangle = 0 = \langle V_i, \text{nullspace}(M) \rangle$. This allows us to write

$$\nabla \ell(w^*) = -\frac{1}{n} \sum_{i=1}^n E_i M^{\frac{1}{2}} \tilde{V}_i, \quad \text{and} \quad \nabla \ell(w^*)^T L^\dagger \nabla \ell(w^*) = \frac{1}{n^2} \sum_{i=1}^n \sum_{\ell=1}^n \tilde{V}_i^T M^{\frac{1}{2}} E_i^T L^\dagger E_\ell M^{\frac{1}{2}} \tilde{V}_\ell.$$

By definition, for every pair $i \neq \ell \in [n]$, \tilde{V}_i is independent of \tilde{V}_ℓ . Moreover, for every $i \in [n]$,

$$\begin{aligned} \mathbb{E}[\tilde{V}_i] &= \mathbb{E}[(M^\dagger)^{\frac{1}{2}} \sum_{j=1}^m \mathbf{1}[y_i = j] R_j \nabla \log F(v) \Big|_{v=(w^*)^T E_i R_j}] \\ &= (M^\dagger)^{\frac{1}{2}} \sum_{j=1}^m F((w^*)^T E_i R_j) R_j \nabla \log F(v) \Big|_{v=(w^*)^T E_i R_j} \\ &= (M^\dagger)^{\frac{1}{2}} \sum_{j=1}^m R_j \nabla F(v) \Big|_{v=(w^*)^T E_i R_j}. \end{aligned}$$

In order to further evaluate this expression, define a function $g : \mathbb{R}^m \rightarrow \mathbb{R}$ as $g(z) = \sum_{j=1}^m F(z^T R_j)$. Then by definition we have $g(z) = 1$. Taking derivatives, we get $0 = \nabla g(z) = \sum_{j=1}^m R_j \nabla F(z^T R_j)$. It follows that $\mathbb{E}[\tilde{V}_i] = 0$, and hence that

$$\begin{aligned} \mathbb{E}[\nabla \ell(w^*)^T L^\dagger \nabla \ell(w^*)] &= \frac{1}{n^2} \mathbb{E}[\sum_{i=1}^n \sum_{\ell=1}^n \tilde{V}_i^T M^{\frac{1}{2}} E_i^T L^\dagger E_\ell M^{\frac{1}{2}} \tilde{V}_\ell] \\ &= \frac{1}{n^2} \mathbb{E}[\sum_{i=1}^n \tilde{V}_i^T M^{\frac{1}{2}} E_i^T L^\dagger E_i M^{\frac{1}{2}} \tilde{V}_i] \\ &\leq \frac{1}{n} \mathbb{E}[\sup_{\ell \in [n]} \|\tilde{V}_\ell\|_2^2] \text{tr}(\frac{1}{n} \sum_{i=1}^n M^{\frac{1}{2}} E_i^T L^\dagger E_i M^{\frac{1}{2}}). \end{aligned}$$

Since $L = \frac{m}{n} \sum_{i=1}^n E_i M E_i^T$, we have $\text{tr}(\frac{1}{n} \sum_{i=1}^n M^{\frac{1}{2}} E_i^T L^\dagger E_i M^{\frac{1}{2}}) = \frac{d-1}{m}$, as well as

$$\|\tilde{V}_\ell\|_2^2 = \sum_{j=1}^m \mathbf{1}[y_i = j] (\nabla \log F(v) \Big|_{v=(w^*)^T E_i R_j})^T R_j^T M R_j \nabla \log F(v) \Big|_{v=(w^*)^T E_i R_j}.$$

Recalling the previously defined matrix M , observe that since R_j is simply a permutation matrix, we have $R_j^T M R_j = M$ for every $j \in [m]$. By chain rule, we have $\langle \nabla \log F(v), 1 \rangle = \frac{1}{F(v)} \langle \nabla F(v), 1 \rangle = 0$, where the last step follows from our previous calculation. It follows that

$$\mathbb{E}[\langle \nabla \ell(w^*), L^\dagger \nabla \ell(w^*) \rangle] \leq \frac{d}{n} \sup_{v \in [-B, B]^m} \|\nabla \log F(v)\|_2^2.$$

Substituting this bound into equation (24) yields the claim.

D.1.1 LOWER BOUND UNDER THE SQUARED L SEMI-NORM

For any pair of quality score vectors w^j and w^k , the KL divergence between the distributions \mathbb{P}_{w^j} and \mathbb{P}_{w^k} is given by

$$D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) = \sum_{i=1}^n \sum_{l=1}^m F(w^{jT} E_i R_l) \log \frac{F(w^{jT} E_i R_l)}{F(w^{kT} E_i R_l)}.$$

Applying the inequality $\log x \leq x - 1$, valid for $x > 0$, we find that

$$D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) \leq \sum_{i=1}^n \sum_{l=1}^m F(w^{jT} E_i R_l) \left(\frac{F(w^{jT} E_i R_l)}{F(w^{kT} E_i R_l)} - 1 \right).$$

Now employing the fact that $\sum_{l=1}^m F(w^{jT} E_i R_l) = \sum_{l=1}^m F(w^{kT} E_i R_l) = 1$ gives

$$\begin{aligned} D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) &\leq \sum_{i=1}^n \sum_{l=1}^m \left(\frac{F(w^{jT} E_i R_l)^2}{F(w^{kT} E_i R_l)} - 2F(w^{jT} E_i R_l) + F(w^{kT} E_i R_l) \right). \\ &= \sum_{i=1}^n \sum_{l=1}^m \frac{(F(w^{jT} E_i R_l) - F(w^{kT} E_i R_l))^2}{F(w^{kT} E_i R_l)} \\ &\leq \frac{1}{F(-B, B, \dots, B)} \sum_{i=1}^n \sum_{l=1}^m (F(w^{jT} E_i R_l) - F(w^{kT} E_i R_l))^2 \\ &\leq \frac{1}{F(-B, B, \dots, B)} \sum_{i=1}^n \sum_{l=1}^m (\langle \nabla F(z_{il}), w^{jT} E_i R_l - w^{kT} E_i R_l \rangle)^2, \end{aligned}$$

for some $z_{il} \in [-B, B]^m$. Letting $\zeta = \frac{\sup_{z \in [-B, B]^m} \|\nabla F(z)\|_{H^\dagger}^2}{F(-B, B, \dots, B)}$ and applying Lemma 16 (noting that $\langle w^{jT} E_i R_l, \text{nullspace}(H) \rangle = 0$ for all i, j, l) gives

$$\begin{aligned} D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) &\leq \sum_{i=1}^n \sum_{l=1}^m \zeta \|w^{jT} E_i R_l - w^{kT} E_i R_l\|_H^2 \\ &\leq \zeta (w^j - w^k)^T \left(\sum_{i=1}^n \sum_{l=1}^m E_i R_l H R_l^T E_i^T \right) (w^j - w^k) \\ &\leq \zeta \lambda_m(H) n \|w^j - w^k\|_L^2, \end{aligned} \tag{26}$$

where the final step is a result of Lemma 11.

Consider the pair of scalars $\alpha \in (0, \frac{1}{4})$ and $\delta > 0$ whose values will be specified later. Let $M(\alpha)$ be as defined in (14). Consider the packing set $\{w^1, \dots, w^{M(\alpha)}\}$ constructed in Appendix A.1. Each of these vectors is of length d , satisfies $\langle w^j, 1 \rangle = 0$, and furthermore, each pair from this set satisfies $\alpha \delta^2 \leq \|w^j - w^k\|_L^2 \leq \delta^2$. Setting $\delta^2 = 0.01 \frac{d}{n \zeta \lambda_m(H)}$ yields

$$D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) \leq 0.01d.$$

Every element from the packing set also satisfies $\|w^j\|_\infty \leq B$ when $n \geq \frac{0.01 \sigma^2 \text{tr}(L^\dagger)}{\zeta B^2 \lambda_m(H)}$, and thus belongs to the class \mathcal{W}_B .

Applying Lemma 6 yields the lower bound

$$\|\widehat{w} - w^*\|_L^2 \geq \frac{\alpha}{2} 0.01 \frac{d}{n\zeta\lambda_m(H)} \left\{ 1 - \frac{0.01d + \log 2}{\log M(\alpha)} \right\}.$$

Setting $\alpha = 0.01$ proves the claim for $d > 9$.

For the case of $d \leq 9$, consider the set of the three d -length vectors $z^1 = [0 \ \cdots \ 0 \ -1]$, $z^2 = [0 \ \cdots \ 0 \ 1]$ and $z^3 = [0 \ \cdots \ 0 \ 0]$. Construct the packing set $\{w^1, w^2, w^3\}$ from these three vectors $\{z^1, z^2, z^3\}$ as done above for the case of $d > 9$. From the calculations made for the general case above, we have for all pairs $\min_{j \neq k} \|w^j - w^k\|_L^2 \geq \frac{\delta^2}{9}$ and $\max_{j,k} \|w^j - w^k\|_L^2 \leq 4\delta^2$, and as a result $\max_{j,k} D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) \leq 4n\zeta\lambda_m(H)\delta^2$. Choosing $\delta^2 = \frac{\log 2}{8n\zeta\lambda_m(H)}$ and applying Lemma 6 proves the claim.

D.1.2 UPPER BOUND UNDER THE SQUARED EUCLIDEAN NORM

The upper bound under the squared ℓ_2 -norm follows directly from the upper bound under the squared L semi-norm in Theorem 4: noting that $(w^* - \widehat{w}) \perp \text{nullspace}(L)$, we get that

$$(w^* - \widehat{w})^T L (w^* - \widehat{w}) \geq \lambda_2(L) \|w^* - \widehat{w}\|_2^2.$$

Substituting this inequality in the upper bound on the minimax risk under the squared L semi-norm in Theorem 4 gives the desired result.

D.1.3 LOWER BOUND UNDER THE SQUARED EUCLIDEAN NORM

Define $\zeta = \frac{\sup_{z \in [-B, B]^m} \|\nabla F(z)\|_{H^\dagger}^2}{F(-B, B, \dots, B)}$. Equation (26) in Appendix D.1.1 shows that for any vectors $w^j, w^k \in \mathcal{W}_B$,

$$D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) \leq \zeta \lambda_m(H) n \|w^j - w^k\|_L^2,$$

Consider the pair of scalars $\alpha \in (0, \frac{1}{4})$ and $\delta > 0$ whose values will be specified later. Let $M(\alpha)$ be as defined in (14). In Appendix B.2 we constructed a set $\{w^1, \dots, w^{M(\alpha)}\}$ of vectors of length d that satisfy $\langle w^j, 1 \rangle = 0$ for every $j \in [M(\alpha)]$, and for every pair of vectors in this set, $\|w^j - w^k\|_2^2 \geq \alpha\delta^2$ and $\frac{1}{\binom{M(\alpha)}{2}} \sum_{j \neq k} \|\tilde{w}^j - \tilde{w}^k\|_L^2 \leq \frac{2\delta^2}{d} \text{tr}(L)$. Applying Lemma 10 gives

$$\frac{1}{\binom{M(\alpha)}{2}} \sum_{j \neq k} \|\tilde{w}^j - \tilde{w}^k\|_L^2 \leq \frac{2\delta^2}{d} m(m-1).$$

Setting $\delta^2 = 0.005 \frac{d^2}{n\zeta\lambda_m(H)m(m-1)}$ yields

$$D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) \leq 0.01d.$$

In a manner similar to Lemma 14 in the pairwise comparison case, one can show that in the general setting of this section, $\text{tr}(L^\dagger) \geq \frac{d^2}{4m(m-1)}$. Then, every element from the packing set also satisfies $\|w^j\|_\infty \leq B$ when $\delta \leq B$, which holds true under our assumption

of $n \geq \frac{c\sigma^2 \text{tr}(L^\dagger)}{\zeta B^2 \lambda_m(H)} \geq \frac{c\sigma^2 d^2}{4m(m-1)\zeta B^2 \lambda_m(H)}$ with $c = 0.01$. Each element of our packing set thus belongs to the class \mathcal{W}_B . Applying Lemma 6 yields the lower bound

$$\|\widehat{w} - w^*\|_L^2 \geq \frac{\alpha}{2} 0.01 \frac{d^2}{n\zeta \lambda_m(H)m(m-1)} \left\{ 1 - \frac{0.01d + \log 2}{\log M(\alpha)} \right\}.$$

Setting $\alpha = 0.01$ proves the claim for $d > 9$.

For the case of $d \leq 9$, consider the set of the three d -length vectors $z^1 = [0 \ \cdots \ 0 \ -1]$, $z^2 = [0 \ \cdots \ 0 \ 1]$ and $z^3 = [0 \ \cdots \ 0 \ 0]$. Construct the packing set $\{w^1, w^2, w^3\}$ from these three vectors $\{z^1, z^2, z^3\}$ as done above for the case of $d > 9$. From the calculations made for the general case above, we have for all pairs $\min_{j \neq k} \|w^j - w^k\|_2^2 \geq \frac{\delta^2}{9}$ and $\max_{j,k} \|w^j - w^k\|_L^2 \leq 4\delta^2$, and as a result $\max_{j,k} D_{\text{KL}}(\mathbb{P}_{w^j} \|\mathbb{P}_{w^k}) \leq 4n\zeta \lambda_m(H)\delta^2$. Choosing $\delta^2 = \frac{\log 2}{8n\zeta \lambda_m(H)}$ and applying Lemma 6 proves the claim.

D.2 Some implied properties of the model

In this section, we prove the two auxiliary lemmas stated at the start of this appendix.

D.2.1 PROOF OF LEMMA 10

From the definition (10) of L , have

$$L1 = \frac{1}{n} \sum_{i=1}^n E_i(mI - 11^T)E_i^T 1 = \frac{1}{n} \sum_{i=1}^n E_i(mI - 11^T)1 = 0,$$

showing that $1 \in \text{nullspace}(L)$.

Now consider any non-zero vector $v := [v_1, \dots, v_d]^T \in \mathbb{R}^d$ such that $v \notin \text{span}(1)$. Then there must exist some $i, j \in [d]$ such that $v_i \neq v_j$. We know that there exists some path from item i to j in the comparison hyper-graph. Thus there must exist some hyper-edge in this path with two items, say i', j' , such that $v_{i'} \neq v_{j'}$. Suppose that hyper-edge corresponds to sample $\ell \in [n]$. Let $v' := E_\ell^T v$. Then $v' \notin \text{span}(1)$. The Cauchy-Schwarz inequality $\langle v', v' \rangle \langle 1, 1 \rangle > (\langle v', 1 \rangle)^2$ thus implies

$$v^T E_\ell(mI - 11^T)E_\ell^T v > 0.$$

Furthermore, for any $v'' \in \mathbb{R}^m$, the Cauchy-Schwarz inequality $\langle v'', v'' \rangle \langle 1, 1 \rangle > (\langle v'', 1 \rangle)^2$ implies that for any $i \in [n]$, we have $v^T E_i(mI - 11^T)E_i^T v \geq 0$. Overall we conclude that have $v^T L v > 0$ for every $v \notin \text{span}(1)$, and hence, $\text{nullspace}(L) = 1$ and $\lambda_2(L) > 0$.

Finally, we have

$$\text{tr}(L) = \frac{1}{n} \sum_{i=1}^n \text{tr}(E_i(mI - 11^T)E_i^T) = \frac{1}{n} \sum_{i=1}^n \left(m \text{tr}(E_i E_i^T) - \text{tr}(E_i 11^T E_i^T) \right). \quad (27)$$

By the definition of the matrices $\{E_i\}_{i \in [n]}$, $\text{tr}(E_i E_i^T) = m$ and $\text{tr}(E_i 11^T E_i^T) = m$. Substituting these values in (27) gives the desired result $\text{tr}(L) = m(m-1)$. \blacksquare

D.2.2 PROOF OF LEMMA 11

Let h_1, \dots, h_m denote the m eigenvectors of H , with $h_1 = \frac{1}{\sqrt{m}}\mathbf{1}$. Then for any vector $v' \in \mathbb{R}^m$,

$$\begin{aligned} v'^T H v' &= \sum_{i=2}^m \lambda_i(H) \langle v', h_i \rangle^2 \geq \lambda_2(H) \sum_{i=2}^m \langle v', h_i \rangle^2 = \lambda_2(H) \left(\sum_{i=1}^m \langle v', h_i \rangle^2 - \frac{1}{m} \langle v', \mathbf{1} \rangle^2 \right) \\ &= \lambda_2(H) v'^T \left(I - \frac{1}{m} \mathbf{1} \mathbf{1}^T \right) v', \end{aligned}$$

where the final step employed the property $\sum_{i=1}^m h_i h_i^T = I$ of the eigenvectors h_1, \dots, h_m of H . A similar argument gives

$$\begin{aligned} v'^T H v' &= \sum_{i=2}^m \lambda_i(H) \langle v', h_i \rangle^2 \leq \lambda_{\max}(H) \sum_{i=2}^m \langle v', h_i \rangle^2 = \lambda_{\max}(H) \left(\sum_{i=1}^m \langle v', h_i \rangle^2 - \frac{1}{m} \langle v', \mathbf{1} \rangle^2 \right) \\ &= \lambda_{\max}(H) v'^T \left(I - \frac{1}{m} \mathbf{1} \mathbf{1}^T \right) v'. \end{aligned}$$

Setting $v' = R_j^T v$ gives

$$\lambda_2(H) v^T R_j \left(I - \frac{1}{m} \mathbf{1} \mathbf{1}^T \right) R_j^T v \leq v^T R_j H R_j^T v \leq \lambda_{\max}(H) v^T R_j \left(I - \frac{1}{m} \mathbf{1} \mathbf{1}^T \right) R_j^T v.$$

Observe that the matrix $I - \frac{1}{m} \mathbf{1} \mathbf{1}^T$ is invariant to permutation of the coordinates, and hence $R_j \left(I - \frac{1}{m} \mathbf{1} \mathbf{1}^T \right) R_j^T = I - \frac{1}{m} \mathbf{1} \mathbf{1}^T$. This gives

$$\frac{\lambda_2(H)}{m} v^T (mI - \mathbf{1} \mathbf{1}^T) v \leq v^T R_j H R_j^T v \leq \frac{\lambda_{\max}(H)}{m} v^T (mI - \mathbf{1} \mathbf{1}^T) v. \quad \blacksquare$$

Appendix E. Some useful tail bounds

In this appendix, we collect a few useful tail bounds for quadratic forms in Gaussian and sub-Gaussian random variables.

Lemma 12 (Tail bound for Gaussian quadratic form) *For any positive semidefinite matrix Q and standard Gaussian vector $g \sim N(0, I_d)$, we have*

$$\mathbb{P}[g^T Q g \geq (\sqrt{\text{tr}(Q)} + \sqrt{\|Q\|_{\text{op}}}) \delta]^2 \leq e^{-\delta/2}. \quad (28)$$

valid for all $\delta \geq 0$.

Proof Note that the function $g \mapsto \|\sqrt{Q}g\|_2$ is Lipschitz with constant $\|\sqrt{Q}\|_{\text{op}}$. Consequently, by concentration for Lipschitz functions of Gaussian vectors (Ledoux, 2001), the random variable $Z = \|\sqrt{Q}g\|_2$ satisfies the upper bound

$$\mathbb{P}[Z \geq \mathbb{E}[Z] + t] \leq \exp\left(-\frac{t^2}{2\|\sqrt{Q}\|_{\text{op}}^2}\right) = \exp\left(-\frac{t^2}{2\|Q\|_{\text{op}}}\right).$$

By Jensen's inequality, we have $\mathbb{E}[Z] = \mathbb{E}[\|\sqrt{Q}g\|_2] \leq \sqrt{\mathbb{E}[g^T Q g]} = \sqrt{\text{tr}(Q)}$. Setting $t = \sqrt{\|Q\|_{\text{op}}} \delta$ completes the proof. \blacksquare

Lemma 13 ((Hanson and Wright, 1971; Rudelson and Vershynin, 2013)) *Let $V \in \mathbb{R}^d$ be a random vector with independent zero-mean components that are sub-Gaussian with parameter K , and let $M \in \mathbb{R}^{d \times d}$ be an arbitrary matrix. Then there is a universal constant $c > 0$ such that*

$$\mathbb{P} [|V^T M V - \mathbb{E}[V^T M V]| > t] \leq 2 \exp \left(-c \min \left\{ \frac{t^2}{K^4 \|M\|_{\text{fro}}^2}, \frac{t}{K^2 \|M\|_{\text{op}}} \right\} \right) \quad \text{for all } t > 0. \quad (29)$$

Appendix F. Properties of Laplacian matrices

By construction, the Laplacian L of the comparison graph is symmetric and positive-semidefinite. By the singular value decomposition, we can write $L = U^T \Lambda U$ where $U \in \mathbb{R}^{d \times d}$ is an orthonormal matrix, and Λ is a diagonal matrix of nonnegative eigenvalues with $\Lambda_{jj} = \lambda_j(L)$ for every $j \in [d]$. Given our assumption of $\lambda_1(L) \leq \dots \leq \lambda_d(L)$, we also have $\Lambda_{11} \leq \dots \leq \Lambda_{dd}$. Also recall that L^\dagger denotes the Moore-Penrose pseudo-inverse of L . In terms of the notation introduced, the Moore-Penrose pseudo-inverse is then given by $L^\dagger = U^T \Lambda^\dagger U$, where Λ^\dagger is a diagonal matrix with entries

$$\Lambda_{jj}^\dagger = \begin{cases} (\Lambda_{jj}^{-1}) & \text{if } \Lambda_{jj} > 0 \\ 0 & \text{otherwise.} \end{cases}$$

The following pair of lemmas establish some useful properties about L .

Lemma 14 *The Laplacian matrix (4) satisfies the trace constraints*

$$\text{tr}(L) = 2, \quad \text{and} \quad \text{tr}(L^\dagger) \geq \frac{d^2}{4}.$$

Proof From the definition (4) of the matrix L , we have $\text{tr}(L) = \frac{1}{n} \sum_{i=1}^n \text{tr}(x_i x_i^T) = 2$. We also know that $\lambda_1(L) = 0$, and hence $\sum_{j=2}^d \lambda_j(L) = 2$. Given the latter constraint, the sum $\sum_{j=2}^d \frac{1}{\lambda_j(L)}$ is minimized when $\lambda_2(L) = \dots = \lambda_d(L)$. Some simple algebra now gives the claimed result. \blacksquare

Lemma 15 *For the matrix L defined in (4), and for a $(n \times d)$ matrix X with x_i^T as its i^{th} row,*

$$\text{tr}\left(\frac{1}{n} x^T L^\dagger x\right) = d - 1, \quad \left\| \frac{1}{n} x^T L^\dagger x \right\|_{\text{fro}} = d - 1, \quad \text{and} \quad \left\| \frac{1}{n} x^T L^\dagger x \right\|_{\text{op}} = 1.$$

Proof Let $Q = \frac{1}{n}x^T L^\dagger x$. Since $L = \frac{1}{n}X^T X = U^T \Lambda U$, the diagonal entries of Λ are the squared singular values of X/\sqrt{n} . Consequently, there must exist an orthonormal matrix V such that $X/\sqrt{n} = V\sqrt{\Lambda}U^T$, and thus we can write $Q = V\sqrt{\Lambda}\Lambda^\dagger\sqrt{\Lambda}V^T$. By definition of the Moore-Penrose pseudo-inverse, the matrix $\sqrt{\Lambda}\Lambda^\dagger\sqrt{\Lambda}$ is a diagonal matrix; since the Laplacian graph is connected, its diagonal contains $(d-1)$ ones and a single zero. Noting that V is an orthonormal matrix gives the claimed result. \blacksquare

For future reference, we state and prove a lemma showing that these two semi-norms satisfy a restricted form of the Cauchy-Schwarz inequality:

Lemma 16 *For any two vectors u and v such that $u \perp \text{nullspace}(L)$ or/and $v \perp \text{nullspace}(L)$, we have*

$$|\langle u, v \rangle| \leq \|u\|_{L^\dagger} \|v\|_L. \quad (30)$$

Proof Since $L = U^T \Lambda U$ and $L^\dagger = U^T \Lambda^\dagger U$, we have

$$\sqrt{v^T L v} \sqrt{u^T L^\dagger u} = \sqrt{v^T U^T \Lambda U v} \sqrt{u^T U^T \Lambda^\dagger U u} = \|\tilde{v}\|_2 \|\tilde{u}\|_2 \geq |\langle \tilde{v}, \tilde{u} \rangle|,$$

where we have defined $\tilde{v} := \sqrt{\Lambda} U v$ and $\tilde{u} := \sqrt{\Lambda^\dagger} U u$. Continuing on,

$$\langle \tilde{v}, \tilde{u} \rangle = v^T U^T \sqrt{\Lambda} \sqrt{\Lambda^\dagger} U u = v^T U U^T u,$$

where we have used the fact that u or/and v are orthogonal to the null space of L . Since U is orthonormal, we conclude that $\langle \tilde{v}, \tilde{u} \rangle = \langle v, u \rangle$, which completes the proof. \blacksquare

Appendix G. Minimax risk without assumptions on quality scores

The setting considered throughout the paper imposes two restrictions (2) on the quality score vector w^* . The first condition is that of shift invariance, that is, $\langle w^*, 1 \rangle = 0$. The necessity of this condition for identifiability under the ORDINAL model is easy to verify. The second condition is that the quality score vectors are B -bounded, that is, $\|w^*\|_\infty \leq B$ for some finite B . In this section, for the sake of completeness, we show that the minimax risk is infinite in the absence of this condition.

Proposition 17 *Any estimator \tilde{w} based on n samples from the ORDINAL model (with unbounded quality score vectors) has error lower bounded as*

$$\sup_{w^* \in \mathcal{W}_\infty} \mathbb{E} \left[\|\tilde{w} - w^*\|_2^2 \right] = \sup_{w^* \in \mathcal{W}_\infty} \mathbb{E} \left[\|\tilde{w} - w^*\|_L^2 \right] = \infty.$$

The remainder of this section is devoted to the formal proof of Proposition 17. Consider the event where for every comparison, the item with the higher quality score in w^* wins. For any $w^* \in \mathcal{W}_\infty \setminus \{0\}$, this event occurs with a probability at least $\frac{1}{2n}$. Under this event, the true w^* is indistinguishable from the quality score vector $cw^* \in \mathcal{W}_\infty$ for every $c \geq 0$, and the error is also unbounded. Since the probability of this event is strictly bounded away from zero, the expected error is also unbounded.

References

- Ammar Ammar and Devavrat Shah. Ranking: Compare, don't score. In *Allerton Conference on Communication, Control, and Computing*, pages 776–783, 2011.
- Donald R Atkinson, Bruce E Wampold, Susana M Lowe, Linda Matthews, and Hyun-Nie Ahn. Asian American preferences for counselor characteristics: Application of the Bradley-Terry-Luce model to paired comparison data. *The Counseling Psychologist*, 26(1):101–123, 1998.
- Hossein Azari Soufiani, David C Parkes, and Lirong Xia. Preference elicitation for general random utility models. In *Uncertainty in Artificial Intelligence: Proceedings of the 29th Conference*. AUAI Press, 2013.
- William Barnett. The modern theory of consumer behavior: Ordinal or cardinal? *The Quarterly Journal of Austrian Economics*, 6(1):41–65, 2003.
- Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, pages 324–345, 1952.
- Tom Bramley. A rank-ordering method for equating tests by expert judgment. *Journal of Applied Measurement*, 6(2):202–223, 2005.
- Mark Braverman and Elchanan Mossel. Noisy sorting without resampling. In *Symposium on Discrete Algorithms*, pages 268–276, 2008.
- Andries E Brouwer and Willem H Haemers. *Spectra of graphs*. Springer, 2011.
- Sourav Chatterjee. Matrix estimation by universal singular value thresholding. *The Annals of Statistics*, 43(1):177–214, 2014.
- Xi Chen, Paul N Bennett, Kevyn Collins-Thompson, and Eric Horvitz. Pairwise ranking aggregation in a crowdsourced setting. In *International Conference on Web Search and Data Mining*, pages 193–202, 2013.
- Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.
- Ofer Gabber and Zvi Galil. Explicit constructions of linear-sized superconcentrators. *Journal of Computer and System Sciences*, 22(3):407–420, 1981.
- Edgar N Gilbert. A comparison of signalling alphabets. *Bell System Technical Journal*, 31(3):504–522, 1952.
- Paul E Green, J Douglas Carroll, and Wayne S DeSarbo. Estimating choice probabilities in multiattribute decision making. *Journal of Consumer Research*, pages 76–84, 1981.

- Bruce Hajek, Sewoong Oh, and Jiaming Xu. Minimax-optimal inference from partial rankings. In *Advances in Neural Information Processing Systems*, pages 1475–1483, 2014.
- David Lee Hanson and Farroll Tim Wright. A bound on tail probabilities for quadratic forms in independent random variables. *The Annals of Mathematical Statistics*, pages 1079–1083, 1971.
- Sandra Heldsinger and Stephen Humphry. Using the method of pairwise comparison to obtain reliable teacher assessments. *The Australian Educational Researcher*, 37(2):1–19, 2010.
- Ralf Herbrich, Tom Minka, and Thore Graepel. Trueskill: A Bayesian skill rating system. In *Advances in Neural Information Processing Systems*, volume 19, page 569, 2007.
- Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *Signal Processing Magazine, IEEE*, 29(6):82–97, 2012.
- Srikanth Jagabathula and Devavrat Shah. Inferring rankings under constrained sensing. In *Advances in Neural Information Processing Systems*, pages 753–760, 2008.
- Kevin G Jamieson and Robert Nowak. Active ranking using pairwise comparisons. In *Advances in Neural Information Processing Systems*, pages 2240–2248, 2011.
- Gabriella Kazai. In search of quality in crowdsourcing for search engine evaluation. In *Advances in Information Retrieval*, pages 165–176. Springer, 2011.
- Zahid Y Khairullah and Stanley Zionts. An approach for preference ranking of alternatives. *European journal of operational research*, 28(3):329–342, 1987.
- Firas Khatib, Frank DiMaio, Seth Cooper, Maciej Kazmierczyk, Mirosław Gilski, Szymon Krzywda, Helena Zabranska, Iva Pichova, James Thompson, Zoran Popović, Mariusz Jaskolski, and David Baker. Crystal structure of a monomeric retroviral protease solved by protein folding game players. *Nature structural & molecular biology*, 18(10):1175–1177, 2011.
- John I Kiger. The depth/breadth trade-off in the design of menu-driven user interfaces. *International Journal of Man-Machine Studies*, 20(2):201–213, 1984.
- Kenneth J Koehler and Harold Ridpath. An application of a biased version of the Bradley-Terry-Luce model to professional basketball results. *Journal of Mathematical Psychology*, 25(3), 1982.
- Paul FM Krabbe. Thurstone scaling as a measurement method to quantify subjective health outcomes. *Medical care*, 46(4):357–365, 2008.
- ASID Lang and Joshua Rio-Ross. Using Amazon Mechanical Turk to transcribe historical handwritten documents. *The Code4Lib Journal*, 2011.

- M. Ledoux. *The Concentration of Measure Phenomenon*. Mathematical Surveys and Monographs. American Mathematical Society, Providence, RI, 2001.
- Michael D Lee, Mark Steyvers, Mindy De Young, and Brent J Miller. A model-based approach to measuring expertise in ranking tasks. In *Proceedings of the 33rd annual conference of the cognitive science society*, 2011.
- E.L. Lehmann and G. Casella. *Theory of Point Estimation*. Springer Texts in Statistics, 1998.
- Peter John Loewen, Daniel Rubenson, and Arthur Spirling. Testing the power of arguments in referendums: A Bradley–Terry approach. *Electoral Studies*, 31(1):212–221, 2012.
- R Duncan Luce. *Individual Choice Behavior: A Theoretical Analysis*. New York: Wiley, 1959.
- Miguel Angel Luengo-Oroz, Asier Arranz, and John Freen. Crowdsourcing malaria parasite quantification: an online game for analyzing images of infected thick blood smears. *Journal of medical Internet research*, 14(6), 2012.
- George A Miller. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological review*, 63(2):81, 1956.
- Sahand Negahban, Sewoong Oh, and Devavrat Shah. Iterative ranking from pair-wise comparisons. In *Advances in Neural Information Processing Systems*, pages 2474–2482, 2012.
- Sahand Negahban, Sewoong Oh, and Devavrat Shah. Rank centrality: Ranking from pair-wise comparisons. *arXiv preprint arXiv:1209.1688v4*, 2015.
- Robert M Nosofsky. Luce’s choice model and Thurstone’s categorical judgment model compared: Kornbrot’s data revisited. *Attention, Perception, & Psychophysics*, 37(1): 89–91, 1985.
- Roberto Imbuzeiro Oliveira. Concentration of the adjacency matrix and of the laplacian in random graphs with independent edges. *arXiv preprint arXiv:0911.0600*, 2009.
- Chris Piech, Jonathan Huang, Zhenghao Chen, Chuong Do, Andrew Ng, and Daphne Koller. Tuned models of peer assessment in MOOCs. In *International Conference on Educational Data Mining*, 2013.
- Robin L Plackett. The analysis of permutations. *Applied Statistics*, pages 193–202, 1975.
- 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*, pages 118–126, 2014.
- Vikas C Raykar, Shipeng Yu, Linda H Zhao, Gerardo Hermsillo Valadez, Charles Florin, Luca Bogoni, and Linda Moy. Learning from crowds. *The Journal of Machine Learning Research*, 99:1297–1322, 2010.

- Daniel Ross. Arpad Elo and the Elo rating system, 2007. <http://en.chessbase.com/post/arpad-elo-and-the-elo-rating-system>.
- Mark Rudelson and Roman Vershynin. Hanson-wright inequality and sub-gaussian concentration. *Electronic Communications in Probability*, 18:1–9, 2013.
- Thomas L Saaty and Mujgan S Ozdemir. Why the magic number seven plus or minus two. *Mathematical and Computer Modelling*, 38(3):233–244, 2003.
- Nihar B Shah, Joseph K Bradley, Abhay Parekh, Martin Wainwright, and Kannan Ramchandran. A case for ordinal peer-evaluation in MOOCs. In *NIPS Workshop on Data Driven Education*, December 2013.
- Nihar B Shah, Sivaraman Balakrishnan, Joseph Bradley, Abhay Parekh, Kannan Ramchandran, and Martin J. Wainwright. When is it better to compare than to score? *arXiv preprint arXiv:1406.6618*, 2014.
- Richard M Shiffrin and Robert M Nosofsky. Seven plus or minus two: a commentary on capacity limitations. *Psychological Review*, 1994.
- Neil Stewart, Gordon DA Brown, and Nick Chater. Absolute identification by relative judgment. *Psychological review*, 112(4):881, 2005.
- John Swets. The relative operating characteristic in psychology. *Science*, 182(4116), 1973.
- Louis L Thurstone. A law of comparative judgment. *Psychological Review*, 34(4):273, 1927.
- Kenneth E Train. *Discrete choice methods with simulation*. Cambridge university press, 2009.
- Kristi Tsukida and Maya R Gupta. How to analyze paired comparison data. Technical report, DTIC Document, 2011.
- A.B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics, 2008.
- RR Varshamov. Estimate of the number of signals in error correcting codes. In *Dokl. Akad. Nauk SSSR*, volume 117, pages 739–741, 1957.
- Luis von Ahn, Benjamin Maurer, Colin McMillen, David Abraham, and Manuel Blum. Recaptcha: Human-based character recognition via web security measures. *Science*, 321(5895):1465–1468, 2008.
- Jing Wang, Panagiotis G Ipeirotis, and Foster Provost. Managing crowdsourcing workers. In *Winter Conference on Business Intelligence*, pages 10–12, 2011.
- Jinfeng Yi, Rong Jin, Shaili Jain, and Anil Jain. Inferring users’ preferences from crowdsourced pairwise comparisons: A matrix completion approach. In *AAAI Conference on Human Computation and Crowdsourcing*, 2013.