Testing and Estimation of Contextual Generalized Thurstone Models

Dongmin Lee
Purdue University
lee4818@purdue.edu

Anuran Makur
Purdue University
amakur@purdue.edu

Japneet Singh
Purdue University
sing1041@purdue.edu

Boyu Xu
Purdue University
xu2202@purdue.edu

Abstract

Many modern machine learning and artificial intelligence methods utilize preference learning models like Bradley-Terry-Luce (BTL) to learn latent scores of items based on pairwise comparison data. Despite the successes of such models in specific applications, such as reinforcement learning from human feedback (RLHF), the theoretical justification for the employed models has often been unclear. In this work, we provide a hypothesis testing procedure to test whether a given dataset of pairwise comparisons follows a contextual generalized Thurstone (CGT) model under certain classes of latent functions. The latent functions determine the latent scores, and can be expressed as weighted sums of basis functions. Our CGT model encompasses a wide class of preference learning models, including the popular BTL model. In particular, we prove bounds on critical thresholds for our tests in the minimax sense, which scale like $1/\sqrt{|\mathcal{E}|^{1/2}k}$ (where $|\mathcal{E}|$ is the number of edges in the comparison graph and k is the number of observed comparisons per pair) under certain important regimes, such as complete graphs and perfect matchings. Furthermore, to establish our testing results, we also derive error bounds for CGT parameter estimation that generalize and improve on known results in the BTL and non-contextual settings. Finally, we conduct experiments on various synthetic and real-world datasets to verify our theoretical results.

1 Introduction

The problem of preference learning from pairwise comparisons has been the subject of significant attention in the machine learning literature, classically in the contexts of econometrics [1], psychometrics [2], and statistics of rankings [3], [4], and most recently in the context of reinforcement learning from human feedback (RLHF) [5]–[7], where human comparisons are used to train a reward function that models human preferences. The prototypical framework used to estimate preference scores from pairwise comparisons is the *Bradley-Terry-Luce* (BTL) or multinomial logit model [8]–[11]. In this framework, we are given n items indexed by $[n] = \{1, \dots, n\}$, and assume that each item $i \in [n]$ is associated with some latent score $w_i \in \mathbb{R}$ such that item j is preferred over item i in a comparison with probability $\exp(w_j)/(\exp(w_i) + \exp(w_j))$.

Although a powerful tool, the vanilla BTL model by itself has limited utility in many machine learning settings, since estimating the scores w_i requires the comparison graph (with items as vertices and edges representing observed comparisons) to be connected [12], which may not be a realistic assumption. For example, in the aforementioned RLHF context for large language model (LLM)

The author order is alphabetical.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

alignment, where each item corresponds to an LLM's response to some prompt, most LLM responses will appear only once in the dataset, making the comparison graph extremely disconnected. It is therefore an interesting technical problem to estimate scores based on general (possibly disconnected) comparison graphs. We can achieve this by considering a *contextual BTL model*, which assumes that there exists some *latent function* (or reward function) that maps each item to its latent score, and focusing on learning this underlying function from comparisons. The function captures the relationship between different scores even in the absence of comparisons between relevant items. It turns out such models were developed in the RLHF literature and have been used to great success (see, e.g., [6], [7], [13], [14]).

Recently, works such as [15]–[17] investigated the error bounds for estimation under such models, providing theoretical guarantees that support their empirical success. However, such error bounds are only applicable when the dataset is truly generated under a valid contextual BTL model, which is not necessarily the case. Although the BTL model arises from natural assumptions such as transitivity and independence of irrelevant alternatives (IIA) [18], evidence suggests that real-life datasets often do not satisfy such assumptions, natural as they may seem (see, e.g., [19]–[21]). Indeed, there exist many competing models to the (contextual) BTL model, such as the Thurstone model [22], [23], which assumes that item j is preferred to item i with probability $\Phi(w_j - w_i)$, where Φ is the cumulative distribution function (CDF) of the standard Gaussian distribution $\mathcal{N}(0, 1)$. These models form a broader class of *(contextual) generalized Thurstone ((C)GT) models* [2], [24].

In practice, in the absence of ground truth parameters, it is often unclear which of the many CGT models is the most appropriate model for a given dataset. For example, it is unclear whether the contextual BTL models utilized in RLHF are the "correct" models (see, e.g., [25]–[27]). To answer this question, a hypothesis testing procedure to determine whether a dataset has been generated under a certain CGT model is desired. Recent developments by [28] and [29] establish test statistics that test for the (non-contextual) BTL and GT models, but to the best of our knowledge, no analogue in the more technically challenging (and perhaps more practically relevant) contextual setting exists in the literature.

In this work, **we develop a hypothesis testing framework for CGT models** to close this gap. In the process, **we also develop estimation results for CGT models**, generalizing and improving upon existing results on the BTL or non-contextual settings. Notably, our main results do not require connectivity of the comparison graph. Although our results hold for general comparison graphs, we give special attention to two specific graph classes, *complete graphs* and *perfect matchings*, which have broad theoretical and practical significance. First, the class of complete graphs, in which the number of edges $|\mathcal{E}|$ scales like n^2, is the traditional setting for the analysis of (non-contextual) GT models (see, e.g., [28]–[32]). This class of graphs also arises often in real-world settings, such as LLM rankings [33] and sports leagues (see Appendix H for an example). On the other hand, the class of perfect matchings (i.e., 1-regular graphs), in which $|\mathcal{E}|$ scales like n, is important because it appears in many machine learning applications such as RLHF, yet cannot be analyzed under the classic framework of non-contextual ranking models due to disconnectivity. Although a rich universe of graphs exists between these two extremes, we focus on these two regimes as a starting point as obtaining tight bounds on general graphs is more difficult (see, e.g., [34]). In this work, we provide explicit bounds on the estimation error and the critical threshold for the testing of CGT models under complete graphs and perfect matchings to characterize the behavior of CGT models under the two extreme ends on the spectrum of graph connectivity.

1.1 Contributions. We begin by formulating the CGT model in Section 2. Then, in Section 3, we analyze the CGT model under the class of linear latent functions (the LCGT model). First, we bound the estimation error in Theorem 1, generalizing results by [15], [16] (which only considers the contextual BTL model) and [29] (which only considers the non-contextual GT model). We find that the H-norm of the estimation error (as defined in Section 3.1) scales like $\sqrt{p/k}$, where p is the number of parameters and k is the number of observed samples per comparison. This recovers the scaling of [16, Theorem 1.1] in the BTL regime and [29, Theorem 3.3] in the non-contextual regime (where $p = n - 1$). Our bounds hold even under model mismatch, i.e., when the dataset is not distributed according to the estimated CGT model.

Next, we introduce the corresponding composite hypothesis testing problem parameterized by the separation parameter ϵ and define the test statistic in (6), which captures how closely the observed comparison data conforms to a valid CGT model with a linear latent function. In Theorem 2, we

derive an upper bound on the critical critical threshold ϵ_c that allows for testing with bounded minimax risk $\mathcal{R} \leq 1/2$. Our bound, which scales like $1/(|\mathcal{E}|^{1/4}k^{1/2}) + p^{1/2}/(|\mathcal{E}|^{1/2}k^{1/2})$, matches the scaling of [29, Theorem 3.2] for dense graphs with $|\mathcal{E}| = \Omega(n^2)$ under the non-contextual model, and improves upon it for sparser datasets. We conclude the section with Proposition 3, which provides a lower bound for the minimax risk under certain graph regimes, establishing matching upper and lower bounds for the critical threshold under complete graphs.

In Section 4, we study the CGT model under the more general class of latent functions that can be represented as a weighted sum of a fixed set of basis functions (the BCGT model). Whereas the bound in Theorem 1 possesses an implicit dependence on the spectrum of H, the corresponding result in the BCGT setting, Theorem 4, makes this dependence explicit by analyzing the estimation error under the ℓ^2-norm. This requires finer control on the estimation error, which is achieved with a covering argument. As in the LCGT setting, our bound is robust under model mismatch. Next, the spectrum of H, in particular its smallest and largest eigenvalues, are investigated in Lemmas 11 and 12 and Corollary 13, where it is shown that the eigenvalues of H scale like $|\mathcal{E}|$ (the number of unique pairs that are compared) under common settings such as complete graphs and perfect matchings when the distribution of context vectors satisfies certain regularity conditions.

Using the bounds on estimation error obtained in Theorem 4, we derive an upper bound on the critical threshold for BCGT model hypothesis testing in Theorem 5, which matches the scaling of Theorem 2 under the eigenvalue regimes of Lemma 12 and Corollary 13. As in the LCGT setting, our analysis of BCGT concludes with Proposition 6, which establishes tightness of our upper bound on the critical threshold in Theorem 5 for complete graphs.

Finally, we present various experiments on synthetic and real-world datasets in Section 5. In particular, we demonstrate that the estimation error and testing threshold scales as our theoretical upper bounds predict in practice. Furthermore, we show that CGT models can provide meaningful insights even for very noisy and disconnected datasets, where traditional GT models like BTL usually struggle.

1.2 Related literature. Here, we briefly summarize the vast literature on preference learning with (contextual) GT and BTL models, especially in the context of machine learning. The problem of modeling the outcome of noisy comparisons was first studied by foundational works such as [22], [23] (which introduced the Thurstone model) and [8]–[11] (which introduced the BTL model). In particular, the theoretical performance of parameter estimation under the non-contextual BTL model has been the subject of intense scrutiny by works such as [30], [31], [34]–[40], which often bound the ℓ^2 or ℓ^∞ error of parameter estimation algorithms such as MLE or spectral methods under various classes of connected comparison graphs including complete graphs and Erdős-Rényi random graphs. Recently, a new line of research has emerged to formally analyze the contextual BTL model, finding bounds for the estimation error under various classes of comparison graphs, such as random multigraphs with edges uniformly sampled with replacement [15], perfect matchings [16], and Erdős-Rényi random graphs [17]. Meanwhile, works that investigate hypothesis testing methods on comparison datasets include [41], [42], which considers two-sample hypothesis testing on pairwise comparison datasets (i.e., determining whether the two datasets share the same distribution), [43], which analyzes testing for the independence of irrelevant alternatives (IIA) condition, and [28], [29], which studies testing for non-contextual BTL and GT models. Finally, there exists a wealth of results in the literature that document the empirical performance of contextual BTL estimation for learning human preferences, including, e.g., [6], [7], [13], [14], [44], [45].

1.3 Notation. We represent integer ranges with $[n] \triangleq \{1, \dots, n\}$. We use $\|\cdot\|_p$ (where $p \in [1, \infty]$) to refer to the usual ℓ^p norm for vectors and to the operator norm induced by the ℓ^p vector norm for matrices. Given a positive (semi)definite matrix $H \in \mathbb{R}^{n\times n}$, we define the H-(semi)norm $\|\boldsymbol{x}\|_H \triangleq \sqrt{\boldsymbol{x}^\mathsf{T} H \boldsymbol{x}}$ for all vectors $\boldsymbol{x} \in \mathbb{R}^p$. Given a symmetric matrix $H \in \mathbb{R}^{p\times p}$, we enumerate its eigenvalues in ascending order as $\lambda_1(H) \leq \cdots \leq \lambda_p(H)$ and define $\lambda_{\min}(H) \triangleq \lambda_1(H)$, $\lambda_{\max}(H) \triangleq \lambda_p(H)$. We use $\mathbf{0}_n, \mathbf{1}_n \in \mathbb{R}^n$ to refer to the all-zeros and all-ones vector, respectively. We use $\boldsymbol{e}_i \in \mathbb{R}^n$ to refer to the i-th standard basis vector of $\mathbb{R}^n$ for all $i \in [n]$. We use standard Bachmann-Landau asymptotic notation such as $O(\cdot)$ and $\Omega(\cdot)$.

2 Problem formulation

We begin with a formal exposition of the pairwise comparison model that we study in this work, the *contextual generalized Thurstone (CGT) model* (or contextual random utility model). Suppose we

have n items indexed by $i \in [n]$. Each item i is associated with a known *context vector* $\boldsymbol{x}_i \in \mathcal{X}$ (where $\mathcal{X}$ is some compact and convex set), which characterizes the properties of item i. We assume that some (unknown) latent function $f\colon [0,1]^p \to \mathbb{R}$ maps each context vector $\boldsymbol{x}_i$ to the latent utility $w_i = f(\boldsymbol{x}_i)$, which measures how strongly item i is preferred over other items. More formally, p_{ij}, the probability of item j being preferred over item i in a pairwise comparison, is given as

$$p_{ij} = F(w_j - w_i) = F(f(\boldsymbol{x}_j) - f(\boldsymbol{x}_i)),$$

where $F\colon \mathbb{R} \to [0,1]$ is the CDF of some real-valued symmetric random variable X (i.e., $F(-x) = 1 - F(x)$). Some common choices of X include $X \sim \text{Logistic}(0,1)$ for the BTL model (where $F(t) = 1/(1+e^{-t})$), $X \sim \mathcal{N}(0,1)$ for the Thurstone model, and $X \sim \text{Laplace}(0,1)$ for the Dawkins model [46]. We note that this is a generalization of the model proposed by [16], [17], which only considers the BTL model.

Next, we define the common setting of the *parameter estimation* and *hypothesis testing* problems on the CGT model. In both problems, we are presented with a dataset of $|\mathcal{E}|k$ comparison results

$$\mathcal{Z} \triangleq \{Z_{ij}^m \sim \text{Bernoulli}(p_{ij}) : (i,j) \in \mathcal{E}, m \in [k]\},$$

where $\mathcal{E}$ is the edge set of the (undirected) comparison graph $\mathcal{G} = ([n], \mathcal{E})$, $k \in \mathbb{Z}_{>0}$ is the number of comparisons between each pair, and each Z_{ij}^m are independent random variables that encode the outcome of the m-th comparison between items i and j. Note that we use the convention that $j > i$ for all $(i,j) \in \mathcal{E}$. Let $\hat{p}_{ij} \triangleq Z_{ij}/k$ be the empirical probability derived from the observed data, and let the true probability matrix $P \in [0,1]^{n\times n}$ and empirical probability matrix $\hat{P} \in [0,1]^{n\times n}$ be defined as $P_{ij} = p_{ij}$ and $\hat{P}_{ij} = \hat{p}_{ij}$ for all $(i,j) \in \mathcal{E}$ and $P_{ij} = \hat{P}_{ij} = 0$ for all $(i,j) \in [n]^2 \setminus \mathcal{E}$.

For the parameter estimation problem, the goal is to find the closest approximation to f given the empirical probabilities $\hat{p}_{ij}$. For the hypothesis testing problem, the goal is to determine whether the data was generated from a CGT model with a latent function f that belongs to a certain class of functions (e.g., linear functions, polynomials with bounded degree, etc.).

2.1 Common assumptions. Next, we lay out the basic assumptions adopted throughout this work.

Assumption 1 (Minimum degree)**.** *Every vertex in $\mathcal{G}$ has degree at least* 1.

Assumption 2 (Dynamic range)**.** *There exists some constant $a \in (0,1)$ such that $p_{ij} \in [a, 1-a]$ for all $(i,j) \in \mathcal{E}$.*

Assumption 3 (Strong log-concavity of F)**.** *There exist some constants $b, \alpha, \beta > 0$ with $F(2b) \in [a, 1-a]$ such that F is α-strongly log-concave and β-Lipschitz on $[-2b, 2b]$.*

Assumption 1 is natural because we get no information from items that are never compared to any other item, thus we can remove such items from consideration. Assumption 2 is standard in the literature when analyzing comparison models, and Assumption 3 holds for common models such as the BTL, Thurstone, and Dawkins models. We also define the quantity $\hat{\beta} \triangleq \beta/(F(2b)F(-2b))$ to simplify some expressions in our results.

3 Linear contextual model

3.1 Preliminaries. In this section, we consider the case where the latent function f is a linear function, i.e., there exists some vector $\boldsymbol{\theta} \in \mathbb{R}^p$ such that $w_i = \boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_i$ for all $i \in [n]$. We call this class of models the *linear contextual generalized Thurstone (LCGT) model*. Although this is a simpler setting than the more general BCGT model considered in Section 4, it is rich enough to provide intuition on the behavior of CGT models, and has practical significance in its own right. For example, as argued in [16], applications that use a neural network to model the latent function can be modeled under the LCGT framework by setting $\boldsymbol{x}_i$ to be the output of the penultimate layer of the neural network. Since one of the main goals of this section is to provide intuition in a simple setting, we present results such as Theorem 1 in a simpler form, while the "full-fat" versions of our results are provided in Section 4.

For convenience, we define $\boldsymbol{x}_{ij} \triangleq \boldsymbol{x}_j - \boldsymbol{x}_i$ for all $(i,j) \in [n]^2$ so that the latent preference probability p_{ij} may be represented as $p_{ij} = F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})$ for all $i \in [n]$, such that the relationship between the parameter vector $\boldsymbol{\theta}$ and the induced probability distribution $P(\boldsymbol{\theta})$ is clear. Let $X \in \mathbb{R}^{|\mathcal{E}|\times p}$ be the matrix obtained by stacking the row vectors $\boldsymbol{x}_{ij}^\mathsf{T}$ for all $(i,j) \in \mathcal{E}$ (any consistent edge ordering may

be used). The feature matrix $H \in \mathbb{R}^{p\times p}$, defined

$$H \triangleq \sum_{(i,j)\in\mathcal{E}} \boldsymbol{x}_{ij}\boldsymbol{x}_{ij}^{\mathsf{T}} = X^{\mathsf{T}}X, \tag{1}$$

plays a critical role in the behavior of LCGT models. In particular, we represent the bounds in Theorem 1 in terms of the $\|\cdot\|_H$ seminorm (which is a norm under Assumption 6).

3.2 Assumptions. We make the following assumptions when analyzing LCGT models.

Assumption 4 (Compact features)**.** *There exists some constant $R > 0$ such that $\forall i \in [n]$, $\|\boldsymbol{x}_i\|_2 \le R$.*

Assumption 5 (Compact parameters)**.** *There exists some constant $b' > 0$ with $b' \le b/R$ (recall that b was defined in Assumption 3) such that $\boldsymbol{\theta} \in \mathbb{R}^p$ falls in the convex set $\Theta \triangleq \{\boldsymbol{\theta} \in \mathbb{R}^p : \|\boldsymbol{\theta}\|_2 \le b'\}$.*

Assumption 6 (Identifiability)**.** *We assume $\lambda_{\min}(H)$ is strictly positive and $p \le n-1$.*

Assumptions 4 and 5 are necessary to ensure that Assumption 2 is not violated. In particular, they imply that $\boldsymbol{\theta}^{\mathsf{T}}\boldsymbol{x}_{ij} \in [-2b, 2b]$ for all $\boldsymbol{\theta} \in \Theta$ and $\boldsymbol{x}_i, \boldsymbol{x}_j$. Such scale regularization conditions arise because CGT models are not scale-invariant due to the nonlinearity of F. In practice, such assumptions can be satisfied by appropriately normalizing the dataset. Assumption 6 ensures that the maximum likelihood estimation (MLE) problem defined in (2) has a unique solution. Note that if one has $p \ge n-1$, the LCGT reduces to a non-contextual GT model when each $\boldsymbol{x}_i$ is set to the one-hot vector $\boldsymbol{e}_i$ for all $i \in [n-1]$ and $\boldsymbol{x}_n$ is set to the all-zeros vector $\mathbf{0}_p$.

3.3 Estimation results. Under the LCGT model, the optimal parameter vector $\boldsymbol{\theta}^* \in \Theta$ is defined to be the vector that minimizes the negative log likelihood function $L\colon \Theta \to \mathbb{R}$ under the true probabilities p_{ij}, i.e.,

$$\boldsymbol{\theta}^* \triangleq \operatorname*{arg\,min}_{\boldsymbol{\theta}\in\Theta} L(\boldsymbol{\theta}) \triangleq \operatorname*{arg\,min}_{\boldsymbol{\theta}\in\Theta} \sum_{(i,j)\in\mathcal{E}} -p_{ij}\log(F(\boldsymbol{\theta}^{\mathsf{T}}\boldsymbol{x}_{ij})) - (1-p_{ij})\log(1-F(\boldsymbol{\theta}^{\mathsf{T}}\boldsymbol{x}_{ij})), \tag{2}$$

and its estimator $\hat{\boldsymbol{\theta}} \in \Theta$ minimizes the negative log likelihood function $\hat{L} : \Theta \to \mathbb{R}$, defined

$$\hat{\boldsymbol{\theta}} \triangleq \operatorname*{arg\,min}_{\boldsymbol{\theta}\in\Theta} \hat{L}(\boldsymbol{\theta}) \triangleq \operatorname*{arg\,min}_{\boldsymbol{\theta}\in\Theta} \sum_{(i,j)\in\mathcal{E}} -\hat{p}_{ij}\log(F(\boldsymbol{\theta}^{\mathsf{T}}\boldsymbol{x}_{ij})) - (1-\hat{p}_{ij})\log(1-F(\boldsymbol{\theta}^{\mathsf{T}}\boldsymbol{x}_{ij})). \tag{3}$$

For completeness, we prove some of the basic properties of $\boldsymbol{\theta}^*$ and $\hat{\boldsymbol{\theta}}$ in Appendix A, such as Lemma 7, which shows for completeness that $\hat{\boldsymbol{\theta}}$ is a valid MLE, and Lemma 8, which establishes the existence and uniqueness of $\boldsymbol{\theta}^*$. Lemma 7 also implies the consistency and asymptotic normality of $\hat{\boldsymbol{\theta}}$ as $k \to \infty$ under certain regularity conditions, as discussed in Appendix A.

Note that our choice of loss function lacks a regularization term, unlike works such as [38]. While the non-contextual GT model benefits from regularization due to its shift-invariance (i.e., shifting every parameter by the same constant does not affect the model), this is not the case for the LCGT model. Although regularization can also be useful to guarantee identifiability even under overparameterization (i.e., when $p > n$), we do not consider this case in this work since testing is vacuous in this regime. That said, we expect most of our analysis to go through without much modification under ℓ^2 regularization (which will only strengthen the strong convexity of the loss function).

We start by analyzing how well we can estimate $\boldsymbol{\theta}^*$ using the MLE $\hat{\boldsymbol{\theta}}$.

Theorem 1 (Error bounds for MLE under LCGT model)**.** *The expectation of the squared H-norm of the estimation error is bounded by $\mathbb{E}[\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H^2] \le \hat{\beta}^2 p/\alpha^2 k$, and for all sufficiently small $\delta \in (0, 2e^{-c_2 p}]$, with probability at least $1-\delta$, we have*

$$\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H \le \frac{\hat{\beta}}{\alpha}\sqrt{\frac{p}{k} + \frac{4c_1\log(2/\delta)}{c_2 k}}, \tag{4}$$

where $c_1, c_2 > 0$ are some constants. Furthermore, for all $s \in \mathbb{Z}_{>0}$, the s-th moment of the error is bounded by

$$\mathbb{E}[\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H^s] \le \left(\frac{2\sqrt{c_1}\hat{\beta}}{\alpha}\sqrt{\frac{p}{k}}\right)^s \left(1 + 2\frac{(2s)!}{(c_2 p)^{2s}}\right). \tag{5}$$

Theorem 1 is proved in Appendix B. Some remarks are in order, especially regarding Theorem 1's relation to other known bounds in the literature and its role in deriving hypothesis testing results such as Theorem 2. First, (4) matches the scaling of [16, Theorem 1.1], which assumes $k = 1$ and bounds the $H/|\mathcal{E}|$-norm of the error of the linear contextual BTL MLE when the comparison graph is

a perfect matching. Furthermore, under the non-contextual setting where $p = n - 1$ and $\boldsymbol{x}_i$ are set according to the discussion under Assumption 6, the tail bound of [29, Theorem 3.3] (which scales like $\sqrt{n \log(1/\delta)/k}$ when converted to a high-probability bound like ours) is looser than (4), and its moment bound matches the scaling of (5). Note that although [29, Theorem 3.3] is stated in terms of the L-seminorm (where L is the Laplacian matrix of $\mathcal{G}$) instead of the H-norm, it is still comparable to our result since our H in the non-contextual setting will be the principal submatrix of L obtained by removing the last row and column, and thus the eigenvalues of H and L are interlaced according to Cauchy's interlacing theorem [47].

Next, note that Theorem 1 holds regardless of whether the true probabilities p_{ij} obey the LCGT model, since the vectors $\boldsymbol{\theta}^*$ and $\hat{\boldsymbol{\theta}}$ remain well-defined even under non-LCGT datasets (i.e., under misspecification). Note that although our error bounds on the MLE of the LCGT model generalize to the misspecified setting rather straightforwardly, this is not necessarily the case for all MLEs in general. Establishing such bounds that are robust under model mismatch is essential to obtain hypothesis testing results such as Theorem 2. This is because our testing procedure first fits an LCGT model on the data then evaluates the closeness of the fit to obtain a test statistic, thus tight bounds on the MLE error under both the null hypothesis and the alternative hypothesis are needed to effectively characterize the behavior of the test statistic. Moment bounds like (5) (up to the fourth moment) are also needed to bound the variance of the test statistic.

3.4 Testing results. Now, we define the hypothesis testing problem for LCGT models. First, let $Q(\boldsymbol{\theta})$ denote the probability matrix under a LCGT model parameterized by $\boldsymbol{\theta}$, i.e., $Q(\boldsymbol{\theta})_{ij} \triangleq F(\boldsymbol{\theta}^\mathsf{T} \boldsymbol{x}_{ij})$ for all $(i,j) \in \mathcal{E}$ and $Q(\boldsymbol{\theta})_{ij} \triangleq 0$ for all $(i,j) \in [n]^2 \setminus \mathcal{E}$. Then, we define the composite hypothesis testing problem for LCGT testing parameterized by the separation radius $\epsilon \in (0,1]$ as follows.

- Null hypothesis (H_0): $\mathcal{Z} \sim P \in P_0$, where $P_0 \triangleq \{P \in [0,1]^{n \times n} : \exists \boldsymbol{\theta}^* \in \Theta, P = Q(\boldsymbol{\theta}^*)\}$.
- Alternative hypothesis ($H_1(\epsilon)$): $\mathcal{Z} \sim P \in P_1(\epsilon)$, where

$$P_1(\epsilon) \triangleq \Big\{ P \in [0,1]^{n \times n} : \tfrac{1}{\sqrt{|\mathcal{E}|}} \inf_{\boldsymbol{\theta} \in \Theta} \|P - Q(\boldsymbol{\theta})\|_\mathrm{F} \geq \epsilon > 0 \text{ and } \forall (i,j) \in [n]^2 \setminus \mathcal{E}, P_{ij} = 0 \Big\}.$$

Note that we normalize the scale of ϵ by scaling the Frobenius norm by $1/\sqrt{|\mathcal{E}|}$, since $|\mathcal{E}|$ is the number of non-zero elements in P and $Q(\cdot)$. Intuitively, P is exactly characterized by a LCGT model under H_0, whereas P is at least ϵ-far from any LCGT model under $H_1(\epsilon)$. The next step is to show that the infimum in $H_1(\epsilon)$, i.e., the distance between P and the closest LCGT model, behaves like the distance between P and the LCGT model parameterized by $\boldsymbol{\theta}^*$ as defined in (2). This is deferred to Proposition 9 in Appendix A. Next, we define the test statistic T for minimax testing as

$$T \triangleq \tfrac{1}{|\mathcal{E}|} \Bigg(\|\hat{P} - Q(\hat{\boldsymbol{\theta}})\|_\mathrm{F}^2 - \underbrace{\sum_{(i,j) \in \mathcal{E}} \hat{p}_{ij} \left(\hat{p}_{ij} - \tfrac{k \hat{p}_{ij} - 1}{k-1} \right)}_{\text{ⓐ}} \Bigg), \tag{6}$$

where the $1/|\mathcal{E}|$ factor ensures that T scales like ϵ^2 under $H_1(\epsilon)$, and the bias correction term ⓐ makes T an unbiased estimator of $\|P - Q(\hat{\boldsymbol{\theta}})\|_\mathrm{F}^2$. Note that we assume that the MLE $\hat{\boldsymbol{\theta}}$ is computed from a separate dataset $\mathcal{Z}_1$ that is independent of the dataset $\mathcal{Z}_2$ that is used for testing (but sampled from the same distribution). In practice, this may be achieved with cross validation techniques such as sample splitting. Next, we define the *minimax risk* $\mathcal{R}(\mathcal{G}, \epsilon)$ and *critical threshold* (or critical radius) ϵ_c of LCGT testing under the framework of minimax testing introduced by [48]. First, we define $\mathcal{R}(\mathcal{G}, \epsilon) \in [0,1]$ as

$$\mathcal{R}(\mathcal{G}, \epsilon) \triangleq \inf_{\hat{H}} \Big(\sup_{P \in P_0} \mathbb{P}_{\mathcal{Z} \sim P}(\hat{H}(\mathcal{Z}) = 1) + \sup_{P \in P_1(\epsilon)} \mathbb{P}_{\mathcal{Z} \sim P}(\hat{H}(\mathcal{Z}) = 0) \Big),$$

where the infimum is over all (possibly randomized) decision rules $\hat{H}(\mathcal{Z}) \in \{0,1\}$, the first supremum is over all probability matrices under H_0, and the second supremum is over all probability matrices under $H_1(\epsilon)$. Intuitively, the minimax risk is the best (i.e., smallest) sum of Type I and Type II errors any decision rule can achieve against worst case inputs (hence the name). Compared to the traditional approach of significance testing (which only considers Type I error), the framework of minimax testing makes it possible to precisely characterize the minimum separation that allows for reliable testing by ensuring that both the Type I and Type II errors are well controlled. This motivates the definition of the critical threshold,

$$\epsilon_c \triangleq \inf \big\{ \epsilon \in (0,1] : \mathcal{R}(\mathcal{G}, \epsilon) \leq \tfrac{1}{2} \big\}. \tag{7}$$

Note that the choice of $1/2$ as the minimax risk threshold is arbitrary; any other constant in $(0, 1)$ may be used instead, but this will not affect ϵ_c beyond some constant factor. Because $\mathcal{R}(\mathcal{G}, \epsilon)$ is monotonically non-increasing with respect to ϵ, the critical threshold ϵ_c represents the phase transition point for hypothesis testing: if $\epsilon > \epsilon_c$, some testing procedure with bounded risk must exist, whereas none can exist when $\epsilon < \epsilon_c$. Intuitively, testing is not possible if ϵ is too small because the probability distribution of $\mathcal{Z}$ conditioned on $H_1(\epsilon)$ may be too close to the distribution conditioned on H_0. In particular, if $\epsilon < \epsilon_c$, the receiver operating characteristic (ROC) curve of any hypothesis test must be upper bounded by the line $y = x + 1/2$ (see Section 5 for further discussion on the interpretation of ROC curves in the context of minimax testing).

We begin with an upper bound on ϵ_c.

Theorem 2 (Upper bound on critical threshold for LCGT testing)**.** *The critical threshold ϵ_c for LCGT testing as defined in (7) satisfies*

$$\epsilon_c \leq c\sqrt{\frac{\sqrt{|\mathcal{E}|}+p}{|\mathcal{E}|k}},$$

where c is some positive constant.

Theorem 2 is proved in Appendix C by showing that the test statistic defined in (6) concentrates around 0 under H_0 and is strictly bounded away from 0 with high probability under $H_1(\epsilon)$ for all sufficiently large ϵ. Compared to [29, Theorem 3.2], which gives an upper bound on the critical threshold that scales like $\sqrt{n/(|\mathcal{E}|k)}$ under the non-contextual model (after accounting for differences in their definition of the critical threshold compared to ours), our bound matches its scaling for complete graphs and improves upon it for sparse graphs with $|\mathcal{E}| = o(n^2)$. In particular, our bound scales like $1/\sqrt{n^{1/2}k}$ in the important regime of complete matchings when p is constant (which is the usual setting for RLHF applications). Finally, we provide an information-theoretic lower bound for the critical threshold.

Proposition 3 (Minimax Lower Bound)**.** *Assume that the comparison graph $\mathcal{G}$ is super-Eulerian with a maximum spanning Eulerian sub-graph $\tilde{\mathcal{G}}$. Then, there exists some constant $c > 0$ such that the minimax risk $\mathcal{R}(\mathcal{G}, \epsilon)$ is lower bounded by*

$$\mathcal{R}(\mathcal{G}, \epsilon) \geq 1 - \tfrac{1}{2}\sqrt{\exp\Big(ck^2\epsilon^4 \sum_{\sigma\in\mathcal{C}} |\sigma|^2\Big) - 1}$$

for all $\epsilon > 0$, where $|\sigma|$ denotes the length of a cycle $\sigma \in \mathcal{C}$ and $\mathcal{C}$ is some cycle decomposition of $\tilde{\mathcal{G}}$. In particular, this implies that the critical threshold ϵ_c is lower bounded by $\Omega(\sqrt{1/(nk)})$ for complete graphs, matching the scaling of the upper bound in Theorem 2.

Proposition 3 is proved in Appendix D.

4 Basis contextual model

4.1 Preliminaries. In this section, we study the *basis contextual generalized Thurstone (BCGT) model*, which generalizes the LCGT model we considered in the previous section. Under the BCGT model, the space of latent functions is the linear span of some fixed set of basis functions $\phi_i : \mathbb{R}^p \to \mathbb{R}$ for all $i \in [D]$, where D is the size of the basis. The basis functions may be stacked to define the vector-valued function $\boldsymbol{\phi}(\cdot) \triangleq (\phi_1(\cdot), \ldots, \phi_D(\cdot))$. Note that the LCGT model can be recovered as a special case by choosing the linear basis $\boldsymbol{\phi}(\boldsymbol{x}) = \boldsymbol{x}$.

Next, given the context vectors $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$, we define the embedding matrix $\Phi \in \mathbb{R}^{n\times D}$ to be the matrix obtained by stacking the row vectors $\boldsymbol{\phi}(\boldsymbol{x}_i)$, i.e., $\Phi \triangleq [\boldsymbol{\phi}(\boldsymbol{x}_1) \quad \boldsymbol{\phi}(\boldsymbol{x}_2) \quad \cdots \quad \boldsymbol{\phi}(\boldsymbol{x}_n)]^\mathsf{T}$. The BCGT is parameterized by the vector $\boldsymbol{\theta} \in \mathbb{R}^D$ such that score function is represented as $f(\boldsymbol{x}) = \boldsymbol{\theta}^\mathsf{T}\boldsymbol{\phi}(\boldsymbol{x})$. Similarly to Section 3, we define $\boldsymbol{x}_{ij} \triangleq \boldsymbol{\phi}(\boldsymbol{x}_j) - \boldsymbol{\phi}(\boldsymbol{x}_i)$ for all $(i, j) \in \mathcal{E}$, let $X \in \mathbb{R}^{|\mathcal{E}|\times p}$ be the matrix obtained by stacking the row vectors $\boldsymbol{x}_{ij}^\mathsf{T}$ for all $(i, j) \in \mathcal{E}$, and define $H \triangleq \sum_{(i,j)\in\mathcal{E}} \boldsymbol{x}_{ij}\boldsymbol{x}_{ij}^\mathsf{T} = X^\mathsf{T}X$. In particular, we are often interested in its extremal eigenvalues, which we refer to as $\lambda_{\min} \triangleq \lambda_{\min}(H)$ and $\lambda_{\max} \triangleq \lambda_{\max}(H)$ for convenience.

4.2 Assumptions. Next, we lay out the assumptions that we adopt under the BCGT model.

Assumption 7 (Centering of basis features)**.** *The column sums of Φ are 0, i.e., $\mathbf{1}_n^\mathsf{T}\Phi = \mathbf{0}_D^\mathsf{T}$.*

Assumption 8 (Bound on basis function)**.** *We bound the ℓ^2-norm of $\boldsymbol{\phi}$, i.e., we assume that $\|\boldsymbol{\phi}(\boldsymbol{x})\|_2 \leq R < \infty$ for all $\boldsymbol{x} \in \mathcal{X}$.*

Assumption 9 (Bound on parameters)**.** *The parameter vector $\boldsymbol{\theta} \in \mathbb{R}^D$ falls in the set $\Theta = \{\boldsymbol{\theta} \in \mathbb{R}^D : \|\boldsymbol{\theta}\|_2 \leq b'\}$, where b' is chosen appropriately to satisfy $b'R \leq b$ (cf. Assumption 5).*

Since the GT model is invariant under translation, Assumption 7 can always be satisfied by translating the basis functions appropriately once the context vectors are known. As in the LCGT setting, the remaining regularity conditions serve to ensure that Assumption 2 is satisfied. Note that unlike Section 3.2, we do not impose an explicit bound on the spectrum of H. Instead, we will use probabilistic arguments to characterize the scaling of H's eigenvalues under an orthonormal basis.

4.3 Main results. Similarly to Section 3.3, $\boldsymbol{\theta}^*$ is defined to be the parameters that minimize the negative log likelihood function L, defined

$$\boldsymbol{\theta}^* \triangleq \arg\min_{\boldsymbol{\theta}\in\Theta} L(\boldsymbol{\theta}) \triangleq \arg\min_{\boldsymbol{\theta}\in\Theta} \sum_{(i,j)\in\mathcal{E}} -p_{ij}\log(F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})) - (1-p_{ij})\log(1-F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})),$$

and the MLE $\hat{\boldsymbol{\theta}}$ is defined as

$$\hat{\boldsymbol{\theta}} \triangleq \arg\min_{\boldsymbol{\theta}\in\Theta} \hat{L}(\boldsymbol{\theta}) \triangleq \arg\min_{\boldsymbol{\theta}\in\Theta} \sum_{(i,j)\in\mathcal{E}} -\hat{p}_{ij}\log(F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})) - (1-\hat{p}_{ij})\log(1-F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})).$$

Next, we present the BCGT counterpart to Theorem 1, presented in terms of the ℓ^2 norm. Note that the BCGT and LCGT models are equivalent in the sense that all LCGT models can be converted to a BCGT model by taking the basis functions to be the linear basis, and all BCGT models can be converted to an LCGT model by using the mapped embeddings themselves as context vectors. However, we present them (and results regarding them) separately to enhance the presentation of our work (by beginning with a simpler model and elaborating on it in the sequel) and because additional interesting results such as Lemmas 11 and 12 can be said for BCGT models that use certain orthonormal bases.

Theorem 4 (Error bounds for MLE under BCGT model)**.** *For all $\delta \in (0,1)$, with probability at least $1-\delta$, we have*

$$\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \leq \tfrac{4}{\alpha\lambda_{\min}}\sqrt{\tfrac{8|\mathcal{E}|R^2\hat{\beta}^2}{k}(\log(2/\delta) + D\log(5))}.$$

Theorem 4 is proved in Appendix E, and corresponding moment bounds are provided in Lemma 10. Our bound can be compared to the estimation bound in [15, Theorem 1] (which considers the linear CBTL model with $k = 1$ under a multigraph with edges selected uniformly with replacement), which scales like $(1/\lambda_{\min})\sqrt{D\log(D/\delta)|\mathcal{E}|}$ (after accounting for differences on how $\lambda_{\min}$ is defined), although it may be difficult to compare the bounds directly due to differences in setting.

Since Theorem 4 presents the estimation error as an ℓ^2 norm, the dependence on the spectrum of H, which was implicit in Theorem 1, now appears explicitly as $\lambda_{\min}$. Thus, to characterize the bound's scaling, the asymptotic behavior of $\lambda_{\min}$ must be determined. To this end, Lemmas 11 and 12, which are deferred to Appendix A, use matrix concentration inequalities to show that $\lambda_{\min}$ and $\lambda_{\max}$ grow like $|\mathcal{E}|$ under complete graphs and perfect matchings when the context vectors are i.i.d. and the basis functions are orthonormal with respect to the distribution of the context vectors. Furthermore, Corollary 13 presents concrete examples of orthonormal bases such as the Fourier basis or Legendre basis [49] that satisfy the assumptions of Lemmas 11 and 12. Thus, under the regime of Lemmas 11 and 12, the bound in Theorem 4 will scale like $\sqrt{(\log(1/\delta)+D)/(|\mathcal{E}|k)}$ with high probability.

Next, we analyze the problem of testing for the BCGT model, where the hypotheses H_0 and $H_1(\epsilon)$ are posed as in Section 3.4. Likewise, we may reuse the definitions of the test statistic T, minimax risk $\mathcal{R}(\mathcal{G},\epsilon)$, and critical threshold ϵ_c in Section 3.4 for the BCGT setting as well. Then, we may use our estimation bounds in Theorem 4 and Lemma 10 to obtain the following upper bound on the critical threshold.

Theorem 5 (Upper bound of critical threshold in basis model)**.** *Under the BCGT model, the critical threshold ϵ_c defined in (7) satisfies*

$$\epsilon_c \leq c\sqrt{\tfrac{1}{k\sqrt{|\mathcal{E}|}} + \tfrac{R^2}{\lambda_{\min}^2}\tfrac{D\lambda_{\max}}{k}},$$

where c is some positive constant.

Theorem 5 is proved in Appendix F. In the regime of Lemmas 11 and 12, the bound scales like $\sqrt{(|\mathcal{E}|^{1/2}+D)/(|\mathcal{E}|k)}$ with high probability, which matches Theorem 2. We conclude with a lower bound on the critical threshold in Proposition 6, which can be proved analogously to Proposition 3.

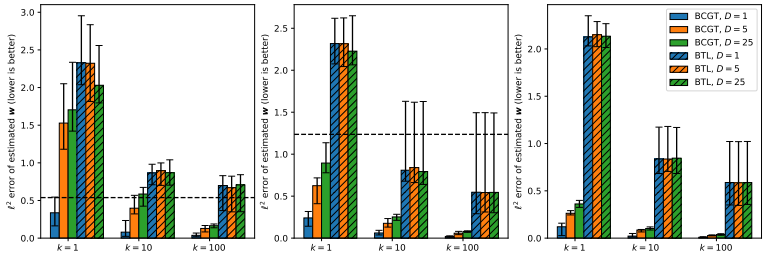

(a) $n = 10$ (b) $n = 50$ (c) $n = 250$

Figure 1: Plot of ℓ^2 estimation error of the BCGT (with polynomial basis) and BTL models on complete graphs with $p = 1$, $n \in \{10, 50, 250\}$, $k \in \{1, 10, 100\}$, and $D \in \{1, 5, 25\}$.

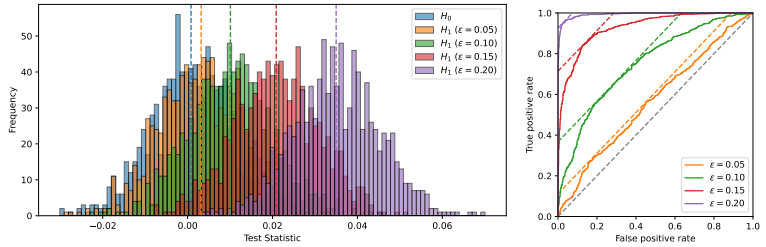

(a) Histograms of T under H_0 and $H_1(\epsilon)$ (b) ROC curves

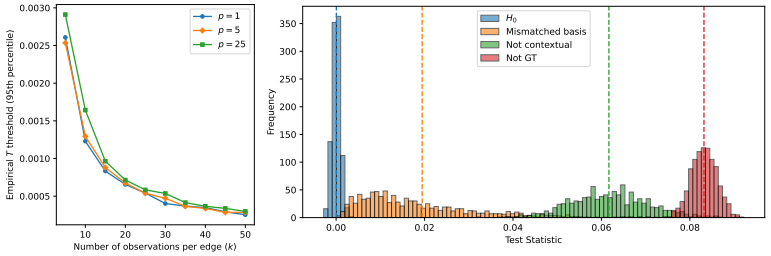

(c) Empirical thresholds (d) Histograms of T under model mismatch

Figure 2: **(a)** Empirical distributions of T under the H_0 and $H_1(\epsilon)$ hypotheses for $\epsilon \in \{0.05, 0.1, 0.15, 0.2\}$. **(b)** ROC curves derived from the test statistics plotted in subfigure (a). **(c)** Plot of the threshold for T to reject H_0 under the LCGT model for $k \in \{5, 10, \ldots, 50\}$ and $p \in \{1, 5, 25\}$. **(d)** Empirical distributions of T under H_0 and various model mismatch settings such as basis mismatch, context mismatch, and general model mismatch.

Proposition 6 (Minimax lower bound in basis model)**.** *For a complete graph $\mathcal{G}$, the critical threshold ϵ_c is lower bounded by $\Omega(\sqrt{1/nk})$, which matches the upper bound in Theorem 5.*

A short statement explaining the argument for Proposition 6 can be found in Appendix G.

5 Experiments

In this section, we present experiment results on synthetic and real-world datasets that illustrate the performance of CGT estimation and testing in various settings. Due to space constraints, we defer most experiments to Appendix H and present only a representative subset here. The experiments were written in Python, and MLEs were solved with the L-BFGS [50] algorithm implemented in PyTorch.

5.1 Synthetic Experiments. First, we compare the ℓ^2 estimation error of the BCGT model (with $F(t) = 1/(1 + e^{-t})$) and the non-contextual BTL model on a BCGT dataset supported on complete comparison graphs. The basis is set to be the polynomial basis with degree $D \in \{1, 5, 25\}$. Each $x_i \in \mathbb{R}$ (i.e., $p = 1$) are independently sampled from $\mathrm{Uniform}(0, 1)$, and $\boldsymbol{\theta} \in \mathbb{R}^D$ are sampled from $\mathcal{N}(0, 1/\sqrt{D})$ (where the scaling ensures that the scale of $\boldsymbol{\phi}(x_i)^\mathsf{T}\boldsymbol{\theta}$ stays invariant over different D).

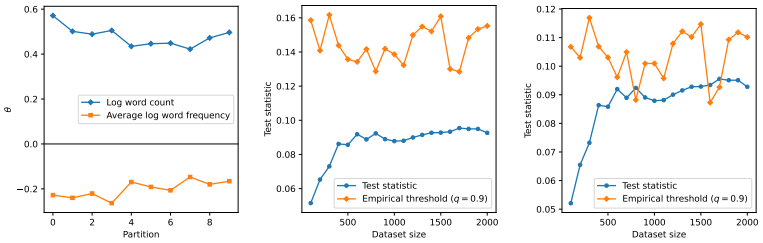

(a) Estimation results (b) T with BTL F (c) T with Thurstone F

Figure 3: **(a)** Result of LCGT estimation on the LLM dataset for each partition of the dataset. **(b)** Plot of the test statistic T and empirical threshold for LCGT testing (with BTL F) for progressively larger subsets of the dataset. **(c)** Plot of the test statistic T and empirical threshold for LCGT testing (with Thurstone F) for progressively larger subsets of the dataset.

Once the comparison outcomes between each pair of items are generated according to the BCGT model, the ℓ^2 error of the score vector $\boldsymbol{w}$ predicted by BCGT and BTL estimation is recorded (note that we measure the error of $\boldsymbol{w}$, not $\boldsymbol{\theta}$, so that the performance of contextual and non-contextual estimation models can be directly compared). We repeat this process 25 times for each combination of the parameters $n \in \{10, 50, 250\}$, $k \in \{1, 10, 100\}$, and $p \in \{1, 5, 25\}$. The results of this experiment are visualized in Figure 1, where the colored bars represent the median and the error bars the first and third quartiles of the 25 trials for each experiment. The horizontal dashed line represents the baseline ℓ^2 error that may be achieved by simply guessing $\boldsymbol{w} = \mathbf{0}_n$. In general, the BCGT model, armed with additional context, tends to outperform the BTL model. Furthermore, as Theorem 4 predicts, the estimation error generally scales as $\sqrt{D/k}$.

Next, we simulate the distribution of the BCGT test statistic T under the H_0 and $H_1(\epsilon)$ hypotheses for various ϵ. Here, the comparison graph is a perfect matching with $n = 500$, and the basis is a quadratic polynomial basis with $p = 2$ and $D = 5$. For each trial, a random dataset under the BCGT model is generated by independently sampling each $\boldsymbol{x}_i \in \mathbb{R}^2$ from $\mathcal{N}(0, 1/2)^2$ and $\boldsymbol{\theta}$ from $\mathcal{N}(0, 1)^5$. Under the $H_1(\epsilon)$ hypothesis, each entry in the probability matrix is perturbed by i.i.d. noise each sampled from $\mathcal{N}(0, \epsilon)$. Then, comparison outcomes are simulated based on the computed probability matrix and the test statistic T is computed. Figure 2(a) visualizes the distributions of T under H_0 and $H_1(\epsilon)$ hypotheses for $\epsilon \in \{0.05, 0.1, 0.15, 0.2\}$ after 1000 trials (where each colored dotted line represents the mean of each empirical distribution), and Figure 2(b) shows the corresponding ROC curves, where each colored dashed line represents the minimum risk corresponding to each curve (also called Youden's J statistic [51]). As theory predicts, the T is centered around 0 under H_0 and scales like ϵ^2 under $H_1(\epsilon)$. Next, Figure 2(c) plots the empirical threshold (in the significance testing sense with $p = 0.05$) for T to reject H_0 as a function of k for various values of p under the LCGT model on a perfect matching with $n = 500$, which shows that the threshold roughly scales like $\sqrt{(|\mathcal{E}|^{1/2} + p)/(|\mathcal{E}|k)}$ as Theorem 2 predicts (with $|\mathcal{E}|$ being the more dominant factor). Finally, Figure 2(d) shows how the distribution of T changes under H_0, basis mismatch (when the data obeys a BCGT model with a Fourier basis, but is tested for a polynomial basis), context mismatch (when a non-contextual dataset is tested for BCGT), and general model mismatch (when the dataset is not GT at all).

5.2 LMArena Human Preferences. The LMArena human preference 55k dataset [33] consists of $n = 57477$ human comparisons between pairs of LLM responses to a given prompt. To construct the context vectors, we computed two simple metrics: log word count and average log word frequency. The log word count is the logarithm of the number of words in the response, and the average log word frequency is the average of the logarithm of the frequency (in the English language) of each word in the response (so that a smaller value corresponds to a higher density of rarer words). Figure 3(a) shows the result of fitting an LCGT model on the dataset partitioned into 10 equal chunks, and Figure 3(b)–(c) show how the test statistic T and empirical threshold evolve for progressively increasing subsets of the dataset under the LCGT model with F taken to be the function that corresponds to the BTL (b) or Thurstone (c) model. Figure 3(a) shows that in general, longer responses with less common words are preferred, while Figure 3(b)–(c) may imply that the LLM dataset may conform closer to a BTL model than a Thurstone model.

Acknowledgment

This work was supported in part by the National Science Foundation (NSF) CAREER Award under Grant CCF-2337808.

References

[1] D. McFadden, "Conditional logit analysis of qualitative choice behavior," in *Frontiers in Econometrics*, P. Zarembka, Ed., New York, NY, USA: Academic Press, 1973, pp. 105–142.

[2] J. I. Yellott Jr, "The relationship between Luce's choice axiom, Thurstone's theory of comparative judgment, and the double exponential distribution," *Journal of Mathematical Psychology*, vol. 15, no. 2, pp. 109–144, Apr. 1977.

[3] A. Jadbabaie, A. Makur, and D. Shah, "Estimation of skill distribution from a tournament," in *Proceedings of the Advances in Neural Information Processing Systems 33 (NeurIPS)*, Vancouver, BC, Canada, Dec. 2020, pp. 8418–8429.

[4] A. Jadbabaie, A. Makur, and D. Shah, "Estimation of skill distributions," *IEEE Transactions on Information Theory*, vol. 70, no. 9, pp. 6447–6480, Sep. 2024.

[5] P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei, "Deep reinforcement learning from human preferences," in *Proceedings of the Advances in Neural Information Processing Systems 30 (NeurIPS)*, Long Beach, CA, USA, Dec. 2017, pp. 4299–4307.

[6] N. Stiennon, L. Ouyang, J. Wu, *et al.*, "Learning to summarize with human feedback," in *Proceedings of the Advances in Neural Information Processing Systems 33 (NeurIPS)*, Dec. 2020, pp. 3008–3021.

[7] L. Ouyang, J. Wu, X. Jiang, *et al.*, "Training language models to follow instructions with human feedback," in *Proceedings of the Advances in Neural Information Processing Systems 35 (NeurIPS)*, New Orleans, LA, USA, Nov. 2022, pp. 27 730–27 744.

[8] R. A. Bradley and M. E. Terry, "Rank analysis of incomplete block designs. I. The method of paired comparisons," *Biometrika*, vol. 39, no. 3/4, pp. 324–345, Dec. 1952.

[9] R. D. Luce, *Individual Choice Behavior: A Theoretical Analysis*. New York, NY, USA: John Wiley & Sons, 1959.

[10] R. L. Plackett, "The analysis of permutations," *Journal of the Royal Statistical Society, Series C (Applied Statistics)*, vol. 24, no. 2, pp. 193–202, 1975.

[11] E. Zermelo, "Die berechnung der turnier-ergebnisse als ein maximumproblem der wahrscheinlichkeitsrechnung," *Mathematische Zeitschrift*, vol. 29, no. 1, pp. 436–460, Dec. 1929.

[12] L. R. Ford, "Solution of a ranking problem from binary comparisons," *The American Mathematical Monthly*, vol. 64, no. 8 Part 2, pp. 28–33, Oct. 1957.

[13] Y. Bai, A. Jones, K. Ndousse, *et al.*, "Training a helpful and harmless assistant with reinforcement learning from human feedback," arXiv:2204.05862 [cs.CL], Apr. 2022. [Online]. Available: https://arxiv.org/abs/2204.05862.

[14] D. Ganguli, L. Lovitt, J. Kernion, *et al.*, "Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned," arXiv:2209.07858v2 [cs.CL], Nov. 2022. [Online]. Available: https://arxiv.org/abs/2209.07858v2.

[15] A. Bower and L. Balzano, "Preference modeling with context-dependent salient features," in *Proceedings of the 37th International Conference on Machine Learning (ICML)*, Jul. 2020, pp. 1067–1077.

[16] B. Zhu, M. Jordan, and J. Jiao, "Principled reinforcement learning with human feedback from pairwise or K-wise comparisons," in *Proceedings of the 40th International Conference on Machine Learning (ICML)*, Honolulu, HI, USA, Jul. 2023, pp. 43 037–43 067.

[17] J. Fan, J. Hou, and M. Yu, "Uncertainty quantification of MLE for entity ranking with covariates," *Journal of Machine Learning Research*, vol. 25, no. 358, pp. 1–83, 2024.

[18] K. J. Arrow, *Social Choice and Individual Values*. New York, NY, USA: John Wiley & Sons, 1951.

[19] D. Davidson and J. Marschak, "Experimental tests of a stochastic decision theory," in *Economic Information, Decision, and Prediction*, Dordrecht, the Netherlands: Springer, 1959, pp. 133–171.

[20] A. Tversky, "Elimination by aspects: A theory of choice," *Psychological Review*, vol. 79, no. 4, pp. 281–299, Jul. 1972.

[21] J. Rieskamp, J. R. Busemeyer, and B. A. Mellers, "Extending the bounds of rationality: Evidence and theories of preferential choice," *Journal of Economic Literature*, vol. 44, no. 3, pp. 631–661, Sep. 2006.

[22] L. L. Thurstone, "A law of comparative judgment," *Psychological Review*, vol. 34, no. 4, pp. 273–286, Jul. 1927.

[23] L. L. Thurstone, "The indifference function," *The Journal of Social Psychology*, vol. 2, no. 2, pp. 139–167, 1931.

[24] J. I. Yellott Jr, "Generalized Thurstone models for ranking: Equivalence and reversibility," *Journal of Mathematical Psychology*, vol. 22, no. 1, pp. 48–69, Aug. 1980.

[25] Q. Zhang, H. Wei, and L. Ying, "Reinforcement learning from human feedback without reward inference: Model-free algorithm and instance-dependent analysis," *Reinforcement Learning Journal*, vol. 3, pp. 1236–1251, 2024.

[26] H. Sun, Y. Shen, and J.-F. Ton, "Rethinking reward modeling in preference-based large language model alignment," in *Proceedings of the Thirteenth International Conference on Learning Representations (ICLR)*, Singapore, Apr. 2025, pp. 46 885–46 922.

[27] Y. Zhang, G. Zhang, Y. Wu, K. Xu, and Q. Gu, "Beyond Bradley-Terry models: A general preference model for language model alignment," in *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, Vancouver, BC, Canada, Jul. 2025, pp. 76 939–76 965.

[28] A. Makur and J. Singh, "Minimax hypothesis testing for the Bradley–Terry–Luce model," *IEEE Transactions on Information Theory*, vol. 71, no. 12, pp. 9163–9202, Dec. 2025.

[29] A. Makur and J. Singh, "Hypothesis testing for generalized Thurstone models," in *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, Vancouver, BC, Canada, Jul. 2025, pp. 42 730–42 764.

[30] G. Simons and Y.-C. Yao, "Asymptotics when the number of parameters tends to infinity in the Bradley-Terry model for paired comparisons," *The Annals of Statistics*, vol. 27, no. 3, pp. 1041–1060, Jun. 1999.

[31] H. Bong and A. Rinaldo, "Generalized results for the existence and consistency of the MLE in the Bradley-Terry-Luce model," in *Proceedings of the 39th International Conference on Machine Learning (ICML)*, Baltimore, MD, USA, Jul. 2022, pp. 2160–2177.

[32] D. Lee, A. Makur, and J. Singh, "Entrywise error bounds for spectral ranking with semi-random adversaries," in *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (KDD '26)*, Jeju Island, Republic of Korea, Aug. 2026, pp. 2449–2460.

[33] W.-L. Chiang, L. Zheng, Y. Sheng, *et al.*, "Chatbot arena: An open platform for evaluating LLMs by human preference," in *Proceedings of the 41st International Conference on Machine Learning (ICML)*, Vienna, Austria, Jul. 2024, pp. 8359–8388.

[34] W. Li, S. Shrotriya, and A. Rinaldo, "ℓ_∞-bounds of the MLE in the BTL model under general comparison graphs," in *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence (UAI)*, Eindhoven, The Netherlands, Aug. 2022, pp. 1178–1187.

[35] S. Negahban, S. Oh, and D. Shah, "Iterative ranking from pair-wise comparisons," in *Proceedings of the Advances in Neural Information Processing Systems 25 (NeurIPS)*, Lake Tahoe, NV, USA, Dec. 2012, pp. 2474–2482.

[36] N. B. Shah, S. Balakrishnan, J. Bradley, A. Parekh, K. Ramchandran, and M. J. Wainwright, "Estimation from pairwise comparisons: Sharp minimax bounds with topology dependence," *Journal of Machine Learning Research*, vol. 17, no. 58, pp. 1–47, 2016.

[37] S. Negahban, S. Oh, and D. Shah, "Rank centrality: Ranking from pairwise comparisons," *Operations Research*, vol. 65, no. 1, pp. 266–287, Jan. 2017.

[38] Y. Chen, J. Fan, C. Ma, and K. Wang, "Spectral method and regularized MLE are both optimal for top-K ranking," *The Annals of Statistics*, vol. 47, no. 4, pp. 2204–2235, Aug. 2019.

[39] P. Chen, C. Gao, and A. Y. Zhang, "Partial recovery for top-k ranking: Optimality of MLE and suboptimality of the spectral method," *The Annals of Statistics*, vol. 50, no. 3, pp. 1618–1652, Jun. 2022.

[40] Y. Yang, A. Chen, L. Orecchia, and C. Ma, "Top-K ranking with a monotone adversary," in *Proceedings of the Thirty Seventh Conference on Learning Theory (COLT)*, Edmonton, AB, Canada, Jun. 2024, pp. 5123–5162.

[41] D. P. McGuire and M. L. Davison, "Testing group differences in paired comparisons data," *Psychological Bulletin*, vol. 110, no. 1, pp. 171–182, 1991.

[42] C. Rastogi, S. Balakrishnan, N. B. Shah, and A. Singh, "Two-sample testing on ranked preference data and the role of modeling assumptions," *Journal of Machine Learning Research*, vol. 23, no. 225, pp. 1–48, 2022.

[43] A. Seshadri and J. Ugander, "Fundamental limits of testing the independence of irrelevant alternatives in discrete choice," in *Proceedings of the 2019 ACM Conference on Economics and Computation*, Phoenix, AZ, USA, Jun. 2019, pp. 65–66.

[44] J. Xu, X. Liu, Y. Wu, *et al.*, "ImageReward: Learning and evaluating human preferences for text-to-image generation," in *Proceedings of the Advances in Neural Information Processing Systems 36 (NeurIPS)*, New Orleans, LA, USA, Dec. 2023, pp. 15 903–15 935.

[45] X. Wu, K. Sun, F. Zhu, R. Zhao, and H. Li, "Human preference score: Better aligning text-to-image models with human preference," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, Paris, France, Oct. 2023, pp. 2096–2105.

[46] R. Dawkins, "A threshold model of choice behaviour," *Animal Behaviour*, vol. 17, no. 1, pp. 120–133, Feb. 1969.

[47] S. Fisk, "A very short proof of Cauchy's interlace theorem for eigenvalues of Hermitian matrices," *The American Mathematical Monthly*, vol. 112, no. 2, p. 118, Feb. 2005.

[48] Y. I. Ingster, "Minimax testing of nonparametric hypotheses on a distribution density in the L_p metrics," *Theory of Probability & Its Applications*, vol. 31, no. 2, pp. 333–337, 1987.

[49] T. M. Dunster, "Legendre and related functions," in *NIST Handbook of Mathematical Functions*, F. W. J. Olver, D. M. Lozier, R. F. Boisvert, and C. W. Clark, Eds., New York, NY, USA: Cambridge University Press, 2010.

[50] D. C. Liu and J. Nocedal, "On the limited memory BFGS method for large scale optimization," *Mathematical Programming*, vol. 45, no. 1, pp. 503–528, Aug. 1989.

[51] W. J. Youden, "Index for rating diagnostic tests," *Cancer*, vol. 3, no. 1, pp. 32–35, 1950.

[52] E. L. Lehmann and G. Casella, *Theory of point estimation*, Second. New York, NY, USA: Springer, 1998.

[53] H. Weyl, "Das asymptotische verteilungsgesetz der eigenwerte linearer partieller differentialgleichungen (mit einer anwendung auf die theorie der hohlraumstrahlung)," *Mathematische Annalen*, vol. 71, no. 4, pp. 441–479, Dec. 1912.

[54] J. A. Tropp, "An introduction to matrix concentration inequalities," *Foundations and Trends in Machine Learning*, vol. 8, no. 1–2, pp. 1–230, May 2015.

[55] R. Vershynin, "Introduction to the non-asymptotic analysis of random matrices," in *Compressed Sensing: Theory and Applications*, Y. C. Eldar and G. Kutyniok, Eds., Cambridge, UK: Cambridge University Press, 2012.

[56] M. Rudelson and R. Vershynin, "Hanson-Wright inequality and sub-Gaussian concentration," *Electronic Communications in Probability*, vol. 18, no. 82, pp. 1–9, 2013.

[57] Y. I. Ingster and I. A. Suslina, *Nonparametric goodness-of-fit testing under Gaussian models*, eng, New York, 2003.

[58] M. G. Kendall, "A new measure of rank correlation," *Biometrika*, vol. 30, no. 1-2, pp. 81–93, Jun. 1938.

A Technical Lemmata

In this appendix, we lay out and prove some technical lemmata for completeness.

Lemma 7 ($\hat{\boldsymbol{\theta}}$ is an MLE)**.** *Suppose that the LCGT model as defined in Section 3 is not misspecified, i.e., there exists some $\boldsymbol{\theta}$ such that $p_{ij} = F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})$ for all $(i,j) \in \mathcal{E}$. Then, $\hat{\boldsymbol{\theta}}$ as defined in (3) is the maximum likelihood estimator of $\boldsymbol{\theta}$ in the sense that $\hat{L}$ is proportional to the negative log likelihood function of $\boldsymbol{\theta}$ conditioned on the observations $\hat{p}_{ij}$ (plus some constant), i.e.,*

$$k\,\hat{L}(\boldsymbol{\theta}) + C = -\log f(\hat{P} \mid \boldsymbol{\theta})$$

where f is the conditional probability density of $\hat{P}$ conditioned on $\boldsymbol{\theta}$ and C is some constant that does not depend on $\boldsymbol{\theta}$. Moreover, if H (as defined in (1)) is not singular (i.e., Assumption 6 holds), the MLE satisfies assumptions (A0)–(A2) of [52, Section 6.3].

Proof. Since $\hat{p}_{ij} \sim \text{Binomial}(k, p_{ij})$, and each $\hat{p}_{ij}$ are independent, we have

$$
\begin{aligned}
-\log f(\hat{P} \mid \boldsymbol{\theta}) &= -\log\left(\prod_{(i,j)\in\mathcal{E}} \binom{k}{\hat{p}_{ij}k} (F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij}))^{\hat{p}_{ij}k}(1-F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij}))^{k-\hat{p}_{ij}k}\right) \\
&= -\sum_{(i,j)\in\mathcal{E}} \log\left(\binom{k}{\hat{p}_{ij}k} (F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij}))^{\hat{p}_{ij}k}(1-F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij}))^{k-\hat{p}_{ij}k}\right) \\
&= -\sum_{(i,j)\in\mathcal{E}} \left(\log\binom{k}{\hat{p}_{ij}k} + \hat{p}_{ij}k\log(F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})) + (k-\hat{p}_{ij}k)\log(1-F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij}))\right) \\
&= k\sum_{(i,j)\in\mathcal{E}} \left(-\hat{p}_{ij}\log(F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})) - (1-\hat{p}_{ij})\log(1-F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij}))\right) \underbrace{-\sum_{(i,j)\in\mathcal{E}} \log\binom{k}{\hat{p}_{ij}k}}_{=C} \\
&= k\hat{L}(\boldsymbol{\theta}) + C.
\end{aligned}
$$

Next, we show that assumptions (A0)–(A2) of [52, Section 6.3] hold if H is not singular. First, in order for (A0) to hold, we must have that for all $\boldsymbol{\theta}_1 \neq \boldsymbol{\theta}_2$, $f(\cdot \mid \boldsymbol{\theta}_1) \neq f(\cdot \mid \boldsymbol{\theta}_2)$. Suppose that $f(\cdot \mid \boldsymbol{\theta}_1) = f(\cdot \mid \boldsymbol{\theta}_2)$, i.e., $\boldsymbol{\theta}_1^\mathsf{T}\boldsymbol{x}_{ij} = \boldsymbol{\theta}_2^\mathsf{T}\boldsymbol{x}_{ij}$ for all $(i,j) \in \mathcal{E}$. This implies that $(\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2)H = \mathbf{0}$, i.e., $\boldsymbol{\theta}_1 = \boldsymbol{\theta}_2$ since H is not singular. This establishes (A0). (A1) is clear from the definition. Finally, (A2) holds with respect to k as the number of observations (not n), as the observations can be regarded as k i.i.d. single-comparison observations of the entire graph. □

In particular, results like Lemma 7 imply that as long as regularity conditions such as those in (A)–(D) of [52, Section 6.5] are met, the MLE $\hat{\boldsymbol{\theta}}$ is consistent and its distribution is asymptotically normal as $k \to \infty$, i.e., $\sqrt{k}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})$ is asymptotically normal with zero mean and covariance $\mathcal{I}(\boldsymbol{\theta})^{-1}$, where $\mathcal{I}(\boldsymbol{\theta})$ is the Fisher information matrix [52, Theorem 5.1]. In particular, the following conditions must be established [52, Section 6.5]:

(A) There must exist an open subset ω of Θ containing $\boldsymbol{\theta}$ such that $f(\hat{P} \mid \boldsymbol{\theta})$ is thrice differentiable (with respect to $\boldsymbol{\theta}$) for all $\boldsymbol{\theta} \in \omega$ and almost all $\hat{P}$.

(B) $\mathbb{E}_{\hat{P}}[\nabla_{\boldsymbol{\theta}} \log f(\hat{P} \mid \boldsymbol{\theta}) \mid \boldsymbol{\theta}] = \mathbf{0}$ and the Fisher information matrix satisfies $\mathcal{I}(\boldsymbol{\theta}) \triangleq \mathbb{E}_{\hat{P}}[(\nabla_{\boldsymbol{\theta}} \log f(\hat{P} \mid \boldsymbol{\theta}))(\nabla_{\boldsymbol{\theta}} \log f(\hat{P} \mid \boldsymbol{\theta}))^\mathsf{T} \mid \boldsymbol{\theta}] = \mathbb{E}_{\hat{P}}[-\nabla^2 \log f(\hat{P} \mid \boldsymbol{\theta}) \mid \boldsymbol{\theta}]$.

(C) $\mathcal{I}(\boldsymbol{\theta})$ is positive definite for all $\boldsymbol{\theta} \in \omega$ (cf. Lemma 8).

(D) There exist functions $M_{ijk} \colon \text{supp}(\hat{P}) \to \mathbb{R}$ such that $\left|\frac{\partial^3}{\partial\theta_i\partial\theta_j\partial\theta_k} \log f(\hat{P} \mid \boldsymbol{\theta})\right| \leq M_{ijk}(\hat{P})$ for all $\boldsymbol{\theta} \in \omega$ and $\mathbb{E}_{\hat{P}}[M_{ijk}(\hat{P}) \mid \boldsymbol{\theta}] < \infty$ for all $(i,j,k) \in [p]^3$.

However, whether these conditions hold depends on the choice of F, and specializing on a particular choice of F is beyond the scope of this work.

On the other hand, the regime of $n \to \infty$ is considerably more complex since observations on different edges of the comparison graph are not identically distributed (and heavily depends on H and the graph structure), so the standard analysis à la [52] does not go through.

Lemma 8 (Existence and Uniqueness of $\boldsymbol{\theta}^*$)**.** *Assume H is positive definite (PD), and the Thurstone function $F : \mathbb{R} \to [0,1]$ satisfying Assumption 3, the optimal solution $\boldsymbol{\theta}^* \in \Theta$ of (2) exists and is unique.*

Proof. For existence, recall that $\boldsymbol{\theta}$ is defined on a compact set $\Theta = \{\boldsymbol{\theta} \in \mathbb{R}^p : \|\boldsymbol{\theta}\|_2 \leq b'\}$ for finite b'. Therefore, $L(\boldsymbol{\theta})$ defined in (2) is a continuous function defined on a compact set, which leads to the existence of optimal $\boldsymbol{\theta}^*$.

For uniqueness, we calculate the Hessian of $L(\boldsymbol{\theta})$:

$$\boldsymbol{\nabla}^2 L(\boldsymbol{\theta}) = \sum_{(i,j)\in\mathcal{E}} -\left(p_{ij} \frac{\mathrm{d}^2}{\mathrm{d}t^2} \log(F(t))\Big|_{t=\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij}} + (1-p_{ij}) \frac{\mathrm{d}^2}{\mathrm{d}t^2} \log(1-F(t))\Big|_{t=\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij}} \right) \boldsymbol{x}_{ij}\boldsymbol{x}_{ij}^\mathsf{T}.$$

Since $F(t)$ and $1-F(t) = F(-t)$ are both α-strongly log-concave on $t \in [-2b, 2b]$ and $\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij} \in [-2b, 2b]$ due to Assumptions 4 and 5, we have that $\boldsymbol{\nabla}^2 L(\boldsymbol{\theta}) \succcurlyeq \alpha H$. Since H is positive definite, the Hessian $\boldsymbol{\nabla}^2 L(\boldsymbol{\theta})$ is also positive definite, which implies the optimal $\boldsymbol{\theta}^*$ must be unique. □

Note that the proof can easily be adapted to the BCGT setting if H (as defined in Section 4) is positive definite by simply changing the definition of $\boldsymbol{x}_{ij}$ from $\boldsymbol{x}_{ij} = \boldsymbol{x}_j - \boldsymbol{x}_i$ to $\boldsymbol{x}_{ij} = \boldsymbol{\phi}(\boldsymbol{x}_j) - \boldsymbol{\phi}(\boldsymbol{x}_i)$.

Proposition 9 (Bounding separation distance)**.** *For some constant $c > 0$, we have*

$$c\|P - Q(\boldsymbol{\theta}^*)\|_\mathrm{F} \le \inf_{\boldsymbol{\theta}\in\Theta} \|P - Q(\boldsymbol{\theta})\|_\mathrm{F} \le \|P - Q(\boldsymbol{\theta}^*)]\|_\mathrm{F}.$$

Proof. The right inequality is trivial since $\boldsymbol{\theta}^* \in \Theta$ and thus $\inf_{\boldsymbol{\theta}\in\Theta} \|P - Q(\boldsymbol{\theta})\|_\mathrm{F} \le \|P - Q(\boldsymbol{\theta}^*)\|_\mathrm{F}$. All that remains is to prove the left inequality. We have

$$\begin{aligned}
&\inf_{\boldsymbol{\theta}\in\Theta} \|P - Q(\boldsymbol{\theta})\|_\mathrm{F}^2 \\
&= \inf_{\boldsymbol{\theta}\in\Theta} \sum_{(i,j)\in\mathcal{E}} \left(p_{ij} - F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})\right)^2 \\
&= \inf_{\boldsymbol{\theta}\in\Theta} \sum_{(i,j)\in\mathcal{E}} \frac{\left(p_{ij} - F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})\right)^2}{F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})\left(1 - F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})\right)} \left(F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})\left(1 - F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})\right)\right) \\
&\ge \inf_{\boldsymbol{\theta}\in\Theta} \sum_{(i,j)\in\mathcal{E}} \frac{\left(p_{ij} - F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})\right)^2}{F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})\left(1 - F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij})\right)} \left(F(-2b)(1 - F(-2b))\right) \\
&\ge \inf_{\boldsymbol{\theta}\in\Theta} \left(F(-2b)(1 - F(-2b))\right) D_\mathrm{KL}\left(\mathrm{Bernoulli}(p_{ij}) \,\middle\|\, \mathrm{Bernoulli}(F(\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij}))\right) \\
&= \inf_{\boldsymbol{\theta}\in\Theta} \left(F(-2b)(1 - F(-2b))\right) L(\boldsymbol{\theta}) \\
&= \left(F(-2b)(1 - F(-2b))\right) L(\boldsymbol{\theta}^*) \\
&= \inf_{\boldsymbol{\theta}\in\Theta} \left(F(-2b)(1 - F(-2b))\right) D_\mathrm{KL}\left(\mathrm{Bernoulli}(p_{ij}) \,\middle\|\, \mathrm{Bernoulli}(F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij}))\right) \\
&\ge 2\left(F(-2b)(1 - F(-2b))\right) \sum_{(i,j)\in\mathcal{E}} \left\|\mathrm{Bernoulli}(p_{ij}) - \mathrm{Bernoulli}(F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij}))\right\|_\mathrm{TV}^2 \\
&= 2\left(F(-2b)(1 - F(-2b))\right) \sum_{(i,j)\in\mathcal{E}} \left(p_{ij} - F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})\right)^2 \\
&= 2\left(F(-2b)(1 - F(-2b))\right) \|P - Q(\boldsymbol{\theta}^*)\|_\mathrm{F}^2.
\end{aligned}$$

We use the inequality $\chi^2(P \,\|\, R) \ge D_\mathrm{KL}(P \,\|\, R) \ge 2\|P - R\|_\mathrm{TV}$ (where $\chi^2(\cdot\,\|\,\cdot)$, $D_\mathrm{KL}(\cdot\,\|\,\cdot)$, and $\|\cdot\|_\mathrm{TV}$ denote the chi-squared divergence, Kullback–Leibler (KL) divergence, and total variation (TV) distance, respectively) for discrete random variables P and R where P is absolutely continuous with respect to R. Choosing $c = 2(F(-2b)(1 - F(-2b))) > 0$ completes the proof. This implies that we can estimate $\inf_{\boldsymbol{\theta}\in\boldsymbol{\Theta}} \|P - Q(\boldsymbol{\theta})\|_\mathrm{F}$ by estimating the term $\|P - Q(\boldsymbol{\theta}^*)\|_\mathrm{F}$ instead. Note that the proof is analogous to the proof of [28, Theorem 3.1]. □

Lemma 10 (Bounding s-th moment of the expectation of estimation bound)**.** *The s-th moment of the expectation of $\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2$ with $s \in \mathbb{Z}$ and $s \ge 1$ has the following bound*

$$\mathbb{E}[\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^s] \le d_s \left(\frac{R}{\lambda_{\min}(H)}\sqrt{\frac{|\mathcal{E}|D}{k}}\right)^s$$

where d_s is a positive p-dependent constant. The above bound leads to an estimation bound on score function

$$\mathbb{E}[\|\hat{\boldsymbol{\theta}}^\mathsf{T}\Phi - \boldsymbol{\theta}^{*\mathsf{T}}\Phi\|_2^s] \le d_s \left(\frac{R^2}{\lambda_{\min}(H)}\sqrt{\frac{n|\mathcal{E}|D}{k}}\right)^s$$

Applying probability matrix operator Q yields

$$\mathbb{E}[\|Q(\hat{\boldsymbol{\theta}}) - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^s] \leq c_s \left(\frac{R}{\lambda_{\min}(H)} \sqrt{\frac{|\mathcal{E}|D\lambda_{\max}(H)}{k}} \right)^s$$

where c_s is a positive p-dependent constant.

Proof. Let $a = \frac{k(\alpha\lambda_{\min})^2}{128|\mathcal{E}|R^2\hat{\beta}^2}$, $b = \log(2 \cdot 5^D)$ and $X = \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2$. The boundary condition $au^2 - \log(2 \cdot 5^D) = 0$ yields $u = u_0 = \sqrt{\frac{b}{a}}$. For $s \geq 2$, we have

$$\begin{aligned}
\mathbb{E}[X^s] &= s \int_0^{+\infty} u^{s-1} \mathbb{P}\left(X > u\right) \mathrm{d}u \\
&= s \int_0^{u_0} u^{s-1} \mathbb{P}\left(X > u\right) \mathrm{d}u + s \int_{u_0}^{+\infty} u^{s-1}\, \mathbb{P}\left(X > u\right) \mathrm{d}u \\
&\leq \left(\frac{b}{a}\right)^{\frac{s}{2}} + s e^b \int_{u_0}^{+\infty} u^{s-1} e^{-au^2}\, \mathrm{d}u.
\end{aligned}$$

By substituting u with $t = au^2 - b$, we get

$$\begin{aligned}
\mathbb{E}[X^s] &= \left(\frac{b}{a}\right)^{\frac{s}{2}} + s \int_0^{+\infty} \frac{1}{2} a^{-\frac{s}{2}} (t+b)^{\frac{s-2}{2}} e^{-t}\, \mathrm{d}t \\
&\leq \left(\frac{b}{a}\right)^{\frac{s}{2}} + 2^{\frac{s-2}{2}} \frac{s}{2} a^{-\frac{s}{2}} \int_0^{+\infty} t^{\frac{s-2}{2}} e^{-t} + b^{\frac{s-2}{2}} e^{-t}\, \mathrm{d}t \\
&= \left(\frac{b}{a}\right)^{\frac{s}{2}} + 2^{\frac{s-2}{2}} \frac{s}{2} a^{-\frac{s}{2}} \left(\tau\left(\frac{s}{2}\right) + b^{\frac{s-2}{2}}\right) \\
&= d_s \left(\frac{\hat{\beta} R}{\alpha \lambda_{\min}} \sqrt{\frac{|\mathcal{E}|D}{k}} \right)^s,
\end{aligned}$$

where d_s is some constant. $s = 1$ holds directly from the tail bound in Theorem 4. Absorbing constant $\frac{\hat{\beta}}{\alpha}$ into d_s yields

$$\mathbb{E}[\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^s] \leq d_s \left(\frac{R}{\lambda_{\min}} \sqrt{\frac{|\mathcal{E}|D}{k}} \right)^s.$$

Since $\|\Phi\|_2 \leq \sqrt{n}R$, we can also derive a bound on scores

$$\mathbb{E}[\|\hat{\boldsymbol{\theta}}^{\mathsf{T}}\Phi - \boldsymbol{\theta}^{*\mathsf{T}}\Phi\|_2^s] \leq d_s \left(\frac{R^2}{\lambda_{\min}} \sqrt{\frac{n|\mathcal{E}|D}{k}} \right)^s$$

Further, in similar to the case of linear regression model, we try to bound the Frobenius norm of the pairwise comparison probability matrix. Referring to Equation 9,

$$\begin{aligned}
\mathbb{E}[\|Q(\hat{\boldsymbol{\theta}}) - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^s] &\leq 2^s \beta^s \mathbb{E}[\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H^s] \\
&\leq 2^s \beta^s \lambda_{\max}^{\frac{s}{2}} \mathbb{E}[\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^s] \\
&\leq c_s \left(\frac{R}{\lambda_{\min}} \sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}} \right)^s,
\end{aligned}$$

by absorbing constants into c_s. □

Lemma 11 (Bounds on minimum and maximum eigenvalues for connected graph)**.** *Suppose each $\boldsymbol{x}_i$ are independent copies of some random variable $\boldsymbol{X}$. Let $\Sigma = \mathbb{E}_{\boldsymbol{X}}[\boldsymbol{\phi}(\boldsymbol{X})\boldsymbol{\phi}(\boldsymbol{X})^{\mathsf{T}}]$ and assume the graph $\mathcal{G}$ is connected. Then for $\delta \in (0,1)$, with probability at least $1 - \delta$,*

$$\lambda_{\min}(H) \geq n\lambda_2(L_{\mathcal{G}}) \left(\lambda_{\min}(\Sigma) - 2R^2 \left(\sqrt{\frac{2\log(2D/\delta)}{n}} + \frac{2\log(2D/\delta)}{3n} \right) \right).$$

Similarly, with probability at least $1-\delta$,

$$\lambda_{\max}(H) \le n\lambda_{\max}(L_{\mathcal{G}}) \left(\lambda_{\max}(\Sigma) + 2R^2 \left(\sqrt{\frac{2\log(2D/\delta)}{n}} + \frac{2\log(2D/\delta)}{3n} \right) \right).$$

Proof. Let $B \in \mathbb{R}^{|\mathcal{E}| \times n}$ be the oriented incidence matrix of the comparison graph $\mathcal{G}$, defined

$$B_{(i,j),k} \triangleq \begin{cases} 1, & \text{if } i = k, \\ -1, & \text{if } j = k, \\ 0 & \text{otherwise} \end{cases}$$

for all $(i,j) \in \mathcal{E}$ and $k \in [n]$ (recall that we use the convention that $(i,j) \in \mathcal{E}$ only if $i < j$). Then, we have $X = B\Phi$. Thus, we can rewrite H as

$$H = X^\mathsf{T} X = (B\Phi)^\mathsf{T}(B\Phi) = \Phi^\mathsf{T} B^\mathsf{T} B \Phi = \Phi^\mathsf{T} L \Phi,$$

where L is the Laplacian matrix of $\mathcal{G}$.

We begin by considering the minimum eigenvalue of H:

$$\begin{aligned} \lambda_{\min} &= \lambda_{\min}(\Phi^\mathsf{T} L \Phi) \\ &= \min_{\|\boldsymbol{v}\|_2 = 1} \boldsymbol{v}^\mathsf{T} \Phi^\mathsf{T} L \Phi \boldsymbol{v} \\ &\ge \lambda_2(L_{\mathcal{G}}) \min_{\|\boldsymbol{v}\|_2 = 1} \|P_\perp \Phi \boldsymbol{v}\|_2^2 \qquad \text{where } P_\perp \triangleq I_n - \frac{1}{n}\mathbf{1}_n \mathbf{1}_n^\mathsf{T} \\ &\ge \lambda_2(L_{\mathcal{G}}) \lambda_{\min}(\Phi^\mathsf{T} P_\perp \Phi). \end{aligned}$$

Since $\Phi^\mathsf{T} P_\perp \Phi = \Phi^\mathsf{T}\Phi - \frac{1}{n}\Phi^\mathsf{T}\mathbf{1}_n\mathbf{1}_n^\mathsf{T}\Phi$, Weyl's inequality [53] implies

$$\lambda_{\min}(\Phi^\mathsf{T} P_\perp \Phi) \ge \lambda_{\min}(\Phi^\mathsf{T}\Phi) - \left\| \frac{1}{n}\Phi^\mathsf{T}\mathbf{1}_n\mathbf{1}_n^\mathsf{T}\Phi \right\|_2 = \lambda_{\min}(\Phi^\mathsf{T}\Phi) - \underbrace{\frac{1}{n}\|\Phi^\mathsf{T}\mathbf{1}_n\|_2^2}_{=0 \text{ (Assumption 7)}}.$$

Thus, it suffices to lower-bound $\lambda_{\min}(\Phi^\mathsf{T}\Phi) = n\lambda_{\min}(\hat{\Sigma})$, where

$$\hat{\Sigma} \triangleq \frac{1}{n}\sum_{i=1}^n \boldsymbol{\phi}(\boldsymbol{x}_i)\boldsymbol{\phi}(\boldsymbol{x}_i)^\mathsf{T} = \frac{1}{n}\Phi^\mathsf{T}\Phi.$$

For each $i \in [n]$, define the random variable $Y_i \triangleq \boldsymbol{\phi}(\boldsymbol{x}_i)\boldsymbol{\phi}(\boldsymbol{x}_i)^\mathsf{T} - \Sigma$ with $(1/n)\sum_{i=1}^n Y_i = \hat{\Sigma} - \Sigma$. By definition, Y_i is independent, symmetric, and zero-mean. Since $\|\boldsymbol{\phi}\|_2$ is bounded by R, $\|Y_i\|_2$ is bounded by $2R^2$. Applying the Hermitian matrix Bernstein inequality [54, Theorem 6.6.1] and a union bound on $\|\hat{\Sigma} - \Sigma\|_2$ with $\|\hat{\Sigma} - \Sigma\|_2 = \max\{\lambda_{\max}(\hat{\Sigma} - \Sigma), \lambda_{\max}(\Sigma - \hat{\Sigma})\}$, we get

$$\mathbb{P}(n\|\hat{\Sigma} - \Sigma\|_2 \ge nt) \le 2D \exp\left(-\frac{n^2 t^2/2}{v + 2R^2 nt/3} \right),$$

where $v = \|\sum_{i=1}^n \mathbb{E}[Y_i^2]\|_2 \le 4nR^4$. Therefore, with probability at least $1-\delta$, we have

$$\begin{aligned} \|\hat{\Sigma} - \Sigma\|_2 &\le \frac{\sqrt{2v\log(2D/\delta)}}{n} + \frac{4R^2\log(2D/\delta)}{3n} \\ &\le \frac{2R^2\sqrt{2n\log(2D/\delta)}}{n} + \frac{4R^2\log(2D/\delta)}{3n}. \end{aligned}$$

According to Weyl's inequality [53], $\lambda_{\min}(\hat{\Sigma}) \ge \lambda_{\min}(\Sigma) - \|\hat{\Sigma} - \Sigma\|_2$. Hence, for $\delta > 0$, with probability at least $1-\delta$,

$$\lambda_{\min} \ge n\lambda_2(L_{\mathcal{G}}) \left(\lambda_{\min}(\Sigma) - 2R^2 \left(\sqrt{\frac{2\log(2D/\delta)}{n}} + \frac{2\log(2D/\delta)}{3n} \right) \right).$$

In a similar way, we can upper-bound the maximum eigenvalue with probability at least $1-\delta$,

$$\lambda_{\max} \le n\lambda_{\max}(L_{\mathcal{G}}) \left(\lambda_{\max}(\Sigma) + 2R^2 \left(\sqrt{\frac{2\log(2D/\delta)}{n}} + \frac{2\log(2D/\delta)}{3n} \right) \right).$$

□

Lemma 12 (Bounds on minimum and maximum eigenvalues under perfect matchings)**.** *Suppose that the degree of the comparison graph $\mathcal{G} = ([n], \mathcal{E})$ (where n is even) is 1. So the graph is a perfect matching with n vertex and $n/2$ edges. Suppose each $\boldsymbol{x}_i$ are independent copies of some random variable $\boldsymbol{X}$. Let $\Sigma = \mathbb{E}_{\boldsymbol{X}}[\boldsymbol{\phi}(\boldsymbol{X})\boldsymbol{\phi}(\boldsymbol{X})^\mathsf{T}]$. Then for $\delta > 0$, with probability at least $1 - \delta$,*

$$\lambda_{\min}(H) \geq n\left(2\lambda_{\min}(\Sigma) - 2R^2\left(3\sqrt{\frac{2\log(2D/\delta)}{n}} + \frac{2\log(2D/\delta)}{n}\right)\right).$$

In a similar way, we can upper-bound the maximum eigenvalue

$$\lambda_{\max}(H) \leq n\left(2\lambda_{\max}(\Sigma) + 2R^2\left(3\sqrt{\frac{2\log(2D/\delta)}{n}} + \frac{2\log(2D/\delta)}{n}\right)\right).$$

Proof. When the graph is a perfect matching, due to the independence in $\boldsymbol{x}_i \sim \boldsymbol{X}$, $\boldsymbol{x}_{ij} = \boldsymbol{\phi}(\boldsymbol{x}_j) - \boldsymbol{\phi}(\boldsymbol{x}_i), (i,j) \in \mathcal{E}$ is also independent with each other. So we can treat $Y_{ij} = \boldsymbol{x}_{ij}\boldsymbol{x}_{ij}^\mathsf{T}, (i,j) \in \mathcal{E}$ as random variables. Calculating the expectation, we get

$$\begin{aligned}
\mathbb{E}[Y_{ij}] &= \mathbb{E}\left[\boldsymbol{x}_{ij}\boldsymbol{x}_{ij}^\mathsf{T}\right] \\
&= \mathbb{E}\left[(\boldsymbol{\phi}(\boldsymbol{x}_j) - \boldsymbol{\phi}(\boldsymbol{x}_i))(\boldsymbol{\phi}(\boldsymbol{x}_j) - \boldsymbol{\phi}(\boldsymbol{x}_i))^\mathsf{T}\right] \\
&= 2\mathbb{E}_{\boldsymbol{X}}\left[\boldsymbol{\phi}(\boldsymbol{X})\boldsymbol{\phi}(\boldsymbol{X})^\mathsf{T}\right] - 2\mathbb{E}_{\boldsymbol{X}}\left[\boldsymbol{\phi}(\boldsymbol{X})\right]\mathbb{E}_{\boldsymbol{X}}\left[\boldsymbol{\phi}(\boldsymbol{X})\right]^\mathsf{T} \\
&= 2\Sigma.
\end{aligned}$$

The last equation follows by $\mathbb{E}_{\boldsymbol{X}}[\boldsymbol{\phi}(\boldsymbol{X})] = 0$ due to feature centering in Assumption 7. Define $\hat{Y}_{ij} \triangleq Y_{ij} - 2\Sigma$ and $Y = \sum_{(i,j)\in\mathcal{E}} \hat{Y}_{ij}$. Random variable $\hat{Y}$ is independent, symmetric, and zero-mean, with $\|\hat{Y}\|_2 \leq \|Y_{ij}\|_2 + 2\|\Sigma\|_2 \leq 6R^2$. Analogous to Lemma 11, applying the Hermitian matrix Bernstein inequality [54, Theorem 6.6.1] and a union bound, we get

$$\mathbb{P}(\|Y\|_2 \geq nt) \leq 2D\exp\left(-\frac{n^2t^2/2}{v(Y) + 6R^2nt/3}\right),$$

where $v(Y) = \|\sum_{(i,j)\in\mathcal{E}}^n \mathbb{E}[\hat{Y}_{ij}^2]\|_2 \leq 36nR^4$. Therefore, with probability at least $1 - \delta$ for $\delta > 0$, we have

$$\begin{aligned}
\left\|\frac{1}{n}Y\right\|_2 &= \left\|\frac{1}{n}Y_{ij} - 2\Sigma\right\|_2 \\
&= \left\|\frac{1}{n}H - 2\Sigma\right\|_2 \\
&\leq \frac{6R^2\sqrt{2n\log(2D/\delta)}}{n} + \frac{4R^2\log(2D/\delta)}{n}.
\end{aligned}$$

By Weyl's inequality [53], $\lambda_{\min}(\frac{1}{n}H) \geq 2\lambda_{\min}(\Sigma) - \|\frac{1}{n}Y\|_2$, which implies with probability at least $1 - \delta$,

$$\lambda_{\min} \geq n\left(2\lambda_{\min}(\Sigma) - 2R^2\left(3\sqrt{\frac{2\log(2D/\delta)}{n}} + \frac{2\log(2D/\delta)}{n}\right)\right).$$

In a similar way, we can upper-bound the maximum eigenvalue

$$\lambda_{\max} \leq n\left(2\lambda_{\max}(\Sigma) + 2R^2\left(3\sqrt{\frac{2\log(2D/\delta)}{n}} + \frac{2\log(2D/\delta)}{n}\right)\right).$$

□

Corollary 13. *Good choices of the basis Φ and the distribution of $\boldsymbol{X}$ will lead to a simple structure on Σ. For simplicity, we assume $\mathcal{G}$ is connected here. In the following two cases, we can simplify the high probability bound in Lemma 11 when n is sufficiently large.*

- *Let Φ be a Fourier basis with maximum frequency K, i.e., $\phi_k(x) = e^{2\pi i\langle k,x\rangle}, k = (k_1, k_2, \ldots, k_p) \in [-K, K]^p$. Suppose $\boldsymbol{X}$ follows a uniform distribution on $[0,1]^p$. Then, we have $\Sigma = I_D$, $R = \sqrt{D}$, and with high probability, $\lambda_{\min}(H) \geq cn\lambda_2(L_{\mathcal{G}})$ and $\lambda_{\max}(H) \leq cn\lambda_{\max}(L_{\mathcal{G}})$ for some constant c.*
- *Let Φ the Legendre basis [49] with degree less or equal to q, i.e., for each multi-index $k = (k_1, k_2, \ldots, k_p) \in \mathbb{N}^p$ such that $\sum_{j=1}^p k_j \leq q$, we have*

$$\phi_k(x) = \prod_{j=1}^p \sqrt{2k_j+1}\, P_{k_j}(x_j) = \prod_{j=1}^p \sqrt{2k_j+1}\frac{1}{2^{k_j}k_j!}\left(\frac{\mathrm{d}^{k_j}}{\mathrm{d}x^{k_j}}(x^2-1)^{-k_j}\Big|_{x=x_j}\right).$$

Assume $\boldsymbol{X} \sim$ Uniform $(-1,1)^p$. We then have $\Sigma = I_D$, $R \leq (2q+1)^{\frac{p}{2}}\sqrt{D}$, and with high probability, $\lambda_{\min}(H) \geq cn\lambda_2(L_{\mathcal{G}})$ and $\lambda_{\max}(H) \leq cn\lambda_{\max}(L_{\mathcal{G}})$ for some constant c.

Proof. For the Fourier basis defined in this corollary, we can simplify the inner product as

$$\begin{aligned}
\mathbb{E}\left[\phi_k(\boldsymbol{x})\phi_{k'}(\boldsymbol{x})\right] &= \int_{[0,1]^p} e^{2\pi i\langle k-k',\boldsymbol{x}\rangle} d\boldsymbol{x} \\
&= \prod_{j=1}^p \int_0^1 e^{2\pi i(k_j-k'_j),x_j} dx_j \\
&= \prod_{j=1}^p \mathbb{1}\{k_j = k'_j\} \\
&= \mathbb{1}\{k = k'\}.
\end{aligned}$$

So $\Sigma = \mathbb{E}_{\boldsymbol{X}}[\boldsymbol{\phi}(\boldsymbol{X})\boldsymbol{\phi}(\boldsymbol{X})^\mathsf{T}] = I_D$, $\lambda_{\min}(\Sigma) = \lambda_{\max}(\Sigma) = 1$. Notice that for all $\boldsymbol{x}$, we have $\|\boldsymbol{\phi}(\boldsymbol{X})\|_2 \leq 1 \cdot \sqrt{D}$, so we can choose $R = \sqrt{D}$.

For the Legendre basis [49], we have

$$\begin{aligned}
\mathbb{E}\left[\phi_k(\boldsymbol{x})\phi_{k'}(\boldsymbol{x})\right] &= \prod_{j=1}^p \frac{1}{2}\int_{-1}^1 \sqrt{2k_j+1}\sqrt{2k'_j+1}P_{k_j}(x_j)P_{k'_j}(x_j)dx_j \\
&= \prod_{j=1}^p \frac{1}{2}\sqrt{2k_j+1}\sqrt{2k'_j+1}\frac{2}{2k_j+1}\mathbb{1}\{k_j = k'_j\} \\
&= \prod_{j=1}^p \mathbb{1}\{k_j = k'_j\} \\
&= \mathbb{1}\{k = k'\}.
\end{aligned}$$

The second equality follows from the following property of Legendre bases [49]:

$$\int_{-1}^1 P_{k_j}(x_j)P_{k'_j}(x_j)dx_j = \frac{2}{2k_j+1}\mathbb{1}\{k_j = k'_j\}.$$

So $\Sigma = \mathbb{E}_{\boldsymbol{X}}[\boldsymbol{\phi}(\boldsymbol{X})\boldsymbol{\phi}(\boldsymbol{X})^\mathsf{T}] = I_D$, $\lambda_{\min}(\Sigma) = \lambda_{\max}(\Sigma) = 1$. Notice that for all $\boldsymbol{x}$, we have $\|\boldsymbol{\phi}(\boldsymbol{X})\|_2 \leq 1 \cdot \sqrt{D}$, so we can choose $R = \sqrt{D}$. Notice that the maximum degree is q, we have

$$\forall \boldsymbol{x}, \|\phi_k(\boldsymbol{x})\|_2 \leq \prod_{j=1}^p \sqrt{2k_j+1} \leq (2q+1)^{\frac{p}{2}}.$$

Hence, for all $\boldsymbol{x}$, $\|\boldsymbol{\phi}(\boldsymbol{x})\|_2 \leq \sqrt{D}(2q+1)^{\frac{p}{2}}$, so we can choose $R = \sqrt{D}(2q+1)^{\frac{p}{2}}$.

In both cases, according to Lemma 11, if $n \geq a(R^2\log(2D/\delta))$ for some constant a such that $1 - 2\sqrt{2/a} - 4/(3a) > 0$, with high probability $1-\delta$, $\lambda_{\min} \geq cn\lambda_2(L)$ and $\lambda_{\max} \leq cn\lambda_{\max}(L)$ hold for constant $c = 1 - 2\sqrt{2/a} - 4/(3a)$. □

B Proof of Theorem 1

Proof. The proof uses some ideas from the proofs of [16, Lemma 3.1] and [29, Theorem 3.3].

Since $\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}\in\Theta} \hat{L}(\boldsymbol{\theta})$ and $\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}\in\Theta} L(\boldsymbol{\theta})$, we have that $\nabla L(\boldsymbol{\theta}^*)^\mathsf{T}(\boldsymbol{\theta} - \boldsymbol{\theta}^*) \geq 0$ and $\hat{L}(\hat{\boldsymbol{\theta}}) \leq \hat{L}(\boldsymbol{\theta}^*)$. These properties imply

$$\begin{aligned}\hat{L}(\hat{\boldsymbol{\theta}}) - \hat{L}(\boldsymbol{\theta}^*) - \nabla\hat{L}(\boldsymbol{\theta}^*)^\mathsf{T}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) &\leq -(\nabla\hat{L}(\boldsymbol{\theta}^*)^\mathsf{T} - \nabla L(\boldsymbol{\theta}^*)^\mathsf{T})(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) \\ &\overset{\text{(a)}}{\leq} \|\nabla\hat{L}(\boldsymbol{\theta}^*) - \nabla L(\boldsymbol{\theta}^*)\|_{H^{-1}}\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H, \end{aligned} \tag{8}$$

where (a) holds due to the generalized Cauchy-Schwarz inequality since $\|\cdot\|_{H^{-1}}$ is a dual of $\|\cdot\|_H$. Next, note that the Hessian of $\hat{L}(\boldsymbol{\theta})$ is

$$\nabla^2\hat{L}(\boldsymbol{\theta}) = \sum_{(i,j)\in\mathcal{E}} \underbrace{-\left(\hat{p}_{ij}\,\frac{\mathrm{d}^2}{\mathrm{d}t^2}\log(F(t))\Big|_{t=\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij}} + (1-\hat{p}_{ij})\,\frac{\mathrm{d}^2}{\mathrm{d}t^2}\log(1-F(t))\Big|_{t=\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij}}\right)}_{\text{ⓐ}} \boldsymbol{x}_{ij}\boldsymbol{x}_{ij}^\mathsf{T}.$$

Since $F(-t) = 1 - F(t)$, by assumption, $F(t)$ and $1 - F(t)$ are both α-strongly log-concave on $t \in [-2b, 2b]$. Therefore, ⓐ $\geq \alpha$ for all $t \in [-2b, 2b]$, and since $\boldsymbol{\theta}^\mathsf{T}\boldsymbol{x}_{ij} \in [-2b, 2b]$ due to Assumptions 4 and 5, we have that $\nabla^2\hat{L}(\boldsymbol{\theta}) \succcurlyeq \alpha H$, which implies that $\hat{L}$ is α-strongly convex under the norm $\|\cdot\|_H$. Thus, $\hat{L}(\hat{\boldsymbol{\theta}}) - \hat{L}(\boldsymbol{\theta}^*) - \nabla\hat{L}(\boldsymbol{\theta}^*)^\mathsf{T}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) \geq \frac{\alpha}{2}\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H^2$, which in combination with (8) leads to

$$\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H^2 \leq \frac{4}{\alpha^2}\|\nabla\hat{L}(\boldsymbol{\theta}^*) - \nabla L(\boldsymbol{\theta}^*)\|_{H^{-1}}^2.$$

By some calculations on the gradient, we can obtain

$$\nabla\hat{L}(\boldsymbol{\theta}^*) - \nabla L(\boldsymbol{\theta}^*) = -\sum_{(i,j)\in\mathcal{E}} \frac{(\hat{p}_{ij} - p_{ij})\;F'(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})}{F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})\,(1 - F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij}))}\,\boldsymbol{x}_{ij} = -X^\mathsf{T}\boldsymbol{v}$$

where the vector $\boldsymbol{v} \in \mathbb{R}^{|\mathcal{E}|}$ consists of the elements

$$v_{ij} \triangleq \frac{(\hat{p}_{ij} - p_{ij})\;F'(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})}{F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})\,(1 - F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij}))}$$

for all $(i, j) \in \mathcal{E}$, which are each independent with zero mean. Since $k\,\hat{p}_{ij} \sim \mathrm{Binomial}(k, p_{ij})$ and

$$\sup_{\boldsymbol{\theta}^*\in\Theta, \boldsymbol{x}_{ij}\in[-1,1]^p} \frac{F'(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})}{F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})\,(1 - F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij}))} \leq \hat{\beta},$$

each v_{ij} has variance $\sigma^2(v_{ij}) \leq \hat{\beta}^2/(4k)$ and is sub-Gaussian with $\|v_{ij}\|_{\psi_2}^2 \leq c_1\hat{\beta}^2/k$ [55] for some constant $c_1 > 0$. Furthermore, we have that $XH^{-1}X^\mathsf{T} = X(X^\mathsf{T}X)^{-1}X^\mathsf{T}$ is an orthogonal projection matrix, thus $\|XH^{-1}X^\mathsf{T}\|_2 \leq 1$ and $\|XH^{-1}X^\mathsf{T}\|_\mathrm{F}^2 = \mathrm{trace}(X(X^\mathsf{T}X)^{-1}X^\mathsf{T}) = \mathrm{trace}(I_p) = p$. Hence, we have

$$\begin{aligned}\mathbb{E}\left[\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H^2\right] &\leq \mathbb{E}\left[\frac{4}{\alpha^2}\|\nabla\hat{L}(\boldsymbol{\theta}^*) - \nabla L(\boldsymbol{\theta}^*)\|_{H^{-1}}^2\right] \\ &= \frac{4}{\alpha^2}\mathbb{E}\left[\boldsymbol{v}^\mathsf{T}XH^{-1}X^\mathsf{T}\boldsymbol{v}\right] \\ &\leq \frac{4}{\alpha^2}\sigma^2(v_{ij})\|XH^{-1}X^\mathsf{T}\|_\mathrm{F}^2 \\ &\leq \frac{\hat{\beta}^2 p}{\alpha^2 k}\end{aligned}$$

and

$$\begin{aligned}&\mathbb{P}\left(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H^2 - \frac{\hat{\beta}^2 p}{\alpha^2 k} > t'\right) \\ &\leq \mathbb{P}\left(\frac{4}{\alpha^2}\|\nabla\hat{L}(\boldsymbol{\theta}^*) - \nabla L(\boldsymbol{\theta}^*)\|_{H^{-1}}^2 - \frac{\hat{\beta}^2 p}{\alpha^2 k} > t'\right)\end{aligned}$$

$$
\begin{aligned}
&\leq \mathbb{P}\left(\left|\boldsymbol{v}^\mathsf{T} X H^{-1} X^\mathsf{T} \boldsymbol{v}\right| - \mathbb{E}\left[\left|\boldsymbol{v}^\mathsf{T} X H^{-1} X^\mathsf{T} \boldsymbol{v}\right|\right] > \frac{\alpha^2 t'}{4}\right) \\
&\overset{\text{(a)}}{\leq} 2\exp\left(-c_2 \min\left\{\frac{\alpha^4 t'^2}{16\left(\frac{c_1\hat{\beta}^2}{k}\right)^2 p}, \frac{\alpha^2 t'}{4\left(\frac{c_1\hat{\beta}^2}{k}\right)}\right\}\right)
\end{aligned}
$$

for some constant $c_2 > 0$ where inequality (a) follows from the Hanson-Wright inequality [56].

Plugging in $t' = 4c_1\hat{\beta}^2 t/(\alpha^2 k)$ produces

$$
\begin{aligned}
&\mathbb{P}\left(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H^2 > \frac{\hat{\beta}^2 p}{\alpha^2 k} + \frac{4c_1\hat{\beta}^2 t}{\alpha^2 k}\right) \\
&\leq 2\exp\left(-c_2 \min\left\{\frac{\alpha^4\left(\frac{4c_1\hat{\beta}^2 t}{\alpha^2 k}\right)^2}{16\left(\frac{c_1\hat{\beta}^2}{k}\right)^2 p}, \frac{\alpha^2 \frac{4c_1\hat{\beta}^2 t}{\alpha^2 k}}{4\left(\frac{c_1\hat{\beta}^2}{k}\right)}\right\}\right) \\
&= 2\exp\left(-c_2 \min\left\{\frac{t^2}{p}, t\right\}\right) \\
&\leq 2e^{-c_2 t}
\end{aligned}
$$

for all $t \geq p$. Thus, for all $\delta \in (0, 2e^{-c_2 p}]$, with probability at least $1-\delta$,

$$
\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H \leq \frac{\hat{\beta}}{\alpha}\sqrt{\frac{p}{k} + \frac{4c_1 \log(2/\delta)}{c_2 k}}.
$$

Furthermore, by plugging in $t' = ((4c_1+1)t - 1)\hat{\beta}^2 p/(\alpha^2 k)$ instead, we may produce the following tail bound which aids in the analysis of the moments.

$$
\begin{aligned}
&\mathbb{P}\left(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H^2 > t\frac{4c_1\hat{\beta}^2 p}{\alpha^2 k}\right) \\
&\leq 2\exp\left(-c_2 \min\left\{\frac{\alpha^4\left(\frac{\hat{\beta}^2 p}{\alpha^2 k}((4c_1+1)t-1)\right)^2}{16\left(\frac{c_1\hat{\beta}^2}{k}\right)^2 p}, \frac{\alpha^2 \frac{\hat{\beta}^2 p}{\alpha^2 k}((4c_1+1)t-1)}{4\left(\frac{c_1\hat{\beta}^2}{k}\right)}\right\}\right) \\
&= 2\exp\left(-c_2 p \min\left\{\frac{((4c_1+1)t-1)^2}{16c_1^2}, \frac{(4c_1+1)t-1}{4c_1}\right\}\right) \\
&= 2\exp\left(-c_2 p\left(\left(1+\frac{1}{4c_1}\right)t - \frac{1}{4c_1}\right)\right) \\
&\leq 2e^{-c_2 p t}
\end{aligned}
$$

for all $t \geq 1$, so if we set $C \triangleq \sqrt{4c_1\hat{\beta}^2 p/(\alpha^2 k)}$ for convenience,

$$
\begin{aligned}
\mathbb{E}\left[\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H^s\right] &= \int_0^\infty \mathbb{P}\left(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H^s > t\right) \mathrm{d}t \\
&= \int_0^\infty s t^{s-1} \mathbb{P}\left(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H > t\right) \mathrm{d}t \\
&= \int_0^C s t^{s-1} \mathbb{P}\left(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H > t\right) \mathrm{d}t + \int_C^\infty s t^{s-1} \mathbb{P}\left(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H > t\right) \mathrm{d}t \\
&\leq \int_0^C s t^{s-1} \,\mathrm{d}t + C\int_1^\infty s(Ct)^{s-1} \mathbb{P}\left(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H > Ct\right) \mathrm{d}t \\
&\leq C^s + C^s \int_1^\infty 2s t^{s-1} \exp(-c_2 p\sqrt{t}) \,\mathrm{d}t
\end{aligned}
$$

$$\leq C^s + C^s \int_0^\infty 4st^{2s-1} \exp(-c_2 pt)\, \mathrm{d}t$$
$$= C^s \left(1 + 2\frac{(2s)!}{(c_2 p)^{2s}}\right).$$

Plugging C back in completes the proof of the moment bound. □

C Proof of Theorem 2

Proof. We first calculate the moment bounds on the probability matrix error $\|Q(\hat{\boldsymbol{\theta}}) - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}$ using the estimation bound in Theorem 1. First, since F is β-Lipchitz, we have

$$\sum_{(i,j)\in\mathcal{E}} \left(F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij}) - F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})\right)^2 \leq \beta^2 \sum_{(i,j)\in\mathcal{E}} \left(\boldsymbol{x}_{ij}^\mathsf{T}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\right)^2 = \beta^2 \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H^2.$$

Consider the s-th moment $\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H$ and take the $\frac{s}{2}$-th power on both sides to obtain

$$\mathbb{E}[\|Q(\hat{\boldsymbol{\theta}}) - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^s] \leq 2^s \beta^s \mathbb{E}[\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H^s] \leq c_s \left(\frac{p}{k}\right)^{\frac{s}{2}}. \tag{9}$$

We use the test statistic on the split dataset $\mathcal{D}_1$ and $\mathcal{D}$. By rewriting (6), we get

$$T = \frac{1}{|\mathcal{E}|} \sum_{(i,j)\in\mathcal{E}} \frac{Z_{ij}(Z_{ij}-1)}{k(k-1)} + F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})^2 - 2F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})\frac{Z_{ij}}{k},$$

where $\hat{\boldsymbol{\theta}}$ is estimated on $\mathcal{D}_1$ and Z_{ij} are computed based on $\mathcal{D}$. We can derive

$$\begin{aligned}
\mathbb{E}[T \mid \mathcal{D}_1] &= \frac{1}{|\mathcal{E}|} \sum_{(i,j)\in\mathcal{E}} \mathbb{E}\left[\frac{Z_{ij}(Z_{ij}-1)}{k(k-1)}\middle|\mathcal{D}_1\right] + F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})^2 - 2\mathbb{E}\left[\frac{Z_{ij}}{k}\middle|\mathcal{D}_1\right] F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij}) \\
&= \frac{1}{|\mathcal{E}|} \sum_{(i,j)\in\mathcal{E}} p_{ij}^2 + F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})^2 - 2p_{ij}F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij}) \\
&= \frac{1}{|\mathcal{E}|} \sum_{(i,j)\in\mathcal{E}} \left(p_{ij} - F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})\right)^2.
\end{aligned}$$

Then the expectation of T can be bounded by

$$\begin{aligned}
\mathbb{E}[T] &= \mathbb{E}\left[\mathbb{E}[T \mid \mathcal{D}_1]\right] \\
&= \frac{1}{|\mathcal{E}|} \sum_{(i,j)\in\mathcal{E}} \mathbb{E}\left[\left(p_{ij} - F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})\right)^2\right] \\
&= \frac{1}{|\mathcal{E}|} \sum_{(i,j)\in\mathcal{E}} \left(p_{ij} - F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})\right)^2 + \mathbb{E}\left[\left(F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij}) - F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})\right)^2\right] \\
&\qquad + 2\mathbb{E}\left[\left(p_{ij} - F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})\right)\left(F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij}) - F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})\right)\right] \\
&\leq \frac{1}{|\mathcal{E}|}\Big(\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^2 + \mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}^2\right] \\
&\qquad + 2\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}\right]\Big).
\end{aligned}$$

According to (9), we have $\mathbb{E}[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}^2] \leq c_2(p/k)$ and $\mathbb{E}[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}] \leq c_1\sqrt{p/k}$, which imply

$$\mathbb{E}[T] \leq \frac{1}{|\mathcal{E}|}\left(\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^2 + c_2\left(\frac{p}{k}\right) + 2c_1\sqrt{\frac{p}{k}}\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}\right).$$

Under the null hypothesis H_0, since $\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}} = 0$, we get $\mathbb{E}_{H_0}[T] \leq c_2(p/(k|\mathcal{E}|))$. Under the alternative hypothesis H_1, since $\mathbb{E}[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}^2] \geq 0$ and $\mathbb{E}[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}] \leq c_1\sqrt{p/k}$, we obtain

$$\mathbb{E}_{H_1}[T] \geq \frac{1}{|\mathcal{E}|}\left(\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^2 - 2c_1\sqrt{\frac{p}{k}}\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}\right).$$

To bound the variance of T, we use the law of total variance to decompose it as

$$\mathrm{var}(T) = \mathbb{E}[\mathrm{var}(T) \mid \mathcal{D}_1] + \mathrm{var}(\mathbb{E}[T \mid \mathcal{D}_1]). \tag{10}$$

The first term in (10) is

$$\begin{aligned}
\mathrm{var}(T \mid \mathcal{D}_1) &= \frac{1}{|\mathcal{E}|^2} \sum_{(i,j)\in\mathcal{E}} \mathrm{var}\left(\frac{Z_{ij}(Z_{ij}-1)}{k(k-1)}\right) + 4F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})^2 \mathrm{var}\left(\frac{Z_{ij}}{k}\right) \\
&\qquad - 4F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})\left(\frac{\mathbb{E}[Z_{ij}^2(Z_{ij}-1)]}{k^2(k-1)} - \frac{\mathbb{E}[Z_{ij}(Z_{ij}-1)]\mathbb{E}[Z_{ij}]}{k^2(k-1)}\right) \\
&= \frac{1}{|\mathcal{E}|^2} \sum_{(i,j)\in\mathcal{E}} \frac{2p_{ij}^2 + 4(k-2)p_{ij}^3 + (6-4k)p_{ij}^4}{k(k-1)} + \frac{4F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})^2 p_{ij}(1-p_{ij})}{k} \\
&\qquad - 4F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})\frac{2(p_{ij}^2 - p_{ij}^3)}{k} \\
&= \frac{1}{|\mathcal{E}|^2} \sum_{(i,j)\in\mathcal{E}} \frac{2p_{ij}^2(1-p_{ij})^2}{k(k-1)} + \frac{4p_{ij}(1-p_{ij})}{k}\left(p_{ij} - F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})\right)^2.
\end{aligned}$$

The second equality is obtained by plugging in

$$\mathbb{E}\left[\frac{Z_{ij}(Z_{ij}-1)}{k(k-1)}\right] = p_{ij}^2, \quad \mathbb{E}\left[\frac{Z_{ij}}{k}\right] = p_{ij}, \quad \mathrm{Cov}\left(\frac{Z_{ij}(Z_{ij}-1)}{k(k-1)}, \frac{Z_{ij}}{k}\right) = \frac{2p_{ij}^2 - 2p_{ij}^3}{k}.$$

Taking expectations on both sides, we get

$$\begin{aligned}
\mathbb{E}[\mathrm{var}(T \mid D_1)] &= \frac{1}{|\mathcal{E}|^2} \sum_{(i,j)\in\mathcal{E}} \frac{2p_{ij}^2(1-p_{ij})^2}{k(k-1)} + \frac{4p_{ij}(1-p_{ij})}{k}\mathbb{E}\left[\left(p_{ij} - F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})\right)^2\right] \\
&\le \frac{1}{8k(k-1)|\mathcal{E}|} + \frac{1}{k|\mathcal{E}|^2}\mathbb{E}\left[\|P - Q(\hat{\boldsymbol{\theta}})\|_\mathrm{F}^2\right] \\
&\le \frac{1}{8k(k-1)|\mathcal{E}|} + \frac{1}{k|\mathcal{E}|^2}\mathbb{E}\left[(\|P - Q(\boldsymbol{\theta}^*)\|_\mathrm{F} + \|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_\mathrm{F})^2\right] \\
&\le \frac{1}{8k(k-1)|\mathcal{E}|} + \frac{1}{k|\mathcal{E}|^2}\left(\|P - Q(\boldsymbol{\theta}^*)\|_\mathrm{F} + \sqrt{\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_\mathrm{F}^2\right]}\right)^2 \\
&\le \frac{1}{8k(k-1)|\mathcal{E}|} + \frac{1}{k|\mathcal{E}|^2}\left(\|P - Q(\boldsymbol{\theta}^*)\|_\mathrm{F} + \sqrt{c_2\frac{p}{k}}\right)^2.
\end{aligned}$$

Next, the second term $\mathrm{var}(\mathbb{E}[T \mid \mathcal{D}_1])$ in (10) can be bounded by

$$\begin{aligned}
\mathrm{var}(\mathbb{E}[T \mid \mathcal{D}_1]) &= \frac{1}{|\mathcal{E}|^2}\mathrm{var}(\|P - Q(\hat{\boldsymbol{\theta}}))\|_\mathrm{F}^2) \\
&= \frac{1}{|\mathcal{E}|^2}\Big(\mathbb{E}[\|P - Q(\hat{\boldsymbol{\theta}}))\|_\mathrm{F}^4] - \mathbb{E}[\|P - Q(\hat{\boldsymbol{\theta}}))\|_\mathrm{F}^2]^2\Big) \\
&\le \frac{1}{|\mathcal{E}|^2}\Big(\mathbb{E}[(\|P - Q(\boldsymbol{\theta}^*)\|_\mathrm{F} + \|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}}))\|_\mathrm{F})^4] \\
&\qquad - \mathbb{E}[(\|P - Q(\boldsymbol{\theta}^*)\|_\mathrm{F} - \|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}}))\|_\mathrm{F})^2]^2\Big).
\end{aligned}$$

Under H_0, $\|P - Q(\boldsymbol{\theta}^*)\|_\mathrm{F} = 0$ implies $|\mathcal{E}|^2\mathrm{var}_{H_0}(\mathbb{E}[T \mid \mathcal{D}_1]) \le \mathbb{E}[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}}))\|_\mathrm{F}^4] \le c_4(p/k)^2$.
Under H_1, we have

$$\begin{aligned}
&\mathrm{var}_{H_1}(\mathbb{E}[T \mid \mathcal{D}_1]) \\
&\le \frac{1}{|\mathcal{E}|^2}\Big(\|P - Q(\boldsymbol{\theta}^*)\|_\mathrm{F}^4 + 4\|P - Q(\boldsymbol{\theta}^*)\|_\mathrm{F}^3\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_\mathrm{F}\right] \\
&\quad + 6\|P - Q(\boldsymbol{\theta}^*)\|_\mathrm{F}^2\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_\mathrm{F}^2\right] + 4\|P - Q(\boldsymbol{\theta}^*)\|_\mathrm{F}\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_\mathrm{F}^3\right]
\end{aligned}$$

$$
\begin{aligned}
&+ \mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}^4\right] - \Big(\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^2 + \mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}^2\right] \\
&- 2\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}\right]\Big)^2\Big) \\
&= \frac{1}{|\mathcal{E}|^2}\Big(8\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^3\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}\right] \\
&+ 4\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^2\Big(\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}^2\right] - \mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}\right]^2\Big) \\
&+ 4\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}\Big(\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}^3\right] \\
&+ \mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}^2\right]\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}\right]\Big) \\
&+ \mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}^4\right] - \mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}^2\right]^2\Big) \\
&\leq \frac{1}{|\mathcal{E}|^2}\Big(8c_1\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^3\sqrt{\frac{p}{k}} + 4c_2\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^2\left(\frac{p}{k}\right) \\
&+ 4(c_3 + c_2c_1)\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}\left(\frac{p}{k}\right)^{3/2} + c_4\left(\frac{p}{k}\right)^2\Big).
\end{aligned}
$$

Combining the two terms, we obtain upper bounds on variance under H_0 and H_1:

$$
\mathrm{var}_{H_0}(T) \leq \frac{1}{8k(k-1)|\mathcal{E}|^2} + c_2\frac{1}{|\mathcal{E}|^2}\left(\left(\frac{p}{k^2}\right) + c_4\left(\frac{p}{k}\right)^2\right),
$$

$$
\begin{aligned}
\mathrm{var}_{H_1}(T) \leq \frac{1}{8k(k-1)|\mathcal{E}|} + \frac{1}{|\mathcal{E}|^2}\Big(&\frac{1}{k}\left(\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}} + \sqrt{c_2\frac{p}{k}}\right)^2 + 8c_1\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^3\sqrt{\frac{p}{k}} \\
&+ 4c_2\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^2\left(\frac{p}{k}\right) + 4(c_3 + c_2c_1)\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}\left(\frac{p}{k}\right)^{3/2} \\
&+ c_4\left(\frac{p}{k}\right)^2\Big).
\end{aligned}
$$

We now can define the decision rule by choosing H_1 if

$$
T > \bar{\gamma}\left(\frac{\sqrt{|\mathcal{E}|} + p}{k|\mathcal{E}|}\right) + c_2\frac{p}{k|\mathcal{E}|}.
$$

By this definition, the Type I error probability can be bounded as

$$
\begin{aligned}
\mathbb{P}_{H_0}\left(T > \bar{\gamma}\left(\frac{\sqrt{|\mathcal{E}|} + p}{k|\mathcal{E}|}\right) + c_2\frac{p}{k|\mathcal{E}|}\right) &\leq \mathbb{P}_{H_0}\left(T - \mathbb{E}_{H_0}[T] > \bar{\gamma}\left(\frac{\sqrt{|\mathcal{E}|} + p}{k|\mathcal{E}|}\right)\right) \\
&\leq \frac{\mathrm{var}_{H_0}(T)}{\mathrm{var}_{H_0}(T) + \bar{\gamma}^2\left(\frac{\sqrt{|\mathcal{E}|}+p}{k|\mathcal{E}|}\right)^2} \\
&\leq \frac{\frac{|\mathcal{E}|}{8k(k-1)} + c_2\left(\frac{p}{k^2}\right) + c_4\left(\frac{p}{k}\right)^2}{\frac{|\mathcal{E}|}{8k(k-1)} + c_2\left(\frac{p}{k^2}\right) + c_4\left(\frac{p}{k}\right)^2 + \bar{\gamma}^2\left(\frac{\sqrt{|\mathcal{E}|}+p}{k}\right)^2} \\
&\leq \frac{\frac{|\mathcal{E}|}{4} + c_2\,p + c_4\,p^2}{\frac{|\mathcal{E}|}{4} + c_2\,p + c_4\,p^2 + \bar{\gamma}^2\left(\sqrt{|\mathcal{E}|} + p\right)^2} \leq \frac{1}{4},
\end{aligned}
$$

where the last inequality can be achieved by large enough $\bar{\gamma}$ where $\bar{\gamma}^2 \geq 1 + 4c_2 + 4c_4$ will suffice. For the Type II error, we have

$$
\mathbb{P}_{H_1}\left(T \leq \bar{\gamma}\left(\frac{\sqrt{|\mathcal{E}|} + p}{k|\mathcal{E}|}\right) + c_2\frac{p}{k|\mathcal{E}|}\right) \leq \mathbb{P}_{H_1}\left(T - \mathbb{E}_{H_1}[T] \leq \bar{\gamma}\left(\frac{\sqrt{|\mathcal{E}|} + p}{k|\mathcal{E}|}\right) + c_2\frac{p}{k|\mathcal{E}|}\right.
$$

$$- \frac{1}{|\mathcal{E}|}\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^2 + 2c_1\frac{1}{|\mathcal{E}|}\sqrt{\frac{p}{k}}\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}\Bigg) \leq \frac{\mathrm{var}_{H_1}(T)}{\mathrm{var}_{H_1}(T) + (A^2 - \Delta)^2},$$

where $A \triangleq (1/\sqrt{|\mathcal{E}|})\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}$ and

$$\Delta \triangleq \bar{\gamma}\left(\frac{\sqrt{|\mathcal{E}|} + p}{k|\mathcal{E}|}\right) + c_2\frac{p}{k|\mathcal{E}|} + 2c_1\frac{1}{\sqrt{|\mathcal{E}|}}\sqrt{\frac{p}{k}}A.$$

In order to ensure that the error is less than or equal to $\frac{1}{4}$, it suffices to make $A^2 \geq 2\sqrt{\mathrm{var}_{H_1}(T)} + \Delta$.

$$A^2 \geq \frac{2}{|\mathcal{E}|}\Bigg(\frac{\sqrt{|\mathcal{E}|}}{2k} + \frac{1}{\sqrt{k}}\left(A\sqrt{|\mathcal{E}|} + \sqrt{c_2\frac{p}{k}}\right) + \sqrt{8c_1}\left(A\sqrt{|\mathcal{E}|}\right)^{\frac{3}{2}}\left(\frac{p}{k}\right)^{\frac{1}{4}} + 2\sqrt{c_2}A\sqrt{|\mathcal{E}|}\sqrt{\frac{p}{k}} + 2\sqrt{(c_3 + c_2c_1)}\sqrt{A\sqrt{|\mathcal{E}|}}\left(\frac{p}{k}\right)^{3/4} + \sqrt{c_4}\left(\frac{p}{k}\right)\Bigg) + \bar{\gamma}\left(\frac{\sqrt{|\mathcal{E}|} + p}{k|\mathcal{E}|}\right) + c_2\frac{p}{k|\mathcal{E}|} + 2c_1\frac{1}{\sqrt{|\mathcal{E}|}}\sqrt{\frac{p}{k}}A.$$

If $A = c_0\sqrt{(\sqrt{|\mathcal{E}|} + p)/(k|\mathcal{E}|)}$ where c_0 is some constant that satisfies

$$c_0^2 \geq 2\left(\frac{1}{2} + (c_0 + \sqrt{c_2})\frac{1}{\sqrt{\sqrt{|\mathcal{E}|} + p}} + \sqrt{8c_1}c_0^{\frac{3}{2}} + 2\sqrt{c_2}c_0 + 2\sqrt{(c_3 + c_2c_1)c_0} + \sqrt{c_4}\right) + \bar{\gamma} + c_2 + 2c_1c_0,$$

we have $A^2 \geq 2\sqrt{\mathrm{var}_{H_1}(T)} + \Delta$. Then, the minimax risk is less or equal to $\frac{1}{2}$. According to Proposition 9, $\inf_{\boldsymbol{\theta}\in\Theta}\|P - Q(\boldsymbol{\theta})\|_{\mathrm{F}} = \Theta(\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}})$. Thus, we may conclude that the critical threshold is bounded by

$$\epsilon_c \leq c\sqrt{\frac{\sqrt{|\mathcal{E}|} + p}{|\mathcal{E}|k}},$$

completing the proof. Note that this proof uses some ideas from [28, Theorem 3.2]. □

D Proof of Proposition 3

Proof. The proof utilizes ideas from [43] and [29, Theorem 3.4] but changes the scaling of the critical threshold to fit our minimax testing framework. Since the comparison graph $\mathcal{G}$ is super-Eulerian, we can find a largest Eulerian spanning sub-graph $\tilde{\mathcal{G}}$ of $\mathcal{G}$. Let P be the pairwise comparison model with equal latent scores for all items, i.e., $p_{ij} = \frac{1}{2}$ for all $(i,j) \in \mathcal{G}$. Here we consider a set of perturbations $\mathcal{B}$ with perturbation matrices $B \in \{-1,0,1\}^{n\times n}$, and the corresponding perturbed pairwise comparison model $R = P + \eta B$ for some constant $\eta \in (0, 1/2 - a]$. Next we demonstrate that the set of perturbations $\mathcal{B}$ satisfies certain properties that are necessary for our analysis. We let $\mathcal{C}$ denote the cycle decomposition of graph $\tilde{\mathcal{G}}$ (which may not be unique) and we define

$$\mathcal{B} \triangleq \Bigg\{B : b_{ij} \in \{-1,1\} \text{ for } (i,j) \in \mathcal{E}_{\tilde{\mathcal{G}}},\ b_{ij} = 0 \text{ for } (i,j) \notin \mathcal{E}_{\tilde{\mathcal{G}}},$$
$$b_{ij} + b_{ji} = 0,\ \sum_{j:(i,j)\in\mathcal{E}_{\tilde{\mathcal{G}}}} b_{ij} = 0,\ \forall i \in [n],$$
$$\forall \sigma \in \mathcal{C}, \exists s \in \{1,-1\},\ s.t.\ b_{v_1v_2} = s, \forall (v_1,v_2) \in \sigma\Bigg\}, \tag{11}$$

where b_{ij} are the entries of B and $\mathcal{E}_{\tilde{\mathcal{G}}}$ is the edge set of $\tilde{\mathcal{G}}$. As shown in (11), for each cycle in $\mathcal{C}$, we orient the edges of the cycle in one direction, clockwise or counterclockwise. We first show that for all $B \in \mathcal{B}$ the model R belongs to $P_1(\epsilon)$ by choosing η appropriately. We will prove that in our LCGT model, setting all scores to be zero is the unique optimal solution to approximate the perturbed pairwise comparison matrix R. Recall that the optimal parameter vector $\boldsymbol{\theta}^*$ is defined as

$$\begin{aligned}\boldsymbol{\theta}^* &= \arg\min_{\boldsymbol{\theta}\in\Theta} L(\boldsymbol{\theta}) \\ &= \arg\min_{\boldsymbol{\theta}\in\Theta} \sum_{(i,j)\in\mathcal{E}_{\mathcal{G}}} -\left(\frac{1}{2}+\eta b_{ij}\right)\log\left(F(\boldsymbol{\theta}^{\mathsf{T}}\boldsymbol{x}_{ij})\right) \\ &\qquad\qquad -\left(\frac{1}{2}-\eta b_{ij}\right)\log\left(1-F(\boldsymbol{\theta}^{\mathsf{T}}\boldsymbol{x}_{ij})\right).\end{aligned}$$

The gradient is given by

$$\nabla L(\boldsymbol{\theta}) = -\sum_{(i,j)\in\mathcal{E}_{\mathcal{G}}} \frac{\left(\frac{1}{2}+\eta\, b_{ij} - F(\boldsymbol{\theta}^{\mathsf{T}}\boldsymbol{x}_{ij})\right) F'(\boldsymbol{\theta}^{\mathsf{T}}\boldsymbol{x}_{ij})}{F(\boldsymbol{\theta}^{\mathsf{T}}\boldsymbol{x}_{ij})\left(1-F(\boldsymbol{\theta}^{\mathsf{T}}\boldsymbol{x}_{ij})\right)}\boldsymbol{x}_{ij}.$$

Notice that if $\boldsymbol{\theta} = \mathbf{0}_n$, we obtain

$$\begin{aligned}\nabla L(\boldsymbol{\theta})|_{\boldsymbol{\theta}=\mathbf{0}_n} &= -4\eta F'(0)\sum_{(i,j)\in\mathcal{E}_{\mathcal{G}}} b_{ij}(\boldsymbol{x}_i-\boldsymbol{x}_j) \\ &= -4\eta F'(0)\left(\sum_i\left(\sum_{j:(i,j)\in\mathcal{E}_{\mathcal{G}}} b_{ij}\right)\boldsymbol{x}_i - \sum_j\left(\sum_{i:(i,j)\in\mathcal{E}_{\mathcal{G}}} b_{ij}\right)\boldsymbol{x}_j\right) \\ &= -4\eta F'(0)\left(\sum_i\left(\sum_{j:(i,j)\in\mathcal{E}_{\tilde{\mathcal{G}}}} b_{ij}\right)\boldsymbol{x}_i - \sum_j\left(\sum_{i:(i,j)\in\mathcal{E}_{\tilde{\mathcal{G}}}} b_{ij}\right)\boldsymbol{x}_j\right).\end{aligned}$$

Due to how we designed the perturbations, $\sum_{j:(i,j)\in\mathcal{E}_{\tilde{\mathcal{G}}}} b_{ij} = \sum_{i:(i,j)\in\mathcal{E}_{\tilde{\mathcal{G}}}} b_{ij} = 0$, which implies $\nabla L(\boldsymbol{\theta})|_{\boldsymbol{\theta}=\mathbf{0}_n} = \mathbf{0}_n$. Since the loss function $L(\boldsymbol{\theta})$ is strongly convex with respect to $\boldsymbol{\theta}$ (see Lemma 8), $\boldsymbol{\theta}^* = \mathbf{0}_n$ is the global optimal solution for the perturbed pairwise comparison matrix. Hence, the score function $\boldsymbol{w}^* = X\boldsymbol{\theta}^* = \mathbf{0}_n$ is the global optimal solution. Proposition 9 yields for some constant $c > 0$,

$$\inf_{\boldsymbol{\theta}\in\Theta}\|R - Q(\boldsymbol{\theta})\|_{\mathrm{F}} \geq c\|R - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}} = c\|P+\eta B - Q(\mathbf{0}_n))\|_{\mathrm{F}} = c\eta\sqrt{|\mathcal{E}_{\tilde{\mathcal{G}}}|}.$$

Thus, we may set η such that $c\eta\sqrt{|\mathcal{E}_{\tilde{\mathcal{G}}}|} = \sqrt{|\mathcal{E}_{\mathcal{G}}|}\epsilon$, which implies every perturbed pairwise comparison matrix R belongs to $P_1(\epsilon)$. Next, we extract the following lemma from the proof of [29, Theorem 3.4]:

Lemma 14. *The set of perturbations $\mathcal{B}$ is defined in (11). Let $\mathbb{P}_0$ be the probability measure corresponding to pairwise comparison model P and let $\mathbb{P}_{\mathcal{B}}$ be the probability measure of the mixture distribution obtained by first sampling B uniformly from $\mathcal{B}$, and then generating the corresponding perturbed model $R = P + \eta B$. We have*

$$\chi^2(\mathbb{P}_{\mathcal{B}}\|\mathbb{P}_0) + 1 \leq \exp\left(128\eta^4 k^2\sum_{\sigma\in\mathcal{C}}|\sigma|^2\right).$$

Le Cam's testing bound [57] states that

$$\mathcal{R}(\mathcal{G},\epsilon) \geq 1 - \|\mathbb{P}_0 - \mathbb{P}_{\mathcal{B}}\|_{\mathrm{TV}} \geq 1 - \frac{1}{2}\sqrt{\chi^2(\mathbb{P}_{\mathcal{B}}\|\mathbb{P}_0)}.$$

Combining Lemma 14 and the relationship $c\eta\sqrt{|\mathcal{E}_{\tilde{\mathcal{G}}}|} = \sqrt{|\mathcal{E}_{\mathcal{G}}|}\epsilon$, we obtain

$$\mathcal{R}(\mathcal{G},\epsilon) \geq 1 - \frac{1}{2}\sqrt{\exp\left(128k^2\frac{\epsilon^4|\mathcal{E}_{\mathcal{G}}|^2}{c^4|\mathcal{E}_{\tilde{\mathcal{G}}}|^2}\sum_{\sigma\in\mathcal{C}}|\sigma|^2\right) - 1}.$$

Furthermore, if we consider the maximal spanning Eulerian subgraph $\tilde{\mathcal{G}}$ of $\mathcal{G}$, the Eulerian property implies $|\mathcal{E}_{\tilde{\mathcal{G}}}| \geq n$ and the maximality of $\tilde{\mathcal{G}}$ implies that there is no cycle in the excluded edges $\mathcal{E}_{\mathcal{G}} \setminus \mathcal{E}_{\tilde{\mathcal{G}}}$. Therefore, the ratio $|\mathcal{E}_{\mathcal{G}}|/|\mathcal{E}_{\tilde{\mathcal{G}}}|$ is naturally bounded by $1 + (n-1)/n < 2$. Therefore, we are able to derive

$$\mathcal{R}(\mathcal{G}, \epsilon) \geq 1 - \frac{1}{2}\sqrt{\exp\left(\frac{512}{c^4}k^2\epsilon^4 \sum_{\sigma \in \mathcal{C}} |\sigma|^2\right) - 1},$$

which yields the lower bound by absorbing the constants into c.

For the special case where the graph $\mathcal{G}$ is complete, if n is odd, the graph is Eulerian since every vertex has an even degree. If n is even, the graph is super-Eulerian since we can remove n edges to get a spanning Eulerian subgraph. Note that the ratio between number of edges in the original graph and the spanning Eulerian subgraph satisfies $\frac{1}{2}n(n-1)/(\frac{1}{2}n(n-1) - n) \leq t$ for some constant t. Moreover, for a complete graph, according to [29, Proposition 3.5], we have $\sum_{\sigma \in \mathcal{C}} |\sigma|^2 = O(n^2)$. If $\epsilon = O(1/\sqrt{nk})$, the risk $\mathcal{R}(\mathcal{G}, \epsilon)$ is lower bounded by some constant, which implies that the testing is impossible. So we conclude for complete graphs, $\epsilon_c = \Omega(1/\sqrt{nk})$. □

E Proof of Theorem 4

Proof. Recall that the gradient and the Hessian of the loss are respectively

$$\nabla \hat{L}(\boldsymbol{\theta}^*) - \nabla L(\boldsymbol{\theta}^*) = \sum_{(i,j)\in\mathcal{E}} \frac{(\hat{p}_{ij} - p_{ij})\, F'(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})}{F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})\left(1 - F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})\right)}\boldsymbol{x}_{ij},$$

$$\nabla^2 \hat{L}(\boldsymbol{\theta}) = -\sum_{(i,j)\in E}\left(\hat{p}_{ij}\frac{\mathrm{d}^2}{\mathrm{d}t^2}\log\left(F(t)\right)\Big|_{t=\boldsymbol{\theta}^{\mathsf{T}}\boldsymbol{x}_{ij}} + (1-\hat{p}_{ij})\frac{\mathrm{d}^2}{\mathrm{d}t^2}\log\left(1-F(t)\right)\Big|_{t=\boldsymbol{\theta}^{\mathsf{T}}\boldsymbol{x}_{ij}}\right)\boldsymbol{x}_{ij}\boldsymbol{x}_{ij}^{\mathsf{T}}.$$

Notice that $F(t)$ and $1 - F(t)$ are both α-strongly log-concave on $t \in [-2b, 2b]$, and thus we have $\nabla^2 \hat{L}(\boldsymbol{\theta}) \succeq \alpha \sum_{(i,j)\in\mathcal{E}} \boldsymbol{x}_{ij}\boldsymbol{x}_{ij}^{\mathsf{T}}$. Similarly to the linear case, we can obtain $\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_H \leq \frac{2}{\alpha}\|\nabla \hat{L}(\boldsymbol{\theta}^*) - \nabla L(\boldsymbol{\theta}^*)\|_{H^{-1}}$. Here we consider the inequality from the perspective of Euclidean norm. Therefore,

$$\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \leq \frac{2\|\hat{L}(\boldsymbol{\theta}^*) - \nabla L(\boldsymbol{\theta}^*)\|_{H^{-1}}}{\alpha\sqrt{\lambda_{\min}(H)}} \leq \frac{2\|\hat{L}(\boldsymbol{\theta}^*) - \nabla L(\boldsymbol{\theta}^*)\|_2}{\alpha\lambda_{\min}}.$$

The remaining part is to derive bound for

$$\|\hat{L}(\boldsymbol{\theta}^*) - \nabla L(\boldsymbol{\theta}^*)\|_2 = \sup_{\|\boldsymbol{u}\|_2 = 1}\left|\boldsymbol{u}^{\mathsf{T}}(\nabla \hat{L}(\boldsymbol{\theta}^*) - \nabla L(\boldsymbol{\theta}^*))\right|.$$

Let $y_{ij} = p_{ij} - \hat{p}_{ij}$ and notice that $\mathbb{E}[y_{ij}] = 0$ and $\mathrm{var}(y_{ij}) = (p_{ij}(1-p_{ij}))/k \leq 1/(4k)$. For any vector $\boldsymbol{u} \in \mathbb{R}^D$ with $\|\boldsymbol{u}\|_2 = 1$,

$$\boldsymbol{u}^{\mathsf{T}}(\nabla \hat{L}(\boldsymbol{\theta}^*) - \nabla L(\boldsymbol{\theta}^*)) = -\frac{F'(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})}{F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})\left(1 - F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})\right)}\sum_{(i,j)\in\mathcal{E}}\frac{\sum_{t=1}^k (Z_{ij}^{(t)} - p_{ij})}{k}(\boldsymbol{u}^{\mathsf{T}}\boldsymbol{x}_{ij}).$$

For convenience, we define the random variable $X_{i,j,t}, (i,j) \in \mathcal{E},\ t \in \{1, 2, \ldots, k\}$ such that

$$X_{i,j,t} \triangleq -\frac{F'(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})}{F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})\left(1 - F(\boldsymbol{\theta}^{*\mathsf{T}}\boldsymbol{x}_{ij})\right)}\frac{\left(Z_{ij}^{(t)} - p_{ij}\right)}{k}(\boldsymbol{u}^{\mathsf{T}}\boldsymbol{x}_{ij}).$$

By our assumption on F, we have

$$\sup_{t\in[-1,1]}\frac{F'(t)}{F(t)\left(1-F(t)\right)} \leq \frac{\beta}{F(-2b)\left(1-F(-2b)\right)} = \hat{\beta}.$$

Due to the independent comparison assumption, $X_{i,j,t}$ are zero-mean, independent random variables with $|X_{i,j,t}| \in [-2R\hat{\beta}/k, 2R\hat{\beta}/k]$. Hoeffding's inequality thus yields

$$\mathbb{P}\left(\left|\boldsymbol{u}^{\mathsf{T}}(\nabla \hat{L}(\boldsymbol{\theta}^*) - \nabla L(\boldsymbol{\theta}^*))\right| \geq \gamma\right) = \mathbb{P}\left(\left|\sum_{i,j,t} X_{i,j,t}\right| \geq \gamma\right) \tag{12}$$

$$\leq 2\exp\left(-\frac{2\gamma^2}{\frac{4}{k^2}\sum_{(i,j)\in\mathcal{E}} k(\boldsymbol{u}^\mathsf{T}\boldsymbol{x}_{ij})^2}\right)$$
$$\leq 2\exp\left(-\frac{k\gamma^2}{8|\mathcal{E}|R^2\hat{\beta}^2}\right).$$

To obtain a union bound, we apply the idea of the covering number. For a unit Euclidean sphere S^{D-1}, define $\mathcal{N}_\epsilon$ as a subset of S^{D-1} such that for every $\boldsymbol{u} \in S^{D-1}$, there exists some $\boldsymbol{v} \in S^{D-1}$ such that the Euclidean distance between $\boldsymbol{u}$ and $\boldsymbol{v}$ is bounded as $\|\boldsymbol{u}-\boldsymbol{v}\|_2 \leq \epsilon$. $\mathcal{N}(S^{D-1},\epsilon)$ denotes the minimal cardinality of $\mathcal{N}_\epsilon$. According to [55, Lemma 5.2], we have

$$\mathcal{N}(S^{D-1},\epsilon) \leq \left(1+\frac{2}{\epsilon}\right)^D.$$

For any $\boldsymbol{u} \in S^{D-1}$, there exists some $\boldsymbol{v} \in \mathcal{N}_\epsilon$ such that $|\boldsymbol{u}^\mathsf{T}(\nabla\hat{L}(\boldsymbol{\theta}^*)-\nabla L(\boldsymbol{\theta}^*))| \leq |\boldsymbol{v}^\mathsf{T}(\nabla\hat{L}(\boldsymbol{\theta}^*)-\nabla L(\boldsymbol{\theta}^*))| + |(\boldsymbol{u}^\mathsf{T}-\boldsymbol{v}^\mathsf{T})(\nabla\hat{L}(\boldsymbol{\theta}^*)-\nabla L(\boldsymbol{\theta}^*))| \leq |\boldsymbol{v}^\mathsf{T}(\nabla\hat{L}(\boldsymbol{\theta}^*)-\nabla L(\boldsymbol{\theta}^*))| + \epsilon(\nabla\hat{L}(\boldsymbol{\theta}^*)-\nabla L(\boldsymbol{\theta}^*))$. Hence, $\|\nabla\hat{L}(\boldsymbol{\theta}^*)-\nabla L(\boldsymbol{\theta}^*)\|_2 \leq \frac{1}{(1-\epsilon)}\max_{\boldsymbol{v}\in\mathcal{N}_\epsilon}|\boldsymbol{v}^\mathsf{T}(\nabla\hat{L}(\boldsymbol{\theta}^*)-\nabla L(\boldsymbol{\theta}^*))|$. To ensure that the bound (12) holds for all vectors $\boldsymbol{u}$ with probability at least $1-\delta$, it suffices to make

$$2\mathcal{N}(S^{D-1},\epsilon)\exp\left(-\frac{k\gamma^2}{8|\mathcal{E}|R^2\hat{\beta}^2}\right) \leq \delta,$$

which can be achieved by

$$\gamma \geq \sqrt{\frac{8|\mathcal{E}|R^2\hat{\beta}^2}{k}\log\frac{2\mathcal{N}(S^{D-1},\epsilon)}{\delta}} = \sqrt{\frac{8|\mathcal{E}|R^2\hat{\beta}^2}{k}\log\frac{2\left(1+\frac{2}{\epsilon}\right)^D}{\delta}}.$$

Using the union bound technique, we obtain

$$\mathbb{P}\left(\exists\,\boldsymbol{u}\in\mathbb{R}^D, \|\boldsymbol{u}\|_2=1 : |\boldsymbol{u}^\mathsf{T}\nabla\hat{L}(\boldsymbol{\theta}^*)| \geq \frac{1}{1-\epsilon}\sqrt{\frac{8|\mathcal{E}|R^2\hat{\beta}^2}{k}\log\frac{2\left(1+\frac{2}{\epsilon}\right)^D}{\delta}}\right) \leq \delta.$$

Without losing of generality, we can assign $\epsilon = \frac{1}{2}$, and thus with probability at least $1-\delta$

$$\|\hat{\boldsymbol{\theta}}-\boldsymbol{\theta}^*\|_2 \leq \frac{4}{\alpha\lambda_{\min}}\sqrt{\frac{8|\mathcal{E}|R^2\hat{\beta}^2}{k}\log\frac{2\cdot 5^D}{\delta}}.$$

□

F Proof of Theorem 5

Proof. The ideas of the proof is similar to Theorem 2. We use the same test statistic (6) on split dataset $\mathcal{D}_1$ and $\mathcal{D}$, and rewrite

$$T = \frac{1}{|\mathcal{E}|}\sum_{(i,j)\in\mathcal{E}}\frac{Z_{ij}(Z_{ij}-1)}{k(k-1)} + F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})^2 - 2F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})\frac{Z_{ij}}{k},$$

where $\hat{\boldsymbol{\theta}}$ is estimated on $\mathcal{D}_1$ and Z_{ij} are computed based on $\mathcal{D}$. We derive

$$\mathbb{E}[T \mid \mathcal{D}_1] = \frac{1}{|\mathcal{E}|}\sum_{(i,j)\in\mathcal{E}}\left(p_{ij}-F(\hat{\boldsymbol{\theta}}^\mathsf{T}\boldsymbol{x}_{ij})\right)^2.$$

Then the expectation of T can be bounded by

$$\mathbb{E}[T] = \mathbb{E}\left[\mathbb{E}[T \mid \mathcal{D}_1]\right]$$
$$\leq \frac{1}{|\mathcal{E}|}\left(\|P-Q(\boldsymbol{\theta}^*)\|_\mathrm{F}^2 + \mathbb{E}\left[\|Q(\boldsymbol{\theta}^*)-Q(\hat{\boldsymbol{\theta}})\|_\mathrm{F}^2\right]\right.$$

$$+ 2\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}} \mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}\right]\Big).$$

Using Lemma 10, we obtain the estimation error bounds

$$\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}^2\right] \leq c_2 \left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^2$$

and

$$\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}\right] \leq c_1 \left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right),$$

which imply

$$\mathbb{E}[T] \leq \frac{1}{|\mathcal{E}|}\left(\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^2 + c_2\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^2 + 2c_1\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}\right).$$

Under the null hypothesis H_0, since $\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}} = 0$, we get $\mathbb{E}_{H_0}[T] \leq c_2(\frac{R}{\lambda_{\min}}\sqrt{\frac{D\lambda_{\max}}{k}})^2$. Under the alternative hypothesis H_1, since $\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}^2\right] \geq 0$ and $\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}\right] \leq c_1(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}})$, we obtain

$$\mathbb{E}_{H_1}[T] \geq \frac{1}{|\mathcal{E}|}\left(\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^2 - 2c_1\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}\right).$$

Next, we bound the variance. We can again use the fact that $\mathrm{var}(T) = \mathbb{E}[\mathrm{var}(T) \mid \mathcal{D}_1] + \mathrm{var}(\mathbb{E}[T \mid \mathcal{D}_1])$. The first term is

$$\mathrm{var}(T \mid \mathcal{D}_1) = \frac{1}{|\mathcal{E}|^2}\sum_{(i,j)\in\mathcal{E}} \frac{2p_{ij}^2(1-p_{ij})^2}{k(k-1)} + \frac{4p_{ij}(1-p_{ij})}{k}\left(p_{ij} - F(\hat{\boldsymbol{\theta}}^{\mathsf{T}}\boldsymbol{x}_{ij})\right)^2.$$

Taking expectations, we get

$$\begin{aligned}\mathbb{E}[\mathrm{var}(T \mid D_1)] &\leq \frac{1}{|\mathcal{E}|^2}\left(\frac{|\mathcal{E}|}{8k(k-1)} + \frac{1}{k}\left(\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}} + \sqrt{\mathbb{E}\left[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}})\|_{\mathrm{F}}^2\right]}\right)^2\right) \\ &\leq \frac{1}{|\mathcal{E}|^2}\left(\frac{|\mathcal{E}|}{8k(k-1)} + \frac{1}{k}\left(\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}} + \sqrt{c_2\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^2}\right)^2\right).\end{aligned}$$

The second term $\mathrm{var}(\mathbb{E}[T \mid \mathcal{D}_1])$ is bounded by

$$\begin{aligned}\mathrm{var}(\mathbb{E}[T \mid \mathcal{D}_1]) \leq \frac{1}{|\mathcal{E}|^2}\Big(&\mathbb{E}\left[\left(\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}} + \|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}}))\|_{\mathrm{F}}\right)^4\right] \\ &- \mathbb{E}\left[\left(\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}} - \|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}}))\|_{\mathrm{F}}\right)^2\right]^2\Big).\end{aligned}$$

Under H_0, $\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}} = 0$ implies $|\mathcal{E}|^2\mathrm{var}_{H_0}(\mathbb{E}[T \mid \mathcal{D}_1]) \leq \mathbb{E}[\|Q(\boldsymbol{\theta}^*) - Q(\hat{\boldsymbol{\theta}}))\|_{\mathrm{F}}^4] \leq c_4(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}})^4$. Under H_1, we have

$$\mathrm{var}_{H_1}(\mathbb{E}[T \mid \mathcal{D}_1]) \leq \frac{1}{|\mathcal{E}|^2}\left(8c_1\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^3\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)\right.$$

$$
\begin{aligned}
&+ 4c_2\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^2 \left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^2 \\
&+ 4(c_3 + c_2c_1)\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}} \left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^3 \\
&+ c_4 \left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^4\Bigg).
\end{aligned}
$$

Combining the two terms, we obtain the following upper bounds on variance under H_0 and H_1:

$$
\mathrm{var}_{H_0}(T) \leq \frac{1}{|\mathcal{E}|^2}\left(\frac{|\mathcal{E}|}{8k(k-1)} + c_2\frac{1}{k}\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^2 + c_4\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^4\right),
$$

$$
\begin{aligned}
&\mathrm{var}_{H_1}(T) \\
&\leq \frac{1}{|\mathcal{E}|^2}\left(\frac{|\mathcal{E}|}{8k(k-1)} + \frac{1}{k}\left(\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}} + \sqrt{c_2\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^2}\right)^2\right. \\
&\quad + 8c_1\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^3\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right) + 4c_2\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^2\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^2 \\
&\quad \left. + 4(c_3 + c_2c_1)\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^3 + c_4\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^4\right).
\end{aligned}
$$

We let the decision rule accept H_1 if

$$
T > \bar{\gamma}\left(\frac{1}{k\sqrt{|\mathcal{E}|}} + \frac{R^2}{\lambda_{\min}^2}\frac{D\lambda_{\max}}{k}\right) + c_2\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{D\lambda_{\max}}{k}}\right)^2.
$$

By this definition, the Type I error can be bounded as

$$
\begin{aligned}
&\mathbb{P}_{H_0}\left(T > \bar{\gamma}\left(\frac{1}{k\sqrt{|\mathcal{E}|}} + \frac{R^2}{\lambda_{\min}^2}\frac{D\lambda_{\max}}{k}\right) + c_2\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{D\lambda_{\max}}{k}}\right)^2\right) \\
&\leq P_{H_0}\left(T - \mathbb{E}_{H_0}[T] > \bar{\gamma}\left(\frac{1}{k\sqrt{|\mathcal{E}|}} + \frac{R^2}{\lambda_{\min}^2}\frac{D\lambda_{\max}}{k}\right)\right) \\
&\leq \frac{\mathrm{var}_{H_0}(T)}{\mathrm{var}_{H_0}(T) + \bar{\gamma}^2(\frac{1}{k\sqrt{|\mathcal{E}|}} + \frac{R^2}{\lambda_{\min}^2}\frac{D\lambda_{\max}}{k})^2} \\
&\leq \frac{\frac{|\mathcal{E}|}{8k(k-1)} + c_2\frac{1}{k}(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}})^2 + c_4(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}})^4}{\frac{|\mathcal{E}|}{8k(k-1)} + c_2\frac{1}{k}\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^2 + c_4\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^4 + \bar{\gamma}^2\left(\frac{\sqrt{|\mathcal{E}|}}{k} + \frac{R^2}{\lambda_{\min}^2}\frac{|\mathcal{E}|D\lambda_{\max}}{k}\right)^2} \\
&\leq \frac{\frac{|\mathcal{E}|}{4} + c_2\left(\frac{R}{\lambda_{\min}}\sqrt{|\mathcal{E}|D\lambda_{\max}}\right)^2 + c_4\left(\frac{R}{\lambda_{\min}}\sqrt{|\mathcal{E}|D\lambda_{\max}}\right)^4}{\frac{|\mathcal{E}|}{4} + c_2\left(\frac{R}{\lambda_{\min}}\sqrt{|\mathcal{E}|D\lambda_{\max}}\right)^2 + c_4\left(\frac{R}{\lambda_{\min}}\sqrt{|\mathcal{E}|D\lambda_{\max}}\right)^4 + \bar{\gamma}^2\left(\sqrt{|\mathcal{E}|} + \frac{R^2}{\lambda_{\min}^2}|\mathcal{E}|D\lambda_{\max}\right)^2} \\
&\leq \frac{1}{4}.
\end{aligned}
$$

The last inequality can be achieved by large enough $\bar{\gamma}$ where

$$\bar{\gamma}^2 \geq \frac{|\mathcal{E}|}{\left(\sqrt{|\mathcal{E}|} + \frac{R^2}{\lambda_{\min}^2}|\mathcal{E}|D\lambda_{\max}\right)^2} + \frac{4c_2}{\sqrt{|\mathcal{E}|} + \frac{R^2}{\lambda_{\min}^2}|\mathcal{E}|D\lambda_{\max}} + 4c_4$$

will suffice. For the Type II error, we have

$$\begin{aligned}
&\mathbb{P}_{H_1}\left(T \leq \bar{\gamma}\left(\frac{1}{k\sqrt{|\mathcal{E}|}} + \frac{R^2}{\lambda_{\min}^2}\frac{D\lambda_{\max}}{k}\right) + c_2\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{D\lambda_{\max}}{k}}\right)^2\right) \\
&\leq \mathbb{P}_{H_1}\left(T - \mathbb{E}_{H_1}[T] \leq \bar{\gamma}\left(\frac{1}{k\sqrt{|\mathcal{E}|}} + \frac{R^2}{\lambda_{\min}^2}\frac{D\lambda_{\max}}{k}\right) + c_2\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{D\lambda_{\max}}{k}}\right)^2 - \frac{1}{|\mathcal{E}|}\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}^2\right. \\
&\qquad \left. + 2c_1\frac{1}{|\mathcal{E}|}\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}\right) \\
&\leq \frac{\mathrm{var}_{H_1}(T)}{\mathrm{var}_{H_1}(T) + (A^2 - \Delta)^2} \leq \frac{1}{4},
\end{aligned}$$

where $A = \frac{1}{\sqrt{|\mathcal{E}|}}\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}}$ and

$$\Delta = \bar{\gamma}\left(\frac{1}{k\sqrt{|\mathcal{E}|}} + \frac{R^2}{\lambda_{\min}^2}\frac{D\lambda_{\max}}{k}\right) + c_2\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{D\lambda_{\max}}{k}}\right)^2 + 2c_1\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{D\lambda_{\max}}{k}}\right)A.$$

In order to make the error less or equal to $\frac{1}{4}$, it suffices to make $A^2 \geq 2\sqrt{\mathrm{var}_{H_1}(T)} + \Delta$.

$$\begin{aligned}
A^2 \geq \frac{2}{|\mathcal{E}|}&\left(\frac{\sqrt{|\mathcal{E}|}}{2k} + \frac{1}{\sqrt{k}}\left(A\sqrt{|\mathcal{E}|} + \sqrt{c_2\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)}\right)\right. \\
&+ \sqrt{8c_1}\left(A\sqrt{|\mathcal{E}|}\right)^{\frac{3}{2}}\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^{\frac{1}{2}} + 2\sqrt{c_2}A\sqrt{|\mathcal{E}|}\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right) \\
&\left. + 2\sqrt{(c_3 + c_2c_1)}\sqrt{A\sqrt{|\mathcal{E}|}}\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^{3/2} + \sqrt{c_4}\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{|\mathcal{E}|D\lambda_{\max}}{k}}\right)^2\right) \\
&+ \bar{\gamma}\left(\frac{1}{k\sqrt{|\mathcal{E}|}} + \frac{R^2}{\lambda_{\min}^2}\frac{D\lambda_{\max}}{k}\right) + c_2\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{D\lambda_{\max}}{k}}\right)^2 + 2c_1\left(\frac{R}{\lambda_{\min}}\sqrt{\frac{D\lambda_{\max}}{k}}\right)A.
\end{aligned}$$

Here we write $A = c\sqrt{\frac{1}{k\sqrt{|\mathcal{E}|}} + \frac{R^2}{\lambda_{\min}^2}\frac{D\lambda_{\max}}{k}}$ for some constant c and if c satisfies

$$\begin{aligned}
c^2 \geq 2\Bigg(&\frac{1}{2} + (c + \sqrt{c_2})\sqrt{\frac{1}{\frac{\sqrt{|\mathcal{E}|}}{k} + \frac{R^2}{\lambda_{\min}^2}|\mathcal{E}|D\lambda_{\max}}} + \sqrt{8c_1}c^{\frac{3}{2}} + 2\sqrt{c_2}c \\
&+ 2\sqrt{(c_3 + c_2c_1)c} + \sqrt{c_4}\Bigg) + \bar{\gamma} + c_2 + 2c_1c,
\end{aligned}$$

we have $A^2 \geq 2\sqrt{\mathrm{var}_{H_1}(T)} + \Delta$. In this case, the minimax risk is less or equal to $\frac{1}{2}$. According to Proposition 9, $\inf_{\boldsymbol{\theta}\in\Theta}\|P - Q(\boldsymbol{\theta})\|_{\mathrm{F}} = \Theta(\|P - Q(\boldsymbol{\theta}^*)\|_{\mathrm{F}})$. Combining the normalization scaling $\frac{1}{\sqrt{|\mathcal{E}|}}$, $\frac{1}{\sqrt{|\mathcal{E}|}}\inf_{\boldsymbol{\theta}\in\Theta}\|P - Q(\boldsymbol{\theta})\|_{\mathrm{F}} = \Theta(A)$. We derive that the critical threshold is bounded by

$$\epsilon_c \leq c\sqrt{\frac{1}{k\sqrt{|\mathcal{E}|}} + \frac{R^2}{\lambda_{\min}^2}\frac{D\lambda_{\max}}{k}},$$

which completes the proof. □

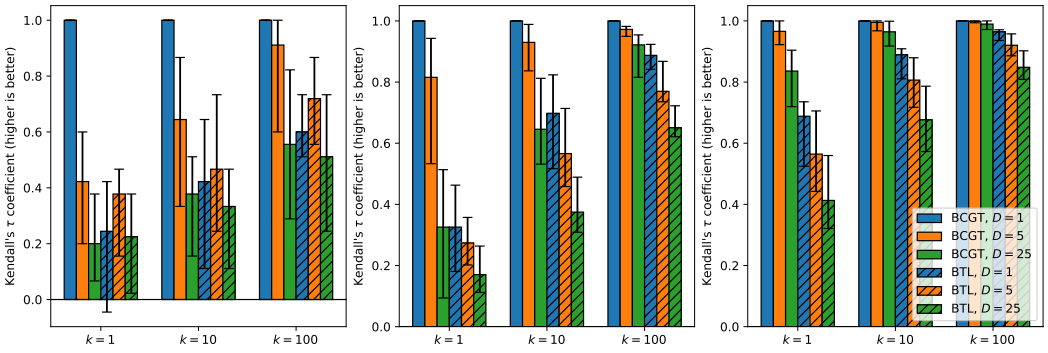

(a) $n = 10$ (b) $n = 50$ (c) $n = 250$

Figure 4: Plot of Kendall's τ coefficient between the true rankings and the rankings estimated by the BCGT (with polynomial basis) and BTL models on complete graphs with $p = 1$, $n \in \{10, 50, 250\}$, $k \in \{1, 10, 100\}$, and $D \in \{1, 5, 25\}$.

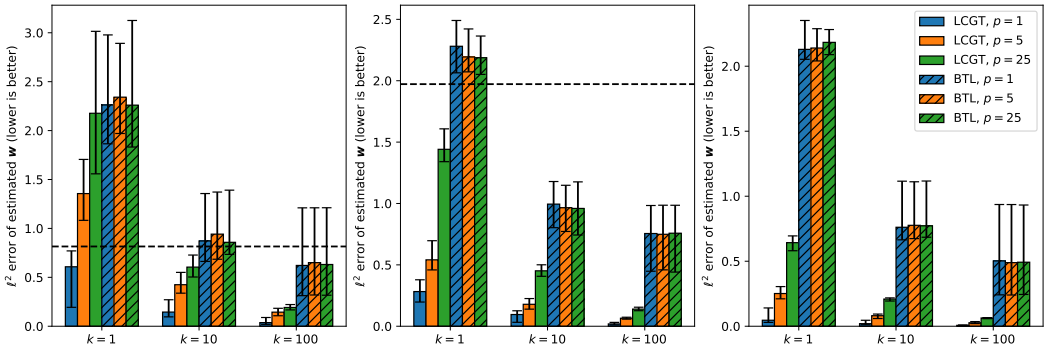

(a) $n = 10$ (b) $n = 50$ (c) $n = 250$

Figure 5: Plot of ℓ^2 estimation error of the LCGT and BTL models on complete graphs with $n \in \{10, 50, 250\}$, $k \in \{1, 10, 100\}$, and $p \in \{1, 5, 25\}$.

G Proof of Proposition 6

Notice that the proof in Appendix D relies on the graph structure and does not depend on the specific structure of the basis. Therefore, Proposition 6 can be proved analogously to Proposition 3 by changing features $\boldsymbol{x}_i$ to $\boldsymbol{\phi}(\boldsymbol{x}_i)$ for all $i \in [n]$.

H Additional experiments

H.1 Numerical simulations. Here we present the results of simulations that could not be included in the main text due to space constraints. The experiment setup is explained in detail in Section 5, so we only provide a brief overview here to accompany each of the plots.

Figure 4 uses the Kendall rank correlation coefficient (also called Kendall's τ coefficient) [58] to compare the accuracy of the rankings estimated by the BCGT (with polynomial basis) and BTL models in the first experiment described in Section 5 (cf. Figure 1, where the ℓ^2 errors of the estimated scores are compared). Note that $\tau = 1$ indicates that the estimated ranking is perfect, whereas a random guess will yield $\tau = 0$ in expectation. Similarly to Figure 1, Figure 4 shows that the rankings estimated by BCGT estimation is usually more accurate than the rankings estimated by BTL estimation. Interestingly, even when the ℓ^2 error of the estimated scores is poor (e.g., when the ℓ^2 error is greater than $\|\boldsymbol{w}\|_2$, which indicates that simply guessing $\boldsymbol{w} = \mathbf{0}$ would have yielded a smaller ℓ^2 error), the rankings still tend to be better than random guesses (i.e., τ tends to be positive).

Figures 5 and 6 show how the estimation error of BCGT (with linear and Fourier bases, respectively) and BTL models compare on synthetic datasets generated on a complete graph.

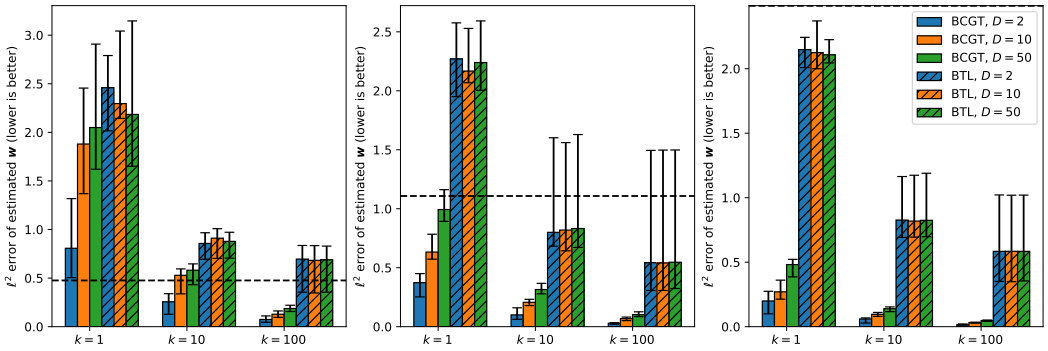

(a) $n = 10$ (b) $n = 50$ (c) $n = 250$

Figure 6: Plot of ℓ^2 estimation error of the BCGT (with Fourier basis) and BTL models on complete graphs with $p = 1$, $n \in \{10, 50, 250\}$, $k \in \{1, 10, 100\}$, and $D \in \{1, 5, 25\}$.

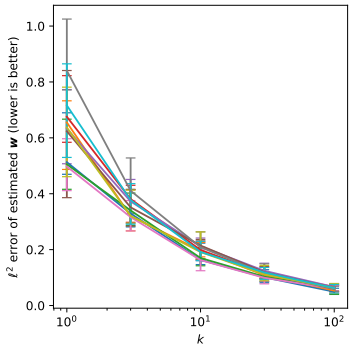

Figure 7: Plot of how the ℓ^2 estimation error of the LCGT model changes with respect to k on complete graphs with $n = 50$ and $d = 5$, where each curve represents a different (fixed) value of $\boldsymbol{\theta}$.

Figure 7 shows how the estimation error of the LCGT model behaves with respect to k for several fixed values of $\boldsymbol{\theta}$. This differs from the above bar plots, which are averaged across many randomly generated instances of $\boldsymbol{\theta}$. Recall that the error bars span the first to third quartiles.

Figure 8 shows how the estimation error of LCGT and BTL models compare on synthetic datasets generated on very sparse Erdős-Rényi graphs (where each edge is independently sampled with probability $q \in \{0.5, 0.05, 0.005\}$). When $q = 0.5$, the graph is relatively dense and the performance of both models is comparable to that in the complete graph setting. On the other hand, under the sparse setting of $q = 0.005$, the graph is disconnected with high probability, which causes a spike in estimation error (especially for the BTL model, since it is impossible to estimate the relationship between disconnected subgraphs without context).

Figure 9 shows how underfitting and overfitting affect the estimation error of the LCGT model. The underlying dataset is generated under an LCGT model with $p = 10$, but estimation is performed for various p that may not match the true data dimension. In general, (except for the outlier $n = 10$, where the small number of items causes overfitting to happen more aggressively when k is small), both underfitting and overfitting hurt performance, but underfitting is more damaging.

Figure 10 shows the distribution of the test statistic T when testing for the LCGT model on complete graphs under H_0 and H_1 when $n \in \{10, 50, 250\}$ and $k \in \{2, 10, 100\}$. Figure 11 shows the corresponding ROC curves. Note that when perfect testing is possible (when there is no overlap between the distributions under H_0 and H_1) the ROC curve may be hidden under the plot axes.

Figure 12 shows the distribution of the test statistic T when testing for the BCGT model (with polynomial basis) on complete graphs under H_0 and H_1 when $n \in \{10, 50, 250\}$ and $k \in \{2, 10, 100\}$. Figure 13 shows the corresponding ROC curves.

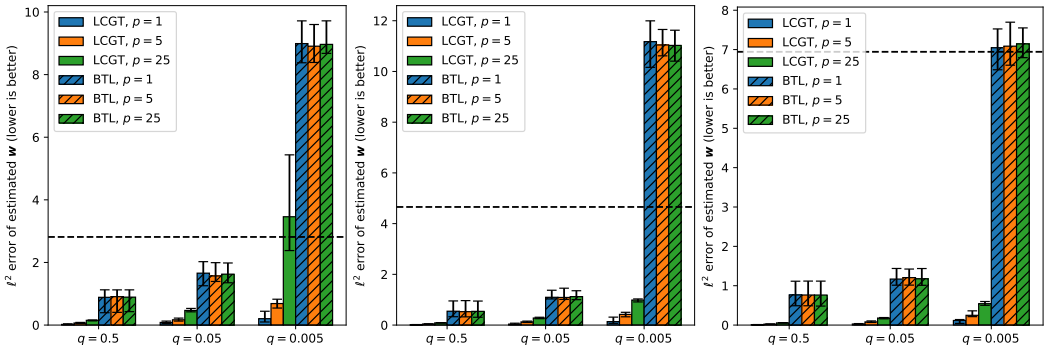

(a) $n = 100$ (b) $n = 250$ (c) $n = 625$

Figure 8: Plot of ℓ^2 estimation error of the LCGT and BTL models on sparse synthetic datasets with $n \in \{100, 250, 625\}$, $k = 100$, $p \in \{1, 5, 25\}$, and Erdös-Rényi sampling probability $q \in \{0.5, 0.05, 0.005\}$.

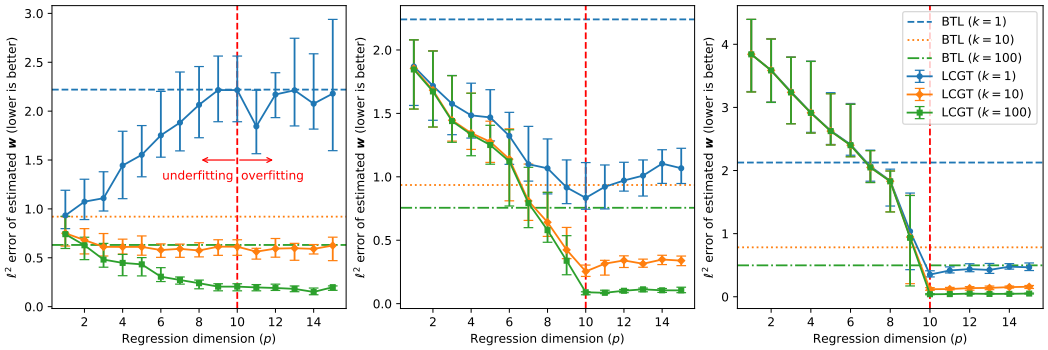

(a) $n = 10$ (b) $n = 50$ (c) $n = 250$

Figure 9: Demonstration of the underfitting and overfitting behavior of LCGT when the linear regression dimension p does not match the true dimensionality of the data $p_{\text{true}} = 10$.

Figure 14 shows the distribution of the test statistic T when testing for the BCGT model (with polynomial basis) on perfect matchings under H_0 and H_1 when $n \in \{100, 500, 2500\}$ and $k \in \{2, 10, 100\}$. Figure 15 shows the corresponding ROC curves. Note how the variance of the test statistic decays more slowly (with respect to n) than in the complete graph setting.

Figure 16 shows the distribution of the biased test statistic (defined as (6) with the bias correction term ⓐ removed) when testing for the BCGT model (with polynomial basis) on perfect matchings under H_0 and H_1 when $n \in \{100, 500, 2500\}$ and $k = 1$ and the corresponding ROC curves.

H.2 Korea Baseball Organization (KBO) League. The Korea Baseball Organization (KBO) league is a baseball league between $n = 10$ teams, where each pair of teams play $k = 16$ games between each other in one season. Win/loss data and OBP (on-base percentage), SLG (slugging percentage), ERA (earned run average) statistics for the years $\{2015, \ldots, 2025\}$ were collected from public data available on `https://www.koreabaseball.com`. Figure 17(a) shows the results of fitting an LCGT model on each year, Figure 17(b) shows how the test statistic T and empirical threshold changes over time as data accumulates, and Figure 17(c) visualizes the result of testing for LCGT using a $\hat{\boldsymbol{\theta}}$ estimate computed from another year's dataset. As expected, the estimation results show that high OBP and SLG (positive statistics for hitters) and low ERA (negative statistic for pitchers) are preferred, Figure 17(b) shows that the test statistic of cumulative data increases over time (due to year-by-year variations that cause the dataset to drift away from an LCGT model), and Figure 17(c) shows that testing using heterogeneous $\mathcal{Z}_1$ and $\mathcal{Z}_2$ usually results in a higher test statistic, suggesting model mismatch.

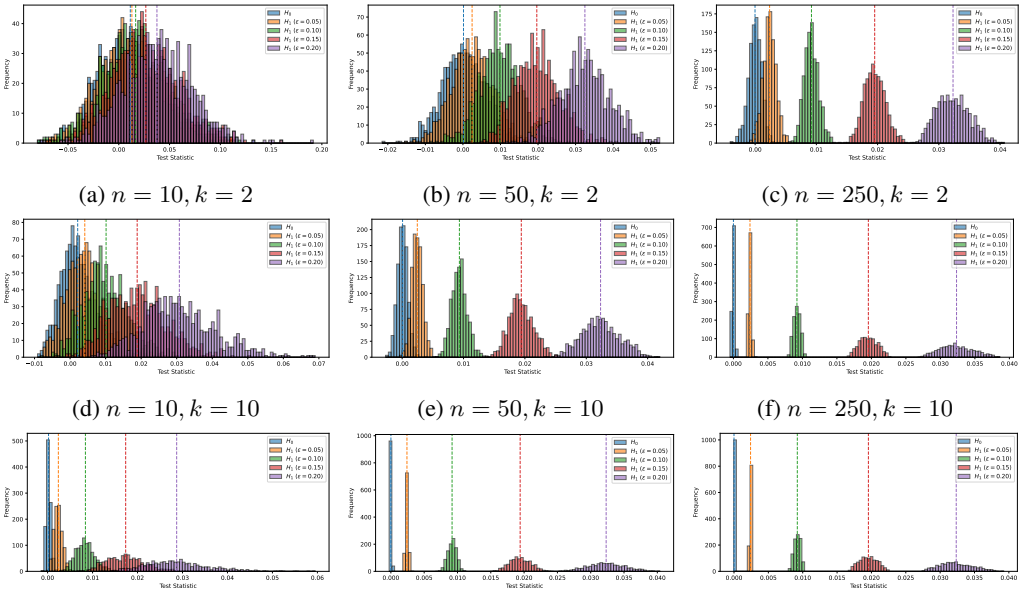

(a) $n = 10, k = 2$ (b) $n = 50, k = 2$ (c) $n = 250, k = 2$

(d) $n = 10, k = 10$ (e) $n = 50, k = 10$ (f) $n = 250, k = 10$

(g) $n = 10, k = 100$ (h) $n = 50, k = 100$ (i) $n = 250, k = 100$

Figure 10: Histograms of the empirical distributions of the test statistic under the H_0 and H_1 hypotheses with various ϵ for the LCGT model, where the comparison graph is a complete graph.

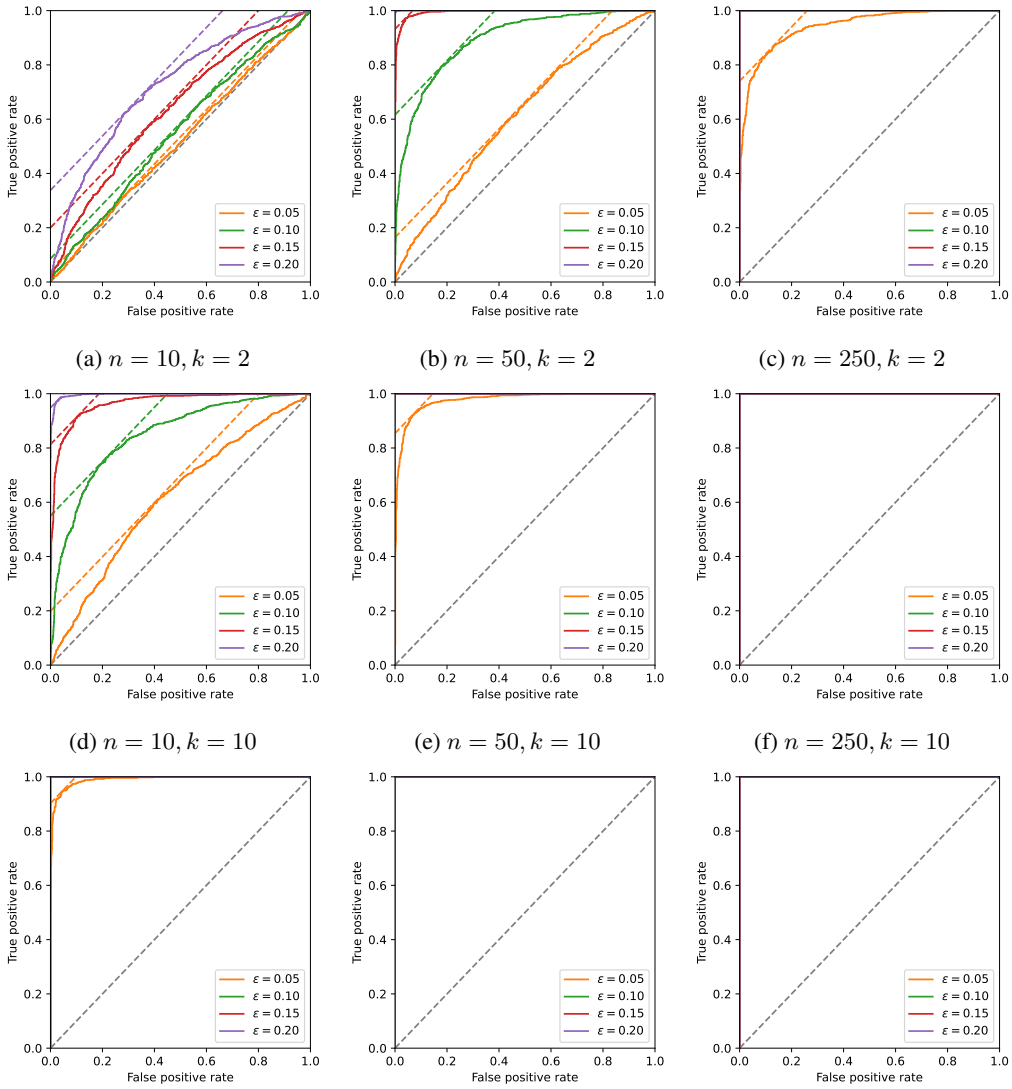

(a) $n = 10, k = 2$ (b) $n = 50, k = 2$ (c) $n = 250, k = 2$

(d) $n = 10, k = 10$ (e) $n = 50, k = 10$ (f) $n = 250, k = 10$

(g) $n = 10, k = 100$ (h) $n = 50, k = 100$ (i) $n = 250, k = 100$

Figure 11: ROC curves derived from the empirical distributions of the test statistic for various ϵ under the LCGT model, where the comparison graph is a complete graph.

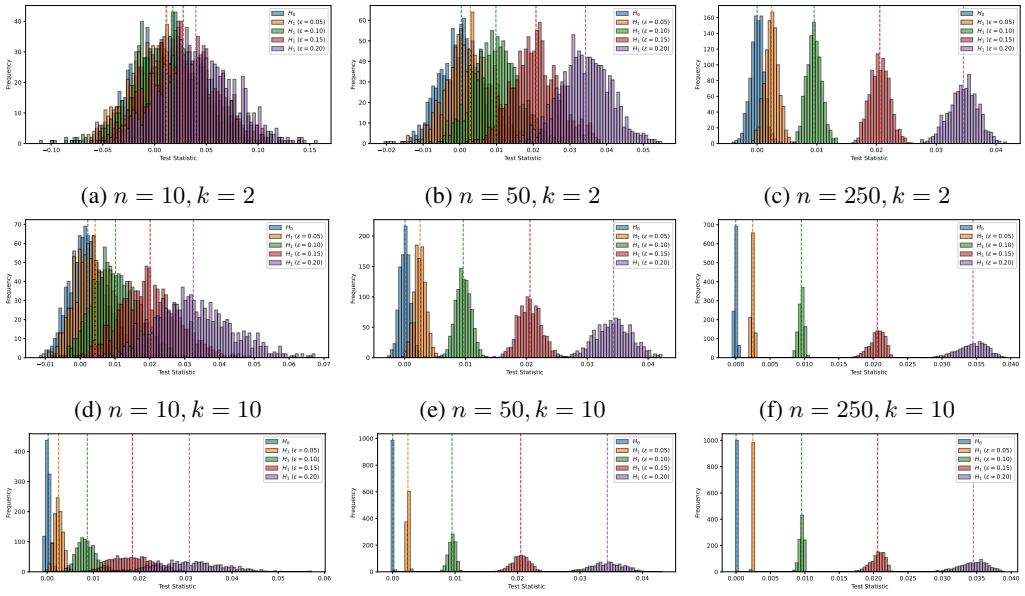

(a) $n = 10, k = 2$ (b) $n = 50, k = 2$ (c) $n = 250, k = 2$

(d) $n = 10, k = 10$ (e) $n = 50, k = 10$ (f) $n = 250, k = 10$

(g) $n = 10, k = 100$ (h) $n = 50, k = 100$ (i) $n = 250, k = 100$

Figure 12: Histograms of the empirical distributions of the test statistic under the H_0 and H_1 hypotheses with various ϵ for the BCGT model with polynomial basis, where the comparison graph is a complete graph.

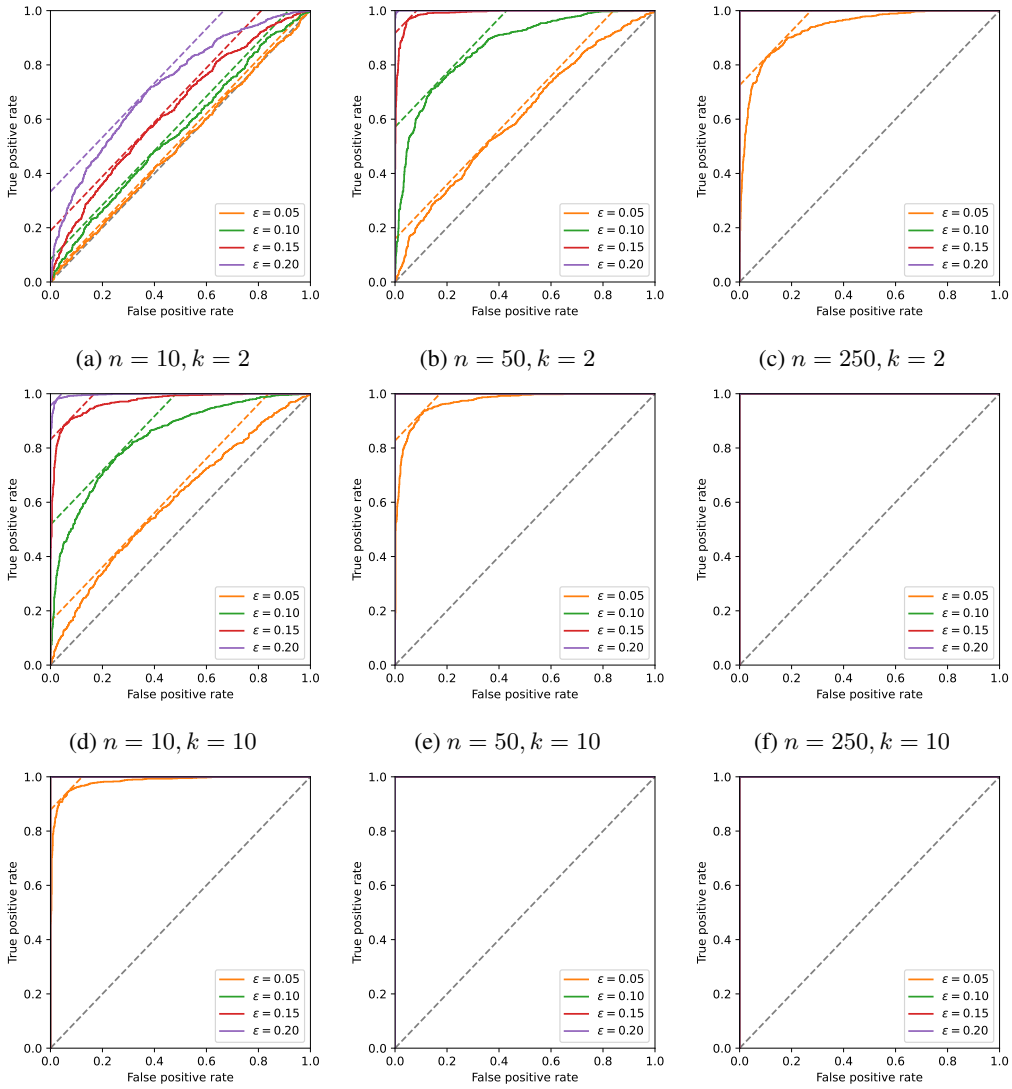

(a) $n = 10, k = 2$ (b) $n = 50, k = 2$ (c) $n = 250, k = 2$

(d) $n = 10, k = 10$ (e) $n = 50, k = 10$ (f) $n = 250, k = 10$

(g) $n = 10, k = 100$ (h) $n = 50, k = 100$ (i) $n = 250, k = 100$

Figure 13: ROC curves derived from the empirical distributions of the test statistic for various ϵ under the BCGT model with polynomial basis, where the comparison graph is a complete graph.

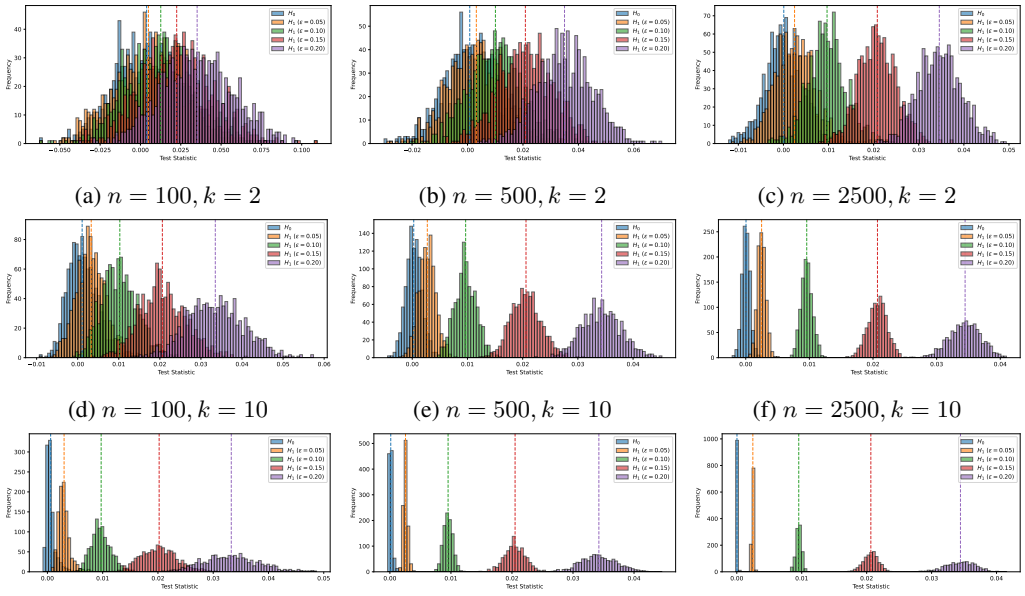

(a) $n = 100, k = 2$ (b) $n = 500, k = 2$ (c) $n = 2500, k = 2$

(d) $n = 100, k = 10$ (e) $n = 500, k = 10$ (f) $n = 2500, k = 10$

(g) $n = 100, k = 100$ (h) $n = 500, k = 100$ (i) $n = 2500, k = 100$

Figure 14: Histograms of the empirical distributions of the test statistic under the H_0 and H_1 hypotheses with various ϵ for the BCGT model with polynomial basis, where the comparison graph is a perfect matching.

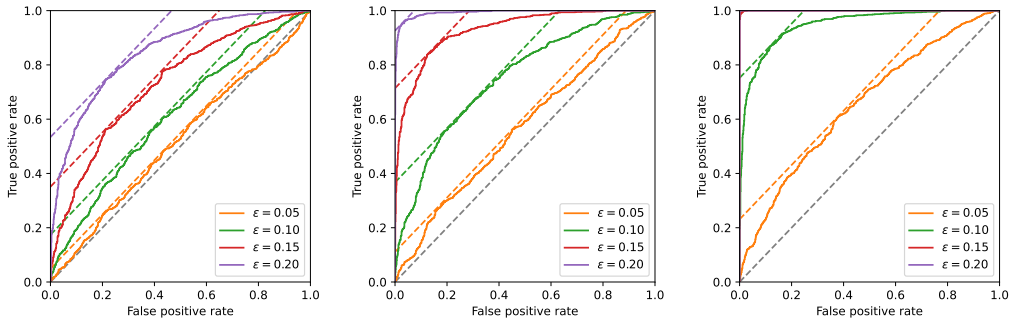

(a) $n = 100, k = 2$ (b) $n = 500, k = 2$ (c) $n = 2500, k = 2$

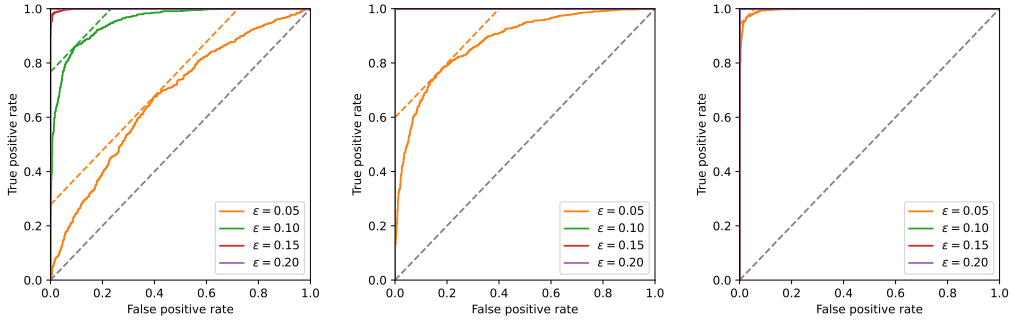

(d) $n = 100, k = 10$ (e) $n = 500, k = 10$ (f) $n = 2500, k = 10$

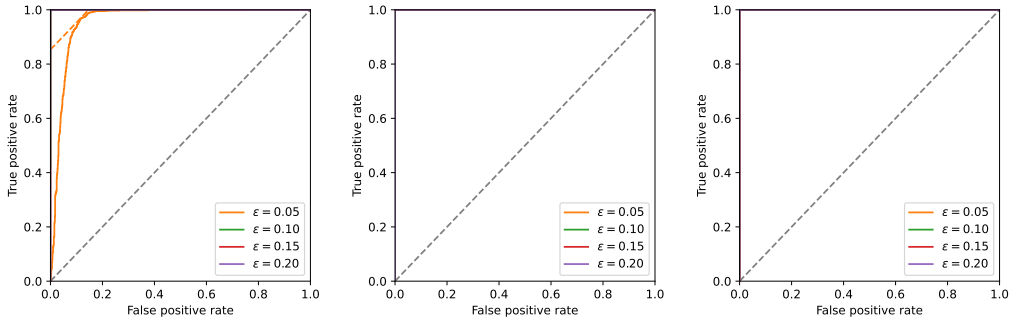

(g) $n = 100, k = 100$ (h) $n = 500, k = 100$ (i) $n = 2500, k = 100$

Figure 15: ROC curves derived from the empirical distributions of the test statistic for various ϵ under the BCGT model with polynomial basis, where the comparison graph is a perfect matching.

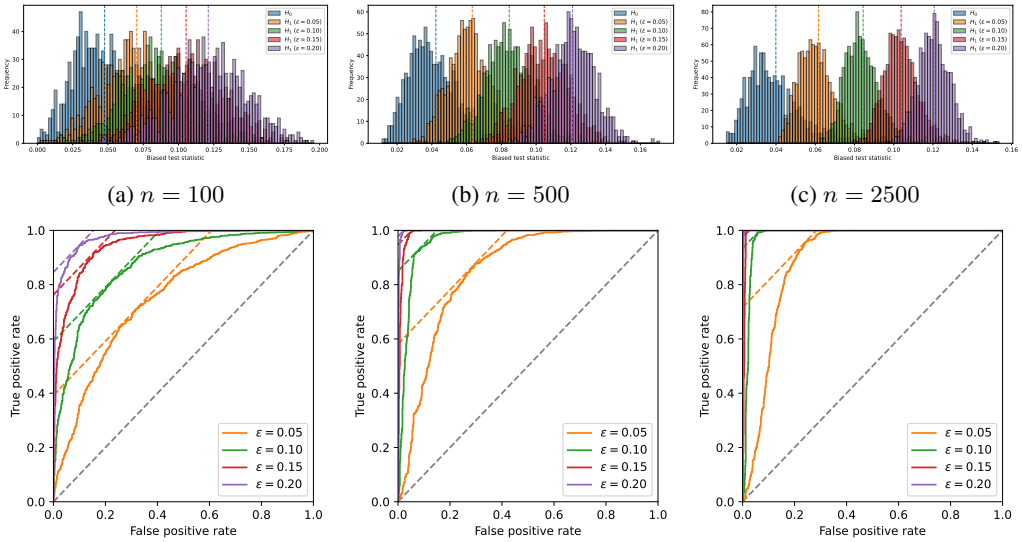

(a) $n = 100$ (b) $n = 500$ (c) $n = 2500$

(d) $n = 100$ (e) $n = 500$ (f) $n = 2500$

Figure 16: Histograms and ROCs of the empirical distributions of the biased test statistic (for $k = 1$) under the H_0 and H_1 hypotheses with various ϵ for the BCGT model with polynomial basis, where the comparison graph is a perfect matching.

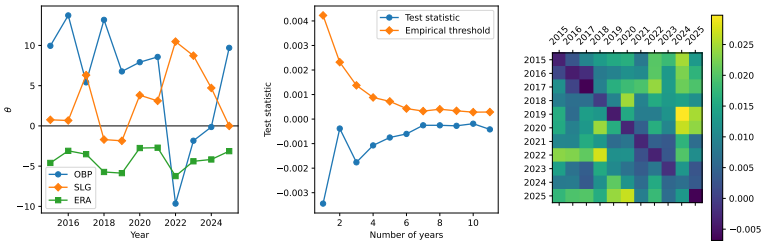

(a) Estimation results (b) Cumulative T (c) Testing with $\mathcal{Z}_1$ and $\mathcal{Z}_2$ from different years

Figure 17: **(a)** Result of LCGT estimation on the KBO dataset for each year, where the OBP, SLG, and ERA statistics are used as the context vectors. **(b)** Plot of the test statistic and empirical threshold computed for the cumulative data over the first $l \in \{1, \ldots, 11\}$ years in the dataset. **(c)** Heatmap of the test statistics computed by taking $\mathcal{Z}_1$ to be the dataset from one year and $\mathcal{Z}_2$ to be the dataset from another.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification:

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification:

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

Justification:

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: The model that we use is described in detail in the main body, and the datasets that we use are cited and readily available.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification:

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification:

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: Error bars are provided for relevant experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [No]

   Justification: The experiments described in this paper can be run on most consumer hardware.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification:

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A]

    Justification: This work is predominantly theoretical in nature.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification:

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All datasets are appropriately cited.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification:

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification:

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification:

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: No LLMs or similar AI systems were used in this research and exposition.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.