A Proof of Theorem 2.1

In this appendix, we present the proof of Theorem 2.1. We first introduce notation and preliminaries in Appendix A.1, to be used subsequently in proving both parts of Theorem 2.1. The proof of Theorem 2.1(b) is presented in Appendix A.2. The proof of Theorem 2.1(a) is presented in Appendix A.3. We first present the proof of Theorem 2.1(b) followed by Theorem 2.1(a), because the proof of Theorem 2.1(a) depends on the proof of Theorem 2.1(b).

In the proof of Theorem 2.1(a), the constants are allowed to depend only on the constant B. In the proof of Theorem 2.1(b), the constants are allowed to depend only on the constants A and B. The proofs for all the lemmas are presented in Appendix A.4.

A.1 Notation and preliminaries

In this appendix, we introduce notation and preliminaries that are used subsequently in the proofs of both Theorem 2.1(b) and Theorem 2.1(a).

(i) Notation

Recall that d denotes the number of items, and k denotes the number of comparisons per pair of items. The d items are associated to a true parameter vector $\theta^* = [\theta_1^*, \ldots, \theta_d^*]$. We have the set $\Theta_B = \{\theta \in \mathbb{R}^d \mid \|\theta\|_\infty \leq B, \sum_{i=1}^d \theta_i = 0\}$ and the set $\Theta_A = \{\theta \in \mathbb{R}^d \mid \|\theta\|_\infty \leq A, \sum_{i=1}^d \theta_i = 0\}$, where A and B are finite constants such that $A > B > 0$. The true parameter vector satisfies $\theta^* \in \Theta_B$.

Denote μ_{ij}^* as the probability that item $i \in [d]$ beats item $j \in [d]$. Under the BTL model, we have

$$\mu_{ij}^* = \frac{1}{1 + e^{-(\theta_i^* - \theta_j^*)}}. \tag{9}$$

For every $r \in [k]$, denote the outcome of the r^{th} comparison between item $i \in [d]$ and item $j \in [d]$ as

$$X_{ij}^{(r)} := \mathbb{1}\{\text{item } i \text{ beats item } j \text{ in their } r^{th} \text{ comparison}\}.$$

We have $X_{ij}^{(r)} \sim \text{Bernoulli}(\mu_{ij}^*)$, independent across all $r \in [k]$ and all $i < j$. Recall that W_{ij} denotes the number of times that item i beats j. We have $W_{ij} = \sum_{r=1}^k X_{ij}^{(r)}$ and therefore $W_{ij} \sim \text{Binom}(k, \mu_{ij}^*)$. Denote μ_{ij} as the fraction of times that item i beats item j. That is,

$$\mu_{ij} := \frac{1}{k} W_{ij} = \frac{1}{k} \sum_{r=1}^k X_{ij}^{(r)}. \tag{10}$$

We have $\mu_{ij} \sim \frac{1}{k}\text{Binom}(k, \mu_{ij}^*)$, independent across all $i < j$.

Finally, we use c, c', c_1, c_2, etc. to denote finite constants whose values may change from line to line. We write $f(n) \lesssim g(n)$ if there exists a constant c such that $f(n) \leq c \cdot g(n)$ for all $n \geq 1$. The notation $f(n) \gtrsim g(n)$ is defined analogously.

(ii) Notion of conditioning

Let E be any event. The conditional bias of any estimator $\widehat{\theta}$ conditioned on the event E is defined as:

$$\beta(\widehat{\theta} \mid E) := \sup_{\theta^* \in \Theta_B} \|\mathbb{E}[\widehat{\theta} \mid E] - \theta^*\|_\infty.$$

We use "w.h.p.$(\frac{1}{dk})$" to denote that an event E happens with probability at least

$$\mathbb{P}(E) > 1 - \frac{c}{dk},$$

for all $d \geq d_0$ and $k \geq k_0$, where d_0, k_0 and c are positive constants.

Similarly, we use "w.h.p.($\frac{1}{dk} \mid E$)" to denote that conditioned on some event E, some other event E' happens with probability at least

$$\mathbb{P}(E' \mid E) \geq 1 - \frac{c}{dk},$$

for all $d \geq d_0$ and $k \geq k_0$, where d_0, k_0 and c are positive constants.

(iii) The negative log-likelihood function and its derivative

Recall that ℓ denotes the negative log-likelihood function. Under the BTL model, we have

$$\begin{aligned}
\ell(\theta) := \ell(\{W_{ij}\};\theta) &= -\sum_{1\leq i<j\leq d}\left[W_{ij}\log\left(\frac{1}{1+e^{-(\theta_i-\theta_j)}}\right) + W_{ji}\log\left(\frac{1}{1+e^{-(\theta_j-\theta_i)}}\right)\right] \\
&= -k\sum_{1\leq i<j\leq d}\left[\mu_{ij}\log\left(\frac{1}{1+e^{-(\theta_i-\theta_j)}}\right) + \mu_{ji}\log\left(\frac{1}{1+e^{-(\theta_j-\theta_i)}}\right)\right] \\
&= k\sum_{1\leq i<j\leq d}\left[\log(e^{\theta_i}+e^{\theta_j}) - \mu_{ij}\theta_i - \mu_{ji}\theta_j\right]. \qquad (11)
\end{aligned}$$

Since $\{\mu_{ij}\}$ is simply a normalized version of $\{W_{ij}\}$, we equivalently denote the negative log-likelihood function as $\ell(\{\mu_{ij}\};\theta)$.

From the expression of ℓ in (11), we compute the gradient $\frac{\partial \ell}{\partial \theta_m}$ for every $m \in [d]$ as

$$\frac{\partial \ell}{\partial \theta_m} = k\sum_{i\neq m}\left(\frac{1}{1+e^{-(\theta_m-\theta_i)}} - \mu_{mi}\right). \qquad (12)$$

Finally, the following lemma from Hunter (2004) shows the strict convexity of the negative log-likelihood function ℓ.

Lemma A.1 (Lemma 2(a) from Hunter (2004)). *The negative log-likelihood function $\ell(\theta)$ is strictly convex in $\theta \in \mathbb{R}^d$.*

(iv) The sigmoid function and its derivatives

Denote the function $f : (-\infty, \infty) \to (0, 1)$ as the sigmoid function $f(x) = \frac{1}{1+e^{-x}}$. It is straightforward to verify that the function f has the following two properties.

- The first derivative f' is positive on $(-\infty, \infty)$. Moreover, on any bounded interval, the first derivative f' is bounded above and below. That is, for any constants $c_1 < c_2$, there exist constants $c_3, c_4 > 0$ such that

$$0 < c_3 < f'(x) < c_4, \qquad \text{for all } x \in (c_1, c_2). \qquad (13a)$$

- The second derivative f'' is bounded on any bounded interval. That is, for any constants $c_1 < c_2$, there exists a constant c_5 such that

$$|f''(x)| < c_5, \qquad \text{for all } x \in (c_1, c_2). \qquad (13b)$$

(v) Existence and uniqueness of MLE

Recall that the MLE (3), the unconstrained MLE (5), and the stretched-MLE (6) are respectively defined as:

$$\widehat{\theta}^{(B)}(\{\mu_{ij}\}) = \operatorname*{argmin}_{\theta\in\Theta_B} \ell(\{\mu_{ij}\};\theta), \qquad (14)$$

$$\widehat{\theta}^{(\infty)}(\{\mu_{ij}\}) = \operatorname*{argmin}_{\theta\in\Theta_\infty} \ell(\{\mu_{ij}\};\theta), \qquad (15)$$

$$\widehat{\theta}^{(A)}(\{\mu_{ij}\}) = \operatorname*{argmin}_{\theta\in\Theta_A} \ell(\{\mu_{ij}\};\theta). \qquad (16)$$

The following lemma shows the existence and uniqueness of the stretched-MLE $\widehat{\theta}^{(A)}$ (16) for any constant $A > 0$, which incorporates the standard MLE $\widehat{\theta}^{(B)}$ by setting $A = B$.

Lemma A.2. *For any finite constant $A > 0$, there always exists a unique solution $\widehat{\theta}^{(A)}$ to the stretched-MLE* (16)*.*

See Appendix A.4.1 for the proof of Lemma A.2.

For the unconstrained MLE, due to the removal of the box constraint in (15), a finite solution $\widehat{\theta}^{(\infty)}$ may not exist. However, the following lemma shows that a unique finite solution exists with high probability.

Lemma A.3. *There exists a unique finite solution $\widehat{\theta}^{(\infty)}$ to the unconstrained MLE* (15) *w.h.p.($\frac{1}{dk}$).*

See Appendix A.4.2 for the proof of Lemma A.3.

In the subsequent proofs of Theorem 2.1(b) and Theorem 2.1(a), we heavily use the unconstrained MLE as an intermediate quantity to analyze the MLE and the stretched-MLE.

A.2 Proof of Theorem 2.1(b)

In this appendix, we present the proof of Theorem 2.1(b). To describe the main steps involved, we first present a proof sketch of a simple case of $d = 2$ items (Appendix A.2.1), followed by the complete proof of the general case (Appendix A.2.2). The reader may pass to the complete proof in Appendix A.2.2 without loss of continuity.

A.2.1 Simple case: 2 items

We first present an informal proof sketch for a simple case where there are $d = 2$ items. The proof for the general case in Appendix A.2.2 follows the same outline. In the case of $d = 2$ items, due to the centering constraint on the true parameter vector θ^*, we have $\theta_2^* = -\theta_1^*$. Similarly, we have $\widehat{\theta}_2 = -\widehat{\theta}_1$ for any estimator that satisfies the centering constraint (in particular, for the stretched-MLE $\widehat{\theta}^{(A)}$ and the unconstrained MLE $\widehat{\theta}^{(\infty)}$). Therefore, it suffices to focus only on item 1. Since there are only two items, for ease of notation, we denote $\mu = \mu_{12}$ and $\mu^* = \mu_{12}^*$. We now present the main steps of the proof sketch.

Proof sketch of the 2-item case (informal):

In the proof sketch, we fix any $\theta^* \in \Theta_B$, and any finite constants A and B such that $A > B > 0$.

Step 1: Establish concentration of μ

By Hoeffding's inequality, we have

$$|\mu - \mu^*| \lesssim \sqrt{\frac{\log k}{k}}, \qquad \text{w.h.p.} \tag{17}$$

Since $|\theta^*| \leq B$, we have that μ^* is bounded away from 0 and 1 by a constant. Hence, for sufficiently large k, there exist constants c_L, c_U where $0 < c_L < c_U < 1$, such that

$$\mu, \mu^* \in (c_L, c_U). \tag{18}$$

Step 2: Write the first-order optimality condition for $\widehat{\theta}^{(\infty)}$

The unconstrained MLE $\widehat{\theta}^{(\infty)}$ minimizes the negative log-likelihood ℓ. If a finite unconstrained MLE $\widehat{\theta}^{(\infty)}$ exists[1], we have $\nabla_{\theta=\widehat{\theta}^{(\infty)}} \ell(\theta) = 0$. Setting $m = 1$ in the gradient expression (12) and plugging in $\widehat{\theta}^{(\infty)}$, we have

$$\begin{aligned} \frac{\partial \ell}{\partial \theta_1}\bigg|_{\theta=\widehat{\theta}^{(\infty)}} &= k\left(\frac{1}{1+e^{-(\widehat{\theta}_1^{(\infty)} - \widehat{\theta}_2^{(\infty)})}} - \mu_{12}\right) \\ &= k\left(\frac{1}{1+e^{-2\widehat{\theta}_1^{(\infty)}}} - \mu\right). \end{aligned} \tag{19}$$

[1] For the proof sketch, we ignore the high-probability nature of Lemma A.3, and assume that a finite $\widehat{\theta}^{(\infty)}$ always exists. It is made precise in the complete proof in Appendix A.2.2.

Setting the derivative (19) to 0, we have

$$\widehat{\theta}_1^{(\infty)} = -\frac{1}{2}\log\left(\frac{1}{\mu} - 1\right). \tag{20}$$

By the definition of $\{\mu^*_{ij}\}$ in (9), we have $\mu^* = \frac{1}{1+e^{-(\theta_1^*-\theta_2^*)}} = \frac{1}{1+e^{-2\theta_1^*}}$, which can be written as

$$\theta_1^* = -\frac{1}{2}\log\left(\frac{1}{\mu^*} - 1\right). \tag{21}$$

Define a function $h : [0, 1] \to \mathbb{R} \cup \{\pm\infty\}$ as

$$h(t) = -\frac{1}{2}\log\left(\frac{1}{t} - 1\right). \tag{22}$$

Subtracting (21) from (20) and using the definition of h from (22), we have

$$\widehat{\theta}_1^{(\infty)} - \theta_1^* = h(\mu) - h(\mu^*). \tag{23}$$

Step 3: Bound the difference between $\widehat{\theta}^{(\infty)}$ and θ^*, by the first-order mean value theorem

It can be verified that h has positive first-order derivative on $(0, 1)$. Moreover, there exists some constant c_1 such that $0 < h'(t) < c_1$ for all $t \in (c_L, c_U)$. Applying the first-order mean value theorem on (23), we have the deterministic relation

$$\widehat{\theta}_1^{(\infty)} - \theta_1^* = h'(\lambda) \cdot (\mu - \mu^*), \tag{24}$$

where λ is a random variable that depends on μ and μ^*, and takes values between μ and μ^*. By (18), we have $\lambda \in (c_L, c_U)$. From (24) we have

$$|\widehat{\theta}_1^{(\infty)} - \theta_1^*| \leq c_1|\mu - \mu^*|. \tag{25}$$

Combining (25) with (17), we have

$$|\widehat{\theta}_1^{(\infty)} - \theta_1^*| \lesssim \sqrt{\frac{\log k}{k}}, \qquad \text{w.h.p.} \tag{26}$$

Step 4: Bound the *expected* difference between $\widehat{\theta}^{(\infty)}$ and θ^*, by the second-order mean value theorem

By the second-order mean value theorem on (23), we have the deterministic relation

$$\widehat{\theta}_1^{(\infty)} - \theta_1^* = h(\mu) - h(\mu^*) = h'(\mu^*) \cdot (\mu - \mu^*) + h''(\widetilde{\lambda}) \cdot (\mu - \mu^*)^2, \tag{27}$$

where $\widetilde{\lambda}$ is a random variable that depends on μ and μ^*, and takes values between μ and μ^*. By (18), we have $\widetilde{\lambda} \in (c_L, c_U)$.

It can be verified that h has bounded second-order derivative. That is, $|h''(t)| < c_2$ for all $t \in (c_L, c_U)$. Taking an expectation over (27), we have

$$\begin{aligned}\mathbb{E}[\widehat{\theta}_1^{(\infty)}] - \theta_1^* &= h'(\mu^*) \cdot (\mathbb{E}[\mu] - \mu^*) + \mathbb{E}[h''(\widetilde{\lambda}) \cdot (\mu - \mu^*)^2] \qquad (28)\\ &\overset{\text{(i)}}{\leq} c_2\mathbb{E}[(\mu - \mu^*)^2]\\ &\overset{\text{(ii)}}{\lesssim} \frac{\log k}{k}, \qquad (29)\end{aligned}$$

where (i) is true because $\mathbb{E}[\mu] = \mu^*$ combined with the fact that $|h''| < c_2$ on (c_L, c_U), and (ii) is true[2] by (17).

[2] For the proof sketch, we ignore the high-probability nature of (17) and treat it as a deterministic relation. It is made precise in the complete proof in Appendix A.2.2.

Step 5: Connect $\widehat{\theta}^{(\infty)}$ back to $\widehat{\theta}^{(A)}$

From (26), we have $|\widehat{\theta}_1^{(\infty)} - \theta_1^*| \leq A - B$ w.h.p. for sufficiently large k. Hence,

$$|\widehat{\theta}_1^{(\infty)}| \leq |\theta_1^*| + |\widehat{\theta}_1^{(\infty)} - \theta_1^*| \leq B + (A - B) = A, \qquad \text{w.h.p.}$$

Moreover, we have $\left|\widehat{\theta}_2^{(\infty)}\right| = \left|\widehat{\theta}_1^{(\infty)}\right| \leq A$. Therefore, with high probability, the unconstrained MLE $\widehat{\theta}^{(\infty)}$ does not violate the box constraint at A, and therefore $\widehat{\theta}^{(\infty)}$ is identical to the stretched-MLE $\widehat{\theta}^{(A)}$. Hence, the bound (29) holds[3] for the stretched-MLE, completing the proof sketch.

A.2.2 Complete Proof

In this appendix, we present the proof of Theorem 2.1(b), by formally extending the 5 steps outlined for the simple case in Appendix A.2.1. In the general case, one notable challenge is that one can no longer write a closed-form solution of the MLE as we did in (20) of Step 2. The first-order optimality condition now becomes a system of equations that describe an implicit relation between θ and μ, requiring more involved analysis.

In the proof, we fix any $\theta^* \in \Theta_B$, and fix any finite constants A and B such that $A > B > 0$.

Step 1: Establish concentration of $\{\mu_{ij}\}$

We first use standard concentration inequalities to establish the following lemma, to be used in the subsequent steps of the proof.

Lemma A.4. *There exists a constant $c > 0$, such that*

$$\left|\sum_{i \neq m} \mu_{mi} - \sum_{i \neq m} \mu_{mi}^*\right| \leq c\sqrt{\frac{d(\log d + \log k)}{k}},$$

simultaneously for all $m \in [d]$ w.h.p.($\frac{1}{dk}$).

See Appendix A.4.3 for the proof of Lemma A.4.

Recall that Lemma A.3 states that a finite unconstrained MLE $\widehat{\theta}^{(\infty)}$ exists w.h.p.($\frac{1}{dk}$). We denote E_0 as the event that Lemma A.3 and Lemma A.4 both hold. For the rest of the proof, we condition on E_0. Since both Lemma A.3 and Lemma A.4 hold w.h.p.($\frac{1}{dk}$), taking a union bound, we have that E_0 holds w.h.p.($\frac{1}{dk}$). That is,

$$\mathbb{P}(E_0) \geq 1 - \frac{c}{dk}, \qquad \text{for some constant } c > 0. \tag{30}$$

Step 2: Write the first-order optimality condition for the unconstrained MLE $\widehat{\theta}^{(\infty)}$

Recall from Lemma A.1 that the negative log-likelihood function ℓ is convex in θ. In this step, we first justify that the whenever a finite unconstrained MLE $\widehat{\theta}^{(\infty)}$ exists, it satisfies the first-order optimality condition $\nabla_{\theta=\widehat{\theta}^{(\infty)}} \ell(\theta) = 0$. (Note that for any optimization problem with constraints, it is in general not true that the derivative of the convex objective equals 0 at the optimal solution.) Then we derive a specific form of the first-order optimality condition, to be used in subsequent steps of the proof.

Given that we have conditioned on E_0 (and therefore on Lemma A.3), a finite solution $\widehat{\theta}^{(\infty)}$ to the unconstrained MLE exists. To show that $\widehat{\theta}^{(\infty)}$ satisfies the first-order optimality condition, we show that $\widehat{\theta}^{(\infty)}$ is also a solution to the following MLE without any constraint at all (that is, we remove the centering constraint too):

$$\operatorname*{argmin}_{\theta \in \mathbb{R}^d} \ell(\theta). \tag{31}$$

If the unconstrained MLE $\widehat{\theta}^{(\infty)}$ is a solution to (31), then it satisfies the first-order condition $\nabla_\theta \ell(\widehat{\theta}^{(\infty)}) = 0$. Now we prove that $\widehat{\theta}^{(\infty)}$ is a solution to (31). Note that the solutions to (31) are shift-invariant. That is, if θ is a

[3] For the proof sketch, we ignore the high-probability nature of the fact that $\widehat{\theta}^{(\infty)} = \widehat{\theta}^{(A)}$, and treat it as a deterministic relation. It is made precise in the complete proof in Appendix A.2.2.

solution to (31), then $\theta + c\mathbf{1}$ is also a solution, where $\mathbf{1}$ is the d-dimensional all-one vector, and c is any constant. Now suppose by contradiction that $\widehat{\theta}^{(\infty)}$ is not a solution to (31). Then there exists some finite $\theta \in \mathbb{R}^d$ such that $\ell(\theta) < \ell(\widehat{\theta}^{(\infty)})$. Now consider $\theta' := \theta - (\frac{1}{d}\sum_{i=1}^d \theta_i)\mathbf{1}$. We have $\theta' \in \Theta_\infty$ because it satisfies the centering constraint, and we have $\ell(\theta') = \ell(\theta) < \ell(\widehat{\theta}^{(\infty)})$ because the solutions to (31) are shift-invariant. The construction of θ' thus contradicts the assumption that $\widehat{\theta}^{(\infty)}$ is optimal for the unconstrained MLE. Hence, $\widehat{\theta}^{(\infty)}$ is a solution to (31), and $\widehat{\theta}^{(\infty)}$ satisfies the first-order optimality condition.

Now we derive a specific form of the first-order optimality condition. Plugging $\widehat{\theta}^{(\infty)}$ into the gradient expression (12) and setting the gradient to 0, we have the deterministic equality

$$\sum_{i \neq m} \frac{1}{1 + e^{-(\widehat{\theta}_m^{(\infty)} - \widehat{\theta}_i^{(\infty)})}} = \sum_{i \neq m} \mu_{mi}, \qquad \text{for every } m \in [d]. \tag{32}$$

In words, the first-order optimality condition (32) means that for any item $m \in [d]$, the probability that item m wins (among all comparisons in which item m is involved) as predicted by the unconstrained MLE $\widehat{\theta}^{(\infty)}$ equals the fraction of wins by item m from the observed comparisons. We now subtract (9) from both sides of (32):

$$\sum_{i \neq m} \left(\frac{1}{1 + e^{-(\widehat{\theta}_m^{(\infty)} - \widehat{\theta}_i^{(\infty)})}} - \frac{1}{1 + e^{-(\theta_m^* - \theta_i^*)}} \right) = \sum_{i \neq m} (\mu_{mi} - \mu_{mi}^*)$$
$$\sum_{i=1}^{d} \left(\frac{1}{1 + e^{-(\widehat{\theta}_m^{(\infty)} - \widehat{\theta}_i^{(\infty)})}} - \frac{1}{1 + e^{-(\theta_m^* - \theta_i^*)}} \right) = \sum_{i \neq m} (\mu_{mi} - \mu_{mi}^*). \tag{33}$$

For ease of notation, we denote the random vector $\delta := \widehat{\theta}^{(\infty)} - \theta^*$. Equivalently, we have $\widehat{\theta}^{(\infty)} = \theta^* + \delta$. Using the definition of δ, we rewrite (33) as:

$$\sum_{i=1}^{d} \left(\frac{1}{1 + e^{-(\theta_m^* - \theta_i^* + \delta_m - \delta_i)}} - \frac{1}{1 + e^{-(\theta_m^* - \theta_i^*)}} \right) = \sum_{i \neq m} (\mu_{mi} - \mu_{mi}^*). \tag{34}$$

Using the definition of the sigmoid function $f(x) = \frac{1}{1+e^{-x}}$, we rewrite (34) as:

$$\sum_{i=1}^{d} [f(\theta_m^* - \theta_i^* + \delta_m - \delta_i) - f(\theta_m^* - \theta_i^*)] = \sum_{i \neq m} (\mu_{mi} - \mu_{mi}^*). \tag{35}$$

In the rest of the proof, we primarily work with the first-order optimality condition in the form of (35).

Step 3: Bound the difference between the unconstrained MLE $\widehat{\theta}^{(\infty)}$ and the true parameter vector θ^*

The first-order optimality condition (35) can be thought of as a system of equations that describes some implicit relation between the unconstrained MLE $\widehat{\theta}^{(\infty)}$ and the observations $\{\mu_{mi}\}$. Intuitively, the concentration of $\{\mu_{mi}\}$ on the RHS of (35) (by Lemma A.4) should imply the concentration of the unconstrained MLE $\widehat{\theta}^{(\infty)}$ on the LHS. The following lemma formalizes this intuition about the concentration of $\widehat{\theta}^{(\infty)}$.

Lemma A.5. *Conditioned on E_0, we have the deterministic relation*

$$|\delta_m| = |\widehat{\theta}_m^{(\infty)} - \theta_m^*| \lesssim \sqrt{\frac{\log d + \log k}{dk}}, \qquad \textit{for every } m \in [d],$$

for all $d \geq d_0$ and $k \geq k_0$, where d_0 and k_0 are constants.

See Appendix A.4.4 for the proof of Lemma A.5.

This lemma provides a deterministic bound on the difference between $\widehat{\theta}^{(\infty)}$ and θ^*. Now we move to analyze the difference between $\widehat{\theta}^{(\infty)}$ and θ^* in expectation.

Step 4: Bound the *expected* difference between the unconstrained MLE $\widehat{\theta}^{(\infty)}$ and the true parameter vector θ^*, using the second-order mean value theorem

In Step 1 we bound the difference between $\{\mu_{mi}\}$ and $\{\mu^*_{mi}\}$ with high-probability. However, if we consider the difference in expectation, we have $\mathbb{E}[\mu_{mi}] = \mu^*_{mi}$. The *expected* difference between $\{\mu_{mi}\}$ and $\{\mu^*_{mi}\}$ is 0, significantly smaller than the high-probability bound in Step 1. Intuitively, we may also expect that the *expected* difference between $\widehat{\theta}^{(\infty)}$ and θ^* is smaller than the deterministic bound in Lemma A.5. In this step, we formalize this intuition.

By the second-order mean value theorem on the LHS of the first-order optimality condition (35), we have the deterministic relation that for every $m \in [d]$,

$$\sum_{i=1}^{d}\left[f'(\theta^*_m - \theta^*_i)\cdot(\delta_m - \delta_i) + \frac{1}{2}f''(\lambda_{mi})\cdot(\delta_m-\delta_i)^2\right] = \sum_{i\neq m}(\mu_{mi} - \mu^*_{mi})$$

$$\sum_{i=1}^{d} f'(\theta^*_m - \theta^*_i)\cdot(\delta_m - \delta_i) = \sum_{i\neq m}(\mu_{mi} - \mu^*_{mi}) - \frac{1}{2}\sum_{i=1}^{d} f''(\lambda_{mi})\cdot(\delta_m - \delta_i)^2, \tag{36}$$

where each λ_{mi} is a random variable that takes values between $\theta^*_m - \theta^*_i$ and $\theta^*_m - \theta^*_i + (\delta_m - \delta_i)$. Taking an expectation over (36) conditional on E_0, we have that for every $m \in [d]$:

$$\sum_{i=1}^{d} f'(\theta^*_m - \theta^*_i)\cdot\mathbb{E}\left[\delta_m - \delta_i \mid E_0\right] = \sum_{i\neq m}(\mathbb{E}[\mu_{mi}\mid E_0] - \mu^*_{mi}) - \frac{1}{2}\sum_{i=1}^{d}\mathbb{E}[f''(\lambda_{mi})(\delta_m-\delta_i)^2 \mid E_0]. \tag{37}$$

Denote the vector $\Delta := \mathbb{E}[\delta \mid E_0] = \mathbb{E}[\widehat{\theta}^{(\infty)} \mid E_0] - \theta^*$. Plugging this definition of Δ into (37) yields

$$\sum_{i=1}^{d} f'(\theta^*_m - \theta^*_i)\cdot(\Delta_m - \Delta_i) = \sum_{i\neq m}(\mathbb{E}[\mu_{mi}\mid E_0] - \mu^*_{mi}) - \frac{1}{2}\sum_{i=1}^{d}\mathbb{E}[f''(\lambda_{mi})(\delta_m-\delta_i)^2 \mid E_0]. \tag{38}$$

We first bound the RHS of (38), and then derive a bound regarding Δ_i on the LHS accordingly.

To bound the RHS of (38), we first consider the term $\mathbb{E}[\mu_{mi} \mid E_0] - \mu^*_{mi}$. In what follows, we state a lemma that is slightly more general than what is needed here. The more general version is used in the subsequent proof of Theorem 2.1(a). To state the lemma, recall the definition that an event E' happens w.h.p.($\frac{1}{dk} \mid E$), if the conditional probability $\mathbb{P}(E' \mid E) \geq 1 - \frac{c}{dk}$, for some constant $c > 0$.

Lemma A.6. *Let E be any event, and let E' be any event that happens w.h.p.($\frac{1}{dk} \mid E$). Then for any $m \neq i$, we have*

$$|\mathbb{E}[\mu_{mi} \mid E', E] - \mathbb{E}[\mu_{mi} \mid E]| \lesssim \frac{1}{dk}. \tag{39}$$

See Appendix A.4.5 for the proof of Lemma A.6.

To apply Lemma A.6, we set E to be the (trivial) event of the entire probability space, and set E' to be E_0 in (39). We have

$$|\mathbb{E}[\mu_{mi} \mid E_0] - \mathbb{E}[\mu_{mi}]| = |\mathbb{E}[\mu_{mi} \mid E_0] - \mu^*_{mi}| \lesssim \frac{1}{dk}. \tag{40}$$

The remaining terms in (38) are handled in the following lemma. This lemma bounds the expected difference between $\widehat{\theta}^{(\infty)}$ and θ^* conditioned on E_0, that is, the quantity $|\Delta_m| = |\mathbb{E}[\widehat{\theta}^{(\infty)}_m \mid E_0] - \theta^*_m|$.

Lemma A.7. *Conditioned on E_0, we have*

$$|\Delta_m| \lesssim \frac{\log d + \log k}{dk}, \qquad \textit{for every } m \in [d],$$

for all $d \geq d_0$ and all $k \geq k_0$, where d_0 and k_0 are constants. Equivalently,

$$\beta(\widehat{\theta}^{(\infty)} \mid E_0) = \|\mathbb{E}[\widehat{\theta}^{(\infty)} \mid E_0] - \theta^*\|_\infty = \|\Delta\|_\infty \lesssim \frac{\log d + \log k}{dk}, \tag{41}$$

for all $d \geq d_0$ and all $k \geq k_0$, where d_0 and k_0 are constants.

See Appendix A.4.6 for the proof of Lemma A.7.

Note that (41) yields the desired rate on the quantity $\beta(\widehat{\theta}^{(\infty)} \mid E_0)$. It remains to show that $\beta(\widehat{\theta}^{(\infty)} \mid E_0)$ is sufficiently close to $\beta(\widehat{\theta}^{(A)})$.

Step 5: Show that the box constraint at A is vacuous for the unconstrained MLE $\widehat{\theta}^{(\infty)}$ and hence $\widehat{\theta}^{(\infty)}$ is the same as the stretched-MLE $\widehat{\theta}^{(A)}$ with high probability, using the deterministic bound in Step 3

To show that $\beta(\widehat{\theta}^{(\infty)} \mid E_0)$ is sufficiently close to $\beta(\widehat{\theta}^{(A)})$, we divide the argument into two parts. First, we show that $\beta(\widehat{\theta}^{(\infty)} \mid E_0) = \beta(\widehat{\theta}^{(A)} \mid E_0)$. Second, we show that $\beta(\widehat{\theta}^{(A)} \mid E_0)$ is close to $\beta(\widehat{\theta}^{(A)})$.

We first show that $\beta(\widehat{\theta}^{(\infty)} \mid E_0) = \beta(\widehat{\theta}^{(A)} \mid E_0)$. Recall that A and B are constants such that $A > B$. Recall from Lemma A.5 that $\|\widehat{\theta}^{(\infty)} - \theta^*\|_\infty \lesssim \frac{\log d + \log k}{dk}$ conditioned on E_0. Hence, there exist constants d_0 and k_0, such that for any $d \geq d_0$ and $k \geq k_0$, we have $\|\widehat{\theta}^{(\infty)} - \theta^*\|_\infty < A - B$ conditioned on E_0. In this case, we have

$$\|\widehat{\theta}^{(\infty)}\|_\infty \leq \|\theta^*\|_\infty + \|\widehat{\theta}^{(\infty)} - \theta^*\|_\infty < B + (A - B) = A, \qquad \text{conditioned on } E_0.$$

Conditioned on E_0, the unconstrained MLE $\widehat{\theta}^{(\infty)}$ obeys the box constraint $\|\widehat{\theta}^{(\infty)}\|_\infty \leq A$. Therefore, $\widehat{\theta}^{(\infty)}$ is also a solution to the stretched-MLE $\widehat{\theta}^{(A)}$. By the uniqueness of $\widehat{\theta}^{(A)}$ from Lemma A.2, we have

$$\widehat{\theta}^{(A)} = \widehat{\theta}^{(\infty)}, \qquad \text{conditioned on } E_0.$$

Hence, we have the relation

$$\beta(\widehat{\theta}^{(\infty)} \mid E_0) = \beta(\widehat{\theta}^{(A)} \mid E_0), \tag{42}$$

completing the first part of the argument.

It remains to show that $\beta(\widehat{\theta}^{(A)} \mid E_0)$ is sufficiently close to $\beta(\widehat{\theta}^{(A)})$. We have

$$\begin{aligned}
\beta(\widehat{\theta}^{(A)}) &= \|\mathbb{E}[\widehat{\theta}^{(A)}] - \theta^*\|_\infty \\
&\overset{\text{(i)}}{=} \|\mathbb{E}[\widehat{\theta}^{(A)} \mid E_0] \cdot \mathbb{P}(E_0) + \mathbb{E}[\widehat{\theta}^{(A)} \mid \overline{E_0}] \cdot \mathbb{P}(\overline{E_0}) - \theta^*\|_\infty \\
&\overset{\text{(ii)}}{\leq} \|\mathbb{E}[\widehat{\theta}^{(A)} \mid E_0] - \theta^*\|_\infty \cdot \mathbb{P}(E_0) + \|\mathbb{E}[\widehat{\theta}^{(A)} \mid \overline{E_0}] - \theta^*\|_\infty \cdot \mathbb{P}(\overline{E_0}) \\
&= \underbrace{\beta(\widehat{\theta}^{(A)} \mid E_0) \cdot \mathbb{P}(E_0)}_{R_1} + \underbrace{\|\mathbb{E}[\widehat{\theta}^{(A)} \mid \overline{E_0}] - \theta^*\|_\infty \cdot \mathbb{P}(\overline{E_0})}_{R_2}.
\end{aligned} \tag{43}$$

where step (i) is true by the law of iterated expectation, and step (ii) is true by the triangle inequality.

Consider the two terms in (43). For R_1, combining (41) and (42) yields

$$\beta(\widehat{\theta}^{(A)} \mid E_0) = \beta(\widehat{\theta}^{(\infty)} \mid E_0) \lesssim \frac{\log d + \log k}{dk}.$$

Therefore,

$$R_1 \lesssim \frac{\log d + \log k}{dk}. \tag{44}$$

Now consider R_2. By the box constraint $\|\widehat{\theta}^{(A)}\|_\infty \leq A$, we have

$$\|\mathbb{E}[\widehat{\theta}^{(A)} \mid \overline{E_0}] - \theta^*\|_\infty \overset{\text{(i)}}{\leq} \|\mathbb{E}[\widehat{\theta}^{(A)} \mid \overline{E_0}]\|_\infty + \|\theta^*\|_\infty \leq A + B, \tag{45}$$

where step (i) is true by the triangle inequality. Recall from (30), the event E_0 happens w.h.p.$(\frac{1}{dk})$. Therefore,

$$\mathbb{P}(\overline{E_0}) \lesssim \frac{1}{dk}. \tag{46}$$

Combining (45) and (46) yields

$$R_2 \lesssim \frac{1}{dk}. \tag{47}$$

Plugging the term R_1 from (44) and the term R_2 from (47) back into (43), we have

$$\beta(\widehat{\theta}^{(A)}) \lesssim \frac{\log d + \log k}{dk},$$

completing the proof of Theorem 2.1(b).

A.3 Proof of Theorem 2.1(a)

Similar to the proof of Theorem 2.1(b), we first present a proof of the simple case of $d = 2$ items. It is important to note that although we present proofs of the 2-item case for both Theorem 2.1(b) and Theorem 2.1(a), their purposes are different. In Theorem 2.1(b) presented in Appendix A.2, the proof sketch of the 2-item case is informal. It serves as a guideline for the general case. Then the main work involved in the general case is to generalize the arguments in the 2-item case step-by-step. On the other hand, in Theorem 2.1(a), the proof of the 2-item case to be presented is formal. It serves as a core sub-problem of the general case. Then the main work involved in the general case is to reduce the problem to the 2-item case, and then the results from the 2-item case directly.

A.3.1 Simple case: 2 items

As in Appendix A.2.1, we first consider the simple case where there are $d = 2$ items. Again, due to the centering constraint, we have $\theta_2^* = -\theta_1^*$ for the true parameter vector θ^*, and we have $\widehat{\theta}_2 = -\widehat{\theta}_1$ for any estimator $\widehat{\theta}$ that satisfies the centering constraint (in particular, for the standard MLE $\widehat{\theta}^{(B)}$ and the unconstrained MLE $\widehat{\theta}^{(\infty)}$). Therefore, it suffices to focus only on item 1. Since there are only two items, for ease of notation, we denote $\mu = \mu_{12}$ and $\mu^* = \mu_{12}^*$.

We consider the true parameter vector $\theta^* = [B, -B]$. By the definition of $\{\mu_{ij}^*\}$ in (9), we have

$$\mu^* = \frac{1}{1 + e^{-(\theta_1^* - \theta_2^*)}} = \frac{1}{1 + e^{-2B}}.$$

The following proposition now lower bounds the bias of the standard MLE $\widehat{\theta}^{(B)}$.

Proposition A.8. *Under $\theta^* = [B, -B]$, the bias of the MLE $\widehat{\theta}^{(B)}$ is bounded as*

$$\beta(\widehat{\theta}^{(B)}) = \|\mathbb{E}[\widehat{\theta}^{(B)}] - \theta^*\|_\infty = |\mathbb{E}[\widehat{\theta}_1^{(B)}] - B| \gtrsim \frac{1}{\sqrt{k}}.$$

Specifically, the bias is negative, that is,

$$\mathbb{E}[\widehat{\theta}_1^{(B)}] - B \leq -\frac{c}{\sqrt{k}}, \tag{48}$$

for some constant $c > 0$.

The rest of this appendix is devoted to proving (48) in Proposition A.8.

For ease of notation, denote $\mu_+ = \mu^* = \frac{1}{1+e^{-2B}}$, and $\mu_- = 1 - \mu^* = \frac{1}{1+e^{2B}}$. In the proof sketch of Theorem 2.1(b) of the case of $d = 2$ items (Appendix A.2.1), we derived the following expression (20) for the unconstrained MLE:

$$\widehat{\theta}_1^{(\infty)}(\mu) = -\frac{1}{2}\log\left(\frac{1}{\mu} - 1\right).$$

Now consider the standard MLE $\widehat{\theta}^{(B)}$. By straightforward analysis, one can derive the following closed-form expression for the standard MLE:

$$\widehat{\theta}_1^{(B)}(\mu) = \begin{cases} -B & \text{if } \mu \in [0, \mu_-] \\ -\frac{1}{2}\log\left(\frac{1}{\mu} - 1\right) & \text{if } \mu \in (\mu_-, \mu_+) \\ B & \text{if } \mu \in [\mu_+, 1]. \end{cases} \tag{49}$$

For ease of notation, we denote a function $h : [0,1] \to [-B, B]$ as

$$h(t) = \begin{cases} -B & \text{if } t \in [0, \mu_-] \\ -\frac{1}{2}\log\left(\frac{1}{t} - 1\right) & \text{if } t \in (\mu_-, \mu_+) \\ B & \text{if } t \in [\mu_+, 1], \end{cases} \tag{50}$$

where $h(t) = \widehat{\theta}_1^{(B)}(\mu = t)$ for any $t \in [0,1]$. Then the standard MLE (49) can be equivalently written as $h(\mu)$. To make the computation of the bias incurred by $\widehat{\theta}^{(B)}$ more tractable, we also define the following auxiliary function $h^+ : [0,1] \to [-B, B]$ as:

$$h^+(t) := \begin{cases} \frac{2B}{\mu_+}(t - \mu_+) + B & \text{if } t \in [0, \mu_+) \\ B & \text{if } t \in [\mu_+, 1]. \end{cases} \tag{51}$$

In words, the function h^+ is piecewise linear. On the interval $[0, \mu_+]$, it is a line passing through the points $(0, -B)$ and (μ_+, B). On the interval $[\mu_+, 1]$, its value equals the constant B. The following lemma now states a relation between $h^+(\mu)$ and $h(\mu)$ in expectation with respect to μ.

Lemma A.9. *Under $\theta^* = [B, -B]$, we have*

$$\mathbb{E}[h(\mu)] \leq \mathbb{E}[h^+(\mu)]. \tag{52}$$

See Appendix A.4.7 for the proof of Lemma A.9.

Now subtracting B from both sides of (52), we have

$$\mathbb{E}[\widehat{\theta}_1^{(B)}] - \theta_1^* = \mathbb{E}[h(\mu)] - B \leq \mathbb{E}[h^+(\mu)] - B. \tag{53}$$

The following lemma states that the bias introduced by $h^+(\mu)$ satisfies the desired rate from Proposition A.8.

Lemma A.10. *Under $\theta^* = [B, -B]$, we have*

$$\mathbb{E}[h^+(\mu)] - B \leq -\frac{c}{\sqrt{k}}, \tag{54}$$

for some constant $c > 0$.

See Appendix A.4.8 for the proof of Lemma A.10.

Combining (53) and (54), we have

$$\mathbb{E}[\widehat{\theta}_1^{(B)}] - \theta_1^* \leq -\frac{c}{\sqrt{k}},$$

completing the proof of (48) in Proposition A.8.

A.3.2 Complete Proof

In this appendix, we present the proof of Theorem 2.1(a). The proof reduces the general case to the 2-item case presented in Appendix A.3.1. In the reduction, we construct an "oracle" MLE, such that the oracle MLE yields identical estimates for item 2 through item d. Specifically, we consider an unconstrained oracle denoted by $\widetilde{\theta}^{(\infty)}$ (without the box constraint), and a constrained oracle denoted by $\widetilde{\theta}^{(B)}$ (with the box constraint at B), to be defined precisely in the proof shortly. Then we derive the closed-form expressions for $\widetilde{\theta}^{(\infty)}$ and $\widetilde{\theta}^{(B)}$, which bear resemblance to the expressions of the the unconstrained MLE and the standard MLE in the 2-item case. Using the proof of the 2-item case, we prove that the constrained oracle $\widetilde{\theta}^{(B)}$ incurs a negative bias of $\Omega(\frac{1}{\sqrt{dk}})$. Given this result, it remains to show that $\widetilde{\theta}^{(B)}$ and $\widehat{\theta}^{(B)}$ differ by $o(\frac{1}{\sqrt{dk}})$ in terms of bias. We decompose the difference between $\widetilde{\theta}^{(B)}$ and $\widehat{\theta}^{(B)}$ into three terms: from $\widetilde{\theta}^{(B)}$ to $\widetilde{\theta}^{(\infty)}$, from $\widetilde{\theta}^{(\infty)}$ to $\widehat{\theta}^{(\infty)}$, and from $\widehat{\theta}^{(\infty)}$ to $\widehat{\theta}^{(B)}$, The second term is bounded by $\widetilde{\mathcal{O}}(\frac{1}{dk})$ by modifying the upper-bound proof of Theorem 2.1(b). The first and the third terms are bounded by carefully analyzing the effect of the box constraint on the oracle MLE and the standard MLE, respectively.

In the proof, we fix any constant $B > 0$, and consider the true parameter vector:

$$\theta^* = \left[B, -\frac{B}{d-1}, -\frac{B}{d-1}, \ldots, -\frac{B}{d-1}\right]. \tag{55}$$

It can be verified that θ^* satisfies both the box constraint at B and the centering constraint, so we have $\theta^* \in \Theta_B$. We prove that the bias on item 1 is negative, and its magnitude is $\Omega(\frac{1}{\sqrt{dk}})$. That is, we prove that

$$\mathbb{E}[\widehat{\theta}_1^{(B)}] - \theta_1^* = \mathbb{E}[\widehat{\theta}_1^{(B)}] - B \leq -\frac{c}{\sqrt{dk}},$$

for some constant $c > 0$. The proof consists of the following 5 steps.

Step 1: Construct oracle estimators $\widetilde{\theta}^{(\infty)}$ (unconstrained) and $\widetilde{\theta}^{(B)}$ (constrained)

Recall that $\mu_{ij} \sim \frac{1}{k}\text{Binom}(k, \mu_{ij}^*)$ is a random variable representing the fraction of times that item i beats item j. We define μ_1 as fraction of wins by item 1, among all comparisons in which item 1 is involved:

$$\mu_1 := \frac{1}{d-1}\sum_{m=2}^{d} \mu_{1m}. \tag{56}$$

We similarly define the true probability $\mu_1^* = \frac{1}{d-1}\sum_{m=2}^{d} \mu_{1m}^*$. With the construction (55) of θ^*, we have $\mu_1^* = \frac{1}{1+e^{-\frac{d}{d-1}B}}$. Now we construct the following random quantities $\{\widetilde{\mu}_{ij}\}_{i\neq j}$ as a function of $\{\mu_{ij}\}_{i\neq j}$:

$$\widetilde{\mu}_{ij} = \begin{cases} \mu_1 & \text{if } i = 1,\ j \in \{2, \ldots, d\} \\ 1 - \mu_1 & \text{if } j = 1,\ i \in \{2, \ldots, d\} \\ \frac{1}{2} & \text{otherwise.} \end{cases} \tag{57}$$

Recall that $\widehat{\theta}^{(\infty)}(\{\mu_{ij}\})$ denotes the unconstrained MLE (15). Now define an "unconstrained oracle" MLE $\widetilde{\theta}^{(\infty)}$ as:

$$\begin{aligned} \widetilde{\theta}^{(\infty)}(\{\mu_{ij}\}) &:= \widehat{\theta}^{(\infty)}(\{\widetilde{\mu}_{ij}\}) \\ &= \operatorname*{argmin}_{\theta\in\Theta_\infty} \ell(\{\widetilde{\mu}_{ij}\}; \theta). \end{aligned} \tag{58a}$$

Similarly, define a "constrained oracle" MLE $\widetilde{\theta}^{(B)}$ as:

$$\begin{aligned} \widetilde{\theta}^{(B)}(\{\mu_{ij}\}) &:= \widehat{\theta}^{(B)}(\{\widetilde{\mu}_{ij}\}) \\ &= \operatorname*{argmin}_{\theta\in\Theta_B} \ell(\{\widetilde{\mu}_{ij}\}; \theta). \end{aligned} \tag{58b}$$

In the subsequent steps, these oracle estimators are used to reduce the general case to the 2-item case.

Step 2: Formalize the oracle information contained in the unconstrained oracle $\widetilde{\theta}^{(\infty)}$ and the constrained oracle $\widetilde{\theta}^{(B)}$

Note that the construction of $\{\widetilde{\mu}_{ij}\}$ in (57) is symmetric with respect to item 2 through item d, that is, for any two items i and i' where $i, i' \in \{2, \ldots, d\}$, we have $\widetilde{\mu}_{ij} = \widetilde{\mu}_{i'j}$ and $\widetilde{\mu}_{ji} = \widetilde{\mu}_{ji'}$ for every $i \in [d] \setminus \{j, j'\}$. Therefore, the construction of $\{\widetilde{\mu}_{ij}\}$ intuitively encodes the "oracle" that item 2 through item d have identical parameters. Formally, define the set $\Theta_{\text{oracle}} := \{\theta \in \mathbb{R}^d \mid \theta_2 = \cdots = \theta_d\}$. The following lemma states that the unconstrained oracle and the constrained oracle incorporate the set Θ_{oracle} into the domain of optimization without altering their solutions.

Lemma A.11. *The unconstrained oracle $\widetilde{\theta}^{(\infty)}$ can be equivalently written as*

$$\widetilde{\theta}^{(\infty)} = \operatorname*{argmin}_{\Theta_\infty\cap\Theta_{\text{oracle}}} \ell(\{\widetilde{\mu}_{ij}\}; \theta). \tag{59a}$$

That is, a solution to (58a) *exists if and only if a solution to* (59a) *exists. Moreover, when the solutions to* (58a) *and* (59a) *exist, they are identical.*

Similarly, the constrained oracle $\widetilde{\theta}^{(B)}$ can be equivalently written as

$$\widetilde{\theta}^{(B)} = \operatorname*{argmin}_{\theta\in\Theta_B\cap\Theta_{\text{oracle}}} \ell(\{\widetilde{\mu}_{ij}\}; \theta). \tag{59b}$$

See Appendix A.4.9 for the proof of Lemma A.11.

Given Lemma A.11 combined with the centering constraint, we parameterize the unconstrained oracle $\widetilde{\theta}^{(\infty)}$ and the constrained oracle $\widetilde{\theta}^{(B)}$ as:

$$\widetilde{\theta}^{(\infty)} = \left[\widetilde{\theta}_1^{(\infty)}, -\frac{1}{d-1}\widetilde{\theta}_1^{(\infty)}, \ldots, -\frac{1}{d-1}\widetilde{\theta}_1^{(\infty)}\right], \tag{60a}$$

$$\widetilde{\theta}^{(B)} = \left[\widetilde{\theta}_1^{(B)}, -\frac{1}{d-1}\widetilde{\theta}_1^{(B)}, \ldots, -\frac{1}{d-1}\widetilde{\theta}_1^{(B)}\right]. \tag{60b}$$

Step 3: Show that the bias of the constrained oracle $\widetilde{\theta}^{(B)}$ on item 1 is bounded by $\mathbb{E}[\widetilde{\theta}_1^{(B)}]-\theta_1^* \leq -\frac{c}{\sqrt{dk}}$, by making a reduction to the 2-item case

In this step, we modify the proof of Proposition A.8 in the 2-item case to lower bound the bias of the constrained oracle $\widetilde{\theta}^{(B)}$. Specifically, we show that given $\theta^* = \left[B, -\frac{B}{d-1}, \ldots, -\frac{B}{d-1}\right]$, the bias on item 1 is bounded as (cf. (48)):

$$\mathbb{E}[\widetilde{\theta}_1^{(B)}] - \theta^* \leq -\frac{c}{\sqrt{dk}},$$

for some constant $c > 0$.

First, we solve for the unconstrained oracle $\widetilde{\theta}^{(\infty)}$ and the constrained oracle $\widetilde{\theta}^{(B)}$ in closed form. Set $m = 1$ in the gradient expression (12). Plugging in the expressions for the unconstrained oracle $\widetilde{\theta}^{(\infty)}$ (60a) and the manipulated observations $\{\widetilde{\mu}_{ij}\}$ (57), we have

$$\left.\frac{\partial \ell}{\partial \theta_1}\right|_{\theta=\widetilde{\theta}^{(\infty)}} = k(d-1)\left[\frac{1}{1+e^{-\frac{d}{d-1}\widetilde{\theta}_1^{(\infty)}}} - \mu_1\right] \tag{61}$$

Setting the derivative (61) to 0, we have

$$\frac{1}{1+e^{-\frac{d}{d-1}\widetilde{\theta}_1^{(\infty)}}} = \mu_1$$

$$\widetilde{\theta}_1^{(\infty)} = -\frac{d-1}{d}\log\left(\frac{1}{\mu_1}-1\right). \tag{62}$$

Denote $\mu_{d,+} = \mu_1^* = \frac{1}{1+e^{-\frac{d}{d-1}B}}$, and $\mu_{d,-} = 1-\mu_{d,+} = \frac{1}{1+e^{\frac{d}{d-1}B}}$. In the notations $\mu_{d,+}$ and $\mu_{d,-}$, the dependency on d is made explicit. When the dependency on d does not need to be emphasized, we also use the shorthand notations μ_+ and μ_-. Now consider the constrained oracle $\widetilde{\theta}^{(B)}$. By straightforward analysis, one can derive the following closed-form expression for the constrained oracle:

$$\widetilde{\theta}_1^{(B)}(\mu_1) = \begin{cases} -B & \text{if } 0 \leq \mu_1 < \mu_{d,-} \\ -\frac{d-1}{d}\log\left(\frac{1}{\mu_1}-1\right) & \text{if } \mu_{d,-} < \mu_1 < \mu_{d,+} \\ B & \text{if } \mu_{d,+} \leq \mu_1 \leq 1. \end{cases} \tag{63}$$

Note the similarity between $\widetilde{\theta}^{(B)}$ in (63) and the 2-item case $\widehat{\theta}_1^{(B)}$ in (49) from Appendix A.3.1. Similar to the function h defined in (50) of the 2-item case, we denote a function $h_d : [0,1] \to [-B, B]$ as:

$$h_d(t) = \begin{cases} -B & \text{if } 0 \leq t < \mu_{d,-} \\ -\frac{d-1}{d}\log\left(\frac{1}{t}-1\right) & \text{if } \mu_{d,-} < t < \mu_{d,+} \\ B & \text{if } \mu_{d,+} \leq t \leq 1, \end{cases}$$

where $h_d(t) = \widetilde{\theta}_1^{(B)}(\mu_1 = t)$ for any $t \in [0,1]$. Then the estimator $\widetilde{\theta}_1^{(B)}(\mu)$ can be equivalently written as $h_d(\mu)$. Similar to the function h^+ defined in (51) of the 2-item case, we define an auxiliary function $h_d^+ : [0,1] \to [-B, B]$ as:

$$h_d^+(t) = \begin{cases} \frac{2B}{\mu_{d,+}}(t-\mu_{d,+}) + B & \text{if } 0 \leq t < \mu_{d,+} \\ B & \text{if } \mu_{d,+} \leq t \leq 1. \end{cases}$$

Note that in the proofs of Lemma A.9 and Lemma A.10, we have only relied on the following two facts:

- There exists a constant c such that

$$\frac{1}{2} < \mu_+ < c < 1.$$

- The random variable μ is sampled as $\mu \sim \frac{1}{k}\text{Binom}(k, \mu_+)$.

In the general case, it can be verified that

- There exists a constant c such that

$$\frac{1}{2} < \mu_{d,+} < c < 1, \qquad \text{for all } d \geq 2.$$

- The random variable μ_1 as defined in (57) is sampled as $\mu_1 \sim \frac{1}{k'}\text{Binom}(k', \mu_+)$, where $k' := (d-1)k$ denotes the total number of comparisons in which item 1 is involved.

To extend the arguments in the 2-item case to the general case, we replace μ by μ_1, replace μ_+ by $\mu_{d,+}$, replace h^+ by h_d^+, and replace k by k' in the proof of Proposition A.8. It can be verified that the arguments in Lemma A.9 and Lemma A.10 still hold after these replacements. Therefore, extending the arguments in Proposition A.8, we have that at $\theta^* = \left[B, -\frac{B}{d-1}, \dots, -\frac{B}{d-1}\right]$,

$$\mathbb{E}[\widetilde{\theta}_1^{(B)}] - \theta_1^* \leq -\frac{c}{\sqrt{k'}} = -\frac{c}{\sqrt{(d-1)k}} \leq -\frac{c'}{\sqrt{dk}}, \tag{64}$$

for some constants $c, c' > 0$.

Step 4: Bound the difference between the unconstrained oracle $\widetilde{\theta}^{(\infty)}$ and the unconstrained MLE $\widehat{\theta}^{(\infty)}$, by modifying the proof of Theorem 2.1(b)

Recall that the random variable μ_1 denotes the fraction of wins by item 1. In this step, we fix any real number $v \in [\frac{1}{2}, \mu_+]$, and denote E_v as the event that we observe $\mu_1 = v$. Then we prove that conditioned on the event E_v, the difference between the unconstrained oracle $\widetilde{\theta}^{(\infty)}$ and the unconstrained MLE $\widehat{\theta}^{(\infty)}$ is small in expectation, by modifying Step 1 to Step 4 in the upper-bound proof of Theorem 2.1(b) in Appendix A.2.2.

We first conceptually explain how to modify the proof of Theorem 2.1(b). Our goal is to bound the difference between $\widetilde{\theta}^{(\infty)}$ and $\widehat{\theta}^{(\infty)}$ in expectation conditioned on the event E_v. By the definition of $\{\widetilde{\mu}_{ij}\}$ in (57), the quantities $\{\widetilde{\mu}_{ij}\}$ are fixed (not random) conditioned on E_v, and hence the unconstrained oracle $\widetilde{\theta}^{(\infty)}$ is fixed conditioned on E_v. We therefore replace the role of the true parameter vector θ^* in the proof of Theorem 2.1(b) by the unconstrained oracle $\widetilde{\theta}^{(\infty)}$. Then we think of the actual observations $\{\mu_{ij}\}$ as a noisy version of $\{\widetilde{\mu}_{ij}\}$, and think of $\widehat{\theta}^{(\infty)}$ as the estimate for $\widetilde{\theta}^{(\infty)}$. Now we modify the proof of Theorem 2.1(b) to bound the expected difference between $\widehat{\theta}^{(\infty)}$ and $\widetilde{\theta}^{(\infty)}$ conditioned on E_v. At the end of this step, we provide more intuition why we need to condition on the event E_v.

Formally, we denote $\{\widetilde{\mu}_{ij}^v\}$ as the values of $\{\widetilde{\mu}_{ij}\}$ conditional on E_v. We denote $\widetilde{\theta}^v$ as the unconstrained oracle $\widetilde{\theta}^{(\infty)}$ conditional on E_v. It can be verified that $\{\widetilde{\mu}_{ij}^v\}$ and $\widetilde{\theta}^v$ are fixed (not random) given any $v \in [\frac{1}{2}, \mu_+]$. Conditioned on E_v, we think of $\widetilde{\theta}^v$ *as if* it is the "true" parameter vector to be estimated (replacing the role of θ^*), and think of $\{\widetilde{\mu}_{ij}^v\}$ *as if* it is the "true" underlying probabilities (replacing the role of $\{\mu_{ij}^*\}$).

Given the definition of $\{\widetilde{\mu}_{ij}\}$ in (57), we have that conditioned on event E_v,

$$\widetilde{\mu}_{ij}^v = \begin{cases} v & \text{if } i = 1,\ j \in \{2, \dots, d\} \\ 1 - v & \text{if } j = 1,\ i \in \{2, \dots, d\} \\ \frac{1}{2} & \text{otherwise.} \end{cases} \tag{65}$$

From the expression (62) of the unconstrained oracle $\widetilde{\theta}^{(\infty)}$, it can be verified that $\widetilde{\theta}^{(\infty)}$ satisfies the deterministic equality

$$\frac{1}{1 + e^{-\left(\widetilde{\theta}_i^{(\infty)} - \widetilde{\theta}_j^{(\infty)}\right)}} = \widetilde{\mu}_{ij}, \qquad \text{for all } i \neq j. \tag{66}$$

Now we start to replicate Step 1 to Step 4 in the proof of Theorem 2.1(b) presented in Appendix A.2.2.

To replicate *Step 1* of Theorem 2.1(b), recall that in the proof of Theorem 2.1(b), we condition on Lemma A.3 and Lemma A.4. We first establish the modified versions of these two lemmas, when conditioned on E_v.

Lemma A.12 (Conditional version of Lemma A.3)**.** *Conditioned on the event E_v, there exists a finite solution $\widehat{\theta}^{(\infty)}$ to the unconstrained MLE* (15) *w.h.p.($\frac{1}{dk} \mid E_v$).*

See Appendix A.4.10 for the proof of Lemma A.12.

Lemma A.13 (Conditional version of Lemma A.4)**.** *Conditioned on the event E_v, there exists a constant $c > 0$, such that*

$$\left| \sum_{i \neq m} \mu_{mi} - \sum_{i \neq m} \widetilde{\mu}_{mi}^v \right| \leq c \sqrt{\frac{d(\log d + \log k)}{k}}, \tag{67}$$

simultaneously for all $m \in [d]$ w.h.p.($\frac{1}{dk} \mid E_v$).

See Appendix A.4.11 for the proof of Lemma A.13.

Recall that we have conditioned on the event E_v. Denote E_0 as the event that Lemma A.12 and Lemma A.13 both hold. (Note that the event E_0 is defined for some fixed v, so to be precise, the event E_0 should be denoted as $E_{0,v}$. For ease of notation, we drop the subscript v.) Taking a union bound of Lemma A.12 and Lemma A.13, we have that E_0 happens w.h.p.($\frac{1}{dk} \mid E_v$). For the rest of the proof, we condition on the events (E_0, E_v).

To replicate *Step 2* of Theorem 2.1(b), we subtract equality (66) from both sides of (32). We obtain the (unconditional) deterministic equality:

$$\sum_{i=1}^{d} \left(\frac{1}{1 + e^{-(\widehat{\theta}_m^{(\infty)} - \widehat{\theta}_i^{(\infty)})}} - \frac{1}{1 + e^{-(\widetilde{\theta}_m^{(\infty)} - \widetilde{\theta}_i^{(\infty)})}} \right) = \sum_{i \neq m} (\mu_{mi} - \widetilde{\mu}_{mi}), \qquad \text{for every } m \in [d]. \tag{68}$$

Conditioning (68) on (E_0, E_v), we have the following deterministic equality, as a modified version of (33):

$$\sum_{i=1}^{d} \left(\frac{1}{1 + e^{-(\widehat{\theta}_m^{(\infty)} - \widehat{\theta}_i^{(\infty)})}} - \frac{1}{1 + e^{-(\widetilde{\theta}_m^{v} - \widetilde{\theta}_i^{v})}} \right) = \sum_{i \neq m} (\mu_{mi} - \widetilde{\mu}_{mi}^v), \qquad \text{conditioned on } (E_0, E_v). \tag{69}$$

To replicate *Step 3* of Theorem 2.1(b), note that v is bounded as $v \in [\frac{1}{2}, \mu_+]$. By the expression (62) of $\widetilde{\theta}^{(\infty)}$ (and hence of $\widetilde{\theta}^v$), it can be verified that $\widetilde{\theta}^v$ is bounded as $|\widetilde{\theta}^v| \leq c$ for some constant c. Denote $\widetilde{\delta} = \widehat{\theta}^{(\infty)} - \widetilde{\theta}^v$. Using the same arguments as in Lemma A.5, we have the deterministic relation that

$$\|\widehat{\theta}^{(\infty)} - \widetilde{\theta}^v\|_\infty = \|\widetilde{\delta}\|_\infty \lesssim \sqrt{\frac{\log d + \log k}{dk}}, \qquad \text{conditioned on } (E_0, E_v). \tag{70}$$

To replicate *Step 4* of Theorem 2.1(b), we first apply the second-order mean value theorem on (69), and then take an expectation conditional on (E_0, E_v). The following equation establishes a modified version of (37):

$$\sum_{i=1}^{d} f'(\widetilde{\theta}_m^v - \widetilde{\theta}_i^v) \cdot \mathbb{E}\left[\widetilde{\delta}_i - \widetilde{\delta}_m \mid E_0, E_v\right] = \sum_{i \neq m} (\mathbb{E}[\mu_{mi} \mid E_0, E_v] - \widetilde{\mu}_{mi}^v) - \frac{1}{2} \sum_{i=1}^{d} \mathbb{E}[f''(\lambda_{mi})(\widetilde{\delta}_m - \widetilde{\delta}_i)^2 \mid E_0, E_v], \tag{71}$$

where each λ_{mi} is a random variable that takes values between $\widetilde{\theta}_m^v - \widetilde{\theta}_i^v$ and $\widetilde{\theta}_m^v - \widetilde{\theta}_i^v + \widetilde{\delta}_m - \widetilde{\delta}_i$. To apply Lemma A.6, we set E as E_v, and set E' as E_0 in (39):

$$|\mathbb{E}[\mu_{ij} \mid E_0, E_v] - \mathbb{E}[\mu_{ij} \mid E_v]| \lesssim \frac{1}{dk}. \tag{72}$$

It can be verified that

$$\mathbb{E}[\mu_{ij} \mid E_v] = \widetilde{\mu}_{ij}^v. \tag{73}$$

Plugging (73) into (72), we have

$$|\mathbb{E}[\mu_{ij} \mid E_0, E_v] - \widetilde{\mu}_{ij}^v| \lesssim \frac{1}{dk}.$$

Using the same arguments as in Lemma A.7 to handle the remaining terms in (71), we have the following upper bound as a modified version of (41):

$$\|\mathbb{E}[\widehat{\theta}^{(\infty)} - \widetilde{\theta}^v \mid E_0, E_v]\|_\infty = \|\mathbb{E}[\widehat{\theta}^{(\infty)} - \widetilde{\theta}^{(\infty)} \mid E_0, E_v]\|_\infty \lesssim \frac{\log d + \log k}{dk}. \tag{74}$$

Now that we have established the desired result (74) of this step, we conclude this step with some intuition why we need to condition on E_v. Without conditioning on E_v, we could still have utilized the proof of Theorem 2.1(b), and could have established a result of the form (cf. (74)):

$$\|\mathbb{E}[\widehat{\theta}^{(\infty)} - \widetilde{\theta}^v \mid E_0]\|_\infty = \|\mathbb{E}[\widehat{\theta}^{(\infty)} - \widetilde{\theta}^{(\infty)} \mid E_0]\|_\infty \lesssim \frac{\log d + \log k}{dk}. \tag{75}$$

Our goal here is to bound the constrained oracle $\widehat{\theta}^{(B)}$ and the constrained MLE $\widetilde{\theta}^{(B)}$ in expectation. However, the fact that two *unconstrained* estimators are close in expectation does not imply that their *constrained* counterparts are close in expectation[4]. Therefore, a bound of the form (75) is not sufficient for our goal, and instead we need to establish some "pointwise" control between $\widehat{\theta}^{(\infty)}$ and $\widetilde{\theta}^{(\infty)}$. That is, whenever the box constraint has little effect on $\widetilde{\theta}^{(\infty)}$, we want to show that the box constraint also has little effect on $\widehat{\theta}^{(\infty)}$. Thus, we condition on the event E_v for any $v \in [\frac{1}{2}, \mu_+]$, and bound the difference between $\widehat{\theta}^{(\infty)}$ and $\widetilde{\theta}^{(\infty)}$ in expectation conditioned on E_v (that is, the bound in (74)). Given this pointwise result, we then integrate over v to establish the desired result that $\widehat{\theta}^{(B)}$ and $\widetilde{\theta}^{(B)}$ are close in expectation, to be presented in the subsequent step of the proof.

Step 5: Bound the expected difference between $\widehat{\theta}^{(B)}$ and $\widetilde{\theta}^{(B)}$, by making a connection between $\widehat{\theta}^{(B)} - \widetilde{\theta}^{(B)}$ and $\widehat{\theta}^{(\infty)} - \widetilde{\theta}^{(\infty)}$

We decompose the bias of the standard MLE $\widehat{\theta}^{(B)}$ as

$$\mathbb{E}[\widehat{\theta}_1^{(B)}] - \theta_1^* = (\mathbb{E}[\widetilde{\theta}_1^{(B)}] - \theta_1^*) + \mathbb{E}[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)}]. \tag{76}$$

Recall from (64) that

$$\mathbb{E}[\widetilde{\theta}_1^{(B)}] - \theta_1^* \leq -\frac{c}{\sqrt{dk}}. \tag{77}$$

In what follows, we prove that

$$\mathbb{E}[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)}] \leq c' \frac{\log d + \log k}{dk}. \tag{78}$$

Then plugging (77) and (78) back into (76) yields

$$\mathbb{E}[\widehat{\theta}_1^{(B)}] - \theta_1^* \leq -\frac{c}{\sqrt{dk}} + c' \frac{\log d + \log k}{dk} \leq -\frac{c''}{\sqrt{dk}},$$

[4] For example, consider the following two univariate estimators. The first estimator always outputs a value within $[-B, B]$. The second estimator sometimes outputs a value within $[-B, B]$, and sometimes outputs a value greater than B. The two estimators could be constructed such that they are close (or equal) in expectation. However, now consider their constrained counterparts. The first estimator is not affected by a box constraint at B, whereas the expected value of second estimator can become significantly smaller due to the box constraint. Therefore, the constrained counterparts of these two estimators may not be close in expectation.

for all $d \geq d_0$ and $k \geq k_0$ where d_0 and k_0 are constants, completing the proof of Theorem 2.1(a).

The rest of this step is devoted to proving (78). To bound $\mathbb{E}[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)}]$, we make a connection between $\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)}$ and $\widetilde{\theta}_1^{(\infty)} - \widehat{\theta}_1^{(\infty)}$, and then we evoke the bound on $\widetilde{\theta}_1^{(\infty)} - \widehat{\theta}_1^{(\infty)}$ from (74) in Step 4.

Recall that μ_1 is a discrete random variable representing the fraction of wins by item 1. By the law of iterated expectation, we have

$$\begin{aligned}\mathbb{E}[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)}] =& \underbrace{\mathbb{E}\left[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)} \mid \frac{1}{2} < \mu_1 < \mu_1^*\right] \cdot \mathbb{P}\left(\frac{1}{2} < \mu_1 < \mu_1^*\right)}_{R_1} \\ &+ \underbrace{\mathbb{E}[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)} \mid \mu_1 \geq \mu_1^*] \cdot \mathbb{P}(\mu_1 \geq \mu_1^*)}_{R_2} + \underbrace{\mathbb{E}\left[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)} \mid \mu_1 < \frac{1}{2}\right] \cdot \mathbb{P}\left(\mu_1 < \frac{1}{2}\right)}_{R_3}. \end{aligned} \tag{79}$$

In what follows, we bound the terms R_1, R_2 and R_3 separately.

Consider the term R_2. From the expression of $\widetilde{\theta}^{(B)}$ in (63), we have $\widetilde{\theta}_1^{(B)} = B$ when $\mu_1 \geq \mu_1^*$. Therefore,

$$\mathbb{E}[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)} \mid \mu_1 \geq \mu_1^*] = \mathbb{E}[\widehat{\theta}_1^{(B)} \mid \mu_1 \geq \mu_1^*] - B \overset{\text{(i)}}{\leq} 0,$$

where (i) is true due to the box constraint $|\widehat{\theta}_1^{(B)}| \leq B$. Hence,

$$R_2 \leq 0. \tag{80}$$

Consider the term R_3, we have $\mathbb{E}[\mu_1] = \mu_1^* = \frac{1}{1+e^{-\frac{d}{d-1}B}}$, and therefore it can be verified that there exists a constant $\tau > 0$, such that $\mu_1^* > \frac{1}{2} + \tau$ for all $d \geq 2$. By Hoeffding's inequality, we have

$$\begin{aligned}\mathbb{P}\left(\mu_1 < \frac{1}{2}\right) &< \mathbb{P}\left(|\mu_1 - \mu_1^*| > \tau\right) \\ &\leq 2\exp\left(-2(d-1)k\tau^2\right) \lesssim \frac{1}{dk}.\end{aligned} \tag{81}$$

Therefore, we have

$$\begin{aligned}R_3 &= \mathbb{E}\left[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)} \mid \mu_1 < \frac{1}{2}\right] \cdot \mathbb{P}\left(\mu_1 < \frac{1}{2}\right) \\ &\overset{\text{(i)}}{\leq} 2B \cdot \mathbb{P}\left(\mu_1 < \frac{1}{2}\right) \\ &\overset{\text{(ii)}}{\lesssim} \frac{1}{dk},\end{aligned} \tag{82}$$

where (i) is true because $|\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)}| \leq |\widehat{\theta}_1^{(B)}| + |\widetilde{\theta}_1^{(B)}| \leq 2B$ by the box constraint, and (ii) is true due to (81).

Now consider the term R_1. Denote $\overline{E_0}$ as the complement of the event E_0. Using the law of iterated expectation again, we have

$$\begin{aligned}R_1 = \mathbb{E}\left[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)} \mid \frac{1}{2} < \mu_1 < \mu_1^*\right] \cdot \mathbb{P}\left(\frac{1}{2} < \mu_1 < \mu_1^*\right) = \\ \underbrace{\mathbb{E}\left[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)} \mid E_0, \frac{1}{2} < \mu_1 < \mu_1^*\right] \cdot \mathbb{P}\left(E_0, \frac{1}{2} < \mu_1 < \mu_1^*\right)}_{R_{11}} \\ + \underbrace{\mathbb{E}\left[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)} \mid \overline{E_0}, \frac{1}{2} < \mu_1 < \mu_1^*\right] \cdot \mathbb{P}\left(\overline{E_0}, \frac{1}{2} < \mu_1 < \mu_1^*\right)}_{R_{12}}\end{aligned} \tag{83}$$

Consider the term R_{12}. We have

$$\begin{aligned}\mathbb{P}\left(\overline{E_0}, \frac{1}{2} < \mu_1 < \mu_1^*\right) &= \sum_{v \in (\frac{1}{2}, \mu_1^*)} \mathbb{P}(\overline{E_0} \mid E_v) \cdot \mathbb{P}(E_v) \\ &\overset{\text{(i)}}{\leq} \frac{c}{dk} \sum_{v \in (\frac{1}{2}, \mu_1^*)} \mathbb{P}(E_v) \\ &\lesssim \frac{1}{dk}, \end{aligned} \tag{84}$$

where (i) is true because E_0 happens w.h.p.$(\frac{1}{dk} \mid E_v)$. Combining (84) with the fact that $|\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)}| \leq 2B$ due to the box constraint, we have

$$R_{12} \lesssim \frac{1}{dk}. \tag{85}$$

Now consider the term R_{11}. We first analyze the constrained oracle $\widetilde{\theta}^{(B)}$. By the expression of $\widetilde{\theta}^{(B)}$ in (63) and the expression of $\widetilde{\theta}^{(\infty)}$ in (62), we have

$$\widetilde{\theta}^{(B)} = \widetilde{\theta}^{(\infty)}, \qquad \text{conditioned on } \frac{1}{2} < \mu_1 < \mu_1^*. \tag{86}$$

Moreover, given $\frac{1}{2} < \mu_1 < \mu_1^*$, by the expression of $\widetilde{\theta}^{(B)}$ in (63), we have

$$0 < \widetilde{\theta}_1^{(B)} < B$$

and therefore by the parameterization of $\widetilde{\theta}^{(B)}$ in (60b),

$$|\widetilde{\theta}_i^{(B)}| \leq \frac{1}{d-1} B \qquad \text{for every } i \in \{2, \ldots, d\}.$$

Hence, there exists a constant $\tau' > 0$ such that

$$\widetilde{\theta}_1^{(B)} > -B + \tau' \tag{88a}$$

and

$$-B + \tau' < \widetilde{\theta}_i^{(B)} < B - \tau' \qquad \text{for every } i \in \{2, \ldots, d\}. \tag{88b}$$

Now we analyze the standard MLE $\widehat{\theta}^{(B)}$. Recall that E_v denotes the event that $\mu_1 = v$. We have that for every $v \in \left(\frac{1}{2}, \mu_1^*\right)$,

$$\|\widehat{\theta}_1^{(\infty)} - \widetilde{\theta}_1^{(B)}\|_\infty \overset{\text{(i)}}{=} \|\widehat{\theta}_1^{(\infty)} - \widetilde{\theta}^{(\infty)}\|_\infty \overset{\text{(ii)}}{\lesssim} \sqrt{\frac{\log d + \log k}{dk}}, \qquad \text{conditioned on } (E_0, E_v), \tag{89}$$

where (i) is true by (86), and (ii) is true by (70) from Step 4. By (89), we have that for every $v \in \left(\frac{1}{2}, \mu_1^*\right)$,

$$\|\widehat{\theta}_1^{(\infty)} - \widetilde{\theta}_1^{(B)}\|_\infty \leq \tau', \qquad \text{conditioned on } (E_0, E_v), \tag{90}$$

for all $d \geq d_0$ and all $k \geq k_0$, where d_0 and k_0 are constants. Combining (90) with (88), if the unconstrained MLE $\widehat{\theta}^{(\infty)}$ violates the box constraint, then only possible case is $\widehat{\theta}_1^{(\infty)} > B$. Then either $\widehat{\theta}_1^{(\infty)} = \widehat{\theta}_1^{(B)}$ (when $\widehat{\theta}^{(\infty)}$ does not violate the box constraint) or $\widehat{\theta}_1^{(\infty)} > B \geq \widehat{\theta}_1^{(B)}$ (when $\widehat{\theta}^{(\infty)}$ violates the box constraint). Hence, for every $v \in (\frac{1}{2}, \mu_1^*)$,

$$\widehat{\theta}_1^{(\infty)} \geq \widehat{\theta}_1^{(B)}, \qquad \text{conditioned on } (E_0, E_v). \tag{91}$$

Combining (86) and (91), we have that for every $v \in (\frac{1}{2}, \mu_1^*)$,

$$\widehat{\theta}^{(B)} - \widetilde{\theta}^{(B)} \leq \widehat{\theta}^{(\infty)} - \widetilde{\theta}^{(\infty)}, \qquad \text{conditioned on } (E_0, E_v). \tag{92}$$

By the law of iterated expectation again, we have

$$\begin{aligned}
R_{11} &= \sum_{v \in (\frac{1}{2}, \mu_1^*)} \mathbb{E}[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)} \mid E_0, \mu_1 = v] \cdot \mathbb{P}(E_0, \mu_1 = v) \\
&= \sum_{v \in (\frac{1}{2}, \mu_1^*)} \mathbb{E}[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)} \mid E_0, E_v] \cdot \mathbb{P}(E_0, E_v) \\
&\overset{\text{(i)}}{\leq} \sum_{v \in (\frac{1}{2}, \mu_1^*)} \mathbb{E}[\widehat{\theta}_1^{(\infty)} - \widetilde{\theta}_1^{(\infty)} \mid E_0, E_v] \cdot \mathbb{P}(E_0, E_v) \\
&\overset{\text{(ii)}}{\lesssim} \frac{\log d + \log k}{dk} \sum_{v \in (\frac{1}{2}, \mu_1^*)} \mathbb{P}(E_0, E_v) \\
&\lesssim \frac{\log d + \log k}{dk},
\end{aligned} \tag{93}$$

where (i) is true due to (92), and (ii) is true due to the bound (74) from Step 4.

Plugging the term R_{11} from (93) and R_{12} from (85) back to (83), we have

$$R_1 = R_{11} + R_{12} \lesssim \frac{\log d + \log k}{dk}. \tag{94}$$

Finally, plugging the terms R_1 from (94), R_2 from (80), and R_3 from (82) back into (79) yields

$$\mathbb{E}[\widehat{\theta}_1^{(B)} - \widetilde{\theta}_1^{(B)}] \lesssim \frac{\log d + \log k}{dk},$$

completing the proof of (78).

A.4 Proofs of Lemmas

In this appendix, we present the proofs of all the lemmas.

A.4.1 Proof of Lemma A.2

We fix any constant $A > 0$.

The stretched-MLE (16) is an optimization over the compact set Θ_A, and the negative log-likelihood function ℓ is continuous. By the Extreme Value Theorem (Rudin, 1976, Theorem 4.16), a solution $\widehat{\theta}^{(A)}$ is guaranteed to exist.

It remains to prove the uniqueness of $\widehat{\theta}^{(A)}$. Assume for contradiction that there exist two solutions $\widehat{\theta}, \widehat{\theta}' \in \Theta_A$ to the stretched-MLE (16) and $\widehat{\theta} \neq \widehat{\theta}'$. By Lemma A.1, the negative log-likelihood function ℓ is strictly convex. Therefore,

$$\frac{1}{2}\left(\ell(\widehat{\theta}) + \ell(\widehat{\theta}')\right) > \ell\left(\frac{\widehat{\theta} + \widehat{\theta}'}{2}\right). \tag{95}$$

It can be verified that $\frac{\widehat{\theta}+\widehat{\theta}'}{2} \in \Theta_A$. Moreover, (95) along with the fact that $\ell(\widehat{\theta}) = \ell(\widehat{\theta}')$ implies that $\frac{\widehat{\theta}+\widehat{\theta}'}{2}$ attains a strictly smaller function value than both $\widehat{\theta}$ and $\widehat{\theta}'$. This contradicts the assumption that $\widehat{\theta}$ and $\widehat{\theta}'$ are both optimal solutions to the stretched-MLE (16).

A.4.2 Proof of Lemma A.3

We first define a "comparison graph" $G(\{W_{ij}\})$ as a function of the pairwise-comparison outcomes $\{W_{ij}\}$. Let each item $i \in [d]$ be a node of the graph. Let there be a directed edge $(i \to j) \in G$, if and only if there exists a comparison where item i beats item j. A directed graph is called strongly-connected if and only if there exists a path from every node i to every other node j.

The following lemma from Ford (1957) relates the existence and uniqueness of a finite unconstrained MLE $\widehat{\theta}^{(\infty)}$ to the strong connectivity of the comparison graph G. This lemma is based on a different parameterization of the BTL model. In this parameterization, each item has a weight $w_i^* > 0$, and the probability that item i beats item j equals $\frac{w_i^*}{w_i^*+w_j^*}$.

Lemma A.14 (Section 2 from Ford (1957))**.** *If the comparison graph $G(\{W_{ij}\})$ is strongly-connected, then there exists a unique solution to the following MLE:*

$$\widehat{w}_{\text{MLE}} = \underset{\substack{w\in\mathbb{R}^d \\ w_i>0,\ \sum_{i=1}^d w_i=1}}{\operatorname{argmin}} \ell_w(\{W_{ij}\}; w),$$

where the negative log-likelihood function ℓ_w is defined as

$$\ell_w(w) = -\sum_{1\le i<j\le d} \left(W_{ij} \log\left(\frac{w_i}{w_i+w_j}\right) + W_{ji}\log\left(\frac{w_j}{w_i+w_j}\right)\right).$$

It can be seen that θ and w are simply different parameterizations of the same problem. There is a one-to-one mapping between θ and w, by taking $\theta_i = \log(w_i)$ and re-centering accordingly (or in the inverse direction, by taking $w_i = e^{\theta_i}$ and normalizing accordingly). Therefore, the existence and the uniqueness of the MLE $\widehat{w}_{\text{MLE}}$ in Lemma A.3 carries over to our unconstrained MLE $\widehat{\theta}^{(\infty)}$ in (15). That is, if the comparison graph G is strongly-connected, then there exists a unique solution $\widehat{\theta}^{(\infty)}$ to the unconstrained MLE. It remains to show that the comparison graph G is strongly-connected w.h.p.($\frac{1}{dk}$).

We first construct an undirected graph $G'(\{W_{ij}\})$ as follows. Let each item $i \in [d]$ be a node of the graph G'. Let there be an undirected edge $(i,j) \in G'$, if and only if in the directed graph G we have both $(i \to j) \in G$ and $(j \to i) \in G$. Equivalently, there exists an undirected edge $(i,j) \in G'$, if and only if $0 < \mu_{ij} < 1$. It can be verified that the connectivity of the undirected graph G' implies the strong connectivity of the directed graph G. Therefore,

$$\mathbb{P}(G \text{ strongly-connected}) \ge \mathbb{P}(G' \text{ connected}). \tag{96}$$

The probability that $(i,j) \in G'$ is $\mathbb{P}(0 < \mu_{ij} < 1)$. By Hoeffding's inequality, we have that for any $t > 0$,

$$\mathbb{P}(|\mu_{ij} - \mu_{ij}^*| > t) < 2e^{-kt^2}, \qquad \text{for all } 1 \le i < j \le d.$$

We have $0 < \frac{1}{1+e^{2B}} \le \mu_{ij}^* \le \frac{1}{1+e^{-2B}} < 1$, for any $i < j$. Since B is a constant, we have that μ_{ij}^* is bounded away from 0 and 1 by a constant. Set $t = \tau$ where τ is any constant such that $0 < \tau < \frac{1}{1+e^{2B}}$. Then for all $1 \le i < j \le d$, we have

$$\begin{aligned}\mathbb{P}(0 < \mu_{ij} < 1) &> \mathbb{P}(\mu_{ij}^* - \tau < \mu_{ij} < \mu_{ij}^* + \tau)\\ &\ge 1 - \mathbb{P}(|\mu_{ij} - \mu_{ij}^*| > \tau)\\ &> 1 - 2e^{-ck},\end{aligned}$$

for some constant $c > 0$.

Recall that the random variables $\{\mu_{ij}\}$ are independent across all $1 \le i < j \le d$. Hence, the probability of the undirected graph G' being connected is at least the probability of an (undirected) Erdős-Rényi random graph being connected, where each edge independently exists with probability $1 - 2e^{-ck}$.

The following lemma from Gilbert (1959) provides an upper bound on the probability of an (undirected) Erdős-Rényi random graph being disconnected (and hence a lower bound on the probability of the graph being connected).

Lemma A.15 (Theorem 1 from Gilbert (1959))**.** *For an (undirected) Erdős-Rényi graph of d nodes, where each edge independently exists with probability p. Let $q := 1 - p$. Then the probability of the graph being disconnected is at most*

$$\left(1 - \frac{d-1}{2} q^{d-1}\right) dq^{d-1}.$$

To apply Lemma A.15, we set $p = 1 - 2e^{-ck}$ and therefore $q = 2e^{-ck}$. Then we have

$$\begin{aligned}\mathbb{P}[G' \text{ disconnected}] &\leq \left(1 - \frac{d-1}{2} q^{d-1}\right) dq^{d-1} \\ &\leq dq^{d-1} \\ &= de^{-ck(d-1)} \\ &\leq \frac{c'}{dk}, \qquad \text{for some constant } c' > 0. \end{aligned} \tag{97}$$

Combining (96) and (97) completes the proof of the lemma.

A.4.3 Proof of Lemma A.4

We first consider any fixed $m \in [d]$. By the definition of $\{\mu_{ij}\}$ in (10), we have

$$\sum_{i \neq m} \mu_{mi} = \frac{1}{k} \sum_{i \neq m} \sum_{r=1}^{k} X_{mi}^{(r)}. \tag{98}$$

There are $(d-1)k$ terms of the form $X_{mi}^{(r)}$ in (98). It can be verified that the terms $X_{mi}^{(r)}$ involved in (98) are independent. Moreover, since $X_{mi}^{(r)} \in \{0, 1\}$, changing the value of a single term $X_{mi}^{(r)}$ changes the value of (98) by $\frac{1}{k}$. By McDiarmid's inequality, we have that for any $t > 0$,

$$\mathbb{P}\left[\left|\sum_{i \neq m} \mu_{mi} - \sum_{i \neq m} \mu_{mi}^*\right| > t\right] \leq 2\exp\left(-\frac{2t^2}{(d-1)k \cdot (\frac{1}{k})^2}\right) = 2\exp\left(-\frac{2kt^2}{(d-1)}\right). \tag{99}$$

Setting $t = c\sqrt{\frac{d(\log d + \log k)}{k}}$ in (99), we have

$$\begin{aligned}\mathbb{P}\left[\left|\sum_{i \neq m} \mu_{mi} - \sum_{i \neq m} \mu_{mi}^*\right| \leq c\sqrt{\frac{d(\log d + \log k)}{k}}\right] &\geq 1 - 2\exp\left(-c' \frac{d}{d-1} (\log d + \log k)\right) \\ &\geq 1 - \frac{c''}{d^2 k}, \end{aligned} \tag{100}$$

for some constants $c', c'' > 0$, provided that the constant $c > 0$ is sufficiently large.

Taking a union bound over $m \in [d]$ on (100) completes the proof.

A.4.4 Proof of Lemma A.5

Denote the random variables $m^+ := \text{argmax}_{i \in [d]} \delta_i$ and $m^- := \text{argmin}_{i \in [d]} \delta_i$. When there are multiple maximizers or minimizers, we arbitrarily choose one.

Setting $m = m^+$ in the first-order optimality condition (35), we have

$$\underbrace{\sum_{i=1}^{d} [f(\theta_{m^+}^* - \theta_i^* + \delta_{m^+} - \delta_i) - f(\theta_{m^+}^* - \theta_i^*)]}_{R^+} = \sum_{i \neq m^+} (\mu_{mi} - \mu_{mi}^*) \overset{\text{(i)}}{\lesssim} \sqrt{\frac{d(\log d + \log k)}{k}}, \tag{101}$$

where (i) is true by Lemma A.4 (recall that the lemma statement is conditioned on the event E_0 that both Lemma A.3 and Lemma A.4 hold).

Denote the function $g(x,t) := f(x+t) - f(x) = \frac{1}{1+e^{-(x+t)}} - \frac{1}{1+e^{-x}}$. The following lemma states three properties for the function g, which are used in later parts of the proof.

Lemma A.16. *We have the following properties for the function g.*

$$g(x,t) = -g(-x,-t), \qquad \text{for all } x,t \in \mathbb{R} \tag{102a}$$

$$g(x,t) \geq g(\tau,t) > 0, \qquad \text{for all } \tau > 0,\ t > 0,\ \text{and all } x \text{ such that } -\tau \leq x \leq \tau \tag{102b}$$

$$g(\tau,t_1) + g(\tau,t_2) \geq g(\tau,t_1+t_2), \qquad \text{for all } \tau > 0, \text{and all } t_1,t_2 \geq 0. \tag{102c}$$

Lemma A.16 can be verified by straightforward algebra. For completeness, we include the proof of Lemma A.16 at the end of this appendix.

By the definition of m^+, we have $\delta_{m^+} = \max_{i\in[d]} \delta_i$, and therefore $\delta_{m^+} - \delta_i \geq 0$ for all $i \in [d]$. Hence, we have

$$\begin{aligned}
R^+ &= \sum_{i=1}^d f(\theta^*_{m^+} - \theta^*_i + \delta_{m^+} - \delta_i) - f(\theta^*_{m^+} - \theta^*_i) \\
&= \sum_{i=1}^d g(\theta^*_{m^+} - \theta^*_i, \delta_{m^+} - \delta_i) \\
&\overset{\text{(i)}}{\geq} \sum_{i=1}^d g(2B, \delta_{m^+} - \delta_i),
\end{aligned} \tag{103}$$

where (i) is true by (102b) combined with the fact that $|\theta^*_i - \theta^*_j| \leq |\theta^*_i| + |\theta^*_j| \leq 2B$ for all $i,j \in [d]$.

Similarly, setting $m = m^-$ in the first-order optimality condition (35), we have

$$\underbrace{\sum_{i=1}^d \left[f(\theta^*_{m^-} - \theta^*_i + \delta_{m^-} - \delta_i) - f(\theta^*_{m^-} - \theta^*_i)\right]}_{R^-} \lesssim \sqrt{\frac{d(\log d + \log k)}{k}}. \tag{104}$$

By the definition of m^-, we have $\delta_{m^-} = \min_{i\in[d]} \delta_i$, and therefore $\delta_i - \delta_{m^-} \geq 0$ for all $i \in [d]$. Hence, we have

$$\begin{aligned}
R^- &= \sum_{i=1}^d f(\theta^*_{m^-} - \theta^*_i + \delta_{m^-} - \delta_i) - f(\theta^*_{m^-} - \theta^*_i) \\
&= \sum_{i=1}^d g(\theta^*_{m^-} - \theta^*_i, \delta_{m^-} - \delta_i) \\
&\overset{\text{(i)}}{=} \sum_{i=1}^d -g(\theta^*_i - \theta^*_{m^-}, \delta_i - \delta_{m^-}) \\
&\overset{\text{(ii)}}{\leq} \sum_{i=1}^d -g(2B, \delta_i - \delta_{m^-}),
\end{aligned} \tag{105}$$

where (i) is true by (102a), and (ii) is true by (102b) combined with the fact that $|\theta^*_i - \theta^*_j| \leq 2B$ for all $i,j \in [d]$.

Combining (103) and (105), we have

$$\begin{aligned}
R^+ - R^- &\geq \sum_{i=1}^d g(2B, \delta_{m^+} - \delta_i) + \sum_{i=1}^d g(2B, \delta_i - \delta_{m^-}) \\
&\overset{\text{(i)}}{\geq} \sum_{i=1}^d g(2B, \delta_{m^+} - \delta_{m^-}) \\
&= d \cdot g(2B, \delta_{m^+} - \delta_{m^-}) \overset{\text{(ii)}}{\geq} 0,
\end{aligned} \tag{106}$$

where (i) is true due to (102c) since $\delta_{m^+} - \delta_i \geq 0$ and $\delta_i - \delta_{m^-} \geq 0$ for all $i \in [d]$, and (ii) is true since $\delta_{m^+} - \delta_{m^-} \geq 0$. On the other hand, combining (101) and (104), we have

$$R^+ - R^- \lesssim \sqrt{\frac{d(\log d + \log k)}{k}}. \tag{107}$$

Combining (106) and (107), we have

$$\begin{aligned} 0 \leq d \cdot g(2B, \delta_{m^+} - \delta_{m^-}) \leq R^+ - R^+ &\lesssim \sqrt{\frac{d(\log d + \log k)}{k}} \\ g(2B, \delta_{m^+} - \delta_{m^-}) &\lesssim \sqrt{\frac{\log d + \log k}{dk}} \\ f(2B + \delta_{m^+} - \delta_{m^-}) - f(2B) &\lesssim \sqrt{\frac{\log d + \log k}{dk}}. \end{aligned} \tag{108}$$

By the first-order mean value theorem on the LHS of (108), we have

$$f(2B + \delta_{m^+} - \delta_{m^-}) - f(2B) = f'(\lambda) \cdot (\delta_{m^+} - \delta_{m^-}) \leq c\sqrt{\frac{\log d + \log k}{dk}}, \tag{109}$$

where λ is a random variable that takes values in the interval $[2B, 2B + \delta_{m^+} - \delta_{m^-}]$.

Let ϵ be any constant such that $0 < \epsilon < 1 - f(2B)$. Then there exists a constant $\tau > 0$ such that $f(2B + \tau) - f(2B) = \epsilon$. On the other hand, there exist constants $d_0 > 0$ and $k_0 > 0$ such that

$$c\sqrt{\frac{\log d + \log k}{dk}} < \epsilon, \qquad \text{for any } d \geq d_0 \text{ and } k \geq k_0. \tag{110}$$

Combining (109) and (110), we have

$$\begin{aligned} f(2B + \delta_{m^+} - \delta_{m^-}) - f(2B) \leq c\sqrt{\frac{\log d + \log k}{dk}} &< \epsilon = f(2B + \tau) - f(2B) \\ f(2B + \delta_{m^+} - \delta_{m^-}) &\leq f(2B + \tau). \end{aligned} \tag{111}$$

By (13a), we have $f' > 0$ on $(-\infty, \infty)$, and hence the function f is monotonically increasing. Hence, from (111), we have $\delta_{m^+} - \delta_{m^-} \leq \tau$, and therefore the interval $[2B, 2B + \delta_{m^+} - \delta_{m^-}]$ is bounded. By the property (13a) of the sigmoid function f, we have $f' > c_3 > 0$ for some constant $c_3 > 0$ in the bounded interval $[2B, 2B + \delta_{m^+} - \delta_{m^-}]$. Recall that λ takes values in the interval $[2B, 2B + \delta_{m^+} - \delta_{m^-}]$. Therefore, we have

$$c_3(\delta_{m^+} - \delta_{m^-}) < f'(\lambda) \cdot (\delta_{m^+} - \delta_{m^-}). \tag{112}$$

Combining (109) and (112), we have

$$\begin{aligned} c_3(\delta_{m^+} - \delta_{m^-}) < f'(\lambda) \cdot (\delta_{m^+} - \delta_{m^-}) &\leq c\sqrt{\frac{\log d + \log k}{dk}} \\ \delta_{m^+} - \delta_{m^-} &\lesssim \sqrt{\frac{\log d + \log k}{dk}}. \end{aligned} \tag{113}$$

By the assumption that $\theta^* \in \Theta_B$, we have $\sum_{i=1}^d \theta^*_i = 0$. Similarly, by the centering constraint on the unconstrained MLE $\widehat{\theta}^{(\infty)}$ in (15), we have $\sum_{i=1}^d \widehat{\theta}_i^{(\infty)} = 0$. Hence, we have the deterministic relation

$$\sum_{i=1}^d \widehat{\theta}_i^{(\infty)} - \sum_{i=1}^d \theta^*_i = \sum_{i=1}^d \delta_i = 0. \tag{114}$$

Hence, $\delta_{m^+} \geq 0$ and $\delta_{m^-} \leq 0$. By (113), we have

$$\delta_{m^+} - \delta_{m^-} = |\delta_{m^+}| + |\delta_{m^-}| \lesssim \sqrt{\frac{\log d + \log k}{dk}}.$$

Hence, $|\delta_{m^+}| \lesssim \sqrt{\frac{\log d + \log k}{dk}}$ and $|\delta_{m^-}| \lesssim \sqrt{\frac{\log d + \log k}{dk}}$. Therefore,

$$|\delta_m| \lesssim \sqrt{\frac{\log d + \log k}{dk}}, \qquad \text{for all } m \in [d],$$

completing the proof of the lemma.

Proof of Lemma A.16: We prove the three parts of the lemma separately.

(a) It can be verified that $f(x) = 1 - f(-x)$. Hence,

$$\begin{aligned} g(x,t) = f(x+t) - f(x) &= [1 - f(-x-t)] - [1 - f(-x)] \\ &= -[f(-x-t) - f(-x)] = -g(-x,-t). \end{aligned}$$

(b) We prove the two parts of the inequality separately.

We first prove that $g(\tau, t) > 0$. By (13a), the function f is strictly increasing. Therefore, for any $t > 0$, we have

$$g(\tau, t) = f(\tau + t) - f(\tau) > 0.$$

Now we prove that $g(x,t) \geq g(\tau,t)$. We have

$$\begin{aligned} g(x,t) - g(\tau,t) &= f(x+t) - f(x) - [f(\tau+t) - f(\tau)] \\ &= \int_x^{x+t} f'(u)\,\mathrm{d}u - \int_\tau^{\tau+t} f'(u)\,\mathrm{d}u \\ &= \int_0^t f'(x+u)\,\mathrm{d}u - \int_0^t f'(\tau+u)\,\mathrm{d}u \\ &= \int_0^t [f'(x+u) - f'(\tau+u)]\,\mathrm{d}u. \end{aligned} \tag{115}$$

By (115), it remains to prove that

$$f'(x+u) \geq f'(\tau+u), \qquad \text{for any } u \in [0,t]. \tag{116}$$

Fix any $u \in [0,t]$. By assumption we have $\tau > 0$. Hence, $\tau + u > 0$. Now we consider the sign of $(x+u)$.

If $x + u \geq 0$, then by the assumption that $x \leq \tau$, we have $0 \leq x + u \leq \tau + u$. It can be verified that f' is decreasing on $[0,\infty)$. Therefore,

$$f'(x+u) \geq f'(\tau+u). \tag{117}$$

If $x + u < 0$, we have

$$0 < -x - u \overset{\text{(i)}}{\leq} \tau - u \overset{\text{(ii)}}{\leq} \tau + u, \tag{118}$$

where (i) is true by the assumption that $x \geq -\tau$, and (ii) is true because $u \in [0,t]$ and therefore $u \geq 0$. We have

$$f'(x+u) \overset{\text{(i)}}{=} f'(-x-u) \overset{\text{(ii)}}{\geq} f'(\tau+u), \tag{119}$$

where (i) holds because it can be verified that $f'(x) = f'(-x)$ for any $x \in \mathbb{R}$, and (ii) is true by combining (118) with the fact that f' is decreasing on $[0,\infty)$.

Combining the two cases of (117) and (119) completes the proof of (116).

(c) We have

$$\begin{aligned} g(\tau, t_1) + g(\tau, t_2) &= f(\tau + t_1) - f(\tau) + f(\tau + t_2) - f(\tau) \\ &= \int_{\tau}^{\tau+t_1} f'(u)\,\mathrm{d}u + \int_{\tau}^{\tau+t_2} f'(u)\,\mathrm{d}u \\ &\overset{\text{(i)}}{\geq} \int_{\tau}^{\tau+t_1} f'(u)\,\mathrm{d}u + \int_{\tau+t_1}^{\tau+t_1+t_2} f'(u)\,\mathrm{d}u \\ &= \int_{\tau}^{\tau+t_1+t_2} f'(u)\,\mathrm{d}u \\ &= f(\tau + t_1 + t_2) - f(\tau) = g(\tau, t_1 + t_2), \end{aligned}$$

where (i) is true because f' is decreasing on $(0, \infty)$, and because $\tau > 0$ and $t_1, t_2 \geq 0$ by assumption.

A.4.5 Proof of Lemma A.6

We fix any $i, j \in [d]$ where $i \neq j$. By the law of iterated expectation, we have

$$\mathbb{E}[\mu_{ij} \mid E] = \mathbb{E}[\mu_{ij} \mid E', E] \cdot \mathbb{P}(E' \mid E) + \mathbb{E}[\mu_{ij} \mid \overline{E}', E] \cdot \mathbb{P}(\overline{E}' \mid E). \tag{120}$$

Subtracting $\mathbb{E}[\mu_{ij} \mid E', E]$ from both sides of (120), we have

$$\begin{aligned} \mathbb{E}[\mu_{ij} \mid E] - \mathbb{E}[\mu_{ij} \mid E', E] &= \mathbb{E}[\mu_{ij} \mid E', E] \cdot [\mathbb{P}(E' \mid E) - 1] + \mathbb{E}[\mu_{ij} \mid \overline{E}', E] \cdot \mathbb{P}(\overline{E}' \mid E) \\ &= (-\mathbb{E}[\mu_{ij} \mid E', E] + \mathbb{E}[\mu_{ij} \mid \overline{E}', E]) \cdot \mathbb{P}(\overline{E}' \mid E). \end{aligned} \tag{121}$$

Taking an absolute value on (121), we have

$$\begin{aligned} |\mathbb{E}[\mu_{ij} \mid E] - \mathbb{E}[\mu_{ij} \mid E', E]| &= \left| -\mathbb{E}[\mu_{ij} \mid E', E] + \mathbb{E}[\mu_{ij} \mid \overline{E}', E] \right| \cdot \mathbb{P}(\overline{E}' \mid E) \\ &\overset{\text{(i)}}{\lesssim} \frac{1}{dk}, \end{aligned}$$

where (i) is true due to the deterministic inequality $0 \leq \mu_{ij} \leq 1$ and the fact that event E' happens w.h.p.$(\frac{1}{dk} \mid E)$.

A.4.6 Proof of Lemma A.7

Denote $m^+ := \text{argmax}_{i \in [d]} \Delta_i$ and $m^- := \text{argmin}_{i \in [d]} \Delta_i$. When there are multiple maximizers or minimizers, we arbitrarily choose one. The proof works similarly in spirit to the proof of Lemma A.5. We first show that $\Delta_{m^+} - \Delta_{m^-}$ satisfies the desired upper bound. Then we show that Δ_{m^+} and Δ_{m^-} have different signs, and therefore the desired upper bound holds on $|\Delta_m|$ uniformly across all $m \in [d]$.

Recall from (38) that for every $m \in [d]$,

$$\sum_{i=1}^{d} f'(\theta_m^* - \theta_i^*) \cdot (\Delta_m - \Delta_i) = \underbrace{\sum_{i \neq m} (\mathbb{E}[\mu_{mi} \mid E_0] - \mu_{mi}^*)}_{R_1} - \underbrace{\frac{1}{2} \sum_{i=1}^{d} \mathbb{E}[f''(\lambda_{mi})(\delta_m - \delta_i)^2 \mid E_0]}_{R_2}, \tag{122}$$

where λ_{mi} is a random variable that takes values between $\theta_m^* - \theta_i^*$ and $\theta_m^* - \theta_i^* + (\delta_m - \delta_i)$.

We consider the two terms on the RHS of (38) separately. For the term R_1, recall from (40) that

$$|\mathbb{E}[\mu_{mi} \mid E_0] - \mu_{mi}^*| \lesssim \frac{1}{dk}.$$

Therefore,

$$|R_1| \lesssim (d - 1) \cdot \frac{1}{dk} \lesssim \frac{1}{k}. \tag{123}$$

Now consider the term R_2. Recall that $\theta^* \in \Theta_B$. Therefore, for every $m \in [d]$, we have $|\theta^*_m| \leq B$. Recall from Lemma A.5 that for every $m \in [d]$, we have

$$|\delta_m| \lesssim \sqrt{\frac{\log d + \log k}{dk}}, \qquad \text{conditioned on } E_0. \tag{124}$$

Let $c > 0$ be any constant. By (124), we have $|\delta_m| \leq c$, for all $d \geq d_0$ and $k \geq k_0$, where d_0 and k_0 are constants which may only depend on c. Hence, conditioned on E_0, the interval between $\theta^*_m - \theta^*_i$ and $\theta^*_m - \theta^*_i + (\delta_m - \delta_i)$ is contained in the interval $[-2B - 2c, 2B + 2c]$. By the property (13b) of the sigmoid function f, we have

$$|f''| < c_5, \qquad \text{on the bounded interval } [-2B - 2c, 2B + 2c].$$

Therefore,

$$\left|\mathbb{E}\left[f''(\lambda_{mi}) \cdot (\delta_m - \delta_i)^2 \mid E_0\right]\right| \leq c_5 \cdot \mathbb{E}[(\delta_m - \delta_i)^2 \mid E_0] \overset{\text{(i)}}{\lesssim} \frac{\log d + \log k}{dk}, \qquad \text{for all } i, m \in [d],$$

where (i) is again by (124). Therefore,

$$|R_2| \lesssim d \cdot \frac{\log d + \log k}{dk} = \frac{\log d + \log k}{k}. \tag{125}$$

Taking an absolute value on (122) and using the triangle inequality, we have

$$\left|\sum_{i=1}^{d} f'(\theta^*_m - \theta^*_i) \cdot (\Delta_m - \Delta_i)\right| \leq |R_1| + |R_2| \overset{\text{(i)}}{\lesssim} \frac{\log d + \log k}{k}, \tag{126}$$

where (i) is true by combining the term R_1 from (123) and the term R_2 from (125). Taking $m = m^+$ in (126), we have

$$\sum_{i=1}^{d} f'(\theta^*_{m^+} - \theta^*_i) \cdot (\Delta_{m^+} - \Delta_i) \leq c\frac{\log d + \log k}{k}. \tag{127}$$

Taking $m = m^-$ in (126), we have

$$\sum_{i=1}^{d} f'(\theta^*_{m^-} - \theta^*_i) \cdot (\Delta_{m^-} - \Delta_i) \geq -c\frac{\log d + \log k}{k}$$

and hence

$$\sum_{i=1}^{d} f'(\theta^*_{m^-} - \theta^*_i) \cdot (\Delta_i - \Delta_{m^-}) \leq c\frac{\log d + \log k}{k}. \tag{128}$$

Adding (127) and (128), we have

$$\underbrace{\sum_{i=1}^{d} f'(\theta^*_{m^+} - \theta^*_i) \cdot (\Delta_{m^+} - \Delta_i) + \sum_{i=1}^{d} f'(\theta^*_{m^-} - \theta^*_i) \cdot (\Delta_i - \Delta_{m^-})}_{R} \leq c\frac{\log d + \log k}{k}. \tag{129}$$

Consider the term R. We have $|\theta^*_m - \theta^*_i| \leq 2B$ for all $i, m \in [d]$. By the property (13a) of the sigmoid function, there exists some constant c_3, such that

$$f'(\theta^*_m - \theta^*_i) > c_3 > 0, \qquad \text{for all } i, m \in [d]. \tag{130}$$

By the definition of m^+ and m^-, we have $\Delta_{m^+} - \Delta_i \geq 0$ and $\Delta_i - \Delta_{m^-} \geq 0$ for every $i \in [d]$. Plugging (130) into (129), combined with the fact that $\Delta_{m^+} - \Delta_i \geq 0$ and $\Delta_i - \Delta_{m^-} \geq 0$, we have

$$\begin{aligned} c_3 \left[\sum_{i=1}^{d} (\Delta_{m^+} - \Delta_i) + \sum_{i=1}^{d} (\Delta_i - \Delta_{m^-}) \right] \leq R \leq c \frac{\log d + \log k}{k} \\ c_3 d \cdot (\Delta_{m^+} - \Delta_{m^-}) \leq c \frac{\log d + \log k}{k} \\ \Delta_{m^+} - \Delta_{m^-} \lesssim \frac{\log d + \log k}{dk}. \end{aligned} \tag{131}$$

By (114) in the proof of Lemma A.5, we have the deterministic relation

$$\sum_{i=1}^{d} \delta_i = 0. \tag{132}$$

Taking an expectation over (132) conditional on E_0, we have

$$\sum_{i=1}^{d} \Delta_i = 0.$$

Hence, $\Delta_{m^+} \geq 0$ and $\Delta_{m^-} \leq 0$. By (131), we have

$$\Delta_{m^+} - \Delta_{m^-} = |\Delta_{m^+}| + |\Delta_{m^-}| \lesssim \frac{\log d + \log k}{dk}.$$

Hence, $|\Delta_{m^+}| \lesssim \frac{\log d + \log k}{dk}$ and $|\Delta_{m^-}| \lesssim \frac{\log d + \log k}{dk}$. Therefore,

$$|\Delta_m| \lesssim \frac{\log d + \log k}{dk}, \qquad \text{for all } m \in [d].$$

A.4.7 Proof of Lemma A.9

To compare the functions h and h^+, we introduce an auxiliary function $h_0 : [0, 1] \to [-B, B]$:

$$h_0(t) = \begin{cases} -B & \text{if } 0 \leq t \leq \mu_- \\ \frac{B}{\mu_+ - \frac{1}{2}}(t - \frac{1}{2}) & \text{if } \mu_- < t < \mu_+ \\ B & \text{if } \mu_+ \leq t \leq 1. \end{cases}$$

In words, the function h_0 is piecewise linear. On the interval $[0, \mu_-]$, its value equals the constant $-B$. On the interval $[\mu_-, \mu_+]$, it is a line passing through the points $(\mu_-, -B)$ and (μ_+, B). On the interval $[\mu_+, 1]$, its value equals the constant B. See Fig. 7 for a comparison of the three functions h, h^+ and h_0.

It can be verified that $h^+(t) \geq h_0(t)$ for any $t \in [0, 1]$. Hence,

$$\mathbb{E}[h^+(\mu)] \geq \mathbb{E}[h_0(\mu)]. \tag{133}$$

Recall that our goal is to prove (52):

$$\mathbb{E}[h(\mu)] \leq \mathbb{E}[h^+(\mu)].$$

Given (133), it suffices to prove that

$$\mathbb{E}[h(\mu)] \leq \mathbb{E}[h_0(\mu)]. \tag{134}$$

The rest of the proof is devoted to proving (134).

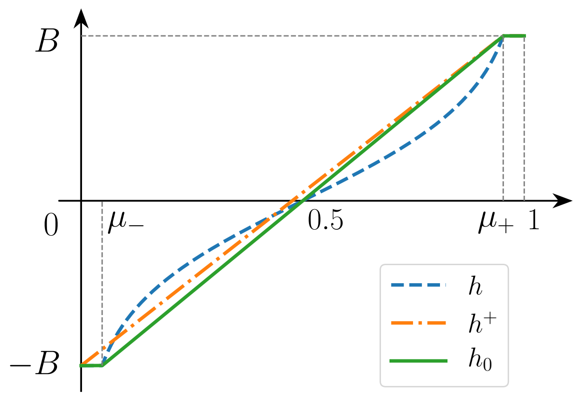

Figure 7: The functions h, h^+ and h_0.

It can be verified that h and h_0 are anti-symmetric around $\frac{1}{2}$. That is, for any $t \in [0,1]$, we have

$$h(t) = -h(1-t) \tag{135a}$$
$$h_0(t) = -h_0(1-t). \tag{135b}$$

In particular, we have

$$h\left(\frac{1}{2}\right) = h_0\left(\frac{1}{2}\right) = 0. \tag{136}$$

It can also be verified that

$$h(t) \geq h_0(t), \qquad \text{for all } t \in \left[0, \frac{1}{2}\right]. \tag{137}$$

Recall the notation of $W = k\mu$ representing the number of times that item 1 beats item 2 among the k comparisons between them. We have $W \sim \text{Binom}(k, \mu_+)$. Therefore,

$$\begin{aligned}
\mathbb{E}[h(\mu)] - \mathbb{E}\left[h_0(\mu)\right] &= \mathbb{E}_W\left[h\left(\frac{W}{k}\right)\right] - \mathbb{E}_W\left[h_0\left(\frac{W}{k}\right)\right] \\
&= \sum_{w=0}^{k}\left[h\left(\frac{w}{k}\right) - h_0\left(\frac{w}{k}\right)\right] \cdot \mathbb{P}(W=w) \\
&\overset{\text{(i)}}{=} \left(\sum_{w=0}^{\lfloor \frac{k}{2} \rfloor} + \sum_{w=\lceil \frac{k}{2} \rceil}^{k}\right)\left[(h-h_0)\left(\frac{w}{k}\right)\right] \cdot \mathbb{P}(W=w) \\
&\overset{\text{(ii)}}{=} \sum_{w=0}^{\lfloor \frac{k}{2} \rfloor}\left[(h-h_0)\left(\frac{w}{k}\right) \cdot \mathbb{P}(W=w) + (h-h_0)\left(1-\frac{w}{k}\right) \cdot \mathbb{P}(W=k-w)\right] \\
&\overset{\text{(iii)}}{=} \sum_{w=0}^{\lfloor \frac{k}{2} \rfloor}(h-h_0)\left(\frac{w}{k}\right) \cdot \left[\mathbb{P}(W=w) - \mathbb{P}(W=k-w)\right],
\end{aligned} \tag{138}$$

where (i) is true by (136). Specifically, when k is even, we double-count the term of $w = \frac{k}{2}$. This term equals $(h-h_0)(\frac{1}{2}) = 0$, so double-counting this term does not affect the equality. Moreover, step (ii) is true by a change of variable $w \leftarrow k-w$ in the second summation, and step (iii) is true by the anti-symmetry (135) of the functions h and h^+.

Now consider the terms in the summation (138). By (137), we have

$$(h-h_0)\left(\frac{w}{k}\right) \geq 0, \qquad \text{for all } 0 \leq w \leq \left\lfloor \frac{k}{2} \right\rfloor. \tag{139}$$

Using the binomial probabilities of $W \sim \text{Binom}(k, \mu_+)$, we also have

$$\begin{aligned}
\mathbb{P}(W = w) - \mathbb{P}(W = k - w) &= \binom{k}{w}[(\mu_+)^w(1-\mu_+)^{k-w} - (\mu_+)^{k-w}(1-\mu_+)^w] \\
&= \binom{k}{w}(\mu_+)^w(1-\mu_+)^w \cdot [(1-\mu_+)^{k-2w} - (\mu_+)^{k-2w}] \\
&\overset{\text{(i)}}{\leq} 0, \qquad \text{for all } 0 \leq w \leq \left\lfloor \frac{k}{2} \right\rfloor, \qquad (140)
\end{aligned}$$

where (i) is true because $\mu_+ = \frac{1}{1+e^{-2B}} > \frac{1}{2}$, combined with the fact that $k - 2w \geq 0$, for all $0 \leq w \leq \lfloor \frac{k}{2} \rfloor$. Plugging (139) and (140) back into (138), we have

$$\mathbb{E}[h(\mu)] - \mathbb{E}[h_0(\mu)] \geq 0,$$

completing the proof of (134).

A.4.8 Proof of Lemma A.10

We have

$$\begin{aligned}
\mathbb{E}[h^+(\mu)] - \theta_1^* &= \mathbb{E}_W\left[h^+\left(\frac{W}{k}\right)\right] - B \\
&= \sum_{w=0}^{k} h^+\left(\frac{w}{k}\right) \cdot \mathbb{P}(W = w) - B \\
&\overset{\text{(i)}}{=} \sum_{w=0}^{\lfloor k\mu_+ \rfloor} \frac{2B}{\mu_+}\left(\frac{w}{k} - \mu_+\right) \cdot \mathbb{P}(W = w) \\
&= c\left(\underbrace{\sum_{w=0}^{\lfloor k\mu_+ \rfloor} \frac{w}{k} \cdot \mathbb{P}(W = w)}_{R_1} - \mu_+ \underbrace{\sum_{w=0}^{\lfloor k\mu_+ \rfloor} \mathbb{P}(W = w)}_{R_2}\right), \qquad (141)
\end{aligned}$$

where (i) is true by plugging in the definition of the function h^+ from (51).

Now we consider the two terms R_1 and R_2 separately. For any integer $n \geq 1$, any integer s such that $0 \leq s \leq n$, and any real number $p \in [0, 1]$, we define $\mathcal{P}_{\text{le}}(n, p, s)$ (resp. $\mathcal{P}_{\text{eq}}(n, p, s)$) as the probability that the value of the random variable $\text{Binom}(n, p)$ is at most (resp. equal to) s. That is,

$$\begin{aligned}
\mathcal{P}_{\text{le}}(n, p, s) &= \mathbb{P}[\text{Binom}(n, p) \leq s], \\
\mathcal{P}_{\text{eq}}(n, p, s) &= \mathbb{P}[\text{Binom}(n, p) = s].
\end{aligned}$$

Then the term R_2 can be written as

$$R_2 = \mathcal{P}_{\text{le}}(k, \mu_+, \lfloor k\mu_+ \rfloor). \qquad (142)$$

For the term R_1, we have

$$\begin{aligned}
R_1 = \sum_{w=0}^{\lfloor k\mu_+\rfloor} \frac{w}{k}\cdot \mathbb{P}(W=w) &= \sum_{w=0}^{\lfloor k\mu_+\rfloor} \frac{w}{k}\cdot \binom{k}{w}\mu_+^w(1-\mu_+)^{(k-w)}\\
&= \sum_{w=1}^{\lfloor k\mu_+\rfloor} \frac{w}{k}\cdot \frac{k!}{w!(k-w)!}\mu_+^w(1-\mu_+)^{(k-w)}\\
&= \mu_+ \sum_{w=1}^{\lfloor k\mu_+\rfloor} \frac{(k-1)!}{(w-1)!(k-w)!}\mu_+^{w-1}(1-\mu_+)^{(k-w)}\\
&\overset{\text{(i)}}{=} \mu_+ \sum_{w=0}^{\lfloor k\mu_+\rfloor-1} \frac{(k-1)!}{(w)!(k-w-1)!}\mu_+^{w}(1-\mu_+)^{(k-1-w)}\\
&= \mu_+ \sum_{w=0}^{\lfloor k\mu_+\rfloor-1} \binom{k-1}{w}\mu_+^{w}(1-\mu_+)^{(k-1-w)}\\
&= \mu_+\cdot \mathcal{P}_{\text{le}}(k-1,\mu_+,\lfloor k\mu_+\rfloor-1),
\end{aligned} \tag{143}$$

where (i) is true by a change of variable $w \leftarrow w-1$. Plugging (142) and (143) back into (141), we have

$$\mathbb{E}[h^+(\mu)] - \theta_1^* = c\mu_+\cdot[\mathcal{P}_{\text{le}}(k-1,\mu_+,\lfloor k\mu_+\rfloor-1) - \mathcal{P}_{\text{le}}(k,\mu_+,\lfloor k\mu_+\rfloor)]. \tag{144}$$

For any integer $n\geq 1$, any integer s such that $0\leq s\leq n$, and any $p\in[0,1]$, we claim the combinatorial equality

$$\mathcal{P}_{\text{le}}(n,p,s) = \mathcal{P}_{\text{le}}(n-1,p,s-1) + (1-p)\cdot\mathcal{P}_{\text{eq}}(n-1,p,s). \tag{145}$$

To prove (145), we use a standard combinatorial argument. Consider n balls, and we select each ball independently with probability p. Then the LHS of (145) equals the probability that at most s balls are selected. This event can be decomposed into two cases. Either there are at most $(s-1)$ balls selected from the first $(n-1)$ balls; or there are exactly s balls selected from the first $(n-1)$ balls, and the last ball is not selected. These two cases correspond to the two terms on the RHS of (145).

Now setting $n=k, p=\mu_+$, and $s=\lfloor k\mu_+\rfloor$ in (145), we have

$$\mathcal{P}_{\text{le}}(k,\mu_+,\lfloor k\mu_+\rfloor) = \mathcal{P}_{\text{le}}(k-1,\mu_+,\lfloor k\mu_+\rfloor-1) + (1-\mu_+)\cdot\mathcal{P}_{\text{eq}}(k-1,\mu_+,\lfloor k\mu_+\rfloor]). \tag{146}$$

Combining (144) and (146), we have

$$\mathbb{E}[h^+(\mu)] - \theta_1^* = -c(1-\mu_+)\cdot\mathcal{P}_{\text{eq}}(k-1,\mu_+,\lfloor k\mu_+\rfloor). \tag{147}$$

It remains to bound the term $\mathcal{P}_{\text{eq}}(k-1,\mu_+,\lfloor k\mu_+\rfloor)$ on the RHS of (147). Writing out the binomial probability, we have

$$\mathcal{P}_{\text{eq}}(k-1,\mu_+,\lfloor k\mu_+\rfloor) = \binom{k-1}{\lfloor k\mu_+\rfloor}\mu_+^{\lfloor k\mu_+\rfloor}(1-\mu_+)^{k-1-\lfloor k\mu_+\rfloor}. \tag{148}$$

By the Stirling's approximation, we have

$$\sqrt{2\pi}\cdot k^{k+\frac{1}{2}}e^{-k} \leq k! \leq e\cdot k^{k+\frac{1}{2}}e^{-k}, \qquad \text{for any integer } k\geq 0.$$

Then for any integer $n\geq 1$, and any integer k such that $0\leq k\leq n$, we have

$$\binom{n}{k} = \frac{n!}{k!(n-k)!} \geq c\frac{n^{n+\frac{1}{2}}}{k^{k+\frac{1}{2}}(n-k)^{n-k+\frac{1}{2}}}. \tag{149}$$

Plugging (149) into (148), we have

$$\begin{aligned}\mathcal{P}_{\text{eq}}(k-1,\mu_+,\lfloor k\mu_+\rfloor) &\geq c\frac{(k-1)^{k-\frac{1}{2}}}{(k-1-\lfloor k\mu_+\rfloor)^{k-\frac{1}{2}-\lfloor k\mu_+\rfloor}\cdot(\lfloor k\mu_+\rfloor)^{\lfloor k\mu_+\rfloor+\frac{1}{2}}}\cdot\mu_+^{\lfloor k\mu_+\rfloor}(1-\mu_+)^{k-1-\lfloor k\mu_+\rfloor}\\ &\geq c\frac{(k-1)^{k-\frac{1}{2}}}{(k-k\mu_+)^{k-\frac{1}{2}-\lfloor k\mu_+\rfloor}\cdot(k\mu_+)^{\lfloor k\mu_+\rfloor+\frac{1}{2}}}\cdot\mu_+^{\lfloor k\mu_+\rfloor}(1-\mu_+)^{k-1-\lfloor k\mu_+\rfloor}\\ &\geq c\frac{(k-1)^{k-\frac{1}{2}}}{k^k\cdot(1-\mu_+)^{k-\frac{1}{2}-\lfloor k\mu_+\rfloor}\cdot(\mu_+)^{\lfloor k\mu_+\rfloor+\frac{1}{2}}}\cdot\mu_+^{\lfloor k\mu_+\rfloor}(1-\mu_+)^{k-1-\lfloor k\mu_+\rfloor}\\ &= c\frac{(k-1)^{k-\frac{1}{2}}}{k^k}\cdot\mu_+^{-\frac{1}{2}}(1-\mu_+)^{-\frac{1}{2}}\\ &\stackrel{\text{(i)}}{=} c\frac{(k-1)^{k-\frac{1}{2}}}{k^k}\geq c\frac{1}{\sqrt{k-1}}(1-\frac{1}{k})^k\gtrsim\frac{1}{\sqrt{k}},\end{aligned}\tag{150}$$

where (i) is true because $\mu_+ = \frac{1}{1+e^{-2B}}$ is bounded away from 0 and 1 by a constant.

Combining (147) and (150), we have

$$\mathbb{E}[h^+(\mu)] - \theta_1^* \leq -\frac{c}{\sqrt{k}}, \qquad \text{for some constant } c > 0.$$

A.4.9 Proof of Lemma A.11

First consider the unconstrained oracle $\widetilde{\theta}^{(\infty)}$. We prove that for any $\theta \notin \Theta_{\text{oracle}}$, there exists some $\theta' \in \Theta_{\text{oracle}}$ such that $\ell(\theta') < \ell(\theta)$, where both θ and θ' satisfy the centering constraint.

Consider any $\theta \notin \Theta_{\text{oracle}}$. By the definition of Θ_{oracle}, there exist some integers i and j where $2 \leq i < j \leq d$, such that $\theta_i \neq \theta_j$. By the symmetry of the manipulated observations $\{\widetilde{\mu}_{ij}\}$ defined in (57) with respect to item 2 through item d, we have that for any $\theta \in \mathbb{R}^d$,

$$\ell(\{\widetilde{\mu}_{i,j};\theta\}) = \ell(\{\widetilde{\mu}_{i,j};\theta_\pi\}),\tag{151}$$

where $\pi : \{2,\ldots,d\} \to \{2,\ldots,d\}$ is any permutation of item 2 through item d, and $\theta_\pi = [\theta_1, \theta_{\pi(2)}, \ldots, \theta_{\pi(d)}]$. For every $s \in \{0, 1, \ldots, d-2\}$, define π_s as the permutation where item 2 through item d are shifted s positions to the left in a circle. That is, for every $i \in \{2,\ldots,d\}$, we have

$$\pi_s(i) = 2 + [(i-2+s) \mod (d-1)].$$

Now define $\theta' = \frac{1}{d-1}\sum_{s=0}^{d-2}\theta_{\pi_s}$. It can be verified that

$$\theta' = \left[\theta_1, \frac{1}{d-1}\sum_{i=2}^{d}\theta_i, \ldots, \frac{1}{d-1}\sum_{i=2}^{d}\theta_i\right] \in \Theta_{\text{oracle}}.\tag{152}$$

Moreover, we have

$$\ell(\theta') = \ell\left(\frac{1}{d-1}\sum_{s=0}^{d-2}\theta_{\pi_s}\right) \stackrel{\text{(i)}}{<} \frac{1}{d-1}\sum_{s=0}^{d-2}\ell(\theta_{\pi_s}) \stackrel{\text{(ii)}}{=} \ell(\theta),$$

where (i) is due to the strict convexity of the negative log-likelihood function ℓ in Lemma A.1, and (ii) is due to (151).

Now we argue the equivalence of the unconstrained oracle $\widetilde{\theta}^{(\infty)}$ defined in (58a) and (59a). If a solution $\widetilde{\theta}^{(\infty)}$ to (58a) exists, then we have $\widetilde{\theta}^{(\infty)} \in \Theta_{\text{oracle}}$ and it is trivially also the solution to (59a). On the other hand, if a solution $\widetilde{\theta}^{(\infty)}$ to (59a) exists, assume for contradiction that $\widetilde{\theta}^{(\infty)}$ is not a solution to (58a). Then either there exists no solution to (58a), or the solution to (58a) is not $\widetilde{\theta}^{(\infty)}$. In either case, there exists some θ such that $\ell(\theta) < \ell(\widehat{\theta}^{(\infty)})$. By (152), we construct some $\theta' \in \Theta_{\text{oracle}}$ such that $\ell(\theta') < \ell(\theta) < \ell(\widehat{\theta}^{(\infty)})$. This contradicts the assumption that $\widehat{\theta}^{(\infty)}$ is the optimal solution to (59a). Hence, Eq. (58a) and (59a) are equivalent definitions of the unconstrained oracle $\widehat{\theta}^{(\infty)}$.

The same argument can be extended to the constrained oracle $\widehat{\theta}^{(B)}$, by noting that if $\theta \in \Theta_B$, then in the construction (152) we have $\theta' \in \Theta_B$.

A.4.10 Proof of Lemma A.12

Note that the lemma statement is conditioned on the event E_v. That is, we observe $\mu_1 = v$ for some $\frac{1}{2} \leq v \leq \mu_+ < 1$. In particular, we have $0 < \mu_1 < 1$. Then there exists at least one directed edge from node 1 to nodes $\{2, \ldots, d\}$, and at least one directed edge from nodes $\{2, \ldots, d\}$ to node 1. Then it suffices to prove that the subgraph consisting of nodes $\{2, \ldots, d\}$ is strongly-connected w.h.p.($\frac{1}{dk}$).

Note that the observations $\{\mu_{ij}\}$ for any $2 \leq i < j \leq d$ are all independent of μ_1, and therefore independent of the event E_v. Using the arguments in Lemma A.3, we have that the subgraph consisting of nodes $\{2, \ldots, d\}$ is strongly-connected w.h.p.($\frac{1}{dk}$).

A.4.11 Proof of Lemma A.13

Note that the lemma statement is conditioned on the event E_v. That is, we observe $\mu_1 = v$ for some $\frac{1}{2} \leq v \leq \mu_+ < 1$.

When $m = 1$, the desired inequality (67) holds trivially, because conditioned on E_v, we have

$$\sum_{i \neq 1} \mu_{1i} - \sum_{i \neq 1} \widetilde{\mu}^v_{1i} = (d-1)v - (d-1)v = 0.$$

Now consider every $m \in \{2 \ldots, d\}$. Consider the (unconditional) McDiarmid's inequality of (100) in the proof of Lemma A.4. Replacing the summation sign $\sum_{i \neq m}$ on the LHS of (100) by the summation sign $\sum_{\substack{i \geq 2 \\ i \neq m}}$ (that is, we further exclude $i = 1$ from the summation) yields the unconditional inequality:

$$\mathbb{P}\left[\left|\sum_{\substack{2 \leq i \leq d \\ i \neq m}} \mu_{mi} - \sum_{\substack{2 \leq i \leq d \\ i \neq m}} \mu^*_{mi}\right| \leq c\sqrt{\frac{d(\log d + \log k)}{k}}\right] \geq 1 - \frac{c'}{d^2 k}, \tag{153}$$

where $c, c' > 0$ are constants. Now we condition (153) on the event E_v. Note that for all $i, m \in \{2, \ldots, d\}$ with $i \neq m$, the terms $\{\mu_{mi}\}$ are independent of E_v. Moreover, by the expression of $\widetilde{\mu}^v_{mi}$ in (65), we have $\mu^*_{mi} = \frac{1}{2} = \widetilde{\mu}^v_{mi}$ conditioned on E_v. Hence, we have

$$\mathbb{P}\left[\left|\sum_{\substack{2 \leq i \leq d \\ i \neq m}} \mu_{mi} - \sum_{\substack{2 \leq i \leq d \\ i \neq m}} \widetilde{\mu}^v_{mi}\right| \leq c\sqrt{\frac{d(\log d + \log k)}{k}} \;\middle|\; E_v\right] \geq 1 - \frac{c'}{d^2 k}. \tag{154}$$

Now we bound the quantity $|\mu_{m1} - \widetilde{\mu}^v_{m1}|$ conditioned on E_v. By the definition of μ_1, we have that among the $(d-1)k$ comparisons $\{X^{(r)}_{1j}\}_{j \in \{2, \ldots, d\}, r \in [k]}$ in which item 1 is involved, there are $(d-1)k\mu_1$ terms that have value 1, and the rest have value 0. Hence, each μ_{1j} can be thought of as the mean of k comparisons sampled without replacement from the $(d-1)k$ comparisons $\{X^{(r)}_{1j}\}_{j \in \{2, \ldots, d\}, r \in [k]}$. By Hoeffding's inequality (sampling without replacement), we have that for every $j \in \{2, \ldots, d\}$,

$$\begin{aligned}\mathbb{P}\left[|\mu_{1j} - \widetilde{\mu}^v_{1j}| \leq c\sqrt{\frac{\log d + \log k}{k}} \;\middle|\; E_v\right] &\geq 1 - 2\exp\left(-c'(\log d + \log k)\right) \\ &\geq 1 - \frac{c''}{d^2 k},\end{aligned}$$

where $c, c', c'' > 0$ are constants. Equivalently, by a change of variables, we have that for every $j \in \{2, \ldots, d\}$,

$$\mathbb{P}\left[|\mu_{m1} - \widetilde{\mu}^v_{m1}| \leq c\sqrt{\frac{\log d + \log k}{k}} \;\middle|\; E_v \right] \geq 1 - \frac{c''}{d^2 k}. \tag{155}$$

Combining (154) and (155) by the triangle inequality, and taking a union bound over $m \in \{2, \ldots, d\}$ completes the proof.

B Proof of Theorem 2.2

In this appendix, we present the proof of Theorem 2.2. Both Theorem 2.2(a) and Theorem 2.2(b) are closely related to Theorem 2 from Shah et al. (2016). Under our setting, the quantity σ defined in Shah et al. (2016) is a universal constant, and the quantities ζ and γ defined in Shah et al. (2016) are constants that depend only on the constant B.

B.1 Proof of Theorem 2.2(a)

Theorem 2.2(a) is a direct consequence of Theorem 2(a) from Shah et al. (2016). We now provide some details on how to apply Theorem 2(a) from Shah et al. (2016). Under our setting, each pair of items is compared k times. Therefore, the sample size n is

$$n = \binom{d}{2} k = \Theta(d^2 k). \tag{156}$$

Moreover, under our setting the underlying topology is a complete graph. Let L denote the scaled Laplacian as defined in Eq. (4) from Shah et al. (2016), and let $L^\dagger$ denote the Moore-Penrose pseudoinverse of L. From Shah et al. (2016), the spectrum of L for a complete graph is $0, \frac{2}{d-1}, \ldots, \frac{2}{d-1}$. Therefore, we have

$$\lambda_2(L) = \frac{2}{d-1}, \tag{157a}$$

$$\mathrm{tr}(L^\dagger) = (d-1) \cdot \frac{d-1}{2} = \frac{(d-1)^2}{2}. \tag{157b}$$

Plugging (156) and (157) into Theorem 2(a) from Shah et al. (2016) shows that the Theorem 2.2(a) holds for all $k \geq k_0$, where k_0 is a constant.

B.2 Proof of Theorem 2.2(b)

The proof of Theorem 2.2(b) closely mimics the proof of Theorem 2(b) from Shah et al. (2016) (which is in turn based on Theorem 1(b) from Shah et al. (2016)). In what follows, we state a minor modification to be made in order to extend the proof from Shah et al. (2016) to Theorem 2.2(b).

In the proof from Shah et al. (2016), the box constraint for the MLE $\widehat{\theta}^{(B)}$ is only used to obtain the following bound (see Appendix A.2 from Shah et al. (2016)):

$$v^T \nabla^2 \ell(w) v \geq \frac{\gamma}{n\sigma^2} \|Xv\|_2^2, \qquad \text{for all } v, w \in \Theta_B. \tag{158}$$

Now we fix any constant A such that $A > B$. It can be verified that (158) still holds when replacing Θ_B by Θ_A, where we now allow γ to depend on both A and B. Since A is assumed to be a constant, we have that γ is still a constant. Then the rest of the arguments from Shah et al. (2016) carry to the proof of Theorem 2.2(b).