

Isotonic Mechanism for Exponential Family Estimation in Machine Learning Peer Review

Yuling Yan* Weijie J. Su[†] Jianqing Fan[‡]

August 16, 2024

Abstract

In 2023, the International Conference on Machine Learning (ICML) required authors with multiple submissions to rank their submissions based on perceived quality. In this paper, we aim to employ these author-specified rankings to enhance peer review in machine learning and artificial intelligence conferences by extending the Isotonic Mechanism to exponential family distributions. This mechanism generates adjusted scores that closely align with the original scores while adhering to author-specified rankings. Despite its applicability to a broad spectrum of exponential family distributions, implementing this mechanism does not require knowledge of the specific distribution form. We demonstrate that an author is incentivized to provide accurate rankings when her utility takes the form of a convex additive function of the adjusted review scores. For a certain subclass of exponential family distributions, we prove that the author reports truthfully *only if* the question involves only pairwise comparisons between her submissions, thus indicating the optimality of ranking in truthful information elicitation. Moreover, we show that the adjusted scores improve dramatically the estimation accuracy compared to the original scores and achieve nearly minimax optimality when the ground-truth scores have bounded total variation. We conclude with a numerical analysis of the ICML 2023 ranking data, showing substantial estimation gains in approximating a proxy ground-truth quality of the papers using the Isotonic Mechanism.

1 Introduction

In recent years, the surge in submissions to artificial intelligence and machine learning conferences, such as ICML, NeurIPS, and ICLR, has made it increasingly challenging to maintain the quality of peer review. For instance, the number of submissions to NeurIPS has grown from 1,673 in 2014 to a record-high of 10,411 in 2022 (Ivanov, 2020; Lawrence and Cortes, 2014; Li, 2022). This increase largely results from the fact that it has become common for an author to submit *multiple* papers to a single artificial intelligence and machine learning conference. For example, one author submitted 32 papers to ICLR 2020 (Sun, 2020). However, the number of qualified reviewers and the time they can offer are limited. This discrepancy has compelled conference organizers to recruit many novice reviewers, which has consequently compromised the quality of peer review (Shah, 2022; Stelmakh et al., 2021).

*Institute for Data, Systems, and Society, Massachusetts Institute of Technology, Cambridge, MA 02142, USA; Email: yulingy@mit.edu.

[†]Department of Statistics and Data Science, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104, USA; Email: suw@wharton.upenn.edu.

[‡]Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544, USA; Email: jqfan@princeton.edu.




Figure 1: Average review scores versus their standard deviations for all papers submitted to ICLR 2022 (<https://github.com/fedebotu/ICLR2022-OpenReviewData>). Each paper typically received three to five reviews, enabling the computation of its average score and standard deviation. Each point in this figure corresponds to a possible (average score, standard deviation) pair, where the size of a point reflects the number of papers with this pair of average review score and standard deviation.

Unlike review processes in most journals, reviewers for machine learning conferences are required to score papers based on quality (for example, on a scale from 1 to 10), and the average score across typically three reviewers serves as the single most important factor in deciding whether to accept or reject a paper. As a manifestation of the decline in review quality, it is increasingly common for review scores to exhibit a significant level of variability (see Figure 1). The arbitrariness of these scores makes it very difficult to make sound decisions regarding whether to accept or reject a paper, or to recommend paper awards. For example, the 2014 NIPS experiment showed that approximately 57% of the accepted papers would be rejected if reviewed twice (Lawrence and Cortes, 2014).

To assess submission quality based on review scores from a statistical perspective, we begin by assuming that the review scores follow an exponential family distribution in the canonical form:

$$p_{\theta}(x) = e^{\theta x - b(\theta)} c(x), \quad (1.1)$$

where θ is a natural parameter, b is a log-partition function, and c is a carrier function. Here, θ or, equivalently, $\mu := \mathbb{E}[X] = b'(\theta)$ can be considered as the ground-truth quality of the paper.¹ Exponential family distributions provide a simple yet flexible approach for modeling the dependence between the mean and the variance of a random variable, as is the case in Figure 1. The most relevant example is the Binomial distribution, whose data type (taking integer values within a certain range) and variance-mean dependence (smaller variance for high and low quality papers) makes it an ideal choice for modeling noisy review scores. For an author who submits $n \geq 2$ papers to a machine learning conference, the average review score X_i of the i th paper is thought of as a realization from the exponential family distribution with parameter θ_i^* :

$$X_i \sim p_{\theta_i^*}(x)$$

for $i = 1, \dots, n$.

¹This ground-truth quality can be understood in practical terms as the expected review score if the paper were to be evaluated by a large number of reviewers.

A readily available estimator of the quality $\mu_i^* = \mathbb{E} X_i = b'(\theta_i^*)$ of the i th paper is simply the raw review score X_i , which is the maximum likelihood estimator (MLE) of μ_i^* . To see this, note that the negative log-likelihood reads:

$$\log(p_{\theta_1}(X_1) \cdots p_{\theta_n}(X_n)) = \sum_{i=1}^n [\theta_i X_i - b(\theta_i) + \log c(X_i)],$$

where the review scores X_1, \dots, X_n are assumed to be independent in light of the fact that the review processes of the n papers are independent.² Minimizing the above negative log-likelihood over θ_i yields $\hat{\theta}_i = (b')^{-1}(X_i)$, which is equivalent to $\hat{\mu}_i = X_i$. The triviality of the MLE in this case arises from treating the n submissions separately, despite the fact that these n papers share a *common* author.

To obtain a better estimation of submission quality, ICML 2023, for the first time in its history, asked any authors who submitted multiple papers to *rank* their submissions from their own perspective, upon submission confirmation (<https://icml.cc/Conferences/2023/CallForPapers>, <https://openrank.cc>). This experiment is based on a recent proposal by Su (2021, 2022) to improve peer review in machine learning conferences. Formally, let an author of n submissions report a ranking of π such that, as claimed by the author, π sorts the ground-truth natural parameters $\theta_1^*, \dots, \theta_n^*$ in descending order.³ It is important to note, however, that the author might report an incorrect ranking, and therefore, π is not necessarily identical to the ground-truth ranking π^* that satisfies

$$\theta_{\pi^*(1)}^* \geq \theta_{\pi^*(2)}^* \geq \cdots \geq \theta_{\pi^*(n)}^*. \quad (1.2)$$

With the author-provided ranking in place, in this paper, we introduce the ranking-constraint MLE to estimate μ_1^*, \dots, μ_n^* :

$$\begin{aligned} \arg \min_{\theta_1, \dots, \theta_n} & \sum_{i=1}^n [-\theta_i X_i + b(\theta_i)] \\ \text{s.t.} & \theta_{\pi(1)} \geq \theta_{\pi(2)} \geq \cdots \geq \theta_{\pi(n)}. \end{aligned} \quad (1.3)$$

This is a convex program because $b(\cdot)$ is convex, and it possesses a unique minimizer $\hat{\boldsymbol{\theta}} = (\hat{\theta}_1, \dots, \hat{\theta}_n)$ in the general setting where $b(\cdot)$ is strictly convex. The estimated scores are

$$\hat{\boldsymbol{\mu}} = b'(\hat{\boldsymbol{\theta}}) = (b'(\hat{\theta}_1), \dots, b'(\hat{\theta}_n)). \quad (1.4)$$

The implementation of the mechanism (1.3) requires knowing the log-partition function $b(\cdot)$, which is also the case for obtaining the adjusted scores in (1.4). Interestingly, we show that this dependence can be circumvented, thereby making the mechanism *robust* to model misspecification. Specifically, by using a classical result from Barlow and Brunk (1972, Theorem 3.1), (1.3) is equivalent to isotonic regression

$$\begin{aligned} \arg \min_{\mu_1, \dots, \mu_n} & \sum_{i=1}^n (X_i - \mu_i)^2 \\ \text{s.t.} & \mu_{\pi(1)} \geq \mu_{\pi(2)} \geq \cdots \geq \mu_{\pi(n)}, \end{aligned} \quad (1.5)$$

²Due to the large numbers of submissions and reviewers, it is unlikely that different papers by the same author would be reviewed by the same reviewer.

³Mathematically speaking, $\pi(1), \dots, \pi(n)$ is a permutation of $1, \dots, n$.

which is referred to as the *Isotonic Mechanism* (Su, 2021), in the sense that the adjusted score vector $\hat{\boldsymbol{\mu}} = b'(\hat{\boldsymbol{\theta}})$ in (1.4) coincides with the solution to (1.5). The feasible set of (1.5) contains the ground-truth parameter $\boldsymbol{\mu}^* = (\mu_1^*, \dots, \mu_n^*)$ when $\pi = \pi^*$, that is, when the author is *truthful*. In particular, the feasible sets of (1.3) and (1.5) are essentially the same in view of the relationship $\mu = b'(\theta)$ and the mapping $b'(\cdot)$ is nondecreasing. This equivalence is convenient because it enables the implementation *without* requiring the knowledge of the probability distribution of the observations X_i 's.

The Isotonic Mechanism can be better understood within the framework of mechanism design (Myerson, 1989). In this context, the author acts as the sole agent possessing private information about the quality of all her papers submitted to a conference. The author provides a ranking that presumably sorts her papers from best to worst to the conference organizers—referred to as the principal. The principal then inputs this ranking, along with the review scores, into the Isotonic Mechanism to obtain adjusted review scores. It is important to note that reviewers assess the papers as they would in any standard machine learning conference, without knowledge of the author-provided rankings. Similarly, the author has no access to the review scores before submitting her ranking.

The Isotonic Mechanism (1.5) is *owner-assisted* (Su, 2022), as it requires authors to provide information pertaining to their own items. For the practical application of this mechanism, several questions immediately arise:

1. First, might the author provide an inaccurate ranking despite being aware of the ground-truth ranking?
2. Is it possible to elicit any other truthful information from the author? For example, can conference organizers inquire how *much* one paper surpasses another of the same author in terms of quality?
3. Lastly, would the estimator derived from the Isotonic Mechanism be more accurate than the raw scores X_i 's? If so, is it optimal in some sense?

A challenge in addressing these questions stems from the *generality* of exponential family distributions, including normal, binomial, and Poisson distributions, among others. In the subsequent part of this section, we present the main results of our paper that respond to these three questions. The proofs of the results in Sections 1.1, 1.2, and 1.3 are given in Sections 2, 3, and 4, respectively, accompanied by additional findings.

1.1 Truthfulness

To analyze the strategy of a *rational* author, it is imperative to specify the author's utility. We assume an additive utility function across all coordinates of the adjusted review score vector $\hat{\boldsymbol{\mu}} = (\hat{\mu}_1, \dots, \hat{\mu}_n)$. This assumption stems from the observation that, given the anonymity of the process and the large number of submissions, the decision-making processes⁴ for each paper are largely independent of one another. This rationale informs the subsequent assumption, as in Su (2022):

Assumption 1 (Utility). *The author's overall utility takes the form*

$$\text{Util} := \mathbb{E} \left[\sum_{i=1}^n U(\hat{\mu}_i) \right] \quad (1.6)$$

⁴ICML 2023 collects rankings but will not utilize adjusted scores in decision-making. Adopting such a significant reform in peer review necessitates time for such a large community.

for some nondecreasing convex function U .

The author is rational, which means that she will report the ranking that maximizes her expected utility, whether truthfully or otherwise. The convexity assumption acknowledges that a higher score may increase the likelihood of winning oral presentations or even best paper awards, thus yielding progressively greater payoffs. This is particularly applicable in the long term, as reviews are often made public at machine learning conferences such as ICLR. Consequently, a higher review score significantly enhances the submission’s impression within the community. Additional justification can be found in Section 4 of [Su \(2021\)](#) and Section 2 of [Su \(2022\)](#).

Assumption 2. *The author knows the ground-truth ranking π^* that fulfills (1.2) or, equivalently, $\mu_{\pi^*(1)}^* \geq \mu_{\pi^*(2)}^* \geq \dots \geq \mu_{\pi^*(n)}^*$.*

For this assumption, the author merely needs to understand the relative quality of her submissions rather than their precise quality.

The subsequent theorem examines a *regular* exponential family distribution $p_\theta(x) = e^{\theta x - b(\theta)} c(x)$, where $b(\cdot)$ is twice differentiable and satisfies $b''(\theta) > 0$ within its domain. This condition implies that X has nonzero and finite variance due to the identity $\text{Var}(X) = b''(\theta)$. Lastly, we presume that X_1, \dots, X_n are independently generated by the exponential family distribution with distinct parameters.

Let $\text{Util}(\pi)$ denote the expected overall utility (1.6) of the solution to the Isotonic Mechanism (1.5) when the input ranking is π .

Theorem 1. *Under Assumptions 1 and 2, the author maximizes her expected overall utility by truthfully reporting the ground-truth ranking π^* . That is,*

$$\text{Util}(\pi^*) \geq \text{Util}(\pi)$$

for any ranking π .

An important feature of this theorem is that it applies to the broad class of exponential family of distributions. In game theoretic terms, this theorem demonstrates that the owner-assisted mechanism (1.5) is incentive-compatible. If the utility function U is strictly convex, it satisfies strict incentive compatibility, which means that $\text{Util}(\pi^*) > \text{Util}(\pi)$ for any ranking π that does not sort $\theta^* = (\theta_1^*, \dots, \theta_n^*)$ or μ^* in descending order.

A numerical illustration of this theorem is presented in Figure 2. Consider a scenario with four submissions whose true scores are $(\mu_1^*, \mu_2^*, \mu_3^*, \mu_4^*) = (8, 7, 6, 4)$. For each submission, three independent review scores are generated from either a binomial (with 10 trials) or Poisson distribution, and their average is used as the original review score. This figure exhibits the author’s expected utility for all $4! = 24$ possible rankings. As observed for both distributions, the expected utility is maximized when the author provides the true ranking $\pi^* = (1, 2, 3, 4)$.

While Theorem 1 offers desired guarantees of truthfulness, the author might not be able to know the complete ranking of her papers. This motivates us to consider the scenario when, instead, the author knows only a *coarse ranking* ([Su, 2022](#)): given positive integers n_1, \dots, n_p satisfying $n_1 + \dots + n_p = n$, let blocks I_1, \dots, I_p be a partition of the indices $1, \dots, n$ such that the size $|I_j| = n_j$ for $1 \leq j \leq p$; a vector $\mathbf{a} \in \mathbb{R}^n$ is said to have coarse ranking $\mathbf{I} = (I_1, \dots, I_p)$, denoted as

$$\mathbf{a}_{I_1} \geq \dots \geq \mathbf{a}_{I_p},$$

if all entries of \mathbf{a}_{I_i} are larger than or equal to all entries of \mathbf{a}_{I_j} as long as $1 \leq i < j \leq p$.

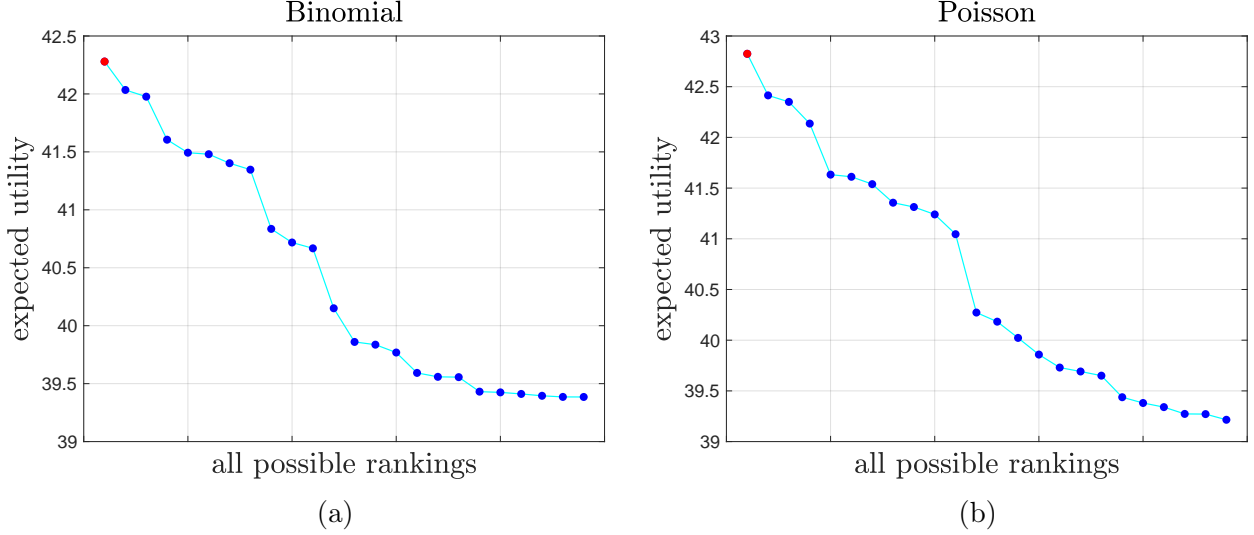


Figure 2: Expected overall utility for all $n! = 4! = 24$ possible rankings in descending order, averaged over 100 runs. The left panel uses review scores simulated from binomial distributions, while the right panel uses review scores simulated from Poisson distributions. In both panels, $(\mu_1^*, \mu_2^*, \mu_3^*, \mu_4^*) = (8, 7, 6, 4)$, and the utility function $U(x) = \max\{x, 0\}^2$ (binomial counts are generated from $\text{Binom}(10, \mu_i^*/10)$). The red dot represents the expected utility when the author reports truthfully. The second to fourth highest expected overall utilities are achieved by the rankings $(2, 1, 3, 4)$, $(1, 3, 2, 4)$, and $(1, 2, 4, 3)$ for binomial distributions, and $(1, 3, 2, 4)$, $(2, 1, 3, 4)$, and $(1, 2, 4, 3)$ for Poisson distributions, respectively. These results suggest that the mechanism may be stable against slightly misspecified rankings.

In a variant of the Isotonic Mechanism, we require any author to provide a coarse ranking of her submissions, which amounts to replacing the second line of (1.3) by

$$\theta_{I_1} \geq \theta_{I_2} \geq \dots \geq \theta_{I_p}.$$

This is equivalent to the following coarse Isotonic Mechanism:

$$\begin{aligned} \arg \min_{\mu_1, \dots, \mu_n} \quad & \sum_{i=1}^n (X_i - \mu_i)^2 \\ \text{s.t.} \quad & \mu_{I_1} \geq \mu_{I_2} \geq \dots \geq \mu_{I_p}. \end{aligned} \tag{1.7}$$

When $n_1 = n_2 = \dots = n_n = 1$, this mechanism reduces to the original Isotonic Mechanism. In general, the coarse Isotonic Mechanism provides the author with certain flexibility in that the author does not need to provide pairwise comparisons between some papers. For example, if $n_1 = 1$ and $n_2 = n - 1$, then the author is asked to report their best paper. In another example, $n_1 = \lfloor n/3 \rfloor$ and $n_2 = \lceil 2n/3 \rceil$, where $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ denote the floor and ceiling functions, respectively. The author is asked to report only the top one-third of her submissions.

The coarse Isotonic Mechanism is truthful, as we show in the following theorem. In this theorem, denote by $\text{Util}(\mathbf{I})$ the expected overall utility of the solution to the coarse Isotonic Mechanism (1.7); define $\mathbf{I}^* = (I_1^*, \dots, I_p^*)$ as

$$I_j^* = \left\{ \pi^*(i) : 1 + \sum_{k=1}^{j-1} n_k \leq i \leq \sum_{k=1}^j n_k \right\}$$

for $j = 1, 2, \dots, p$. This is the ground-truth coarse ranking with respect to $\boldsymbol{\mu}^*$, which satisfies

$$\boldsymbol{\mu}_{I_1^*}^* \geq \boldsymbol{\mu}_{I_2^*}^* \geq \dots \geq \boldsymbol{\mu}_{I_p^*}^*$$

Theorem 2. *Under Assumptions 1 and 2, the author maximizes her expected overall utility by truthfully reporting the coarse ground-truth ranking \mathbf{I}^* . That is,*

$$\text{Util}(\mathbf{I}^*) \geq \text{Util}(\mathbf{I})$$

for any coarse ranking \mathbf{I} of size (n_1, \dots, n_p) .

It is worth mentioning that in the peer review process, the author decides both the coarseness (n_1, \dots, n_p) and the reported coarse ranking. The above theorem asserts that for a given coarseness, the author maximizes her utility by truthfully reporting the coarse ranking, but it says nothing about the decision of the coarseness (e.g., the author might claim that the quality of her n papers are indistinguishable). In order to elicit as much truthful information as possible, in practice, the conference organizer might set a lower bound for p for an author that depends on her number of submissions n .

1.2 Is ranking optimal?

While we have shown truthfulness of eliciting a ranking or coarse ranking from the author, we aim to investigate the feasibility of truthfully eliciting other types of information for estimation. To achieve this, it is necessary to first formalize the types of information that conference organizers can elicit. Specifically, the n -dimensional Euclidean space \mathbb{R}^n , where the ground-truth vector $\boldsymbol{\mu}^*$ resides, is partitioned into $\mathcal{S} := \{S_1, S_2, \dots\}$ such that

$$\bigcup_{S \in \mathcal{S}} S = \mathbb{R}^n, \quad \text{and } S \cap S' = \emptyset$$

for any $S, S' \in \mathcal{S}$ but $S \neq S'$. The cardinality of \mathcal{S} can be either finite or infinite. As a collection of sets, \mathcal{S} is referred to as a knowledge partition and a set $S \in \mathcal{S}$ is called a knowledge element (Su, 2022). By definition, there must exist a unique knowledge element S^* that contains the ground-truth vector $\boldsymbol{\mu}^*$.

We focus on a class of knowledge partitions that are cut by hyperplanes. Specifically, a knowledge partition is cut by $(n-1)$ -dimensional hyperplanes H_1, \dots, H_K if any knowledge element is enclosed by some of H_1, \dots, H_K and no other hyperplane passes through its interior. Such a knowledge partition is called *linear*.

In the following variant of Assumption 2, we assume that the author knows which is the *true* knowledge element S^* . This assumption reduces to Assumption 2 when \mathcal{S} is the collection of all rankings.

Assumption 2'. *The author possesses sufficient understanding of her submissions to identify the specific knowledge element within the linear knowledge partition \mathcal{S} that contains the ground-truth $\boldsymbol{\mu}^*$.*

In this mechanism, the author is mandated to choose an element, say, S , from the linear knowledge partition \mathcal{S} and claim that S contains the ground-truth vector. This step is done before she observes the review scores $\mathbf{X} = (X_1, \dots, X_n)$. As earlier, S may or may not be the true knowledge

element S^* . The conference organizers solve the following shape-restricted regression:⁵

$$\begin{aligned} \arg \min_{\mu_1, \dots, \mu_n} \quad & \sum_{i=1}^n (X_i - \mu_i)^2 \\ \text{s.t.} \quad & \boldsymbol{\mu} \in S. \end{aligned} \tag{1.8}$$

Denote by $\hat{\boldsymbol{\mu}}^S$ the unique minimizer of this convex optimization problem. As in Assumption 1, the author seeks to maximize her expected overall utility

$$\text{Util}(S) = \mathbb{E} \left[\sum_{i=1}^n U(\hat{\mu}_i^S) \right]$$

by reporting any knowledge element S in the linear knowledge partition \mathcal{S} . We are interested in the class of linear knowledge partitions that incentivize the author to tell the truth.

Definition 1. *A linear knowledge partition \mathcal{S} is said to be truthful if, under Assumptions 1 and 2', the author could always maximize her expected overall utility by reporting the ground-truth knowledge element. That is,*

$$\text{Util}(S^*) \geq \text{Util}(S)$$

for any $S \in \mathcal{S}$.

Remark 1. This truthfulness condition requires that it holds for any possible values of the parameters $\theta_1, \dots, \theta_n$ in the exponential family distribution (1.1).

In a nutshell, Theorems 1 and 2 say that the collections of all rankings and, more generally, of all coarse rankings, are truthful. As a defining feature of these knowledge partitions, any (coarse) ranking is enclosed by *pairwise-comparison* hyperplanes $x_i - x_j = 0$ for some $1 \leq i < j \leq n$. For better estimation, however, more fine-grained knowledge partitions are preferred as they narrow down the “search space” S in (1.8). Specifically, does there exist a truthful linear knowledge partition that is more fine-grained than rankings?

Our next main result shows that under certain regularity conditions, which are satisfied by many exponential family distributions as we will show in Section 3, any linear truthful knowledge partition must be cut by pairwise-comparison hyperplanes up to an intercept.

Theorem 3. *Assume that the image of $b'(\cdot)$ is \mathbb{R} , and*

$$\lim_{\theta \rightarrow -\infty} \frac{b'(\theta)}{\sqrt{b''(\theta)}} \rightarrow -\infty, \quad \lim_{\theta \rightarrow \infty} \frac{b'(\theta)}{\sqrt{b''(\theta)}} \rightarrow \infty. \tag{1.9}$$

Consider any linear knowledge partition \mathcal{S} cut by hyperplanes H_1, \dots, H_K . If \mathcal{S} is truthful, then for each $1 \leq k \leq K$, H_k must take the form $x_{i_k} - x_{j_k} = \text{const}$ for some $1 \leq i_k < j_k \leq n$.

This theorem does not specify the intercepts of these hyperplanes. Although we hypothesize that these intercepts should be 0, presenting a proof for general exponential family distributions seems difficult, and the analysis might require examination on a case-by-case basis. For instance, the following proposition illustrates that in the case of normal distributions, the intercepts must be 0.

⁵This optimization may be convex or not, depending on whether the set S is convex or not.

Proposition 1. *Suppose $X_i \sim \mathcal{N}(\mu_i^*, \sigma^2)$ independently for $i = 1, \dots, n$, with $\sigma^2 > 0$. Consider any linear knowledge partition \mathcal{S} cut by hyperplanes H_1, \dots, H_K . If \mathcal{S} is truthful, then for each $1 \leq k \leq K$, H_k must be a pairwise-comparison hyperplane that takes the form $x_{i_k} - x_{j_k} = 0$ for some $1 \leq i_k < j_k \leq n$.*

These results demonstrate that pairwise comparisons are essential for truthfulness, and the author must consider the principle that “you can’t have your cake and eat it too” when providing information about her submission quality. The author can make pairwise comparisons by knowing any arbitrary monotone transformation of the coordinates of the ground-truth $\boldsymbol{\mu}^*$, without needing to know the individual values.

All $n(n-1)/2$ pairwise-comparison hyperplanes produce $n!$ isotonic cones, each of which is parameterized by a ranking π and takes the form $\boldsymbol{\mu} : \mu_{\pi(1)} \geq \dots \geq \mu_{\pi(n)}$. No isotonic cone can be further divided by pairwise-comparison hyperplanes. Consequently, this theorem implies that the collection of rankings is the most fine-grained truthful linear knowledge partition, at least in the case of normally distributed scores. Whenever the author knows the ground-truth ranking or, equivalently, Assumption 2 is satisfied, the Isotonic Mechanism is optimal among all estimators of the form (1.8).

1.3 Optimal estimation

It remains unclear whether the Isotonic Mechanism, along with its truthfulness, enhances estimation. If so, quantifying the extent it surpasses the original review scores \mathbf{X} is vital. We begin addressing these questions by presenting the following simple result.

Proposition 2. *Under Assumptions 1 and 2, the solution $\hat{\boldsymbol{\mu}}$ to the Isotonic Mechanism (1.5) satisfies*

$$\mathbb{E} \|\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}^*\|^2 \leq \mathbb{E} \|\mathbf{X} - \boldsymbol{\mu}^*\|^2,$$

where, throughout the paper, $\|\cdot\|$ denotes the ℓ_2 norm.

Remark 2. Due to the truthfulness of the mechanism, $\hat{\boldsymbol{\mu}}$ is the solution to (1.5) given the ground-truth ranking $\pi = \pi^*$. The other results in this subsection are also under Assumptions 1 and 2, but this fact will not be reiterated for the sake of simplicity. This theorem is also applicable to the coarse Isotonic Mechanism.

Moreover, the improvement can be considerable in a certain regime. We assume that all coordinates μ_1^*, \dots, μ_n^* are lower and upper bounded by V_{\min} and V_{\max} , respectively. This supplementary assumption is not restrictive when implementing the Isotonic Mechanism in machine learning conferences. For instance, all review scores must range from 1 to 10 for NeurIPS 2022 and 1 to 8 for ICML 2023.

Proposition 3. *Let the ground-truth scores μ_1^*, \dots, μ_n^* be arbitrary and satisfy $V_{\min} \leq \mu_1^*, \dots, \mu_n^* \leq V_{\max}$ for constants $V_{\min} < V_{\max}$. As $n \rightarrow \infty$, the Isotonic Mechanism satisfies*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \|\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}^*\|^2 = 0.$$

This proposition demonstrates that the (truthful) ranking constraint of the Isotonic Mechanism allows for estimation consistency. In stark contrast, the original review scores generally obey

$$\frac{1}{n} \mathbb{E} \|\mathbf{X} - \boldsymbol{\mu}^*\|^2 = \frac{\sum_{i=1}^n b''(\theta_i^*)}{n} \not\rightarrow 0$$

in the limit $n \rightarrow \infty$.

Next, we establish that the Isotonic Mechanism not only surpasses the original scores but also attains nearly minimax estimation among all estimators under certain conditions. We employ $\sup_{V_{\min} \leq \boldsymbol{\mu}^* \leq V_{\max}}$ as a shorthand for $\sup_{V_{\min} \leq \mu_i^* \leq V_{\max}, \forall 1 \leq i \leq n}$. Let

$$\sigma^2 := \max_{\theta_{\min} \leq \theta \leq \theta_{\max}} b''(\theta),$$

where $\theta_{\max} = (b')^{-1}(V_{\max})$ and $\theta_{\min} = (b')^{-1}(V_{\min})$. The quantity σ^2 represents the highest variance of review scores when the ground-truth score resides within $[V_{\min}, V_{\max}]$.

Theorem 4. *Under certain regularity conditions, there exist two constants $C_1, C_2 > 0$ such that*

$$\sup_{V_{\min} \leq \boldsymbol{\mu}^* \leq V_{\max}} \mathbb{E} \|\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}^*\|^2 \leq C_1 n \sigma^2 \min \left\{ 1, \left(\frac{V_{\max} - V_{\min}}{n\sigma} \right)^{2/3} + \frac{\log n}{n} \right\}, \quad (1.10)$$

where $\hat{\boldsymbol{\mu}}$ is the optimal solution to (1.5), and

$$\inf_{\tilde{\boldsymbol{\mu}}} \sup_{V_{\min} \leq \boldsymbol{\mu}^* \leq V_{\max}} \mathbb{E} \|\tilde{\boldsymbol{\mu}} - \boldsymbol{\mu}^*\|^2 \geq C_2 n \sigma^2 \min \left\{ 1, \left(\frac{V_{\max} - V_{\min}}{n\sigma} \right)^{2/3} + \frac{1}{n} \right\}, \quad (1.11)$$

where the infimum is taken over all estimators of the true mean vector $\boldsymbol{\mu}^*$ using data X_1, \dots, X_n .

The formal statement of Theorem 4 is provided in Theorem 5 in Section 4. The regularity condition is mild and can be satisfied by many exponential family distributions, including Gaussian, binomial, Poisson, and Gamma distributions. The upper bound (1.10) is developed in Zhang (2002). We obtain the lower bound (1.11) by extending the proof technique for the additive Gaussian noise setting in Bellec and Tsybakov (2015) to general exponential family distributions.

Theorem 4 demonstrates that the Isotonic Mechanism achieves minimax-optimal estimation, at most up to a logarithmic factor. This logarithmic factor would not be in effect when the total variation $V_{\max} - V_{\min}$ remains constant and n is large. The boundedness of the total variation indeed holds for peer review in machine learning conferences, as the scores must be integers between 1 and, for example, 8 or 10. In this case, the minimax estimation risk is $O((V_{\max} - V_{\min})^{2/3} n^{1/3} \sigma^{4/3})$, and this mechanism precisely achieves this rate. For comparison, the ℓ_2 risk of using the original scores \mathbf{X} is $n\sigma^2$. This shows that the Isotonic Mechanism significantly outperforms the original scores in terms of estimation when the number of submissions n is large and the noise level σ is high.

Figure 3 presents a numerical experiment comparing the estimation accuracy of the Isotonic Mechanism and the averaged review scores, using the same setting as in Section 1.1. The Isotonic Mechanism is significantly more accurate than the review scores.

1.4 Related work

The declining quality of peer review in machine learning conferences has been increasingly recognized, resulting in a flurry of research activity toward improving reviewer assignments (Jecmen et al., 2020; Kobren et al., 2019; Leyton-Brown et al., 2024), and bias reduction (MacKay et al., 2017; Paul, 1981; Wang and Shah, 2018), among other things; see Shah (2022) for a survey.

The most related work is Su (2021) and Su (2022), which consider the same problem as the present paper, but in the setting where the review scores follow

$$X_i = \mu_i^* + z_i \quad (1.12)$$




Figure 3: Estimation errors $\|\hat{\mu} - \mu^*\|_2^2/n$ for the Isotonic Mechanism and $\|\mathbf{X} - \mu^*\|_2^2/n$ for the original review scores, with varying numbers of submissions. The left panel presents results when the review scores follow binomial distributions, while the right panel exhibits results when the review scores follow Poisson distributions, in the same setting as the numerical experiment in Section 1.1. For any n , we set $\mu_i^* = 9 - 6(i - 1)/(n - 1)$ for $i = 1, \dots, n$.

with i.i.d. noise terms z_1, \dots, z_n . Specifically, Su (2021) proposes the Isotonic Mechanism and proves its truthfulness when applied to (1.12). In Su (2022), the author further shows that pairwise comparisons are necessary for the truthfulness of the general mechanism (1.8) in the setting of the additive-noise model (1.12).

However, it is worth noting that the techniques developed in Su (2021) and Su (2022) are not sufficient to cope with general exponential family distributions. In more detail, the proof of truthfulness in Su (2021) for the additive-noise model (1.12) heavily relies on a coupling technique based on the exchangeability of the noise terms. In the case of general exponential family distributions, the noise distribution and particularly, the variance of the noise, vary with the mean. Our proof of Theorems 1 and 2 employs a novel technique that addresses the mean-variance dependence.

Moreover, the proof of Theorem 3 is much more sophisticated for exponential family distributions than that for the additive-noise model (1.12) in Su (2022). This distinction arises from the mean-variance dependence in exponential family distributions that prevents us from constructing certain solutions to the general mechanism (1.8), essential to the proof in Su (2022). To tackle this challenge, we leverage a technique based on symmetric distributions, which may be a contribution of independent interest.

From a different angle, the present work connects to cheap talk in game theory and mechanism design (Chakraborty and Harbaugh, 2007). Broadly speaking, cheap talk is a type of game where providing information is free and does not directly affect the payoffs. In the Isotonic Mechanism, it does not cost the author/sender to provide the ranking information, thus broadly falling under cheap talk. Furthermore, Battaglini (2002) shows that common ground can be attained between a privately informed sender and a conference organizer/receiver in multidimensional cheap talk. More recently, Lipnowski and Ravid (2020) characterizes the class of information that can be truthfully elicited from the sender. However, the major distinction between this line of literature and the present paper is that the Isotonic Mechanism assumes the setting where the receiver has private

information in the form of review scores.

2 Truthfulness

In this section, we prove Theorems 1 and 2. As we shall see soon, our proofs make use of several technique novelties that are not found in Su (2021, 2022) in order to cope with the dependence between mean and variance for exponential family distributions.

2.1 Preliminaries

We start by collecting some elements that will be useful for the proofs.

For any vector $\mathbf{a} \in \mathbb{R}^n$, let

$$a_{[1]} \geq a_{[2]} \geq \cdots \geq a_{[n]}$$

denote the components of \mathbf{a} in descending order. Let \mathcal{C} denote the standard isotonic cone

$$\mathcal{C} := \{\mathbf{x} \in \mathbb{R}^n : x_1 \geq x_2 \geq \cdots \geq x_n\}. \quad (2.1)$$

As is clear, \mathcal{C} is a convex cone. Let $\mathcal{P}_{\mathcal{C}} : \mathbb{R}^n \rightarrow \mathcal{C}$ denote the Euclidean projection onto the standard isotonic cone. That is, for any $\mathbf{x} \in \mathbb{R}^n$,

$$\mathcal{P}_{\mathcal{C}}(\mathbf{x}) := \arg \min_{\mathbf{y} \in \mathcal{C}} \|\mathbf{x} - \mathbf{y}\|^2. \quad (2.2)$$

Next, we introduce two types of majorization.

Definition 2 (Marshall et al. (1979)). *For any two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, we say that \mathbf{a} majorizes \mathbf{b} , denoted $\mathbf{a} \succeq \mathbf{b}$, if*

$$\sum_{i=1}^k a_{[i]} \geq \sum_{i=1}^k b_{[i]}$$

holds for all $1 \leq k \leq n$, with equality when $k = n$.

Definition 3 (Su (2022)). *For any two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, we say that \mathbf{a} majorizes \mathbf{b} in the natural order, denoted $\mathbf{a} \succeq_{\text{no}} \mathbf{b}$, if*

$$\sum_{i=1}^k a_i \geq \sum_{i=1}^k b_i$$

holds for all $1 \leq k \leq n$, with equality when $k = n$.

The following two lemmas show how the two majorization definitions relate to each other and showcase the connection between the first definition and the convex utility (1.6) in Assumption 1.

Lemma 1. *For any two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, if $\mathbf{a} \succeq_{\text{no}} \mathbf{b}$, then $\mathcal{P}_{\mathcal{C}}(\mathbf{a}) \succeq \mathcal{P}_{\mathcal{C}}(\mathbf{b})$.*

Proof. See Su (2022, Lemma 6.5). □

Lemma 2 (Hardy–Littlewood–Pólya inequality). *For any two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, the inequality*

$$\sum_{i=1}^n h(a_i) \geq \sum_{i=1}^n h(b_i)$$

holds for any convex function $h : \mathbb{R} \rightarrow \mathbb{R}$ if and only if $\mathbf{a} \succeq \mathbf{b}$.

Proof. See Marshall et al. (1979, Section 4.B.2). □

2.2 Proof of Theorem 1

Without loss of generality, we assume that the ground-truth ranking $\pi^*(i) = i$. That is,

$$\mu_1^* \geq \mu_2^* \geq \cdots \geq \mu_n^* \quad \text{and} \quad \theta_1^* \geq \theta_2^* \geq \cdots \geq \theta_n^*.$$

For two different rankings π and ν , we say that π is an upward swap of ν if there exist $1 \leq i < j \leq n$ such that

$$\pi(i) = \nu(j) < \pi(j) = \nu(i)$$

and $\pi(k) = \nu(k)$ for all $k \neq i, j$. Roughly speaking, π and ν are identical except for their i th and j th entries, and π is “close” to π^* .

Our proof is a stone’s throw once we show that

$$\text{Util}(\pi) \geq \text{Util}(\nu) \tag{2.3}$$

as long as π is an upward swap of ν . To see this point, note the following well-known fact concerning permutations:

Lemma 3. *For any permutation π of $\{1, \dots, n\}$, there exist an integer $m \geq 1$ and a sequence of permutations π_1, \dots, π_m such that $\pi_1 = \pi^*$, $\pi_m = \pi$, and for each $1 \leq i \leq m - 1$, π_i is an upward swap of π_{i+1} .*

Remark 3. For completeness, a proof of this fact is given in the appendix.

To prove Theorem 1, we can get

$$\text{Util}(\pi^*) = \text{Util}(\pi_1) \geq \text{Util}(\pi_2) \geq \cdots \geq \text{Util}(\pi_m) = \text{Util}(\pi),$$

by applying (2.3) repeatedly, where $\{\pi_i\}_{1 \leq i \leq m}$ are constructed in Lemma 3 that satisfies: $\pi_1 = \pi^*$, $\pi_m = \pi$, and for each $1 \leq i \leq m - 1$, π_i is an upward swap of π_{i+1} .

The remainder of the proof aims to establish (2.3). To begin with, write

$$\text{Util}(\pi) = \mathbb{E} \left[\sum_{k=1}^n U([\mathcal{P}_{\mathcal{C}}(\pi \circ \mathbf{X})]_k) \right]$$

and

$$\text{Util}(\nu) = \mathbb{E} \left[\sum_{k=1}^n U([\mathcal{P}_{\mathcal{C}}(\nu \circ \mathbf{X})]_k) \right],$$

where \mathcal{C} and $\mathcal{P}_{\mathcal{C}}$ are defined in (2.1) and (2.2), respectively. Next, define

$$\Delta(x, y) = \mathbb{E} \left[\sum_{k=1}^n U([\mathcal{P}_{\mathcal{C}}(\pi \circ \mathbf{X})]_k) - \sum_{k=1}^n U([\mathcal{P}_{\mathcal{C}}(\nu \circ \mathbf{X})]_k) \mid X_{\pi(i)} = x, X_{\pi(j)} = y \right].$$

It is straightforward to check the following properties for this function:

(a) Recall that $\pi(i) = \nu(j) < \pi(j) = \nu(i)$ and $\pi(k) = \nu(k)$ for all $k \neq i, j$. Therefore, we have

$$\Delta(x, y) = -\Delta(y, x) \tag{2.4}$$

for any x, y . This immediately implies that

$$\Delta(x, x) = 0. \tag{2.5}$$

(b) If $x > y$, conditional on $X_{\pi(i)} = x$ and $X_{\pi(j)} = y$, we have $\pi \circ \mathbf{X} \succeq_{\text{no}} \nu \circ \mathbf{X}$. By Lemma 1, we get

$$\mathcal{P}_{\mathcal{C}}(\pi \circ \mathbf{X}) \succeq \mathcal{P}_{\mathcal{C}}(\nu \circ \mathbf{X}).$$

Together with Lemma 2, the Hardy–Littlewood–Pólya inequality, the majorization relationship above yields

$$\sum_{i=1}^n U([\mathcal{P}_{\mathcal{C}}(\pi \circ \mathbf{X})]_i) \geq \sum_{i=1}^n U([\mathcal{P}_{\mathcal{C}}(\nu \circ \mathbf{X})]_i)$$

for any (nondecreasing) convex function U . Therefore we know that

$$\begin{aligned} \Delta(x, y) &\geq 0 \quad \text{for } x > y, \\ \Delta(x, y) &\leq 0 \quad \text{for } x < y. \end{aligned} \tag{2.6}$$

Taken together, the elements above show that

$$\begin{aligned} \text{Util}(\pi) - \text{Util}(\nu) &\stackrel{\text{(i)}}{=} \mathbb{E} \left[\mathbb{E} \left[\sum_{k=1}^n U([\mathcal{P}_{\mathcal{C}}(\pi \circ \mathbf{X})]_k) - \sum_{k=1}^n U([\mathcal{P}_{\mathcal{C}}(\nu \circ \mathbf{X})]_k) \mid X_{\pi(i)}, X_{\pi(j)} \right] \right] \\ &= \mathbb{E} [\Delta(X_{\pi(i)}, X_{\pi(j)})] \\ &= \int \Delta(x, y) P_{\theta_{\pi(i)}}(dx) P_{\theta_{\pi(j)}}(dy), \end{aligned}$$

where (i) indicates the use of the tower property of conditional expectation and $P_{\theta}(dx)$ denotes the probability measure induced by the exponential family distribution (1.1). To proceed, we have

$$\begin{aligned} &\int \Delta(x, y) P_{\theta_{\pi(i)}}(dx) P_{\theta_{\pi(j)}}(dy) \\ &\stackrel{\text{(ii)}}{=} \int_{x>y} \Delta(x, y) P_{\theta_{\pi(i)}}(dx) P_{\theta_{\pi(j)}}(dy) + \int_{x<y} \Delta(x, y) P_{\theta_{\pi(i)}}(dx) P_{\theta_{\pi(j)}}(dy) \\ &\stackrel{\text{(iii)}}{=} \int_{x>y} \Delta(x, y) P_{\theta_{\pi(i)}}(dx) P_{\theta_{\pi(j)}}(dy) - \int_{x>y} \Delta(x, y) P_{\theta_{\pi(j)}}(dx) P_{\theta_{\pi(i)}}(dy) \\ &= \int_{x>y} \Delta(x, y) c(x) c(y) \exp(\theta_{\pi(i)}x + \theta_{\pi(j)}y - b(\theta_{\pi(i)}) - b(\theta_{\pi(j)})) \rho(dx) \rho(dy) \\ &\quad - \int_{x>y} \Delta(x, y) c(x) c(y) \exp(\theta_{\pi(i)}y + \theta_{\pi(j)}x - b(\theta_{\pi(i)}) - b(\theta_{\pi(j)})) \rho(dx) \rho(dy) \\ &= e^{-b(\theta_{\pi(i)})-b(\theta_{\pi(j)})} \int_{x>y} \Delta(x, y) c(x) c(y) \left[e^{\theta_{\pi(i)}x+\theta_{\pi(j)}y} - e^{\theta_{\pi(i)}y+\theta_{\pi(j)}x} \right] \rho(dx) \rho(dy) \\ &\stackrel{\text{(iv)}}{\geq} 0, \end{aligned}$$

where ρ denotes the base measure with respect to the exponential family distribution.

To be complete, (ii) follows from (2.5); (iii) utilizes (2.4); and (iv) holds since when $x > y$, we have $\Delta(x, y) \geq 0$, which is ensured by (2.6), and that $\theta_{\pi(i)}x + \theta_{\pi(j)}y \geq \theta_{\pi(i)}y + \theta_{\pi(j)}x$ as $\theta_{\pi(i)} \geq \theta_{\pi(j)}$. Therefore, we have proved (2.3) when π is an upward swap of ν .

This concludes the proof of Theorem 1.

2.3 Proof of Theorem 2

Without loss of generality, as earlier, we assume that $\pi^*(i) = i$:

$$\mu_1^* \geq \mu_2^* \geq \cdots \geq \mu_n^* \quad \text{and} \quad \theta_1^* \geq \theta_2^* \geq \cdots \geq \theta_n^*.$$

For any coarse ranking $\mathbf{I} := (I_1, \dots, I_p)$ of sizes n_1, \dots, n_p , let $\pi_{\mathbf{I}, \mathbf{X}}$ be the permutation such that

$$I_1 = \{\pi_{\mathbf{I}, \mathbf{X}}(i)\}_{1 \leq i \leq n_1}, \quad I_2 = \{\pi_{\mathbf{I}, \mathbf{X}}(i)\}_{n_1 < i \leq n_1 + n_2}, \dots, \quad I_p = \{\pi_{\mathbf{I}, \mathbf{X}}(i)\}_{n_1 + \dots + n_{p-1} < i \leq n}$$

and

$$X_{\pi_{\mathbf{I}, \mathbf{X}}(i)} \geq X_{\pi_{\mathbf{I}, \mathbf{X}}(j)} \tag{2.7}$$

for all $i < j \in I_k, 1 \leq k \leq p$. As is clear, $\pi_{\mathbf{I}, \mathbf{X}}$ is a random permutation.

Henceforth, we omit the subscript \mathbf{X} and use $\pi_{\mathbf{I}}$ to denote this permutation as long as it is clear from the context. It follows from Su (2022, Lemma 6.14) that the solution to the coarse Isotonic Mechanism (1.7) with the coarse ranking \mathbf{I} is equivalent to the Isotonic Mechanism with ranking $\pi = \pi_{\mathbf{I}}$.

We say that a coarse ranking \mathbf{I} is a coarse upward swap of another coarse ranking $\mathbf{J} := (J_1, \dots, J_p)$ of the same sizes if there exist $1 \leq k_1 < k_2 \leq p$ and $a \in I_{k_1}, b \in I_{k_2}$ such that

$$J_{k_1} = (I_{k_1} \setminus \{a\}) \cup \{b\}, \quad J_{k_2} = (I_{k_2} \setminus \{b\}) \cup \{a\}, \quad \text{and} \quad a < b.$$

Now, we will show that if \mathbf{I} is a coarse upward swap of \mathbf{J} , then

$$\text{Util}(\pi_{\mathbf{I}}) \geq \text{Util}(\pi_{\mathbf{J}}). \tag{2.8}$$

Let $j = \pi_{\mathbf{J}}^{-1}(a)$ and $i = \pi_{\mathbf{J}}^{-1}(b)$. By definition, $i < j$. Let $\pi_{\mathbf{J} \rightarrow \mathbf{I}}$ be derived from $\pi_{\mathbf{J}}$ by changing only the indices i and j . That is,

$$\pi_{\mathbf{J} \rightarrow \mathbf{I}}(i) = a, \quad \pi_{\mathbf{J} \rightarrow \mathbf{I}}(j) = b, \quad \text{and} \quad \pi_{\mathbf{J} \rightarrow \mathbf{I}}(l) = \pi_{\mathbf{J}}(l)$$

for all $l \neq i, j$. Then, $\pi_{\mathbf{J} \rightarrow \mathbf{I}}$ is an upward swap of $\pi_{\mathbf{J}}$, which allows us to use (2.3) to conclude that

$$\text{Util}(\pi_{\mathbf{J} \rightarrow \mathbf{I}}) \geq \text{Util}(\pi_{\mathbf{J}}). \tag{2.9}$$

In addition, it is straightforward to check that

$$\begin{aligned} I_1 &= \{\pi_{\mathbf{I}}(i)\}_{1 \leq i \leq n_1} = \{\pi_{\mathbf{J} \rightarrow \mathbf{I}}(i)\}_{1 \leq i \leq n_1}, \quad \dots \\ I_p &= \{\pi_{\mathbf{I}}(i)\}_{n_1 + \dots + n_{p-1} < i \leq n} = \{\pi_{\mathbf{J} \rightarrow \mathbf{I}}(i)\}_{n_1 + \dots + n_{p-1} < i \leq n}, \end{aligned}$$

which together with (2.7) establishes

$$\pi_{\mathbf{I}} \circ \mathbf{X} \succeq_{\text{no}} \pi_{\mathbf{J} \rightarrow \mathbf{I}} \circ \mathbf{X}.$$

By applying Lemma 1, this immediately gives

$$\mathcal{P}_{\mathcal{C}}(\pi_{\mathbf{I}} \circ \mathbf{X}) \succeq \mathcal{P}_{\mathcal{C}}(\pi_{\mathbf{J} \rightarrow \mathbf{I}} \circ \mathbf{X}).$$

In view of the Hardy–Littlewood–Pólya inequality (Lemma 2), we get

$$\sum_{i=1}^n U([\mathcal{P}_{\mathcal{C}}(\pi_{\mathbf{I}} \circ \mathbf{X})]_i) \geq \sum_{i=1}^n U([\mathcal{P}_{\mathcal{C}}(\pi_{\mathbf{J} \rightarrow \mathbf{I}} \circ \mathbf{X})]_i)$$

for any (nondecreasing) convex function U , which immediately gives

$$\text{Util}(\pi_{\mathbf{I}}) \geq \text{Util}(\pi_{\mathbf{J} \rightarrow \mathbf{I}}). \quad (2.10)$$

Taking (2.9) and (2.10) collectively proves (2.8).

As with Lemma 3 for rankings, for any coarse ranking \mathbf{I} that is not truthful, we can show that there exist a sequence of coarse rankings $\mathbf{I}_1, \dots, \mathbf{I}_m$ of the same sizes such that $\mathbf{I}_1 = \mathbf{I}^*$ (the truthful coarse ranking), $\mathbf{I}_m = \mathbf{I}$, and for each $1 \leq j \leq m-1$, \mathbf{I}_j is a coarse upward swap of \mathbf{I}_{j+1} . This fact allows us to apply (2.8) repeatedly and obtain

$$\text{Util}(\pi_{\mathbf{I}^*}) = \text{Util}(\pi_{\mathbf{I}_1}) \geq \text{Util}(\pi_{\mathbf{I}_2}) \geq \dots \geq \text{Util}(\pi_{\mathbf{I}_m}) = \text{Util}(\pi_{\mathbf{I}}),$$

thereby finishing the proof.

3 Ranking is (almost) optimal for truthfulness

In this section, we present the proof for Theorem 3 and Proposition 1. We begin by gathering some preliminary facts in Section 3.1 and then outline the proof in Section 3.2. The proofs of Theorem 3 and Proposition 1 are provided in Sections 3.3 and 3.4, respectively.

3.1 Preliminaries

Definition 4 (Marshall et al. (1979)). *For any two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, we say that \mathbf{a} weakly majorizes \mathbf{b} , denoted $\mathbf{a} \succeq_w \mathbf{b}$, if*

$$\sum_{i=1}^k a_{[i]} \geq \sum_{i=1}^k b_{[i]}$$

holds for all $1 \leq k \leq n$.

Lemma 4 (Hardy–Littlewood–Pólya inequality). *For any two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, the inequality*

$$\sum_{i=1}^n h(a_i) \geq \sum_{i=1}^n h(b_i)$$

holds for all nondecreasing convex function $h : \mathbb{R} \rightarrow \mathbb{R}$ if and only if $\mathbf{a} \succeq_w \mathbf{b}$.

Proof. See Marshall et al. (1979, Section 4.B.2). □

Suppose that the linear knowledge partition in \mathbb{R}^n is constructed by a finite set of separating hyperplanes, where a hyperplane is defined as $\{\mathbf{x} : \mathbf{a}^\top \mathbf{x} = b\}$ for some $\mathbf{a} \in \mathbb{R}^n$ and $b \in \mathbb{R}$. We group these hyperplanes into K groups $\mathcal{H}_1, \dots, \mathcal{H}_K$, where each group \mathcal{H}_k consists of l_k parallel hyperplanes. To be precise, we let

$$\mathcal{H}_k = \{H_{k,1}, \dots, H_{k,l_k}\}, \quad \text{where} \quad H_{k,l} = \{\mathbf{x} : \mathbf{a}_k^\top \mathbf{x} = b_{k,l}\}, \quad \forall 1 \leq k \leq K, 1 \leq l \leq l_k.$$

In addition, we assume that $\|\mathbf{a}_k\|_2 = 1$ and $b_{k,1} > \dots > b_{k,l_k}$ for all $1 \leq k \leq K$ for identifiability. We assume that the hyperplanes in different groups are not parallel, namely for each $1 \leq k_1 \neq k_2 \leq K$, $\mathbf{a}_{k_1} \neq \mathbf{a}_{k_2}$. We also define a set of hyperplanes

$$I_{i,j} := \{\mathbf{x} : x_i - x_j = 0\}, \quad \forall 1 \leq i \neq j \leq n.$$

We also collect a few basic facts in geometry, which will be useful throughout the proof. For any point $\mathbf{x}_0 \in \mathbb{R}^n$ and a hyperplane $H = \{\mathbf{x} : \mathbf{a}^\top \mathbf{x} = b\}$, the Euclidean projection of \mathbf{x}_0 onto H is

$$\mathcal{P}_H(\mathbf{x}_0) = \mathbf{x}_0 - \frac{\mathbf{a}^\top \mathbf{x}_0 - b}{\|\mathbf{a}\|^2} \mathbf{a}. \quad (3.1)$$

The Euclidean distance between \mathbf{x}_0 and H is

$$\text{dist}(\mathbf{x}_0, H) = \frac{|\mathbf{a}^\top \mathbf{x}_0 - b|}{\|\mathbf{a}\|}. \quad (3.2)$$

3.2 Proof architecture

In this section we present the outline of the proof. For simplicity we only focus on the case where each group \mathcal{H}_k only contains one hyperplane (denoted by H_k), and the review scores $X_i \sim \mathcal{N}(\mu_i^*, \sigma^2)$. The general setting will be treated in the formal proof.

We first show that each hyperplane H_k is parallel to some $I_{i,j}$. In order to better appreciate our proof idea, we first sketch the proof of (Su, 2022, Theorem 1), which can be viewed as the noiseless case ($\sigma = 0$). That paper proves the result by contradiction: if H_k is not parallel to any $I_{i,j}$, then we will be able to find a point $\boldsymbol{\omega} \in H_k$ such that each entry of $\boldsymbol{\omega}$ is different. For some sufficiently small $\varepsilon > 0$, assuming that $\boldsymbol{\omega} + \varepsilon \mathbf{a}_k$ is the true score, then $\boldsymbol{\omega} + \varepsilon \mathbf{a}_k$ should weakly majorize its projection onto the knowledge element on the other side of H_k , which is given by $\boldsymbol{\omega}$; similarly one can also deduce that $\boldsymbol{\omega} - \varepsilon \mathbf{a}_k$ also weakly majorizes $\boldsymbol{\omega}$. In view of (Su, 2022, Lemma 6.3), this leads to a contradiction. However in the noisy case ($\sigma > 0$), the analysis becomes challenging since the noisy score can take value in the entire \mathbb{R}^n , and there is no simple expression for the projection of the noisy score onto other knowledge element. This issue can be resolved by moving $\boldsymbol{\omega}$ and $\boldsymbol{\omega} \pm \varepsilon \mathbf{a}_k$ far away from the origin and changing the scale, which can effectively be viewed as reducing the magnitude of noise. More concretely, we can find a vector \mathbf{u} that is parallel to H_k and has distinct entries. Then for any $t > 0$, we can consider two possible ground truths $\boldsymbol{\omega} + t\mathbf{u} \pm \varepsilon t \mathbf{a}_k$, and the noisy scores look like $\boldsymbol{\omega} + t\mathbf{u} \pm \varepsilon t \mathbf{a}_k + \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_n)$. By rescaling everything with $1/t$, as t goes to infinity we can show that both $\mathbf{u} \pm \varepsilon \mathbf{a}_k$ weakly majorize \mathbf{u} , which (similar to the noiseless case) leads to a contradiction.

Then we show that each hyperplane H_k passes through the origin. If this is not true, we will find a point $\mathbf{y} \in H_k$ that is sufficiently far away from the origin and other hyperplanes $\{H_{k'} : k' \neq k\}$. Consider the case when the ground truth is \mathbf{y} and the utility function takes the form $U(\mathbf{x}) = \|\mathbf{x}\|_2^2$, then we can compute the expected utility when one reports the knowledge element on both sides of H_k that contains \mathbf{y} respectively, which turns out to be different. This leads to a contradiction, which in turn shows that each hyperplane must pass through the origin. Taking the above two results collectively shows that each hyperplane H_k must be a pairwise comparison hyperplane.

3.3 Step 1: determining the normal vectors $\{\mathbf{a}_k\}_{1 \leq k \leq K}$ (Proof of Theorem 3)

In Step 1, we demonstrate that for any given $1 \leq k \leq K$, there must exist $1 \leq i < j \leq n$ such that \mathbf{a}_k is a constant multiple of the vector $\mathbf{e}_i - \mathbf{e}_j$. We will establish this through a proof by contradiction. Suppose this statement is false; then, we can identify a direction $\mathbf{u} \in \mathbb{R}^d$ satisfying: (i) \mathbf{u} is parallel to the hyperplanes in \mathcal{H}_k , (ii) \mathbf{u} is not parallel to the hyperplanes in other $\mathcal{H}_{k'}$ for $k' \neq k$, and (iii) \mathbf{u} is not parallel to the hyperplane $I_{i,j}$ for any $1 \leq i < j \leq n$. To see why this is true,

we can express these conditions equivalently using the following systems of equations/inequalities:

$$\begin{cases} \mathbf{a}_k^\top \mathbf{u} = 0, \\ \mathbf{a}_{k'}^\top \mathbf{u} \neq 0, & \text{for all } k' \neq k \\ u_i \neq u_j, & \text{for all } 1 \leq i \neq j \leq n. \end{cases} \quad (3.3)$$

The existence of such \mathbf{u} is evident given that $\mathbf{a}_{k'} \neq \mathbf{a}_k$ for all $k' \neq k$ and that \mathbf{a}_k is not a constant multiple of $\mathbf{e}_i - \mathbf{e}_j$ for all $1 \leq i \neq j \leq n$. We can then select an arbitrary \mathbf{u} satisfying (3.3) and assume $\|\mathbf{u}\| = 1$ for identifiability.

Next, we fix an arbitrary point \mathbf{y}_1 on $H_{k,1}$, and let \mathbf{y}_2 be its Euclidean projection onto H_{k,l_k} . Define

$$\mathbf{y}_1(t) := \mathbf{y}_1 + t\mathbf{u} \quad \text{and} \quad \mathbf{y}_2(t) := \mathbf{y}_2 + t\mathbf{u}.$$

We know that

$$\text{dist}(\mathbf{y}_1(t), \mathbf{y}_2(t)) = \frac{b_{k,1} - b_{k,l_k}}{\|\mathbf{a}_k\|} = b_{k,1} - b_{k,l_k}, \quad \forall t \geq 0 \quad (3.4)$$

We can use (3.2) to compute, for $s = 1, 2$

$$\begin{aligned} \text{dist}(\mathbf{y}_s(t), H_{k',l}) &= \left| \mathbf{a}_{k'}^\top (\mathbf{y}_s + t\mathbf{u}) - b_{k',l} \right| = \left| \mathbf{a}_{k'}^\top \mathbf{y}_s - b_{k',l} + t\mathbf{a}_{k'}^\top \mathbf{u} \right| \\ &\geq t \left| \mathbf{a}_{k'}^\top \mathbf{u} \right| - \left| \mathbf{a}_{k'}^\top \mathbf{y}_s - b_{k',l} \right| \end{aligned}$$

for any $k' \neq k$ and $1 \leq l \leq l_{k'}$. In view of (3.3), there exists a constant

$$c_1 := \frac{1}{4} \min_{k' \neq k} \left| \mathbf{a}_{k'}^\top \mathbf{u} \right| > 0$$

and a sufficiently large constant T_1 satisfying

$$T_1 > \frac{1}{2c_1} \max_{k' \neq k, s \in \{1,2\}, 1 \leq l \leq l_{k'}} \left| \mathbf{a}_{k'}^\top \mathbf{y}_s - b_{k',l} \right|,$$

such that for any $t \geq T_1$ and $s = 1, 2$

$$\text{dist}(\mathbf{y}_s(t), H_{k',l}) \geq 2c_1 t \quad (3.5)$$

holds for any $k' \neq k$, $1 \leq l \leq l_{k'}$.

Define

$$\boldsymbol{\omega}_1 := \mathbf{y}_1(T_1) \quad \text{and} \quad \boldsymbol{\omega}_2 = \mathbf{y}_2(T_1).$$

For any $t \geq 0$ and $\varepsilon, \delta > 0$, we further define

$$\begin{aligned} \boldsymbol{\omega}_1(t) &:= \mathbf{y}_1(T_1 + t) = \boldsymbol{\omega}_1 + t\mathbf{u}, \\ \boldsymbol{\omega}_2(t) &:= \mathbf{y}_2(T_1 + t) = \boldsymbol{\omega}_2 + t\mathbf{u}, \end{aligned}$$

and

$$\begin{aligned} \mathbf{z}_1(t, \varepsilon) &:= \mathbf{y}_1(T_1 + t) + \varepsilon t \mathbf{a}_k = \boldsymbol{\omega}_1(t) + \varepsilon t \mathbf{a}_k = \boldsymbol{\omega}_1 + t\mathbf{u} + \varepsilon t \mathbf{a}_k, \\ \mathbf{z}_2(t, \varepsilon) &:= \mathbf{y}_2(T_1 + t) - \varepsilon t \mathbf{a}_k = \boldsymbol{\omega}_2(t) - \varepsilon t \mathbf{a}_k = \boldsymbol{\omega}_2 + t\mathbf{u} - \varepsilon t \mathbf{a}_k, \end{aligned}$$

as well as

$$\tilde{\mathbf{z}}_1(t, \delta) := \mathbf{y}_1(T_1 + t) - \delta \mathbf{a}_k = \boldsymbol{\omega}_1(t) - \delta \mathbf{a}_k = \boldsymbol{\omega}_1 + t\mathbf{u} - \delta \mathbf{a}_k,$$

$$\tilde{\mathbf{z}}_2(t, \delta) := \mathbf{y}_2(T_1 + t) + \delta \mathbf{a}_k = \boldsymbol{\omega}_2(t) + \delta \mathbf{a}_k = \boldsymbol{\omega}_2 + t\mathbf{u} + \delta \mathbf{a}_k.$$

It is easy to check that $\mathbf{z}_1(t, \varepsilon)$, $\boldsymbol{\omega}_1(t)$, $\boldsymbol{\omega}_2(t)$ and $\mathbf{z}_1(t, \varepsilon)$ lie on a line sequentially, and the Euclidean projection of $\mathbf{z}_1(t, \varepsilon)$ (resp. $\mathbf{z}_2(t, \varepsilon)$) onto $H_{k,1}$ (resp. H_{k,l_k}) is $\boldsymbol{\omega}_1(t)$ (resp. $\boldsymbol{\omega}_2(t)$). Fix any constant $\varepsilon_0 \leq c_1$. For any $t \geq 0$ and any $0 < \varepsilon \leq \varepsilon_0$, we can show that the line segment between $\boldsymbol{\omega}_1(t)$ (resp. $\boldsymbol{\omega}_2(t)$) and $\mathbf{z}_1(t, \varepsilon)$ (resp. $\mathbf{z}_2(t, \varepsilon)$) does not cross any separating hyperplane $H_{k,l}$ (except that its endpoint $\boldsymbol{\omega}_1(t)$ (resp. $\boldsymbol{\omega}_2(t)$) is on $H_{k,1}$ (resp. H_{k,l_k}). To see why this is true, for $s = 1, 2$, we have $\text{dist}(\mathbf{z}_s(t, \varepsilon), \boldsymbol{\omega}_s(t)) = \varepsilon t$, and from (3.5) we know that

$$\text{dist}(\boldsymbol{\omega}_s(t), H_{k',l}) = \text{dist}(\boldsymbol{\omega}_s(T_1 + t), H_{k',l}) \geq 2c_1(T_1 + t)$$

for any $k' \neq k$ and $1 \leq l \leq l_{k'}$. Therefore the two line segments does not cross any $H_{k',l}$ for $k' \neq k$ and $1 \leq l \leq l_{k'}$, as long as $\varepsilon \leq \varepsilon_0 \leq c_1$. Recall that we assume $b_{k,1} > \dots > b_{k,l_k}$, and as a result the two line segments does not cross any $H_{k,l}$ for $1 < l < l_k$. This gives two immediate consequences:

1. For any $t \geq 0$ and any $0 < \varepsilon \leq \varepsilon_0$ we have

$$\text{dist}(\mathbf{z}_s(t, \varepsilon), H_{k',l}) \geq \text{dist}(\boldsymbol{\omega}_s(t), H_{k',l}) - \text{dist}(\boldsymbol{\omega}_s(t), \mathbf{z}_s(t, \varepsilon)) \geq 2c_1T + c_1t$$

for any $k' \neq k$ and $1 \leq l \leq l_{k'}$, and

$$\text{dist}(\mathbf{z}_s(t, \varepsilon), H_{k,l}) \geq \varepsilon t$$

for any $1 \leq l \leq l_k$. Taking the above two inequalities collectively yields

$$\min_{1 \leq k \leq K, 1 \leq l \leq l_k} \text{dist}(\mathbf{z}_s(t, \varepsilon), H_{k,l}) \geq \min\{c_1, \varepsilon\}t = \varepsilon t \quad (3.6)$$

as long as $\varepsilon \leq \varepsilon_0 \leq c_1$.

2. As an immediate consequence of (3.6), for $s = 1, 2$, the set $\{\mathbf{z}_1(t, \varepsilon) : t \geq 0, 0 < \varepsilon \leq \varepsilon_0\}$ (resp. $\{\mathbf{z}_2(t, \varepsilon) : t \geq 0, 0 < \varepsilon \leq \varepsilon_0\}$) is a subset of an element S_1 (resp. S_2) from the knowledge partition. This is because the only intersection of this set with the separating hyperplanes $\{H_{k,l} : 1 \leq k \leq K, 1 \leq l \leq l_k\}$ is $\mathbf{z}_1(0, 0) = \boldsymbol{\omega}_1$ (resp. $\mathbf{z}_2(0, 0) = \boldsymbol{\omega}_2$).

By similar arguments, we can also show that there exists a constant

$$\delta_0 := \begin{cases} \frac{1}{2} \min \left\{ \frac{b_{k,1} - b_{k,2}}{\|\mathbf{a}_k\|_2}, \frac{b_{k,l_k-1} - b_{k,l_k}}{\|\mathbf{a}_k\|_2}, c_1T \right\}, & \text{if } l_k \geq 2, \\ \frac{1}{2}cT, & \text{if } l_k = 1, \end{cases}$$

such that for $s = 1, 2$, the set $\{\tilde{\mathbf{z}}_1(t, \delta) : t \geq 0, 0 < \delta \leq \delta_0\}$ (resp. $\{\tilde{\mathbf{z}}_2(t, \delta) : t \geq 0, 0 < \delta \leq \delta_0\}$) is a subset of an element S'_1 (resp. S'_2) from the linear knowledge partition.

Fix any $0 < \varepsilon \leq \varepsilon_0$. For any $t \geq 0$, when the ground truth is $\mathbf{z}_1(t, \varepsilon)$, let the noisy score be \mathbf{X}_t . By the truth telling assumption we know that for any nondecreasing convex function $U : \mathbb{R} \rightarrow \mathbb{R}$ we have

$$\mathbb{E}[U(\mathcal{P}_{S_1}(\mathbf{z}_1(t, \varepsilon) + \mathbf{e}))] \geq \mathbb{E}[U(\boldsymbol{\omega}_1(t) + \mathcal{P}_{H_{k,1}}(\mathbf{e}))], \quad (3.7)$$

where $\mathbf{e} = \mathbf{X}_t - \mathbf{z}_1(t, \varepsilon)$ is the noise, and $\mathcal{P}_{S_1}(\mathbf{z}_1(t, \varepsilon) + \mathbf{e})$ (resp. $\boldsymbol{\omega}_1(t) + \mathcal{P}_{H_{k,1}}(\mathbf{e})$) is the Euclidean projection of the noisy score \mathbf{X}_t onto S_1 (resp. S'_1). Then we claim that this leads to $\mathbf{u} + \varepsilon \mathbf{a}_k \succeq_w \mathbf{u}$. Actually if this does not hold, then Lemma 2 asserts the existence of a nondecreasing continuous convex function $h : \mathbb{R} \rightarrow \mathbb{R}$ such that

$$h(\mathbf{u} + \varepsilon \mathbf{a}_k) < h(\mathbf{u}). \quad (3.8)$$

Let $U(x) = h(tx)$, which is also a nondecreasing convex function. Then (3.7) gives

$$\mathbb{E} \left[h \left(\frac{\mathcal{P}_{S_1}(\mathbf{z}_1(t, \varepsilon) + \mathbf{e})}{t} \right) \right] \geq \mathbb{E} \left[h \left(\frac{\boldsymbol{\omega}_1(t) + \mathcal{P}_{H_{k,1}}(\mathbf{e})}{t} \right) \right]. \quad (3.9)$$

In view of the assumption (1.9), we know that

$$\frac{\mathbf{e}}{t} = \frac{\mathbf{X}_t - \mathbf{z}_1(t, \varepsilon)}{t} \xrightarrow{\text{a.s.}} \mathbf{0}$$

We also learn from (3.6) that the distance between $\mathbf{z}_1(t, \varepsilon)$ and the boundary of S_1 (which must be one of the separating hyperplanes) is at least εt , therefore as $t \rightarrow \infty$ the sequence of random vector

$$\frac{(\mathcal{I} - \mathcal{P}_{S_1})(\mathbf{z}_1(t, \varepsilon) + \mathbf{e})}{t} \xrightarrow{\text{a.s.}} \mathbf{0},$$

where \mathcal{I} is the identity operator. This further gives

$$\begin{aligned} \frac{\mathcal{P}_{S_1}(\mathbf{z}_1(t, \varepsilon) + \mathbf{e})}{t} &= \frac{\mathbf{z}_1(t, \varepsilon) + \mathbf{e}}{t} - \frac{(\mathcal{I} - \mathcal{P}_{S_1})(\mathbf{z}_1(t, \varepsilon) + \mathbf{e})}{t} \\ &= \frac{\boldsymbol{\omega}_1 + t\mathbf{u} + \varepsilon t\mathbf{a}_k + \mathbf{e}}{t} - \frac{(\mathcal{I} - \mathcal{P}_{S_1})(\mathbf{z}_1(t, \varepsilon) + \mathbf{e})}{t} \\ &\xrightarrow{\text{a.s.}} \mathbf{u} + \varepsilon\mathbf{a}_k \end{aligned}$$

as $t \rightarrow \infty$. It is also straightforward to check that

$$\frac{\boldsymbol{\omega}_1(t) + \mathcal{P}_{H_{k,1}}(\mathbf{e})}{t} = \frac{\boldsymbol{\omega}_1 + t\mathbf{u} + \mathcal{P}_{H_{k,1}}(\mathbf{e})}{t} \xrightarrow{\text{a.s.}} \mathbf{u}.$$

as $t \rightarrow \infty$. Since $h(\cdot)$ is continuous, by the continuous mapping theorem, we know that

$$\mathbb{E} \left[h \left(\frac{\mathbf{z}_1(t, \varepsilon) + \mathbf{e}}{t} \right) \right] \rightarrow h(\mathbf{u} + \varepsilon\mathbf{a}_k) \quad \text{and} \quad \mathbb{E} \left[h \left(\frac{\boldsymbol{\omega}_1(t) + \mathcal{P}_{H_{k,1}}(\mathbf{e})}{t} \right) \right] \rightarrow h(\mathbf{u}),$$

as $t \rightarrow \infty$. Since (3.9) holds for any t , we can take the limit $t \rightarrow \infty$ to achieve $h(\mathbf{u} + \varepsilon\mathbf{a}_k) \geq h(\mathbf{u})$, which contradicts with (3.8). This shows that $\mathbf{u} + \varepsilon\mathbf{a}_k \succeq_w \mathbf{u}$. Similarly we can also show that $\mathbf{u} - \varepsilon\mathbf{a}_k \succeq_w \mathbf{u}$. Since each element of \mathbf{u} is different, and $\mathbf{u} \pm \varepsilon\mathbf{a}_k \succeq_w \mathbf{u}$ both hold for any $\varepsilon \in (0, \varepsilon_0]$, we can apply (Su, 2022, Lemma 6.3) to show that $\varepsilon = 0$, which is a contradiction. This means that for any given $1 \leq k \leq K$, there must exist $1 \leq i < j \leq n$ such that \mathbf{a}_k is a constant multiple of the vector $\mathbf{e}_i - \mathbf{e}_j$, which finishes Step 1 as well as the proof of Theorem 3.

3.4 Step 2: determining the intercepts $\{b_{k,l}\}_{1 \leq k \leq K, 1 \leq l \leq l_k}$ (Proof of Proposition 1)

Based on what we have shown in Step 1, we can further assume, without loss of generality, that for any $1 \leq k \leq K$

$$\mathbf{a}_k = \frac{1}{\sqrt{2}}(\mathbf{e}_{i_k} - \mathbf{e}_{j_k})$$

for some $1 \leq i_k < j_k \leq n$. We also require that $(i_k, j_k) \neq (i_{k'}, j_{k'})$ for $k \neq k'$. In Step 2 we aim to show that for any $1 \leq k \leq K$, we must have $l_k = 1$. We prove this by contradiction. Suppose there exists $1 \leq k \leq K$ such that $l_k \geq 2$, and assume without loss of generality that $b_{k,1} > 0$ (the case when $l_k = 1$ is much easier, and the proof can be obtained by taking $b_{k,2} = -\infty$ in the following proof).

For any $t \geq 0$, let \mathbf{v} be a unit vector satisfying

$$\begin{cases} \mathbf{a}_k^\top \mathbf{v} = 0, \\ \mathbf{a}_{k'}^\top \mathbf{v} \neq 0, & \text{for all } k' \neq k, \\ v_i > 0, & \text{for all } 1 \leq i \leq n. \end{cases} \quad (3.10)$$

Fix any $\mathbf{y}(0) \in H_{k,1}$, and define

$$\mathbf{y}(t) = \mathbf{y}(0) + t\mathbf{v}, \quad \forall t \geq 0.$$

Note that for any $1 \leq l \leq l_k$ and any $t \geq 0$, we have

$$\text{dist}(\mathbf{y}(t), H_{k,l}) = \frac{b_{k,1} - b_{k,l}}{\|\mathbf{a}_k\|_2} = b_{k,1} - b_{k,l},$$

Akin to Step 1, we can check that $\mathbf{y}_t \in H_{k,1}$ and there exists some sufficiently large constant $T > 0$ and some sufficiently small constant $c_2 > 0$, such that for any $t \geq T$, we have

$$\min_{k' \neq k, 1 \leq l \leq l_{k'}} \text{dist}(\mathbf{y}(t), H_{k',l}) \geq 4c_2t, \quad (3.11)$$

and

$$\max_{1 \leq l \leq l_k} \text{dist}(\mathbf{y}(t), H_{k,l}) \leq c_2t, \quad (3.12)$$

and

$$\min_{1 \leq i \leq n} y_i(t) \geq 4c_2t. \quad (3.13)$$

Define a Euclidean ball

$$\mathbf{B}(t) := \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x} - \mathbf{y}(t)\|_2 \leq 2c_2t\},$$

and its three subsets separated by $H_{k,1}$ and $H_{k,2}$:

$$\begin{aligned} \mathbf{B}_1(t) &:= \mathbf{B}(t) \cap \left\{ \mathbf{x} : \mathbf{a}_k^\top \mathbf{x} \geq b_{k,1} \right\}, \\ \mathbf{B}_2(t) &:= \mathbf{B}(t) \cap \left\{ \mathbf{x} : b_{k,1} \geq \mathbf{a}_k^\top \mathbf{x} \geq b_{k,2} \right\}, \\ \mathbf{B}_3(t) &:= \mathbf{B}(t) \cap \left\{ \mathbf{x} : \mathbf{a}_k^\top \mathbf{x} \leq b_{k,2} \right\}. \end{aligned}$$

It is straightforward to check that for all $t \geq T$, $\mathbf{B}_1(t)$ (resp. $\mathbf{B}_2(t)$) lives in the same element, denoted by S_1 (resp. S_2), of the linear knowledge partition. For any $t \geq T$, when the ground truth is $\mathbf{y}(t)$, by the truth-telling assumption we know that for any nondecreasing convex function $U : \mathbb{R} \rightarrow \mathbb{R}$ we have

$$\mathbb{E}[U(\mathcal{P}_{S_1}(\mathbf{y}(t) + \mathbf{e}))] = \mathbb{E}[U(\mathcal{P}_{S_2}(\mathbf{y}(t) + \mathbf{e}))] \quad (3.14)$$

with $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_n)$, because $\mathbf{y}(t) \in S_1$ and $\mathbf{y}(t) \in S_2$ hold simultaneously. We take a test function of the form $U(x) = \max\{x^2, 0\}$, which is nondecreasing and convex. We first decompose both sides of (3.14) into

$$\begin{aligned} \mathbb{E}[U(\mathcal{P}_{S_1}(\mathbf{y}(t) + \mathbf{e}))] &= \underbrace{\mathbb{E}[U(\mathcal{P}_{S_1}(\mathbf{y}(t) + \mathbf{e})) \mathbb{1}\{\mathbf{y}(t) + \mathbf{e} \in \mathbf{B}_1(t)\}]}_{=: \alpha_1(t)} \\ &\quad + \underbrace{\mathbb{E}[U(\mathcal{P}_{S_1}(\mathbf{y}(t) + \mathbf{e})) \mathbb{1}\{\mathbf{y}(t) + \mathbf{e} \in \mathbf{B}_2(t) \cup \mathbf{B}_3(t)\}]}_{=: \alpha_2(t)} \end{aligned}$$

$$+ \underbrace{\mathbb{E}[U(\mathcal{P}_{S_1}(\mathbf{y}(t) + \mathbf{e})) \mathbf{1}\{\mathbf{y}(t) + \mathbf{e} \notin \mathbf{B}(t)\}]}_{=:\alpha_3(t)}$$

and

$$\begin{aligned} \mathbb{E}[U(\mathcal{P}_{S_2}(\mathbf{y}(t) + \mathbf{e}))] &= \underbrace{\mathbb{E}[U(\mathcal{P}_{S_2}(\mathbf{y}(t) + \mathbf{e})) \mathbf{1}\{\mathbf{y}(t) + \mathbf{e} \in \mathbf{B}_1(t)\}]}_{=:\beta_1(t)} \\ &+ \underbrace{\mathbb{E}[U(\mathcal{P}_{S_2}(\mathbf{y}(t) + \mathbf{e})) \mathbf{1}\{\mathbf{y}(t) + \mathbf{e} \in \mathbf{B}_2(t)\}]}_{=:\beta_2(t)} \\ &+ \underbrace{\mathbb{E}[U(\mathcal{P}_{S_2}(\mathbf{y}(t) + \mathbf{e})) \mathbf{1}\{\mathbf{y}(t) + \mathbf{e} \in \mathbf{B}_3(t)\}]}_{=:\beta_3(t)} \\ &+ \underbrace{\mathbb{E}[U(\mathcal{P}_{S_2}(\mathbf{y}(t) + \mathbf{e})) \mathbf{1}\{\mathbf{y}(t) + \mathbf{e} \notin \mathbf{B}(t)\}]}_{=:\beta_4(t)}. \end{aligned}$$

Then we analyze $\alpha_1(t)$, $\alpha_2(t)$, $\alpha_3(t)$, $\beta_1(t)$, $\beta_2(t)$, $\beta_3(t)$, and $\beta_4(t)$, respectively.

- In view of (3.11) and (3.12), we know that when $\mathbf{y}(t) + \mathbf{e} \in \mathbf{B}_1(t)$,

$$\mathcal{P}_{S_2}(\mathbf{y}(t) + \mathbf{e}) = \mathcal{P}_{H_{k,1}}(\mathbf{y}(t) + \mathbf{e}) = \mathbf{y}(t) + \mathcal{P}_{H_{k,1}}(\mathbf{e}), \quad (3.15)$$

when $\mathbf{y}(t) + \mathbf{e} \in \mathbf{B}_2(t) \cup \mathbf{B}_3(t)$,

$$\mathcal{P}_{S_1}(\mathbf{y}(t) + \mathbf{e}) = \mathcal{P}_{H_{k,1}}(\mathbf{y}(t) + \mathbf{e}) = \mathbf{y}(t) + \mathcal{P}_{H_{k,1}}(\mathbf{e}), \quad (3.16)$$

By symmetry of $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_n)$, we know that

$$\alpha_2(t) = \beta_1(t). \quad (3.17)$$

- In view of (3.13), we know that for any $\mathbf{x} \in \mathbf{B}(t)$, \mathbf{x} has positive entries and therefore $U(\mathbf{x}) = \|\mathbf{x}\|_2^2$. Let $\phi(\mathbf{x}) = (2\pi)^{-n/2} \exp(-\|\mathbf{x}\|_2^2/2)$ be the probability density function of $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_n)$. Recall that $\mathbf{B}_1(t) \subset S_1$, $\mathbf{B}_2(t) \subset S_2$, we know that

$$\begin{aligned} \alpha_1(t) &= \mathbb{E}[U(\mathbf{y}(t) + \mathbf{e}) \mathbf{1}\{\mathbf{y}(t) + \mathbf{e} \in \mathbf{B}_1(t)\}] = \mathbb{E}[\|\mathbf{y}(t) + \mathbf{e}\|_2^2 \mathbf{1}\{\mathbf{y}(t) + \mathbf{e} \in \mathbf{B}_1(t)\}] \\ &= \int_{\mathbf{a}_k^\top \mathbf{x} \geq 0, \|\mathbf{x}\|_2 \leq 2c_2 t} \|\mathbf{y}(t) + \mathbf{x}\|_2^2 \phi(\mathbf{x}) d\mathbf{x} = \int_{\mathbf{a}_k^\top \mathbf{x} \leq 0, \|\mathbf{x}\|_2 \leq 2c_2 t} \|\mathbf{y}(t) - \mathbf{x}\|_2^2 \phi(\mathbf{x}) d\mathbf{x}, \end{aligned}$$

where the last line holds since $\phi(\cdot)$ is symmetric around $\mathbf{0}$, and

$$\begin{aligned} \beta_2(t) &= \mathbb{E}[U(\mathbf{y}(t) + \mathbf{e}) \mathbf{1}\{\mathbf{y}(t) + \mathbf{e} \in \mathbf{B}_2(t)\}] = \mathbb{E}[\|\mathbf{y}(t) + \mathbf{e}\|_2^2 \mathbf{1}\{\mathbf{y}(t) + \mathbf{e} \in \mathbf{B}_2(t)\}] \\ &= \int_{0 \geq \mathbf{a}_k^\top \mathbf{x} \geq b_{k,2} - b_{k,1}, \|\mathbf{x}\|_2 \leq 2c_2 t} \|\mathbf{y}(t) + \mathbf{x}\|_2^2 \phi(\mathbf{x}) d\mathbf{x} \end{aligned}$$

In view of (3.11) and (3.12), we know that when $\mathbf{y}(t) + \mathbf{e} \in \mathbf{B}_3(t)$,

$$\mathcal{P}_{S_2}(\mathbf{y}(t) + \mathbf{e}) = \mathcal{P}_{H_{k,2}}(\mathbf{y}(t) + \mathbf{e}) = \mathbf{y}(t) + \mathbf{e} - \left\{ \mathbf{a}_k^\top [\mathbf{y}(t) + \mathbf{e}] - b_{k,2} \right\} \mathbf{a}_k = \mathbf{y}(t) + \mathbf{e} - c_t(\mathbf{e}) \mathbf{a}_k,$$

where we define

$$c_t(\mathbf{x}) := \mathbf{a}_k^\top [\mathbf{y}(t) + \mathbf{x}] - b_{k,2} = \mathbf{a}_k^\top \mathbf{x} + b_{k,1} - b_{k,2}.$$

Therefore we know that

$$\begin{aligned}\beta_3(t) &= \mathbb{E} \left[\|\mathbf{y}(t) + \mathbf{e} - c_t(\mathbf{e}) \mathbf{a}_k\|_2^2 \mathbf{1} \{ \mathbf{y}(t) + \mathbf{e} \in \mathbf{B}_3(t) \} \right] \\ &= \int_{\mathbf{a}_k^\top \mathbf{x} \leq b_{k,2} - b_{k,1}, \|\mathbf{x}\|_2 \leq 2c_2 t} \|\mathbf{y}(t) + \mathbf{x} - c_t(\mathbf{x}) \mathbf{a}_k\|_2^2 \phi(\mathbf{x}) \, d\mathbf{x}.\end{aligned}$$

• Then we have

$$\begin{aligned}\alpha_1(t) - \beta_2(t) - \beta_3(t) &= \underbrace{\int_{0 \geq \mathbf{a}_k^\top \mathbf{x} \geq b_{k,2} - b_{k,1}, \|\mathbf{x}\|_2 \leq 2c_2 t} \left[\|\mathbf{y}(t) - \mathbf{x}\|_2^2 - \|\mathbf{y}(t) + \mathbf{x}\|_2^2 \right] \phi(\mathbf{x}) \, d\mathbf{x}}_{=:\gamma_1(t)} \\ &\quad + \underbrace{\int_{\mathbf{a}_k^\top \mathbf{x} \leq b_{k,2} - b_{k,1}, \|\mathbf{x}\|_2 \leq 2c_2 t} \left[\|\mathbf{y}(t) - \mathbf{x}\|_2^2 - \|\mathbf{y}(t) + \mathbf{x} - c_t(\mathbf{x}) \mathbf{a}_k\|_2^2 \right] \phi(\mathbf{x}) \, d\mathbf{x}}_{=:\gamma_2(t)}.\end{aligned}$$

Regarding $\gamma_1(t)$, we have

$$\begin{aligned}\gamma_1(t) &= \int_{0 \geq \mathbf{a}_k^\top \mathbf{x} \geq b_{k,2} - b_{k,1}, \|\mathbf{x}\|_2 \leq 2c_2 t} \left[-4\mathbf{y}(t)^\top \mathbf{x} \right] \phi(\mathbf{x}) \, d\mathbf{x} \\ &= \int_{0 \geq \mathbf{a}_k^\top \mathbf{x} \geq b_{k,2} - b_{k,1}, \|\mathbf{x}\|_2 \leq 2c_2 t} \left[-4\mathbf{y}(0)^\top \mathbf{x} + t\mathbf{v}^\top \mathbf{x} \right] \phi(\mathbf{x}) \, d\mathbf{x} \\ &= \int_{0 \geq \mathbf{a}_k^\top \mathbf{x} \geq b_{k,2} - b_{k,1}, \|\mathbf{x}\|_2 \leq 2c_2 t} \left[-4\mathbf{y}(0)^\top \mathbf{x} \right] \phi(\mathbf{x}) \, d\mathbf{x}\end{aligned}$$

Here the penultimate step follows from $\mathbf{a}_k^\top \mathbf{v} = 0$, which is guaranteed by (3.10). Similarly, we also have

$$\begin{aligned}\gamma_2(t) &= \int_{\mathbf{a}_k^\top \mathbf{x} \leq b_{k,2} - b_{k,1}, \|\mathbf{x}\|_2 \leq 2c_2 t} \left\{ -4\mathbf{y}(t)^\top \mathbf{x} + 2c_t(\mathbf{x}) \mathbf{a}_k^\top [\mathbf{y}(t) + \mathbf{x}] - c_t(\mathbf{x})^2 \right\} \phi(\mathbf{x}) \, d\mathbf{x} \\ &= \int_{\mathbf{a}_k^\top \mathbf{x} \leq b_{k,2} - b_{k,1}, \|\mathbf{x}\|_2 \leq 2c_2 t} \left\{ -4\mathbf{y}(0)^\top \mathbf{x} + 2c_t(\mathbf{x}) \left(\mathbf{a}_k^\top \mathbf{x} + b_{k,1} \right) - c_t(\mathbf{x})^2 \right\} \phi(\mathbf{x}) \, d\mathbf{x} \\ &= \int_{\mathbf{a}_k^\top \mathbf{x} \leq b_{k,2} - b_{k,1}, \|\mathbf{x}\|_2 \leq 2c_2 t} \left[-4\mathbf{y}(0)^\top \mathbf{x} + 2 \left(\mathbf{a}_k^\top \mathbf{x} + b_{k,1} - b_{k,2} \right) b_{k,2} + \left(\mathbf{a}_k^\top \mathbf{x} + b_{k,1} - b_{k,2} \right)^2 \right] \phi(\mathbf{x}) \, d\mathbf{x} \\ &= \int_{\mathbf{a}_k^\top \mathbf{x} \leq b_{k,2} - b_{k,1}, \|\mathbf{x}\|_2 \leq 2c_2 t} \left[-4\mathbf{y}(0)^\top \mathbf{x} + \left(\mathbf{a}_k^\top \mathbf{x} + b_{k,1} \right)^2 - b_{k,2}^2 \right] \phi(\mathbf{x}) \, d\mathbf{x}.\end{aligned}$$

It is straightforward to check that

$$\begin{aligned}\lim_{t \rightarrow \infty} \gamma_1(t) &= \int_{0 \geq \mathbf{a}_k^\top \mathbf{x} \geq b_{k,2} - b_{k,1}} \left[-4\mathbf{y}(0)^\top \mathbf{x} \right] \phi(\mathbf{x}) \, d\mathbf{x} =: \gamma_1, \\ \lim_{t \rightarrow \infty} \gamma_2(t) &= \int_{\mathbf{a}_k^\top \mathbf{x} \leq b_{k,2} - b_{k,1}} \left[-4\mathbf{y}(0)^\top \mathbf{x} + \left(\mathbf{a}_k^\top \mathbf{x} + b_{k,1} \right)^2 - b_{k,2}^2 \right] \phi(\mathbf{x}) \, d\mathbf{x} =: \gamma_2.\end{aligned}$$

Let $Z_1 = \mathbf{a}_k^\top \mathbf{e}$, $Z_2 = \mathbf{v}^\top \mathbf{e}$, and $Z_3 = [(\mathbf{I}_n - \mathbf{a}_k \mathbf{a}_k^\top) \mathbf{y}(0)]^\top \mathbf{e}$. It is clear that Z_1 and Z_3 are independent. Recall that $\mathbf{a}_k^\top \mathbf{v} = 0$, we know that $Z_1, Z_2 \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$. We also know that

$$\mathbf{y}(0)^\top \mathbf{e} = \left[\mathbf{a}_k \mathbf{a}_k^\top \mathbf{y}(0) + \left(\mathbf{I}_n - \mathbf{a}_k \mathbf{a}_k^\top \right) \mathbf{y}(0) \right]^\top \mathbf{e} = b_{k,1} Z_1 + Z_3.$$

Then we know that

$$\begin{aligned}\gamma_1 &= \mathbb{E}[-4(b_{k,1}Z_1 + Z_3) \mathbb{1}\{b_{k,2} - b_{k,1} \leq Z_1 \leq 0\}] \\ &= -4b_{k,1}\mathbb{E}[Z_1 \mathbb{1}\{b_{k,2} - b_{k,1} \leq Z_1 \leq 0\}],\end{aligned}$$

where the last step uses the fact that Z_1 and Z_3 are independent. We also have

$$\begin{aligned}\gamma_2 &= \mathbb{E}\left[\left[-4(b_{k,1}Z_1 + Z_3) + (Z_1 + b_{k,1})^2 - b_{k,2}^2\right] \mathbb{1}\{Z_1 \leq b_{k,2} - b_{k,1}\}\right] \\ &= \mathbb{E}\left[\left[-4b_{k,1}Z_1 + (Z_1 + b_{k,1})^2 - b_{k,2}^2\right] \mathbb{1}\{Z_1 \leq b_{k,2} - b_{k,1}\}\right].\end{aligned}$$

Therefore

$$\begin{aligned}\gamma_1 + \gamma_2 &= -4b_{k,1}\mathbb{E}[Z_1 \mathbb{1}\{Z_1 \leq 0\}] + \mathbb{E}\left[\left[(Z_1 + b_{k,1})^2 - b_{k,2}^2\right] \mathbb{1}\{Z_1 \leq b_{k,2} - b_{k,1}\}\right] \\ &= \frac{4}{\sqrt{2\pi}}b_{k,1} + \mathbb{E}\left[\left[(Z_1 + b_{k,1})^2 - b_{k,2}^2\right] \mathbb{1}\{Z_1 \leq b_{k,2} - b_{k,1}\}\right].\end{aligned}$$

Define a function

$$\begin{aligned}f(x) &:= \mathbb{E}\left[\left[(Z_1 + b_{k,1})^2 - x^2\right] \mathbb{1}\{Z_1 \leq x - b_{k,1}\}\right] \\ &= \int_{-\infty}^{x-b_{k,1}} \left[(z + b_{k,1})^2 - x^2\right] \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{z^2}{2}\right) dz,\end{aligned}$$

and compute its derivative

$$\begin{aligned}f'(x) &= \left[(x - b_{k,1} + b_{k,1})^2 - x^2\right] \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{(x - b_{k,1})^2}{2}\right] - \int_{-\infty}^{x-b_{k,1}} \frac{2x}{\sqrt{2\pi}} \exp\left(-\frac{z^2}{2}\right) dz \\ &= -x \int_{-\infty}^{x-b_{k,1}} \frac{2}{\sqrt{2\pi}} \exp\left(-\frac{z^2}{2}\right) dz.\end{aligned}$$

We can see that $f(x)$ is increasing when $x \leq 0$, and is decreasing when $x > 0$. As a result,

$$\begin{aligned}\inf_{x \leq b_{k,1}} f(x) &= \min\left\{\inf_{x \rightarrow -\infty} f(x), f(b_{k,1})\right\} = \min\left\{0, \int_{-\infty}^0 (z^2 + 2b_{k,1}z) \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{z^2}{2}\right) dz\right\} \\ &= \min\left\{0, \frac{1}{2} - \frac{2}{\sqrt{2\pi}}b_{k,1}\right\}.\end{aligned}$$

Therefore we reach

$$\begin{aligned}\gamma_1 + \gamma_2 &= \frac{4}{\sqrt{2\pi}}b_{k,1} + f(b_{k,2}) \geq \frac{4}{\sqrt{2\pi}}b_{k,1} + \inf_{x \leq b_{k,1}} f(x) = \frac{4}{\sqrt{2\pi}}b_{k,1} + \min\left\{0, \frac{1}{2} - \frac{2}{\sqrt{2\pi}}b_{k,1}\right\} \\ &= \min\left\{\frac{4}{\sqrt{2\pi}}b_{k,1}, \frac{1}{2} + \frac{2}{\sqrt{2\pi}}b_{k,1}\right\} > 0,\end{aligned}$$

namely we have

$$\lim_{t \rightarrow \infty} \alpha_1(t) - \beta_2(t) - \beta_3(t) = \gamma_1 + \gamma_2 > 0. \quad (3.18)$$

- Lastly, we will analyze $\alpha_3(t)$ and $\beta_4(t)$, which is relatively straightforward. We have

$$\alpha_3(t) = \mathbb{E}[U(\mathcal{P}_{S_1}(\mathbf{y}(t) + \mathbf{e})) \mathbb{1}\{\mathbf{y}(t) + \mathbf{e} \notin \mathbf{B}(t)\}]$$

$$\begin{aligned}
&= \mathbb{E} [U(\mathcal{P}_{S_1}(\mathbf{y}(t) + \mathbf{e})) \mathbf{1}\{\|\mathbf{e}\|_2 \geq 2c_2t\}] \\
&\leq \mathbb{E} \left[\|\mathcal{P}_{S_1}(\mathbf{y}(t) + \mathbf{e})\|_2^2 \mathbf{1}\{\|\mathbf{e}\|_2 \geq 2c_2t\} \right],
\end{aligned}$$

where the last step follows from $U(x) = \max\{0, x^2\} \leq x^2$. By the non-expansiveness of Euclidean projection, we have

$$\|\mathcal{P}_{S_1}(\mathbf{y}(t) + \mathbf{e}) - \mathcal{P}_{S_1}(\mathbf{0})\|_2 \leq \|\mathbf{y}(t) + \mathbf{e} - \mathbf{0}\|_2 = \|\mathbf{y}(t) + \mathbf{e}\|_2.$$

Therefore we know that

$$0 \leq \alpha_3(t) \leq \mathbb{E} [(\|\mathbf{y}(t) + \mathbf{e}\|_2 + \|\mathcal{P}_{S_1}(\mathbf{0})\|_2) \mathbf{1}\{\|\mathbf{e}\|_2 \geq 2c_2t\}].$$

By taking $t \rightarrow \infty$, we immediately have

$$\lim_{t \rightarrow \infty} \alpha_3(t) = 0. \quad (3.19)$$

Similarly we can also show that

$$\lim_{t \rightarrow \infty} \beta_4(t) = 0. \quad (3.20)$$

From (3.14) we know that

$$\alpha_1(t) + \alpha_2(t) + \alpha_3(t) - \beta_1(t) - \beta_2(t) - \beta_3(t) - \beta_4(t) = 0$$

for all $t \geq T$. But from (3.17), (3.18), (3.19) and (3.20) we know that

$$\lim_{t \rightarrow \infty} \alpha_1(t) + \alpha_2(t) + \alpha_3(t) - \beta_1(t) - \beta_2(t) - \beta_3(t) - \beta_4(t) > 0,$$

which is a contradiction. Therefore we have shown that $l_k = 1$ and $b_{k,1} = 0$ for all $1 \leq k \leq K$, which finishes Step 2 and the proof of Proposition 1.

4 Minimax optimality

We now present a formal version of Theorem 4. Recall that for $i = 1, \dots, n$, the independent review scores $X_i \sim p_{\theta_i^*}(x)$, where $p_{\theta}(x) = e^{\theta x - b(\theta)} c(x)$. We have $\mu_i^* = \mathbb{E}[X_i] = b'(\theta_i^*)$ and $\text{Var}(X_i) = b''(\theta_i^*)$. We assume the true scores lie within a closed interval $[V_{\min}, V_{\max}]$, where V_{\max} and V_{\min} can be problem-specific, and define the parameter space

$$\mathcal{I}_n(V_{\min}, V_{\max}) := \left\{ \boldsymbol{\mu}^* \in \mathbb{R}^n : V_{\max} \geq \mu_{\pi^*(1)}^* \geq \dots \geq \mu_{\pi^*(n)}^* \geq V_{\min} \right\}.$$

Without loss of generality, we consider only positive scores, i.e., $V_{\min} \geq 0$. As before, we assume that $b''(\theta) > 0$ on its domain, which essentially requires that the noisy scores have nonzero variance. Then, we have $\theta_{\max} \geq \theta_{\pi^*(1)}^* \geq \dots \geq \theta_{\pi^*(n)}^* \geq \theta_{\min}$, where $\theta_{\max} = (b')^{-1}(V_{\max})$ and $\theta_{\min} = (b')^{-1}(V_{\min})$. Let

$$\sigma^2 := \max_{\theta_{\min} \leq \theta \leq \theta_{\max}} b''(\theta)$$

represent the largest variance of any random variable $X \sim p_{\theta}(x)$ for any $\theta \in [\theta_{\min}, \theta_{\max}]$. We assume that $V_{\max} - V_{\min} \geq \sigma/\sqrt{n}$ to ensure the estimation problem is non-trivial. We proceed by assuming the following regularity condition.

Assumption 3. *There exist two constants $C_{\text{var}}, C_{\text{int}} > 0$ such that, there exists $[\tilde{V}_{\min}, \tilde{V}_{\max}] \subseteq [V_{\min}, V_{\max}]$ satisfying $\tilde{V}_{\max} - \tilde{V}_{\min} \geq C_{\text{int}}(V_{\max} - V_{\min})$, such that*

$$\text{Var}_{X \sim p_{\theta}(x)}(X) = b''(\theta) \geq C_{\text{var}}\sigma^2 \quad \text{where } \theta = (b')^{-1}(\mu)$$

for any $\mu \in [\tilde{V}_{\min}, \tilde{V}_{\max}]$.

Assumption 3 is mild, as it essentially requires that the variance of the noisy score not change dramatically when the true score lies within a non-negligible part of the $[V_{\min}, V_{\max}]$ interval. After presenting the following minimax-optimal estimation guarantees, we will verify that it is satisfied by common exponential family distributions.

Theorem 5 (Formal version of Theorem 4). *Suppose that Assumption 3 holds. Then there exist constants $C_1, C_2 > 0$ depending on C_{var} and C_{int} such that*

$$\sup_{V_{\min} \leq \mu^* \leq V_{\max}} \mathbb{E} \|\hat{\mu} - \mu^*\|_2^2 \leq C_1 n \sigma^2 \min \left\{ 1, \left(\frac{V_{\max} - V_{\min}}{n\sigma} \right)^{2/3} + \frac{\log n}{n} \right\}, \quad (4.1)$$

where $\hat{\mu}$ is the optimal solution to (1.5), and

$$\inf_{\tilde{\mu}} \sup_{V_{\min} \leq \mu^* \leq V_{\max}} \mathbb{E} \|\tilde{\mu} - \mu^*\|_2^2 \geq C_2 n \sigma^2 \min \left\{ 1, \left(\frac{V_{\max} - V_{\min}}{n\sigma} \right)^{2/3} + \frac{1}{n} \right\}, \quad (4.2)$$

where the infimum is taken over all estimators $\tilde{\mu}$ for the true mean vector μ^* based on observed data X_1, \dots, X_n .

Before delving into the proof, let us first discuss the implication of Theorem 5 when applied to the following commonly used exponential family distributions.

1. Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$ with fixed variance σ^2 is an exponential family distribution with $\theta = \mu/\sigma^2$ and $b(\theta) = \sigma^2\theta^2/2$. It can be verified that Assumption 3 holds with $\tilde{V}_{\max} = V_{\max}$, $\tilde{V}_{\min} = V_{\min}$, $C_{\text{int}} = 1$ and $C_{\text{var}} = 1$. In comparison to the standard MLE (which simply uses X_i to estimate μ_i^* for each $i \in \{1, \dots, n\}$), with an estimation error of $\mathbb{E}\|\mathbf{X} - \mu^*\|^2 = n\sigma^2$, the estimation error of isotonic regression is

$$O \left(\min \left\{ 1, \left(\frac{V_{\max} - V_{\min}}{n\sigma} \right)^{2/3} + \frac{\log n}{n} \right\} \right)$$

times smaller, which is significantly better when $V_{\max} - V_{\min} \ll n\sigma$.

2. Binomial distribution $\text{Binom}(m, p)$ with success probability p and fixed number of trials m is an exponential family distribution with $\theta = \log[p/(1-p)]$, $b(\theta) = m \log[1 + \exp(\theta)]$ and $\mu = mp$. Since the mean is upper bounded by m in this setting, we consider $V_{\min} = 0$ and $V_{\max} = m$. Then $\sigma^2 = m/4$ and it can be confirmed that Assumption 3 holds with, e.g. $\tilde{V}_{\max} = 3m/4$, $\tilde{V}_{\min} = m/4$, $C_{\text{int}} = 1/2$ and $C_{\text{var}} = 3/4$. We can check that for binomial distribution, both the estimation error bound (4.1) and the minimax lower bound (4.2) scales as $\min\{mn, m^{4/3}n^{1/3}\}$, suggesting that isotonic regression is minimax-optimal (no need to exclude logarithmic factors). The estimation error of the Isotonic Mechanism is $O(\min\{1, (m/n^2)^{1/3}\})$ times smaller than the standard MLE (which scales as mn), which is much better when $n \gg \sqrt{m}$.

3. The Poisson distribution $\text{Poisson}(\lambda)$ with rate $\lambda \in \mathbb{R}^+$ is an exponential family distribution characterized by $\theta = \log \lambda$, $b(\theta) = \exp(\theta)$, and $\mu = \lambda$. We can verify that $\sigma^2 = V_{\max}$, and Assumption 3 is satisfied with $\tilde{V}_{\max} = V_{\max}$, $\tilde{V}_{\min} = \max\{V_{\min}, V_{\max}/2\}$, $C_{\text{int}} = 1/2$, and $C_{\text{var}} = 1/2$. The estimation error for isotonic regression is

$$O\left(\min\left\{1, \left(\frac{V_{\max} - V_{\min}}{n\sqrt{V_{\max}}}\right)^{2/3} + \frac{\log n}{n}\right\}\right)$$

times smaller than the estimation error of the standard MLE, particularly when $V_{\max} - V_{\min} \ll n\sqrt{V_{\max}}$.

4. The Gamma distribution $\text{Gamma}(m, \beta)$ with probability density function

$$f(x) = \frac{1}{\Gamma(m)\beta^m} x^{m-1} \exp\left(-\frac{x}{\beta}\right), \quad \forall x \in \mathbb{R}^+$$

is an exponential family distribution with $\theta = -1/\beta$, $b(\theta) = -m \log(-\theta)$ and $\mu = m\beta$. We can check that $\sigma^2 = V_{\max}^2/m$, and Assumption 3 holds with $\tilde{V}_{\max} = V_{\max}$, $\tilde{V}_{\min} = \max\{V_{\min}, V_{\max}/2\}$, $C_{\text{int}} = 1/2$ and $C_{\text{var}} = 1/4$. When $V_{\max} - V_{\min} \ll nV_{\max}$, the estimation error of isotonic regression is at least $O(\min\{1, (m/n^2)^{1/3}\})$ times smaller than the standard MLE, which is significantly better when $n \gg \sqrt{m}$.

For the remainder of this section, we provide the proof of Theorem 5.

4.1 Establishing the upper bound (4.1)

The upper bound follows from Zhang (2002, Theorem 2.2). We first identify the quantities in Zhang (2002) as

$$y_i = X_i, \quad \varepsilon_i = X_i - \mathbb{E}[X_i] = X_i - \mu_i^*,$$

and it is straightforward to check that $V = V_{\max} - V_{\min}$,

$$\mathbb{E}[\varepsilon_i] = 0, \quad \mathbb{E}[\varepsilon_i^2] = \text{var}(X_i) \leq \sigma^2.$$

Then we can apply Zhang (2002, Theorem 2.2) to achieve

$$R_n \lesssim \sigma \min\left\{1, \left[\left(\frac{V_{\max} - V_{\min}}{n\sigma}\right)^{2/3} + \frac{\log n}{n}\right]^{1/2}\right\}$$

and as a result

$$\mathbb{E}\|\boldsymbol{\mu} - \boldsymbol{\mu}^*\|^2 = nR_n^2(f) \leq C_1 n\sigma^2 \min\left\{1, \left(\frac{V_{\max} - V_{\min}}{n\sigma}\right)^{2/3} + \frac{\log n}{n}\right\}$$

for some universal constant $C_1 > 0$, which finishes the proof for the upper bound (4.1).

4.2 Establishing the lower bound (4.2)

We follow the proof idea of Bellec and Tsybakov (2015, Corollary 5), which focuses on the case when the observations are Gaussian, to establish the minimax lower bound for general exponential family distributions. It suffices to consider the true scores in the interval $[\tilde{V}_{\min}, \tilde{V}_{\max}]$, which is a subset of $[V_{\min}, V_{\max}]$. This motivates one to define

$$\tilde{V} := \tilde{V}_{\max} - \tilde{V}_{\min} \geq C_{\text{int}}(V_{\max} - V_{\min}).$$

We look at the following two cases separately.

Case 1: when n is not too small. We first consider the case when

$$\left(\frac{\tilde{V}}{n\sigma}\right)^{2/3} \geq \frac{4}{n}. \quad (4.3)$$

We would like to show that

$$\inf_{\hat{\boldsymbol{\mu}}} \sup_{\boldsymbol{\mu}^* \in \mathcal{I}_n(V_{\min}, V_{\max})} \mathbb{E} \|\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}^*\|^2 \geq C_2 n \sigma^2 \min \left\{ 1, \left(\frac{V_{\max} - V_{\min}}{n\sigma} \right)^{2/3} \right\} \quad (4.4)$$

for some universal constant $C_2 > 0$. Let

$$k := \min \left\{ \left\lfloor \left(\frac{n\tilde{V}^2}{c^2\sigma^2} \right)^{1/3} \right\rfloor, n \right\} \quad (4.5)$$

for some constant $c = C_{\text{var}}/16$. We know from (4.3) that $k \geq 8$ as long as $n \geq 8$. Without loss of generality, we assume that n is a multiple of k for simplicity; the proof for the general case is similar. In view of the Varshamov-Gilbert bound (Tsybakov, 2008, Lemma 2.9), there exists $\Omega := \{\boldsymbol{\omega}^{(0)}, \dots, \boldsymbol{\omega}^{(M)}\} \subset \{0, 1\}^k$ satisfying $M \geq 2^{k/8}$, $\boldsymbol{\omega}^{(0)} = \mathbf{0}$ and

$$d_{\text{H}}(\boldsymbol{\omega}^{(i)}, \boldsymbol{\omega}^{(j)}) \geq \frac{k}{8}, \quad \forall 0 \leq i < j \leq M, \quad (4.6)$$

where $d_{\text{H}}(\mathbf{x}, \mathbf{y})$ is the Hamming distance $\sum_{i=1}^n \mathbf{1}\{x_i \neq y_i\}$ between two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$. Let

$$\gamma = c \sqrt{\frac{\sigma^2 k}{n}}. \quad (4.7)$$

Then for each $\boldsymbol{\omega} \in \Omega$, we define $\boldsymbol{\mu}^{\boldsymbol{\omega}} \in \mathbb{R}^n$ as

$$\mu_i^{\boldsymbol{\omega}} = \tilde{V}_{\min} + \left\lfloor \frac{i-1}{n} k \right\rfloor \frac{1}{k} \tilde{V} + \gamma \omega_{\lfloor \frac{i-1}{n} k \rfloor + 1}$$

for $1 \leq i \leq n$. Note that in view of (4.5), we have $\gamma \leq \tilde{V}/k$, and it is straightforward to check that

$$\tilde{V}_{\min} \leq \mu_1^{\boldsymbol{\omega}} \leq \mu_i^{\boldsymbol{\omega}} \leq \mu_n^{\boldsymbol{\omega}} \leq \tilde{V}_{\max}$$

holds for any $1 \leq i \leq n$. Therefore for any $\boldsymbol{\omega} \in \Omega$ we have $\boldsymbol{\mu}^{\boldsymbol{\omega}} \in \mathcal{I}_n(V_{\min}, V_{\max})$. In addition, recall that n is a multiple of k , therefore for any $\boldsymbol{\omega}, \boldsymbol{\omega}' \in \Omega$

$$\begin{aligned} \|\boldsymbol{\mu}^{\boldsymbol{\omega}} - \boldsymbol{\mu}^{\boldsymbol{\omega}'}\|^2 &= \gamma^2 \sum_{i=1}^n (\omega_{\lfloor \frac{i-1}{n} k \rfloor + 1} - \omega'_{\lfloor \frac{i-1}{n} k \rfloor + 1})^2 = \frac{\gamma^2 n}{k} d_{\text{H}}(\boldsymbol{\omega}, \boldsymbol{\omega}') \stackrel{(i)}{\geq} \frac{\gamma^2 n}{8} = \frac{c^2}{8} \sigma^2 k \\ &= \frac{c^2}{8} \sigma^2 \min \left\{ \left\lfloor \left(\frac{n\tilde{V}^2}{c^2\sigma^2} \right)^{1/3} \right\rfloor, n \right\} \\ &\stackrel{(ii)}{\geq} \frac{c^2}{8} C_{\text{int}}^{2/3} \sigma^2 \min \left\{ \left\lfloor \left(\frac{n(V_{\max} - V_{\min})^2}{c^2\sigma^2} \right)^{1/3} \right\rfloor, n \right\} \end{aligned} \quad (4.8)$$

as long as $\boldsymbol{\omega} \neq \boldsymbol{\omega}'$. Here (i) follows from (4.6), and (ii) holds since $\tilde{V} \geq C_{\text{int}} V$.

Consider any $\mu_1, \mu_2 \in [\tilde{V}_{\min}, \tilde{V}_{\max}]$, and let $\theta_1 = (b')^{-1}(\mu_1)$ and $\theta_2 = (b')^{-1}(\mu_2)$. Then the Kullback-Leibler (KL) divergence between two distributions p_{θ_1} and p_{θ_2} satisfies

$$\begin{aligned} \text{KL}(p_{\theta_1} \parallel p_{\theta_2}) &= \int \log \frac{p_{\theta_1}(x)}{p_{\theta_2}(x)} p_{\theta_1}(x) dx = \int [(\theta_1 - \theta_2)x - b(\theta_1) + b(\theta_2)] p_{\theta_1}(x) dx \\ &= (\theta_1 - \theta_2) b'(\theta_1) - b(\theta_1) + b(\theta_2) \leq \frac{1}{2} \max_{\theta_{\min} \leq \theta \leq \theta_{\max}} b''(\theta) (\theta_1 - \theta_2)^2 \\ &= \frac{\sigma^2}{2} (\theta_1 - \theta_2)^2 \leq \frac{(\mu_1 - \mu_2)^2}{2C_{\text{var}}^2 \sigma^2}, \end{aligned} \quad (4.9)$$

where the last relation follows from

$$(\mu_1 - \mu_2)^2 = [b'(\theta_1) - b'(\theta_2)]^2 \geq \min_{\theta_{\min} \leq \theta \leq \theta_{\max}} [b''(\theta)]^2 (\theta_1 - \theta_2)^2 \geq C_{\text{var}}^2 \sigma^4 (\theta_1 - \theta_2)^2,$$

which is a direct consequence of Assumption 3. For each $\omega \in \Omega$, let P_ω be the data distribution when $\mu^* = \mu^\omega$. Then we have

$$\begin{aligned} \text{KL}(P_\omega \parallel P_0) &\stackrel{(i)}{=} \sum_{i=1}^n \text{KL}(p_{\theta_i^\omega} \parallel p_{\theta_i^0}) \stackrel{(ii)}{\leq} \sum_{i=1}^n \frac{(\mu_i^\omega - \mu_i^0)^2}{2C_{\text{var}}^2 \sigma^2} = \frac{1}{2C_{\text{var}}^2 \sigma^2} \|\mu^\omega - \mu^0\|^2 \\ &\stackrel{(iii)}{=} \frac{1}{2C_{\text{var}}^2 \sigma^2} \frac{\gamma^2 n}{k} d_{\text{H}}(\omega, \mathbf{0}) \stackrel{(iv)}{\leq} \frac{\gamma^2 n}{2C_{\text{var}}^2 \sigma^2} \stackrel{(v)}{\leq} \frac{c^2}{2C_{\text{var}}^2} k \\ &\leq \frac{8c^2}{C_{\text{var}}^2} \log_2 M \stackrel{(vi)}{\leq} \frac{1}{16} \log_2 M < \frac{1}{8} \log M. \end{aligned} \quad (4.10)$$

Here (i) follows from the additivity property of the KL divergence (cf. [Tsybakov \(2008, page 85\)](#)) and for each $i \in [n]$, $\theta_i^\omega = (b')^{-1}(\mu_i^\omega)$; (ii) follows from (4.9); (iii) follows from the first line of (4.8); (iv) holds since the Hamming distance between k -dimensional vectors is at most k ; (v) follows from (4.7); and (vi) holds when $c \leq C_{\text{var}}/16$. Taking collectively (4.8) and (4.10), we can utilize [Tsybakov \(2008, Theorem 2.7\)](#) to establish the minimax lower bound (4.4).

Case 2: when n is small. We are left with the case when

$$\left(\frac{\tilde{V}}{n\sigma}\right)^{2/3} \leq \frac{4}{n}, \quad (4.11)$$

and we would like to show that

$$\inf_{\hat{\mu}} \sup_{\mu^* \in \mathcal{I}_n(V_{\min}, V_{\max})} \mathbb{E} \|\hat{\mu} - \mu^*\|^2 \geq C_2 \sigma^2. \quad (4.12)$$

We can consider the special case when all μ_i^* are identical and this information is known a priori. Then the problem reduces to estimating $\mu^* \in [V_{\min}, V_{\max}]$ from i.i.d. data $X_1, \dots, X_n \sim p_{\theta^*}(x)$ where $\theta^* = (b')^{-1}(\mu^*)$, and it suffices to show that

$$\inf_{\hat{\mu}} \max_{\mu^* \in [V_{\min}, V_{\max}]} |\hat{\mu} - \mu^*|^2 \geq C_2 \frac{\sigma^2}{n}. \quad (4.13)$$

for some $C_2 > 0$. Let $\mu_1 = \tilde{V}_{\max}$ and $\mu_2 = \tilde{V}_{\max} - c\sigma/\sqrt{n}$ for some sufficiently small constant $c > 0$. Since $\tilde{V}_{\max} - \tilde{V}_{\min} \geq C_{\text{int}}(V_{\max} - V_{\min})$ and $V_{\max} - V_{\min} \geq \sigma/\sqrt{n}$, we know that $\mu_2 \geq \tilde{V}_{\min}$ as long as

$c \leq C_{\text{int}}$. Denote by P_μ the data distribution when the sample average is μ , and let $\theta_1 = (b')^{-1}(\mu_1)$ and $\theta_2 = (b')^{-1}(\mu_2)$. We have

$$\text{KL}(P_{\mu_1} \parallel P_{\mu_2}) \stackrel{\text{(i)}}{=} n \text{KL}(p_{\theta_1} \parallel p_{\theta_2}) \stackrel{\text{(ii)}}{\leq} n \frac{(\mu_1 - \mu_2)^2}{2C_{\text{var}}^2 \sigma^2} = \frac{c}{2C_{\text{var}}^2} \stackrel{\text{(iii)}}{\leq} \frac{1}{8}.$$

Here (i) follows from the additivity property of the KL divergence (cf. [Tsybakov \(2008, Page 85\)](#)); (ii) follows from (4.9); and (iii) holds as long as $c \leq C_{\text{var}}^2/4$. Then we can apply [Tsybakov \(2008, Theorem 2.2\)](#) to obtain the desired minimax lower bound (4.13), which in turn establishes (4.12).

Finally, taking (4.4) and (4.12) collectively gives the minimax lower bound (4.2) as claimed.

5 Numerical experiments

In this section, we conduct numerical experiments to illustrate the performance of the Isotonic Mechanism on a real data example from ICML 2023 and synthetic data.

5.1 The ICML 2023 data

To test the Isotonic Mechanism, the second author of the present paper led an experiment at ICML 2023 through a collaboration between [openreview.net](#) and [openrank.cc](#). The latter website was recently developed specifically for this experiment. Being the second largest conference in machine learning, ICML received a total of 6,538 submissions from 18,515 authors in 2023. Among 4,505 authors who submitted at least two papers, 1,331 authors provided the rankings of their papers for our experiment.

In evaluating the performance of the Isotonic Mechanism on the ICML 2023 ranking data, one challenge arises from the absence of ground-truth scores of the submissions, namely, μ_1^*, \dots, μ_n^* . Recognizing that reviewers also provided their levels of confidence for their review scores, our approach is to use scores given by confident reviewers as a *surrogate* for the ground-truth scores. More precisely, let a submission be reviewed by m reviewers, with scores $X^{(1)}, \dots, X^{(m)}$. Without loss of generality, suppose that the first $m - 1$ reviewers were more confident than the last reviewer.⁶ Formally, we treat the average from the $m - 1$ relative confident scores,

$$\tilde{\mu}^* = \frac{1}{m-1} \sum_{j=1}^{m-1} X^{(j)}, \tag{5.1}$$

as the surrogate for the ground-truth quality of this submission and use the least confident review score, $\tilde{X} = X^{(m)}$, as the “observed” review score of the same submission. To facilitate data processing, we discarded 195 submissions that only had one review score. However, it is worth mentioning that, in practice, the average would be taken over all m review scores instead of some $m - 1$ review scores. As such, these 195 submissions could be included when implementing the Isotonic Mechanism in actual usage.

Consider an author A who submitted n papers with surrogate ground-truth $\tilde{\mu}_1^*, \dots, \tilde{\mu}_n^*$ and noisy scores $\tilde{X}_1, \dots, \tilde{X}_n$.⁷ We run the Isotonic Mechanism using the ranking provided by this author to

⁶In the case of ties, we pick the last reviewer from the least confident reviewers uniformly at random.

⁷We excluded 60 authors in our experiment whose rankings were not informative in the sense that all papers of the same author were ranked first.

#Submissions	2	3	4	5	6	7	8	9
Sample size	786	284	88	52	28	11	9	2
$\widehat{\text{MSE}}_{\text{raw}}$	2.0908	1.9414	2.4904	2.0741	1.9552	1.6054	1.5789	1.1605
$\widehat{\text{MSE}}_{\text{IM}}$	1.8290	1.7375	1.9842	1.6365	1.5720	1.2875	0.9502	0.9420
Improvement	12.52%	10.50%	20.33%	21.10%	19.60%	19.80%	39.82%	18.83%
#Submissions	10	11	12	13	14	15	16	17
Sample size	2	1	1	2	0	1	2	2
$\widehat{\text{MSE}}_{\text{raw}}$	2.6139	2.0783	2.1991	1.8750	NA	1.2741	2.2474	1.5613
$\widehat{\text{MSE}}_{\text{IM}}$	1.3556	1.8294	1.3935	0.9850	NA	1.1574	1.7899	0.7040
Improvement	48.14%	11.98%	36.63%	47.47%	NA	9.16%	20.36%	54.91%

Table 1: Experimental results on the ICML 2023 ranking data. The improvement of the Isotonic Mechanism over the original review scores is defined as $(\widehat{\text{MSE}}_{\text{raw}} - \widehat{\text{MSE}}_{\text{IM}})/\widehat{\text{MSE}}_{\text{raw}}$. For each $2 \leq n \leq 17$, the comparison of the two approaches is evaluated for authors of n submissions.

produce the adjusted scores $\hat{\mu}_1, \dots, \hat{\mu}_n$. With the surrogate ground-truth in place, we evaluate the performance of the unadjusted review scores and the Isotonic Mechanism as follows:

$$\widehat{\text{MSE}}_{\text{raw},A} = \frac{1}{n} \sum_{j=1}^n (\tilde{X}_j - \tilde{\mu}_j^*)^2 \quad \text{and} \quad \widehat{\text{MSE}}_{\text{IM},A} = \frac{1}{n} \sum_{j=1}^n (\hat{\mu}_j - \tilde{\mu}_j^*)^2. \quad (5.2)$$

Next, we calculate the average across all authors with the same number of submissions for the two aforementioned errors:

$$\widehat{\text{MSE}}_{\text{raw},n} = \frac{1}{|\mathcal{A}_n|} \sum_{A \in \mathcal{A}_n} \widehat{\text{MSE}}_{\text{raw},A} \quad \text{and} \quad \widehat{\text{MSE}}_{\text{IM},n} = \frac{1}{|\mathcal{A}_n|} \sum_{A \in \mathcal{A}_n} \widehat{\text{MSE}}_{\text{IM},A},$$

where \mathcal{A}_n is the set of all authors who submitted n papers. The results are presented in Table 1, which shows that the Isotonic Mechanism outperforms the original review scores consistently. In general, the gain becomes more significant as the number of submissions increases. In particular, it is nearly 20% or more for authors with more than three submissions.

5.2 Synthetic data

In the analysis of the ICML 2023 data in Section 5.1, there is a mismatch that reduces the effectiveness of the Isotonic Mechanism. The error criterion $\widehat{\text{MSE}}_{\text{IM},A}$ utilizes the surrogate ground-truth $\tilde{\mu}_1^*, \dots, \tilde{\mu}_n^*$, whereas the Isotonic Mechanism uses the author-provided ranking, which is presumably about μ_1^*, \dots, μ_n^* . In real-world applications, however, this mismatch would not occur. Thus, in practical scenarios, the improvement gained from the Isotonic Mechanism might be even more significant than what is observed in Section 5.1.

To make a comparison from this viewpoint, we resort to synthetic data that mimic the ICML 2023 data. Let $\mathcal{X}_{\text{ICML2023}}$ denote the collection of average review scores of all 6,345 submissions to ICML 2023.⁸ For any $2 \leq n \leq 17$, we generate a synthetic review instance as follows:

- An author submits n papers with ground-truth quality μ_1^*, \dots, μ_n^* that are sampled with replacement from $\mathcal{X}_{\text{ICML2023}}$.

⁸Among the total 6,538 submission to ICML 2023, 193 were desk rejected or withdrawn before review scores were posted.

#Submissions	2	3	4	5	6	7	8	9
Mean($\widehat{\text{MSE}}_{\text{IM}}$)	0.6408	0.6180	0.5619	0.4968	0.4558	0.4384	0.3942	0.3853
Std($\widehat{\text{MSE}}_{\text{IM}}$)	0.6490	0.5172	0.4285	0.3561	0.3187	0.2825	0.2382	0.2379
Mean($\widehat{\text{MSE}}_{\text{raw}}$)	0.7467	0.8005	0.8339	0.7911	0.7918	0.8195	0.7749	0.7884
Std($\widehat{\text{MSE}}_{\text{raw}}$)	0.7241	0.6362	0.5739	0.5038	0.4570	0.4351	0.3651	0.3668
Improvement	14.18%	22.80%	32.61%	37.20%	42.43%	46.50%	49.12%	51.13%
#Submissions	10	11	12	13	14	15	16	17
Mean($\widehat{\text{MSE}}_{\text{IM}}$)	0.3639	0.3406	0.3230	0.3176	0.2932	0.2910	0.2736	0.2677
Std($\widehat{\text{MSE}}_{\text{IM}}$)	0.2245	0.1829	0.1756	0.1698	0.1507	0.1496	0.1339	0.1338
Mean($\widehat{\text{MSE}}_{\text{raw}}$)	0.7925	0.7932	0.7934	0.7996	0.7900	0.7899	0.7884	0.8001
Std($\widehat{\text{MSE}}_{\text{raw}}$)	0.3724	0.3329	0.3183	0.3007	0.2804	0.2886	0.2751	0.2738
Improvement	54.09%	57.06%	59.29%	60.29%	62.89%	63.16%	65.30%	66.55%

Table 2: Experimental results on the synthetic data. The mean and standard deviation of the estimation errors are computed from 1,000 independent runs for each number of submissions.

- For each $1 \leq i \leq n$, let the average review score X_i of the i th paper be the average of 3 i.i.d. $\text{Binomial}(10, \mu_i^*/10)$ random variables.⁹
- Run isotonic regression to obtain $\widehat{\mu}_1, \dots, \widehat{\mu}_n$ for the synthetic review scores X_1, \dots, X_n and the ranking of μ_1^*, \dots, μ_n^* . Calculate $\widehat{\text{MSE}}_{\text{raw}}$ and $\widehat{\text{MSE}}_{\text{IM}}$ as (5.2), with \widetilde{X}_j and $\widetilde{\mu}_j^*$ replaced by X_j and μ_j^* , respectively.

We run 1,000 independent trials for each n to compute the mean and the standard deviation of the two estimation errors. Our results, shown in Table 2, indeed demonstrates more significant gain of the Isotonic Mechanism than Table 1 does. Even for a small number of submissions such as $n = 2$ or 3, the gain of using the Isotonic Mechanism is already noticeable.

5.3 More simulations on nonconvex utility functions

In this section, we design numerical experiments to study whether the theory established in this paper still hold beyond the current model assumptions. Our truthfulness result (cf. Theorem 1) relies on the convex utility assumption (cf. Assumption 1). While this assumption is satisfied if the author hopes to achieve a higher score to increase the chance of winning oral/spotlight presentation or best paper awards, it does not applies to the case where an author only cares about the acceptance of the paper. In this case, the utility of such an author can be given by an indicator function $U(x) = \mathbf{1}\{x \geq c\}$ for some threshold $c > 0$. For example, when the review scores are between 1 and 10, taking $c = 6$ reflects the utility that the author’s utility only depends on the acceptance of the paper. A natural question is, under such a nonconvex utility function, whether it is possible for the author to get a higher overall utility by dishonestly manipulating the reported ranking.

To answer this question, we consider the setting when an author submits two papers with true scores $\mu_1^*, \mu_2^* \in [0, 10]$, and they receive noisy review scores $X_i \sim \text{Bernoulli}(10, \mu_i^*/10)$ for $i = 1, 2$ respectively. In this case, there are only two possible rankings (the ground truth ranking and the reversed ranking). We use Monte Carlo methods to evaluate the difference of the overall utility of reporting the true and reversed rankings for all possible pairs $(\mu_1^*, \mu_2^*) \in [0, 10]^2$, namely

$$\mathbb{E}[U(\widehat{\mu}_1) + U(\widehat{\mu}_2)] - \mathbb{E}[U(\bar{\mu}_1) + U(\bar{\mu}_2)],$$

⁹Here, the number 3 was chosen because it is the median number of reviewers assigned to an ICML 2023 submission.




Figure 4: The difference of utility between reporting the correct and incorrect ranking, for two different types of utility functions. The difference is positive (resp. negative) means that the author has higher utility when reporting the correct (resp. incorrect) ranking, and is represented by the green (resp. red) color.

where $(\hat{\mu}_1, \hat{\mu}_2)$ (resp. $(\bar{\mu}_1, \bar{\mu}_2)$) are the output of the isotonic mechanism with ground truth (resp. reversed) ranking. The result for the utility function $U(x) = \mathbf{1}\{x \geq 6\}$ is presented in the left panel of Figure 4, which shows that even if one only cares the acceptance of the papers, for most (μ_1^*, μ_2^*) pairs, the utility of reporting truthfully is still higher, with the exception when the difference between μ_1^* and μ_2^* is very large. Namely, it is possible to sacrifice the score of a very good paper a bit to gain a better chance for a much worse paper to get accepted. However we believe that this can hardly happen in practice, because of the following reasons.

- It is reasonable to believe that an author capable of writing a very good paper (e.g., with true score 8 – 9) will prefer to seek for higher chance of winning paper awards or oral presentation with such a good paper, rather than using it to save a low-quality paper.
- Given the review scores are highly noisy, it is actually risky to use this strategy in practice. Although the expected utility might be higher, there is a significant chance that both papers can end up being rejected, and we believe that most authors may not want to take this risk.
- Most ML papers are written by several co-authors, who will submit their own rankings. It is not likely that they all wish to save a low-quality paper using the good one they co-authored. If they do not reach such a consensus and only one co-author is trying to manipulate his or her rank, we believe it is not difficult to detect this outlier from other co-authors' ranking.

Inspired by the above discussion, we also consider another nonconvex utility function

$$U(x) = \mathbf{1}\{x \geq 6\} + 0.5 \mathbf{1}\{x \geq 8\} + 0.5 \mathbf{1}\{x \geq 9\},$$

which also incorporates the utility from winning paper awards or oral presentations. The right panel of Figure 4 depicts the result for this utility function, which suggests that there is almost no gain in increasing the utility function by manipulating the ranking.

6 Discussion

In this paper, we examine the problem of eliciting private information from an owner to estimate a score vector generated from an exponential family distribution with varying parameters. Our primary contribution demonstrates that the Isotonic Mechanism (Su, 2021, 2022) maintains truthfulness in this extended regime, as the owner would optimize her convex utility by providing the ground-truth ranking for estimation. Additionally, we establish that the mechanism’s truthfulness necessitates pairwise comparisons between the score means under certain conditions, which directly illustrates the optimality of ranking in truthful private information elicitation. Our analysis reveals that the Isotonic Mechanism achieves near-optimal estimation properties from a minimax perspective in a regime of bounded total variation, with the estimator produced by the mechanism significantly outperforming the original scores. This extension is made possible through the introduction of several new technical elements for analyzing the Isotonic Mechanism.

The extended Isotonic Mechanism is particularly well-suited for applications such as peer review in large machine learning and artificial intelligence conferences (e.g., ICML and NeurIPS), where review scores are integer-valued¹⁰ and exhibit a strong dependence between mean and variance. The flexibility and richness of exponential family distributions aptly address this dependence. Furthermore, a key feature of the extended mechanism is its independence from the specific form of the exponential family distribution that generates the observations, rendering it easy to implement in practice.

Several important directions for future research emerge from our current work. One immediate question concerns further investigation of the necessary conditions for truthfulness by examining whether the intercepts in Theorem 3 are equal to zero beyond the Gaussian case of Proposition 1. Additionally, extending our mechanism to accommodate multivariate scores would be a welcome advance, which would allow for the evaluation of distinct aspects of a paper (e.g., novelty and writing quality). This multidimensional approach could provide a more nuanced and comprehensive evaluation framework. Another promising direction involves developing a ranking method for exponential family distributions in the context of multiple co-authors (Wu et al., 2023). From an empirical standpoint, assessing the robustness of the Isotonic Mechanism in scenarios where the author might have uncertainty in ranking the submissions would be valuable. This extension would explore situations where the author-provided ranking is only approximately truthful. More broadly, for journal submissions where authors typically submit one paper at a time, a “credit” system incorporating submission history with self-ratings could extend the Isotonic Mechanism to the sequential setting. Such a system could balance short-term evaluations with long-term reputational effects, presenting a worthwhile avenue for future research in peer review mechanisms.

Acknowledgements

J. Fan was supported by ONR grant N00014-19-1-2120 and the NSF grants DMS-2053832 and DMS-2210833. W.J. Su was supported in part by NSF grants CCF-1934876 and CAREER DMS-1847415. Y. Yan was supported in part by the Charlotte Elizabeth Procter Honorific Fellowship from Princeton University. We would like to thank Bart Lipman for suggesting relevant references.

¹⁰In our setting, we consider the average score over three or four review scores, resulting in an integer multiple of $\frac{1}{3}$ or $\frac{1}{4}$.

References

- Richard E Barlow and Hugh D Brunk. The isotonic regression problem and its dual. *Journal of the American Statistical Association*, 67(337):140–147, 1972.
- Marco Battaglini. Multiple referrals and multidimensional cheap talk. *Econometrica*, 70(4):1379–1401, 2002.
- Pierre C Bellec and Alexandre B Tsybakov. Sharp oracle bounds for monotone and convex regression through aggregation. *Journal of Machine Learning Research*, 16(56):1879–1892, 2015.
- Archishman Chakraborty and Rick Harbaugh. Comparative cheap talk. *Journal of Economic Theory*, 132(1):70–94, 2007.
- Sergei Ivanov. NeurIPS 2020. comprehensive analysis of authors, organizations, and countries. <https://medium.com/criteo-engineering/neurips-2020-comprehensive-analysis-of-authors-organizations-and-countries-a1b55a08132e>, 2020. [Online].
- Steven Jecmen, Hanrui Zhang, Ryan Liu, Nihar Shah, Vincent Conitzer, and Fei Fang. Mitigating manipulation in peer review via randomized reviewer assignments. *Advances in Neural Information Processing Systems*, 33:12533–12545, 2020.
- Ari Kobren, Barna Saha, and Andrew McCallum. Paper matching with local fairness constraints. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1247–1257, 2019.
- Neil Lawrence and Corinna Cortes. The NIPS experiment, 2014.
- Kevin Leyton-Brown, Yatin Nandwani, Hedayat Zarkoob, Chris Cameron, Neil Newman, Dinesh Raghu, et al. Matching papers and reviewers at large conferences. *Artificial Intelligence*, 331:104119, 2024.
- Xin Li. Acceptance rates for the major top-tier ai-related conferences. <https://github.com/lixin4ever/Conference-Acceptance-Rate>, 2022. [Online].
- Elliot Lipnowski and Doron Ravid. Cheap talk with transparent motives. *Econometrica*, 88(4):1631–1660, 2020.
- Robert S MacKay, Ralph Kenna, Robert J Low, and Sarah Parker. Calibration with confidence: a principled method for panel assessment. *Royal Society open science*, 4(2):160760, 2017.
- Albert W Marshall, Ingram Olkin, and Barry C Arnold. Inequalities: theory of majorization and its applications. 1979.
- Roger B Myerson. *Mechanism design*. Springer, 1989.
- SR Paul. Bayesian methods for calibration of examiners. *British Journal of Mathematical and Statistical Psychology*, 34(2):213–223, 1981.
- Nihar B Shah. Challenges, experiments, and computational solutions in peer review. *Communications of the ACM*, 65(6):76–87, 2022.

Ivan Stelmakh, Nihar B Shah, Aarti Singh, and Hal Daumé III. A novice-reviewer experiment to address scarcity of qualified reviewers in large conferences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 4785–4793, 2021.

Weijie J Su. You are the best reviewer of your own papers: An owner-assisted scoring mechanism. *Advances in Neural Information Processing Systems*, 34:27929–27939, 2021.

Weijie J Su. A truthful owner-assisted scoring mechanism. *arXiv preprint arXiv:2206.08149*, 2022.

Shao-Hua Sun. ICLR2020-openreviewdata. <https://github.com/shaohua0116/ICLR2020-OpenReviewData>, 2020. [Online].

A.B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer New York, 2008. ISBN 9780387790527. URL <https://books.google.com/books?id=mwB8rUBsbqoC>.

Jingyan Wang and Nihar B Shah. Your 2 is my 1, your 3 is my 9: Handling arbitrary miscalibrations in ratings. *arXiv preprint arXiv:1806.05085*, 2018.

Jibang Wu, Haifeng Xu, Yifan Guo, and Weijie Su. A truth serum for eliciting self-evaluations in scientific reviews. *arXiv preprint arXiv:2306.11154*, 2023.

Cun-Hui Zhang. Risk bounds in isotonic regression. *The Annals of Statistics*, 30(2):528–555, 2002.

A Appendix

A.1 Proof of Lemma 3

We prove the lemma by induction. When $n = 1$, the result is trivially true. Suppose that the result holds for all $1 \leq n \leq N - 1$. Now we show that the result also holds for $n = N$. Let $j = \pi^{-1}(N)$, namely j is the index such that $\pi(j) = N$. We define $\tilde{\pi}$ by swapping the value of $\pi(j)$ and $\pi(N)$ in π :

$$\tilde{\pi}(i) = \begin{cases} N & \text{if } i = N, \\ \pi(N) & \text{if } i = j, \\ \pi(i) & \text{otherwise.} \end{cases}$$

It is straightforward to check that $\tilde{\pi}$ is an upward swap of π when $j < N$, and $\tilde{\pi} = \pi$ when $j = N$. Since $\tilde{\pi}(N) = N$ and $\pi^*(N) = N$, by restricting π^* and $\tilde{\pi}$ to $[N - 1]$, we obtain two permutations $\bar{\pi}^*$ and $\bar{\pi}$ over $[N - 1]$, namely

$$\bar{\pi}^*(i) := \pi^*(i), \quad \bar{\pi}(i) = \tilde{\pi}(i), \quad \forall i \in [N - 1].$$

Then we can apply the induction hypothesis to show that there exist $m \geq 1$ and a sequence of permutations over $[N - 1]$: $\bar{\pi}_1, \dots, \bar{\pi}_m$ such that $\bar{\pi}_1 = \bar{\pi}^*$, $\bar{\pi}_m = \bar{\pi}$, and for all $1 \leq i \leq m - 1$, $\bar{\pi}_i$ is an upward swap of $\bar{\pi}_{i+1}$. Then we extend $\bar{\pi}_1, \dots, \bar{\pi}_m$ to be permutations $\tilde{\pi}_1, \dots, \tilde{\pi}_m$ over $[N]$ by defining

$$\tilde{\pi}_k(i) = \begin{cases} \bar{\pi}_k(i) & \text{if } i \leq N - 1, \\ N & \text{if } i = N, \end{cases}$$

for all $k \in [m]$ and $i \in [N]$. When $j < N$, let

$$\pi_1 = \tilde{\pi}_1, \quad \pi_2 = \tilde{\pi}_2, \quad \dots \quad \pi_m = \tilde{\pi}_m, \quad \pi_{m+1} = \pi.$$

It is straightforward to check that $\pi_1 = \pi^*$, $\pi_m = \pi$, and for each $1 \leq i \leq m$, π_i is an upward swap of π_{i+1} . When $j = N$, let

$$\pi_1 = \tilde{\pi}_1, \quad \pi_2 = \tilde{\pi}_2, \quad \dots \quad \pi_m = \tilde{\pi}_m.$$

We can also check that $\pi_1 = \pi^*$, $\pi_m = \pi$, and for each $1 \leq i \leq m - 1$, π_i is an upward swap of π_{i+1} . This concludes the proof.