A Truthful Owner-Assisted Scoring Mechanism

Weijie J. Su

University of Pennsylvania

June 2022

Abstract

Alice (owner) has knowledge of the underlying quality of her items measured in grades. Given the noisy grades provided by an independent party, can Bob (appraiser) obtain accurate estimates of the ground-truth grades of the items by asking Alice a question about the grades? We address this when the payoff to Alice is additive convex utility over all her items. We establish that if Alice has to truthfully answer the question so that her payoff is maximized, the question must be formulated as *pairwise comparisons* between her items. Next, we prove that if Alice is required to provide a ranking of her items, which is the most fine-grained question via pairwise comparisons, she would be truthful. By incorporating the ground-truth ranking, we show that Bob can obtain an estimator with the optimal squared error in certain regimes based on any possible way of truthful information elicitation. Moreover, the estimated grades are substantially more accurate than the raw grades when the number of items is large and the raw grades are very noisy. Finally, we conclude the paper with several extensions and some refinements for practical considerations.

1 Introduction

An owner has a number of items, which are evaluated by an independent party. The raw grades provided by the party are, however, noisy. Suppose that the owner is authoritative regarding the knowledge of her items' quality. Instead of directly using the raw grades, can an appraiser improve on the raw grades of the items, for example, by eliciting hopefully useful information from the owner?

Given below are three examples of practical scenarios that can be modeled by the setting above.

- *Peer review.* One author submitting many papers to a single machine learning conference, such as NeurIPS, ICML, AAAI, or ICLR, is a common practice [29]. The ratings provided by the reviewers are perhaps the most important factor in rejecting and accepting the submissions but, unfortunately, suffer from surprisingly high variability [25, 27]. On the contrary, the authors often have a good understanding of the quality of their papers.

- *Player valuation.* Every player on a soccer team is rated by sports performance analysis agencies, such as the FIFA Index and InStat. However, the team manager has additional information regarding the strengths and health conditions of the players in the team.

Email: suw@wharton.upenn.edu. A preliminary form of part of the results was presented at the Conference on Neural Information Processing Systems (NeurIPS), 2021 [28].

- *Second-hand market.* A car leasing company is selling cars and they have been rated by a vehicle valuation agency. In addition to the ratings, the leasing company has some private information about the reliability of these used cars.

In its simplest form, the situation of the owner and the appraiser can be formulated as follows. Imagine that an appraiser named Bob observes a noisy vector $\boldsymbol{y} = (y_1, \ldots, y_n)$ from the model

$$y_i = R_i + z_i$$

for $i = 1, \ldots, n$, where $\boldsymbol{R} = (R_1, \ldots, R_n)$ is the ground truth and $\boldsymbol{z} = (z_1, \ldots, z_n)$ denotes the noise vector. Suppose the ground truth $\boldsymbol{R}$ is known to the owner named Alice. We partition the space $\mathbb{R}^n$ into disjoint sets $S_1, \ldots, S_m$. These sets are presented as a question for Alice, and she must inform Bob of exactly *one* set that, she alleges, contains the ground truth $\boldsymbol{R}$, *before* seeing the observation $\boldsymbol{y}$. Ideally, we wish to design a partition $\{S_1, \ldots, S_m\}$ that incentivizes Alice to tell the truth, thereby offering useful side information to Bob for possibly better estimation of the ground truth.

However, Alice does not necessarily need to honestly report the set that truly contains the ground truth. Indeed, Alice can pick any element from $\{S_1, \ldots, S_m\}$ in her own interest. For example, if asked directly, "what are the exact grades of your items?" Alice would have an incentive to report higher values than what she actually knows.

To understand whether Alice would be truthful in relation to Bob's approach to estimation, it is imperative to recognize that the two parties are driven by different objectives. Bob's goal is to obtain an estimator of the ground truth $\boldsymbol{R}$ as accurate as possible. Given Alice's selection S from the partition $\{S_1, \ldots, S_m\}$, Bob has to trust her and consider the constraint $\boldsymbol{R} \in S$, though the constraint itself may be incorrect. Perhaps the simplest way of estimating $\boldsymbol{R}$ is to solve the optimization problem

$$\min_{\boldsymbol{r}} \ \|\boldsymbol{y} - \boldsymbol{r}\|^2 \tag{1.1}$$

subject to the constraint $\boldsymbol{r} = (r_1, \ldots, r_n) \in S$, where $\|\cdot\|$ denotes the Euclidean/ℓ_2 norm throughout the paper. This is constrained maximum likelihood estimation when the noise variables $z_1, \ldots, z_n$ are independent and identically distributed (i.i.d.) normal random variables with mean zero [1]. The solution serves as an appraisal of the grades of n items. Arguably, a fine-grained partition $\{S_1, \ldots, S_m\}$ is preferred because it enables Bob to better narrow down the search space for estimating the ground truth.

From Alice's perspective, however, the estimation accuracy means little as she might already know the exact values of $\boldsymbol{R}$. Instead, her objective is to maximize her payoff as a function of the solution $\widehat{\boldsymbol{R}}$ to (1.1). To formalize this viewpoint, letting U be any nondecreasing *convex* function, we assume that Alice is rational and strives to maximize the expected overall utility

$$\mathbb{E}\left[U(\widehat{R}_1) + \cdots + U(\widehat{R}_n)\right]$$

as much as she can by reporting any element S from the given partition $\{S_1, \ldots, S_m\}$, either truthfully or not.

As the main result, we address the needs of both Alice and Bob simultaneously in the following theorem.

Theorem 1.1 (informal). *The most precise information that Bob can assure to truthfully elicit from Alice is the ranking of her items in descending (or, equivalently, ascending) order of the ground-truth* *grades $R_1, \ldots, R_n$. Moreover, the estimator* (1.1) *provided the true ranking is more accurate than the raw observation $\boldsymbol{y}$.*

This mechanism requires Alice to provide an arbitrary ranking of her items. To do so in a possibly truthful manner, it suffices for Alice to know the relative magnitudes instead of the exact values of $R_1, \ldots, R_n$. With the ranking serving as the constraint, (1.1) is a convex quadratic program and is equivalent to isotonic regression [4]. As such, we call it the *Isotonic Mechanism*.

The optimality of the Isotonic Mechanism lies in its truthfulness as well as the most informative location it offers for estimation. The combination of the two appeals is established by taking together the following two theorems. The first one provides a necessary condition, showing that Alice would be truthful *only if* the questions are based on *pairwise comparisons*.

Theorem 1.2 (formal statement in Theorem 1)**.** *If Alice is always truthful under the aforementioned assumptions, then the partition $\{S_1, \ldots, S_m\}$ must be separated by several pairwise-comparison hyperplanes $x_i - x_j = 0$ for some pairs $1 \le i < j \le n$.*

If a partition is pairwise-comparison-based, then to determine whether or not $\boldsymbol{x} \in S$ for any element S in the partition $\{S_1, \ldots, S_m\}$, it suffices to check whether or not $x_i \ge x_j$ for some pairs $i < j$. Conversely, this is also true. For example, the collection of $\{\boldsymbol{x} \in \mathbb{R}^3 : \min(x_i, x_{i+1}) \ge x_{i+2}\}$ for $i = 1, 2, 3$ is such a partition in three dimensions,[1] where we adopt the cyclic convention $x_{i+3} = x_i$. To see this, note that $\boldsymbol{x} \in S_i$ if $x_i \ge x_{i+2}$ and $x_{i+1} \ge x_{i+2}$. On the contrary, for instance, the collection of all spheres $\{\boldsymbol{x} \in \mathbb{R}^n : \|\boldsymbol{x}\| = c\}$ for all $c \ge 0$ cannot be generated from pairwise comparisons.

However, there are examples of pairwise-comparison-based partitions for which Alice would not be truthful (see Section 4). Consequently, it seems on the surface that one needs to prove or disprove for every pairwise-comparison-based partition. From a practical viewpoint, however, we can bypass this cumbersome step since it is plausible to expect that the solution to (1.1) would become better if the partition $\{S_1, \ldots, S_m\}$ becomes finer-grained. In this respect, the best possible hope is to show truthfulness for the most fine-grained partition induced by pairwise-comparison hyperplanes, which is the collection of all $n!$ rankings of the n items.

This result is confirmed by the following theorem.

Theorem 1.3 (formal statement in Theorem 2 and Proposition 2.8)**.** *If Alice is required to provide a ranking of her items as the constraint for the estimation problem* (1.1)*, her expected overall utility would be maximized if she reports the ground-truth ranking of $R_1, \ldots, R_n$, and Bob can improve the estimation accuracy by making use of this ranking.*

Taken together, Theorems 1.2 and 1.3 give our main result, Theorem 1.1.

Considering the above, the Isotonic Mechanism is remarkable in that it satisfies two desiderata: first, it renders Alice honest; consequently, it provides Bob with the most fine-grained information that is achievable by pairwise comparisons. Moreover, the accuracy improvement of the Isotonic Mechanism over the raw observation becomes especially substantial if the noise in the observation $\boldsymbol{y}$ is significant and the number of items n is large.

Interestingly, this is very much the case with peer review in some major machine learning and artificial intelligence conferences: while the high variability of review ratings may be owing to the large number of novice reviewers [8], an increasingly common trend is that an author often submits

[1] We ignore the overlaps in the boundaries, which are of Lebesgue measure zero.

a host of papers to a single conference. For instance, one researcher submitted as many as 32 papers to ICLR in 2020 [29].

The remainder of the paper is structured as follows. In Section 2, we lay out the setting and introduce precise assumptions for a formal version of Theorem 1.2, briefly discussing the estimation properties. Next, Section 3 states the formal results that the owner would be truthful under the Isotonic Mechanism. In addition, we show by examples in Section 4 that some pairwise-comparison-based partitions are truthful while some are not. In Section 5, we present several extensions demonstrating that honesty yields the highest payoff in more general settings. The proofs of our main results are given in Section 6. Section 7 concludes the paper by discussing several directions for future research.

2 When is honesty possible?

To develop a mechanism that enables information elicitation from the owner, we need to specify the class of questions that the appraiser can ask. In our setting, a question is represented by a partition $\mathcal{S} := \{S_1, \ldots, S_m\}$ of the Euclidean space $\mathbb{R}^n$, to which the ground truth $\boldsymbol{R}$ belongs.[2] We call $\mathcal{S}$ a *knowledge partition* and the sets $S_1, \ldots, S_m$ *knowledge elements*.

The owner is required to pick a knowledge element, say, S, from the knowledge partition $\mathcal{S}$ and sends the message "the ground truth is in the set S" to the appraiser. She is *not* allowed to observe $\boldsymbol{y}$ while making the decision or, equivalently, her decision is made before the independent party rates the items. On the other hand, the appraiser knows nothing about the ground truth $\boldsymbol{R}$ but can observe $\boldsymbol{y} = \boldsymbol{R} + \boldsymbol{z}$. Given the message that "the true vector is in S" provided by the owner, the appraiser solves the following optimization program

$$\begin{aligned} \min_{\boldsymbol{r}} \quad & \|\boldsymbol{y} - \boldsymbol{r}\|^2 \\ \text{s.t.} \quad & \boldsymbol{r} \in S \end{aligned} \tag{2.1}$$

and uses its solution as an estimator of the ground truth $\boldsymbol{R}$. This program is equivalent to projecting $\boldsymbol{y}$ onto the knowledge element S.

Formally, knowledge elements are closed sets with nonempty interiors. Their union $\cup_{k=1}^m S_k = \mathbb{R}^n$ and their interiors are disjoint from each other, that is, $\mathring{S}_k \cap \mathring{S}_l = \emptyset$ for any $1 \le k < l \le m$.[3] In addition, we assume that the boundary between any two adjacent knowledge elements is a piecewise smooth surface.[4] We call a surface smooth if any point on the surface is locally defined by the equation $f(x_1, \ldots, x_n) = 0$ for some continuously differentiable function f with nondegenerate gradient ∇f.

We make the following assumptions to investigate under what conditions this type of information elicitation would incentivize the owner to truthfully report the knowledge element that contains the ground truth.

Assumption 2.1. *The owner has sufficient knowledge of the ground truth $\boldsymbol{R}$ of her $n \ge 2$ items to allow her to determine which knowledge element of $\mathcal{S}$ contains $\boldsymbol{R}$.*

[2] The number of elements m in a partition can be finite or infinite. Notably, it is infinite in Proposition 7.1.

[3] The boundaries have Lebesgue measure zero. Thus, although some knowledge elements might overlap in their boundaries, for a generic ground truth vector $\boldsymbol{R}$, there is a unique knowledge element that contains it.

[4] Here we do not say two sets are adjacent if the two share merely a corner where three or more sets meet, which is in the same spirit as the Four-Color Theorem.

Remark 2.1. This assumption is weaker than assuming that the owner knows the exact values of the ground-truth grades $R_1, \ldots, R_n$.

Assumption 2.2. *The noise variables $z_1, \ldots, z_n$ are i.i.d. draws from a probability distribution.*

Remark 2.2. The condition of independence can be relaxed to exchangeability for most results in the paper, unless otherwise specified. That is, we assume that the vector $(z_1, \ldots, z_n)$ has the same probability distribution as $(z_{\pi(1)}, \ldots, z_{\pi(n)})$ for any permutation $\pi(1), \ldots, \pi(n)$ of the indices $1, \ldots, n$. This generalization is useful when the noise terms are influenced by a latent factor. Notably, the noise distribution can have nonzero mean in this assumption.

Let $\widehat{\boldsymbol{R}} := (\widehat{R}_1, \ldots, \widehat{R}_n)$ be the appraiser's estimate of the ground-truth vector. The last assumption is concerned with the overall utility that the owner strives to maximize.

Assumption 2.3. *Given estimates $\widehat{R}_1, \ldots, \widehat{R}_n$, the overall utility of the owner takes the form*

$$U(\widehat{\boldsymbol{R}}) := \sum_{i=1}^{n} U(\widehat{R}_i), \tag{2.2}$$

where U is a nondecreasing convex function. The owner attempts to maximize the expected overall utility $\mathbb{E}\, U(\widehat{\boldsymbol{R}})$.

To put the convexity assumption differently, the marginal utility U' is nondecreasing. Convex utility is often assumed in the literature [14], and in particular, it does not contradict the economic law of diminishing marginal utility [13] since the grade measures quality as opposed to quantity. In peer review in machine learning conferences, for example, high ratings largely determine whether an accepted paper would be presented as a poster, an oral presentation, or granted a best paper award. While an oral presentation draws slightly more attention than a poster, a best paper award would drastically enhance the impact of the paper. Accordingly, the marginal utility tends to be larger when the ratings are higher. Another example is the diamond-quality grading system used by the Gemological Institute of America. Typically, the price of a high-grade diamond increases faster with its grade than a low-grade diamond.

In Section 3.2 and Section 5, we consider two relaxations of this assumption.

2.1 A necessary condition

This subsection characterizes the class of truthful knowledge partitions. To highlight its dependence on the knowledge element S reported by the owner, call $\widehat{\boldsymbol{R}}_S$ the solution to the optimization program (2.1).

Definition 2.4. A knowledge partition $\mathcal{S}$ is said to be *truthful* with utility function U if the owner maximizes her expected overall utility (2.2) by reporting the knowledge element from $\mathcal{S}$ that truly contains the ground truth. That is,

$$\mathbb{E}\, U(\widehat{\boldsymbol{R}}_S) \geq \mathbb{E}\, U(\widehat{\boldsymbol{R}}_{S'})$$

for all $S, S' \in \mathcal{S}$ such that S contains the ground truth $\boldsymbol{R}$.

For convenience, the following theorem precludes the trivial case where $\mathcal{S}$ contains only one element $\mathbb{R}^n$. By definition, this trivial knowledge partition is truthful.

Theorem 1. *If the owner tells the truth whenever Assumptions 2.1, 2.2, and 2.3 are satisfied, then the boundary between any two adjacent knowledge elements is piecewise-flat and each flat surface must be part of a pairwise-comparison hyperplane that is defined by*

$$x_i - x_j = 0$$

for some $1 \le i < j \le n$.

Remark 2.3. This characterization of truthful knowledge partitions is obtained by taking an arbitrary convex utility. For a specific utility function, however, the partition might not be based on pairwise comparisons. See Proposition 7.1 in Section 7.

Remark 2.4. The squared ℓ_2 loss in (2.1) can be replaced by the sum of Bregman divergences. Let ϕ be a twice continuously differentiable, strictly convex function and denote by $D_\phi(y, r) = \phi(y) - \phi(r) - (y - r)\phi'(r)$ its Bregman divergence. Then, Theorem 1 remains true if the appraiser uses the solution to the following program:

$$\begin{aligned} \min_{\boldsymbol{r}} \quad & \sum_{i=1}^{n} D_\phi(y_i, r_i) \\ \text{s.t.} \quad & \boldsymbol{r} \in S. \end{aligned} \tag{2.3}$$

This program reduces to (2.1) when $\phi(x) = x^2$. Another example is the Kullback–Leibler divergence $D_\phi(y, r) = y \log \frac{y}{r} + (1 - y) \log \frac{1-y}{1-r}$ for $0 < y, r < 1$, which is generated by the negative entropy $\phi(x) = x \log x + (1 - x) \log(1 - x)$.

This necessary condition is equivalent to the following: for any point $\boldsymbol{x}$, one can determine which knowledge element contains $\boldsymbol{x}$ by performing pairwise comparisons of some coordinates of this point. For example, consider the collection of

$$S_i = \{\boldsymbol{x} \in \mathbb{R}^n : x_i \text{ is the largest among } x_1, x_2, \ldots, x_n\} \tag{2.4}$$

for $i = 1, \ldots, n$. This is a pairwise-comparison-based knowledge partition because one can conclude that $x \in S_i$ if and only if $x_i \ge x_j$ for all $j \ne i$. However, it is important to note that the converse of Theorem 1 is not true.[5] Indeed, we will show that some pairwise-comparison-based knowledge partitions are truthful, while some are not. This is the subject of Section 4.

Since all pairwise-comparison hyperplanes pass through the origin, an immediate consequence of Theorem 1 is the following result.

Corollary 2.5. *If a knowledge partition* $\mathcal{S}$ *is truthful whenever Assumptions 2.1, 2.2, and 2.3 are satisfied, then any knowledge element* S *of* $\mathcal{S}$ *is a cone. That is, if* $\boldsymbol{x} \in S$, *then* $\lambda \boldsymbol{x} \in S$ *for all* $\lambda \ge 0$.

An important example of a pairwise-comparison-based knowledge partition is the collection of the standard isotonic cone $\{\boldsymbol{x} \equiv (x_1, \ldots, x_n) : x_1 \ge x_2 \ge \cdots \ge x_n\}$ under all $n!$ permutations. Explicitly, an element of this knowledge partition takes the form

$$S_\pi = \{\boldsymbol{x} : x_{\pi(1)} \ge x_{\pi(2)} \ge \cdots \ge x_{\pi(n)}\}$$

[5]However, Theorem 4 shows that the knowledge partition formed by (2.4) is truthful.

for some permutation π of $1, 2, \ldots, n$. As is evident, this is generated by pairwise-comparison hyperplanes and is the *most fine-grained* in the sense that the element S_π cannot be further split by pairwise-comparison hyperplanes. Henceforth, we refer to it as the *isotonic knowledge partition.*

This knowledge partition is fundamental because any truthful partition can be derived from it. To see this point, note that any knowledge partition satisfying the necessary condition in Theorem 1 can be obtained by merging the $n!$ isotonic cones into several groups. Each group corresponds to a knowledge element in the resulting knowledge partition.

The discussion above formally implies the following consequence. We say that a knowledge partition $\mathcal{S}_1$ is coarser than another $\mathcal{S}_2$ if any element of $\mathcal{S}_1$ is a union of several elements of $\mathcal{S}_2$.

Corollary 2.6. *If a knowledge partition $\mathcal{S}$ is truthful whenever Assumptions 2.1, 2.2, and 2.3 are satisfied, then $\mathcal{S}$ is coarser than the isotonic knowledge partition*

$$\{S_\pi : \pi \textit{ is a permutation of } 1, 2, \ldots, n\}.$$

In particular, the cardinality of $\mathcal{S}$ is no more than $n!$.

Remark 2.5. Since the union of cones is also a cone, this result implies Corollary 2.5. However, a knowledge element is not necessarily convex and can even be noncontiguous.

Recognizing that it requires at least two items to compare pairwise, it seems necessary to have $n \geq 2$ for the existence of a truthful knowledge partition. Indeed, this intuition is confirmed by the following proposition. A proof of this result is provided in the Appendix, which however does not directly follow from Theorem 1 since the theorem assumes $n \geq 2$ from the beginning.

Proposition 2.7. *Other than the trivial knowledge partition $\mathcal{S} = \{\mathbb{R}\}$, there does not exist a truthful knowledge partition when the dimension $n = 1$ under Assumptions 2.1*[6]*, 2.2, and 2.3.*

There is a copious body of research using pairwise comparisons to recover a ranking of the items [10, 31, 7, 20, 26, 21]. Our results offer new reflections on the use of comparative measurements for estimation, albeit from a different perspective.

2.2 An estimation perspective

Although it is unclear whether a pairwise-comparison-based knowledge partition would really ensure truthfulness, we can say something about its estimation properties. If the reported knowledge element is a convex set, then the estimated grades by the appraiser are more accurate than the raw grades.

Recall that $\widehat{\boldsymbol{R}}_S$ is the solution to (2.1).

Proposition 2.8. *Suppose that the knowledge element S in the optimization program* (2.1) *is convex and contains the ground truth $\boldsymbol{R}$. Then,* (2.1) *improves the estimation accuracy of the ground-truth grades in the sense that*

$$\mathbb{E}\left\|\widehat{\boldsymbol{R}}_S - \boldsymbol{R}\right\|^2 \leq \mathbb{E}\left\|\boldsymbol{y} - \boldsymbol{R}\right\|^2 = \mathbb{E}\left\|\boldsymbol{z}\right\|^2.$$

Proof of Proposition 2.8. Consider the (possibly degenerate) triangle formed by $\boldsymbol{y}, \boldsymbol{R}, \widehat{\boldsymbol{R}}_S$. Assuming the angle $\measuredangle(\boldsymbol{y}, \widehat{\boldsymbol{R}}_S, \boldsymbol{R}) \geq 90°$ for the moment, we immediately conclude that $\|\boldsymbol{y} - \boldsymbol{R}\| \geq \|\widehat{\boldsymbol{R}}_S - \boldsymbol{R}\|$, thereby proving the proposition. To finish the proof, suppose the contrary that $\measuredangle(\boldsymbol{y}, \widehat{\boldsymbol{R}}_S, \boldsymbol{R}) < 90°$.

[6]For Proposition 2.7, the requirement $n \geq 2$ is dropped.

Then there must exist a point $\boldsymbol{R}'$ on the segment between $\widehat{\boldsymbol{R}}_S$ and $\boldsymbol{R}$ such that $\|\boldsymbol{y}-\boldsymbol{R}'\| < \|\boldsymbol{y}-\widehat{\boldsymbol{R}}_S\|$. Since both $\widehat{\boldsymbol{R}}_S$ and $\boldsymbol{R}$ belong to the (convex) isotonic cone $\{\boldsymbol{x} : x_{\pi^\star(1)} \geq \cdots \geq x_{\pi^\star(n)}\}$, the point $\boldsymbol{R}'$ must be in the isotonic cone as well. However, this contradicts the fact that $\widehat{\boldsymbol{R}}_S$ is the (unique) point of the isotonic cone with the minimum distance to $\boldsymbol{y}$.

□

A crucial point we wish to make here, however, is that there are good reasons to choose a fine-grained knowledge partition over a coarse one, provided that both are truthful. To show this point, we investigate how the accuracy depends on the coarseness of the knowledge partitions. This question is addressed by the following proposition.

Proposition 2.9. *Suppose that the noise vector $\boldsymbol{z}$ in the observation $\boldsymbol{y} = \boldsymbol{R} + \boldsymbol{z}$ consists of i.i.d. copies of normal random variables $\mathcal{N}(0, \sigma^2)$. Let S_1 and S_2 be two cones such that $S_2 \subset S_1$ and both contain the ground truth $\boldsymbol{R}$. Then, we have*

$$\limsup_{\sigma \to 0} \frac{\mathbb{E}\,\|\widehat{\boldsymbol{R}}_{S_2} - \boldsymbol{R}\|^2}{\mathbb{E}\,\|\widehat{\boldsymbol{R}}_{S_1} - \boldsymbol{R}\|^2} \leq 1, \quad \limsup_{\sigma \to \infty} \frac{\mathbb{E}\,\|\widehat{\boldsymbol{R}}_{S_2} - \boldsymbol{R}\|^2}{\mathbb{E}\,\|\widehat{\boldsymbol{R}}_{S_1} - \boldsymbol{R}\|^2} \leq 1.$$

The proof of this proposition is deferred to the Appendix.

This result is consistent with the intuition that a correct constraint with a smaller feasible region might lead to better estimation. We conjecture that Proposition 2.9 remains true for any noise level and leave it for future research.

3 The Isotonic Mechanism

Proposition 2.9 shows that a truthful knowledge partition would yield better estimation if it becomes finer-grained in certain regimes. Hence, the most promising prospect for a truthful mechanism is the most fine-grained knowledge partition induced by pairwise-comparison hyperplanes. This is just the aforementioned isotonic knowledge partition:

$$\{S_\pi : \pi \text{ is a permutation of } 1, 2, \ldots, n\}. \tag{3.1}$$

Letting π denote the ranking supplied by the owner, this mechanism asks the appraiser to solve the following optimization program:

$$\begin{aligned} \min_{\boldsymbol{r}} \quad & \|\boldsymbol{y} - \boldsymbol{r}\|^2 \\ \text{s.t.} \quad & \boldsymbol{r} \in S_\pi. \end{aligned} \tag{3.2}$$

This program is simply isotonic regression, hence the name the Isotonic Mechanism. The estimator of the ground truth can be obtained by projecting the raw grades $\boldsymbol{y}$ onto the isotonic cone S_π. A salient feature of this approach to information elicitation is that it is computationally tractable since (3.2) is a convex quadratic programming problem. Indeed, this optimization problem can be solved very efficiently by the pool adjacent violators algorithm [15, 3, 5].

The main result of this section shows that, under the Isotonic Mechanism, the optimal strategy for the owner is to report the ground-truth ranking $\pi^\star$, which satisfies

$$R_{\pi^\star(1)} \geq R_{\pi^\star(2)} \geq \cdots \geq R_{\pi^\star(n)}.$$

Under Assumption 2.1 for the present setting, it is sufficient for the owner to know the relative magnitudes of R_i's, which enables the truthful ranking of the items.

Theorem 2. *Under Assumptions 2.1, 2.2, and 2.3, the Isotonic Mechanism is truthful. That is, the owner can maximize her expected overall utility by truthfully reporting the ranking $\pi^\star$.*

Remark 3.1. Write $\widehat{\boldsymbol{R}}_\pi$ for the solution to (3.2). In the absence of ties in the true grades $\boldsymbol{R}$, if U is strictly convex, then honesty is strictly optimal in the sense that $\mathbb{E}\, U(\widehat{\boldsymbol{R}}_\pi) < \mathbb{E}\, U(\widehat{\boldsymbol{R}}_{\pi^\star})$ for any ranking π that is not identical to $\pi^\star$.

Remark 3.2. This theorem holds under a slightly more general version of Assumption 2.3. The utility function U can be any convex function and is not necessarily nondecreasing.

To get a handle on the truthfulness of the Isotonic Mechanism, it is instructive to recognize a crucial fact concerning isotonic regression: on top of the mean-preserving constraint

$$\sum_{i=1}^{n} \widehat{R}_{\pi,i} = \sum_{i=1}^{n} y_i,$$

loosely speaking, the solution to isotonic regression tends to exhibit less variability across its entries if an incorrect ranking is provided, as opposed to the true ranking [15]. Consequently, Jensen's inequality suggests that the overall *convex* utility $\sum_{i=1}^n U(\widehat{R}_{\pi,i})$ tends to be small in the case of an incorrect ranking.[7]

When moving to the noiseless setting, $\boldsymbol{y} = \boldsymbol{R}$, we can illustrate this point in a more concrete manner. In this setting, $\boldsymbol{R}$ is a feasible point to (3.2) when π is truthfully set to $\pi^\star$, and thus the optimal solution found by the Isotonic Mechanism is simply $\widehat{\boldsymbol{R}}_{\pi^\star} = \boldsymbol{R}$. In contrast, when $\pi \neq \pi^\star$, the pool adjacent violators algorithm would keep averaging over certain entries of the observation vector $\boldsymbol{y}$ until it obeys the (incorrect) ranking π. The averaging effect tends to lower $\sum_{i=1}^n U(\widehat{R}_{\pi,i})$ owing to the convex nature of the utility function U.

In passing, as with Theorem 1, Theorem 2 remains true if the squared ℓ_2 risk is replaced by the sum of Bregman divergences $\sum_{i=1}^n D_\phi(y_i, r_i)$, that is, the objective of (2.3). In fact, minimizing $\sum_{i=1}^n D_\phi(y_i, r_i)$ over the isotonic cone S_π leads to the same solution as the Isotonic Mechanism, no matter the choice of ϕ, as long as it is continuously differentiable and strictly convex [22].

3.1 Estimation properties

In addition to being truthful, we show that the Isotonic Mechanism improves estimation accuracy significantly, especially in the case of a large number of items and significant noise in the raw grades. Denote by

$$\mathrm{TV}(\boldsymbol{R}) := \inf_{\pi} \sum_{i=1}^{n-1} |R_{\pi(i)} - R_{\pi(i+1)}| = \sum_{i=1}^{n-1} |R_{\pi^\star(i)} - R_{\pi^\star(i+1)}| = R_{\pi^\star(1)} - R_{\pi^\star(n)}$$

the total variation of $\boldsymbol{R}$.

A more refined result on estimation on top of Proposition 2.8 is as follows.

Proposition 3.1. *Let $z_1, \ldots, z_n$ be i.i.d. normal random variables $\mathcal{N}(0, \sigma^2)$. For fixed $\sigma > 0$ and $V > 0$, the Isotonic Mechanism taking as input the ground-truth ranking $\pi^\star$ satisfies*

$$0.4096 + o_n(1) \le \frac{\sup_{\mathrm{TV}(\boldsymbol{R}) \le V} \mathbb{E} \left\| \widehat{\boldsymbol{R}}_{\pi^\star} - \boldsymbol{R} \right\|^2}{n^{\frac{1}{3}} \sigma^{\frac{4}{3}} V^{\frac{2}{3}}} \le 7.5625 + o_n(1),$$

where both $o_n(1)$ terms tend to 0 as $n \to \infty$.

[7] The proof of this theorem does not, however, use Jensen's inequality.

Remark 3.3. This proposition is adapted from an existing result on isotonic regression. See [33, 6].

This result says that the squared error risk of the Isotonic Mechanism is $O(n^{1/3}\sigma^{4/3})$. In contrast, for comparison, the risk of using the raw grades $\boldsymbol{y}$ is

$$\|\boldsymbol{y} - \boldsymbol{R}\|^2 = \mathbb{E}\left[\sum_{i=1}^{n} z_i^2\right] = n\sigma^2.$$

The ratio between the two risks is $O(n^{1/3}\sigma^{4/3})/(n\sigma^2) = O(n^{-2/3}\sigma^{-2/3})$. Therefore, the Isotonic Mechanism is especially favorable when both n and σ are large. While interpreting Proposition 3.1, however, it is important to notice that the total variation of the ground truth is fixed. Otherwise, when $R_{\pi^\star(i)} \gg R_{\pi^\star(i+1)}$ for all i, the solution of the Isotonic Mechanism is roughly the same as the raw-grade vector $\boldsymbol{y}$ because it satisfies the constraint $\boldsymbol{y} \in S_{\pi^\star}$ with high probability. Accordingly, the Isotonic Mechanism has a risk of about $n\sigma^2$ in this extreme case. That said, the Isotonic Mechanism in general is superior to using the raw grades, according to Proposition 2.8.

3.2 True-grade-dependent utility

The utility of an item might depend on its ground-truth grade. In light of this, we consider a relaxation of Assumption 2.3 by taking into account heterogeneity in the utility function.

Assumption 3.2. *Given estimates $\widehat{R}_1, \ldots, \widehat{R}_n$, the overall utility of the owner takes the form*

$$U(\widehat{\boldsymbol{R}}) := \sum_{i=1}^{n} U(\widehat{R}_i; R_i),$$

where $U(x; R)$ is convex in its first argument and satisfies

$$\frac{\mathrm{d}U(x; R)}{\mathrm{d}x} \geq \frac{\mathrm{d}U(x; R')}{\mathrm{d}x}$$

whenever $R > R'$.

The inequality in this assumption amounts to saying that the marginal utility increases with respect to the true grade of the item. For instance, an owner might prefer a high-quality item being rated higher over a low-quality item being rated higher.

An example of true-grade-dependent utility takes the form $U(x; R) = g(R)h(x)$, where $g \geq 0$ is nondecreasing and h is a nondecreasing convex function. Taking any nondecreasing $g_1, \ldots, g_L \geq 0$ and nondecreasing convex $h_1, \ldots, h_L$, more generally, the following function

$$U(x; R) = g_1(R)h_1(x) + g_2(R)h_2(x) + \cdots + g_L(R)h_L(x)$$

satisfies Assumption 3.2.

Theorem 2 remains true in the presence of heterogeneity in the owner's utility, as we show below.

Theorem 3. *Under Assumptions 2.1, 2.2, and 3.2, the Isotonic Mechanism is truthful.*

This theorem is proved in Section 6.3.

4 Incomplete knowledge

While we have determined perhaps the most important truthful knowledge partition, it is tempting to find other truthful pairwise-comparison-based partitions. From a practical viewpoint, another motivation for doing so is that the owner might not precisely know the ground-truth ranking, and only has partial knowledge of it.

To begin with, we present a counterexample to show that the converse of Theorem 1 is not true. Consider $\mathcal{S} = \{S_1, S_2\}$, where $S_1 = \{\boldsymbol{x} : x_1 \geq x_2 \geq \cdots \geq x_n\}$, $S_2 = \mathbb{R}^n \setminus S_1$, and $\boldsymbol{R} = (n\epsilon, (n-1)\epsilon, \ldots, 2\epsilon, \epsilon) \in S_1$ for some small $\epsilon > 0$. Note that S_1 and S_2 are separated by pairwise-comparison hyperplanes. Taking utility $U(x) = x^2$ or $\max\{x, 0\}^2$ and letting the noise terms $z_1, \ldots, z_n$ be i.i.d. standard normal random variables, we show in the Appendix that the owner would be better off reporting S_2 instead of S_1, the set that truly contains the ground truth. Thus, this pairwise-comparison-based knowledge partition is not truthful.

In the remainder of this section, we introduce two useful knowledge partitions and show their truthfulness.

Local ranking. Other than the isotonic knowledge partition (3.1), perhaps the simplest nontrivial truthful knowledge partitions are induced by local rankings: first partition $\{1, \ldots, n\}$ into several subsets of sizes, say, $n_1, n_2, \ldots, n_p$ such that $n_1 + \cdots + n_p = n$; then the owner is asked to provide a ranking of the n_q items indexed by each subset for $q = 1, \ldots, p$, but does not make any between-subset comparisons.

This is captured by the following practical scenario.

Scenario 4.1. *The owner leads a team of p subordinates. For $1 \leq q \leq p$, the q^{th} subordinate of her team produces n_q items and informs the owner of a ranking of the n_q items according to their values. However, no pairwise comparisons are provided to the owner between items made by different subordinates.*

Formally, letting $\mathcal{S}_{\text{IM}}(n)$ be a shorthand for the isotonic knowledge partition in n dimensions, we can write the resulting knowledge partition as

$$\mathcal{S}_{\text{IM}}(n_1) \times \mathcal{S}_{\text{IM}}(n_2) \times \cdots \times \mathcal{S}_{\text{IM}}(n_p),$$

which has a cardinality of $n_1!n_2!\cdots n_p!$. Recognizing that the overall utility is additively separable, we readily conclude that this knowledge partition is truthful and the owner will report the ground-truth local ranking for each subset.

Coarse ranking. Another example is induced by a coarse ranking: given $n_1, n_2, \ldots, n_p$ such that $n_1 + n_2 + \cdots + n_p = n$, the owner partitions $\{1, 2, \ldots, n\}$ into p ordered subsets $I_1, I_2, \ldots, I_p$ of sizes $n_1, n_2, \ldots, n_p$, respectively; but she does not reveal any comparisons within each subset at all. The appraiser wishes that the owner would report the ground-truth coarse ranking $(I_1^\star, I_2^\star, \ldots, I_p^\star)$, which satisfies

$$\boldsymbol{R}_{I_1^\star} \geq \boldsymbol{R}_{I_2^\star} \geq \cdots \geq \boldsymbol{R}_{I_p^\star} \tag{4.1}$$

(here we write $\boldsymbol{a} \geq \boldsymbol{b}$ for two vectors of possibly different lengths if any component of $\boldsymbol{a}$ is larger than or equal to any component of $\boldsymbol{b}$).

For instance, taking $n_q = 1$ for $q = 1, \ldots, p-1$ and $n_p = n - p + 1$, the owner is required to rank only the top $p - 1$ items. Another example is to consider $p = 10$ and $n_1 = \cdots = n_{10} = 0.1n$

(assume n is a multiple of 10), in which case the owner shall identify which items are the top 10%, which are the next top 10%, and so on.

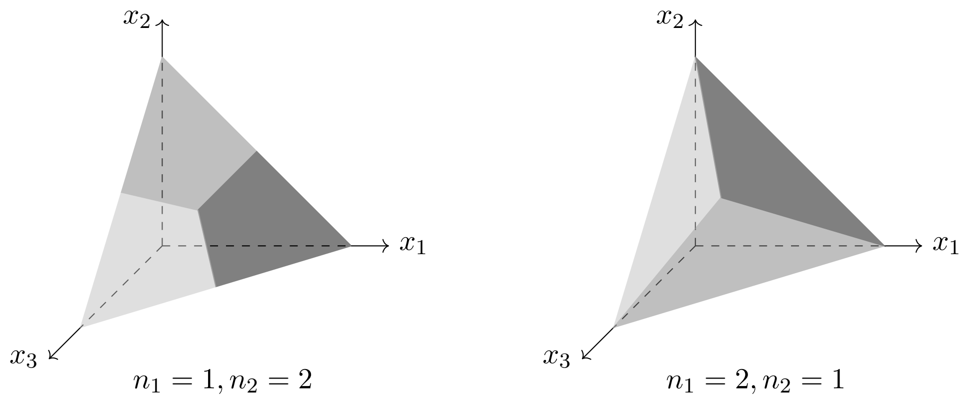

Figure 1: Knowledge partitions induced by coarse rankings in $n = 3$ dimensions. The illustration shows a slice of the partitions restricted to the nonnegative orthant.

Writing $\boldsymbol{I} := (I_1, \ldots, I_p)$, we denote by

$$S_{\boldsymbol{I}} := \{\boldsymbol{x} : \boldsymbol{x}_{I_1} \geq \boldsymbol{x}_{I_2} \geq \cdots \geq \boldsymbol{x}_{I_p}\}$$

the knowledge element indexed by $\boldsymbol{I}$. There are in total $\frac{n!}{n_1! \cdots n_p!}$ knowledge elements, which together form a knowledge partition. As is evident, any two adjacent knowledge elements are separated by pairwise-comparison hyperplanes. Figure 1 illustrates two such knowledge partitions in the case $n = 3$. The coarse ranking $\boldsymbol{I}$ of the owner's choosing may or may not be correct. Nevertheless, this is what the appraiser would incorporate into the estimation of the ground truth:

$$\begin{aligned} \min_{\boldsymbol{r}} \quad & \|\boldsymbol{y} - \boldsymbol{r}\|^2 \\ \text{s.t.} \quad & \boldsymbol{r} \in S_{\boldsymbol{I}}, \end{aligned} \tag{4.2}$$

which is a convex optimization program since the knowledge element $S_{\boldsymbol{I}}$ is convex. We call (4.2) a coarse Isotonic Mechanism.

A use case of this mechanism can be found in the following scenario.

Scenario 4.2. *The owner makes n_q products in grade q, for $q = 1, \ldots, p$. Products of different grades have significantly different values, but the owner cannot tell the difference between products of the same grade. The products are shuffled so that only the owner knows the grade information of each product.*

The following result shows that this new knowledge partition is truthful. Although this knowledge partition is pairwise-comparison-based, Theorem 4 does not follow from Theorem 1. Indeed, the proof of Theorem 4 given in Section 6.4 relies on some different ideas.

Theorem 4. *Under Assumptions 2.1, 2.2, and 2.3, the expected overall utility is maximized if the owner truthfully reports the coarse ranking that fulfills* (4.1).

Remark 4.1. Taking $n_1 = 1$ and $n_2 = n - 1$, Theorem 4 shows that the collection of knowledge elements taking the form (2.4) is a truthful knowledge partition.

One can construct other truthful knowledge partitions by integrating these two types of partitions. Instead of giving a complete ranking of items from each subset as in the local ranking setting, for example, one can provide a coarse ranking for each subset. It is evident that the resulting knowledge partition is truthful. An interesting problem for future investigation is to identify other truthful knowledge partitions based on these two prototypes.

We believe that the coarse Isotonic Mechanism gives inferior estimation performance compared with the vanilla Isotonic Mechanism. This conjecture would be true, for example, if one can prove Proposition 2.9 for any fixed distribution of the noise variables. We leave this for future research.

5 Extensions

In this section, we show that truthfulness continues to be the optimal strategy for the owner in more general settings.

Robustness to inconsistencies. The owner might not have complete knowledge of the true ranking in some scenarios, but is certain that some rankings are more consistent than others. More precisely, consider two rankings π_1 and π_2 such that neither is the ground-truth ranking, but the former can be obtained by swapping two entries of the latter in an upward manner in the sense that

$$R_{\pi_1(i)} = R_{\pi_2(j)} > R_{\pi_1(j)} = R_{\pi_2(i)}$$

for some $1 \le i < j \le n$ and $\pi_1(k) = \pi_2(k)$ for all $k \ne i, j$. In general, π_1 is said to be more consistent than π_2 if π_1 can be sequentially swapped from π_2 in an upward manner.

If the owner must choose between π_1 and π_2, she would be better off reporting the more consistent ranking, thereby being truthful in a relative sense. This shows the robustness of the Isotonic Mechanism against inconsistencies in rankings. A proof of this result is presented in the Appendix.

Proposition 5.1. *Suppose π_1 is more consistent than π_2 with respect to the ground truth $\boldsymbol{R}$. Under Assumptions 2.2 and 2.3, reporting π_1 yields higher or equal overall utility in expectation under the Isotonic Mechanism than reporting π_2.*

Intuitively, one might expect that a more consistent ranking would also lead to better estimation performance. If this intuition were true, it would lead to an extension of Proposition 2.8. We leave this interesting question to future research.

Multiple knowledge partitions. Given several truthful knowledge partitions, say, $\mathcal{S}_1, \ldots, \mathcal{S}_K$, one can offer the owner the freedom of choosing any knowledge element from these partitions. The resulting mechanism remains truthful. Formally, we have the following result.

Proposition 5.2. *Let $\mathcal{S}_1, \ldots, \mathcal{S}_K$ be truthful knowledge partitions. If the owner is required to report one knowledge element from any of these knowledge partitions, then she must be truthful in order to maximize her expected overall utility as much as possible.*

That is, if the owner chooses some knowledge element $S \in \mathcal{S}_k$ for $1 \le k \le K$ such that S does not contain the ground truth $\boldsymbol{R}$, she can always improve her overall utility in expectation by reporting the knowledge element in $\mathcal{S}_k$ truly containing $\boldsymbol{R}$. She can randomly pick a truthful knowledge element from any of $\mathcal{S}_1, \ldots, \mathcal{S}_K$ when it is unclear which knowledge partition leads to the highest overall utility. As Proposition 2.8 still holds in the case of multiple knowledge partitions, honesty would always lead to better estimation accuracy than using the raw observation as long as all knowledge elements are convex.

This result allows for certain flexibility in truthfully eliciting information, especially when we are not sure which knowledge partition satisfies Assumption 2.1. An immediate application is to take several knowledge partitions induced by coarse rankings (4.1) in the hope that, for at least one knowledge partition, the owner can determine the truthful knowledge element. For example, it seems plausible to take approximately equal sizes for the subsets:

$$n_1 \approx n_2 \approx \cdots \approx n_p \approx \frac{n}{p}.$$

However, the owner might not have sufficient knowledge about her items to provide the true coarse ranking, thereby violating Assumption 2.1. To circumvent this issue, we can let the owner pick any coarse ranking such that the number of subsets p is not smaller than, say, $\sqrt{n}$, and the largest subset size $\max_{1 \le i \le p} n_i$ is not greater than, say, $n/10$.

Nonseparable utility functions. The overall utility in Assumption 2.3 can be generalized to certain nonseparable functions. Explicitly, let the overall utility function $U(\boldsymbol{x})$ be symmetric in its n coordinates and satisfy

$$(x_i - x_j)\left(\frac{\partial U(\boldsymbol{x})}{\partial x_i} - \frac{\partial U(\boldsymbol{x})}{\partial x_j}\right) \ge 0 \tag{5.1}$$

for all $\boldsymbol{x}$. The following result shows that the owner's optimal strategy continues to be honesty.

Proposition 5.3. *Under Assumptions 2.1 and 2.2, the Isotonic Mechanism is truthful if the owner aims to maximize this type of overall utility in expectation.*

Theorem 2 is implied by this proposition when the (univariate) utility function U is differentiable. To see this point, note that a separable overall utility $U(\boldsymbol{x}) = U(x_1) + \cdots + U(x_n)$ in Assumption 2.3 satisfies

$$(x_i - x_j)\left(\frac{\partial U(\boldsymbol{x})}{\partial x_i} - \frac{\partial U(\boldsymbol{x})}{\partial x_j}\right) = (x_i - x_j)\left(U'(x_i) - U'(x_j)\right).$$

Since U' is a nondecreasing function, we get $(x_i - x_j)(U'(x_i) - U'(x_j)) \ge 0$.

On the other hand, the applicability of Proposition 5.3 is broader than that of Theorem 2 as there are symmetric functions that satisfy (5.1) but are not separable. A simple example is $U(\boldsymbol{x}) = \max\{x_1, x_2, \ldots, x_n\}$, and an owner with this overall utility is only concerned with the highest-rated item.[8] More generally, letting $x_{(1)} \ge x_{(2)} \ge \cdots \ge x_{(k)}$ be the $k \le n$ largest entries of $\boldsymbol{x}$, this proposition also applies to

$$U(\boldsymbol{x}) = h(x_{(1)}) + h(x_{(2)}) + \cdots + h(x_{(k)})$$

for any nondecreasing convex function h.

Proposition 5.3 follows from the proof of Theorem 2 in conjunction with Remark 6.5 in Section 6.3.

Multiple owners. An item can be shared by multiple owners while applying the Isotonic Mechanism. For example, a machine learning paper is often written by multiple authors. We introduce a variant of the Isotonic Mechanism that can tackle the case of multiple owners.

[8]This overall utility obeys (5.1), but partial derivatives shall be replaced by partial subdifferentials. This is also the case for the next example.

To set up the problem, imagine that there are n items and M owners. Let $\texttt{ind}_{ij} = 1$ if the i^{th} item is shared by the j^{th} owner for $1 \le i \le n$ and $1 \le j \le M$, and otherwise $\texttt{ind}_{ij} = 0$. This results in an $n \times M$ matrix that indicates the ownership information. Taking the ownership matrix as input, Algorithm 1 partitions the n items into several disjoint groups such that the items in each group are shared by the same owner and different groups correspond to different owners. The Isotonic Mechanism is invoked in each group. Owing to the independence between different groups, the Isotonic Mechanism is truthful across all groups in the partition.

Algorithm 1 Partition for the Isotonic Mechanism

Input: Set of items $\mathcal{I}$, set of owners $\mathcal{O}$, and the ownership matrix $\{\texttt{ind}_{ij}\}_{1\le i\le n, 1\le j\le M}$
while $\mathcal{I}$ is *not* empty **do**
  for Owner $o \in \mathcal{O}$ **do**
    Let w_o be the number of items in $\mathcal{I}$ shared by o
  Find the owner $o^\star$ such that $o^\star = \operatorname{argmax}_{o\in\mathcal{O}} w_o$ ▷ *Randomly choose $o^\star$ in the presence of ties*
  Apply the Isotonic Mechanism to $o^\star$ and the items in $\mathcal{I}$ shared by $o^\star$
  Reduce the sets: $O \leftarrow O \setminus \{o^\star\}$ and $\mathcal{I} \leftarrow \mathcal{I} \setminus \{\text{items shared by } o^\star\}$

In this algorithm, a plausible criterion is to prefer a partition with many large groups. In the case of papers and authors, this criterion is equivalent to giving priority to authors who submit a large number of papers. However, some groups may be singletons.

6 Proofs

Here, we prove Theorems 1, 2, 3, and 4. Proofs of other technical results in the paper are relegated to the Appendix.

6.1 Proof of Theorem 1

We prove this theorem in a slightly more general setting where (2.1) is replaced by (2.3). That is, the squared error loss is replaced by the sum of Bregman divergences. We start by introducing the following definition.

Definition 6.1 ([18])**.** We say that a vector $\boldsymbol{a} \in \mathbb{R}^n$ *weakly majorizes* another vector $\boldsymbol{b} \in \mathbb{R}^n$, denoted $\boldsymbol{a} \succeq_{\mathrm{w}} \boldsymbol{b}$, if

$$\sum_{i=1}^{k} a_{(i)} \ge \sum_{i=1}^{k} b_{(i)} \tag{6.1}$$

for all $1 \le k \le n$, where $a_{(1)} \ge \cdots \ge a_{(n)}$ and $b_{(1)} \ge \cdots \ge b_{(n)}$ are sorted in descending order from $\boldsymbol{a}$ and $\boldsymbol{b}$, respectively. If (6.1) reduces to an equality for $k = n$ while the rest $n-1$ inequalities remain the same, we say $\boldsymbol{a}$ *majorizes* $\boldsymbol{b}$ and write as $\boldsymbol{a} \succeq \boldsymbol{b}$.

The following lemma characterizes majorization via convex functions.

Lemma 6.2 (Hardy–Littlewood–Pólya inequality)**.** *Let $\boldsymbol{a}$ and $\boldsymbol{b}$ be two vectors in $\mathbb{R}^n$.*

(a) *The inequality*

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

holds for all nondecreasing convex functions h *if and only if* $\boldsymbol{a} \succeq_{\mathrm{w}} \boldsymbol{b}$.

(b) *The same inequality holds for all convex functions* h *if and only if* $\boldsymbol{a} \succeq \boldsymbol{b}$.

Remark 6.1. This is a well-known result in theory of majorization. For a proof of Lemma 6.2, see [18, 2]. For the proof of Theorem 1, however, only part (a) is needed. Part (b) will be used in the proofs of Theorems 2 and 4.

The following lemma is instrumental to the proof of Theorem 1. Its proof is presented later in this subsection.

Lemma 6.3. *Let* $\boldsymbol{a}$ *be a vector such that each element is different. There exists* $\delta > 0$ *such that if* $\boldsymbol{b}_1 + \boldsymbol{b}_2 = 2\boldsymbol{a}$ *and* $\|\boldsymbol{b}_1 - \boldsymbol{b}_2\| < \delta$, *then* $\boldsymbol{b}_1 \succeq_{\mathrm{w}} \boldsymbol{a}$ *and* $\boldsymbol{b}_2 \succeq_{\mathrm{w}} \boldsymbol{a}$ *cannot hold simultaneously unless* $\boldsymbol{b}_1 = \boldsymbol{b}_2 = \boldsymbol{a}$.

Proof of Theorem 1. Let S and S' be two neighboring knowledge elements in the knowledge partition $\mathcal{S}$. By assumption, the boundary between S and S' is a piecewise smooth surface. Pick an arbitrary point $\boldsymbol{x}$ on the boundary where the surface is locally smooth. Let $\epsilon > 0$ be small and $\boldsymbol{R} = \boldsymbol{x} + \epsilon\boldsymbol{v}$ and $\boldsymbol{R}' = \boldsymbol{x} - \epsilon\boldsymbol{v}$ for some unit-norm vector $\boldsymbol{v}$ that will be specified later. Assume without loss of generality that $\boldsymbol{R} \in S$ and $\boldsymbol{R}' \in S'$. For simplicity, we consider the noiseless setting where $\boldsymbol{y} = \boldsymbol{R}$ and $\boldsymbol{y}' = \boldsymbol{R}'$.

When the ground truth is $\boldsymbol{R}$, by assumption, the owner would truthfully report S as opposed to S'. Put differently, the overall utility by reporting S is higher than or equal to that by reporting S'. As is evident, the mechanism would output $\boldsymbol{y}$ if the owner reports S; if the owner reports S', then it would output the point, say, $\widehat{\boldsymbol{r}}$, that minimizes the sum of Bregman divergences $\sum_{i=1}^{n} D_\phi(y_i, r_i)$ over the boundary between S and S'. Assuming $\widehat{\boldsymbol{r}} = \boldsymbol{x} + o(\epsilon)$ for any sufficiently small ϵ as given for the moment, we get

$$U(\boldsymbol{x} + \epsilon\boldsymbol{v}) = U(\boldsymbol{y}) \geq U(\widehat{\boldsymbol{r}}) = U(\boldsymbol{x} + o(\epsilon))$$

for any nondecreasing convex function U. By Lemma 6.2, then, we must have $\boldsymbol{x} + \epsilon\boldsymbol{v} \succeq_{\mathrm{w}} \boldsymbol{x} + o(\epsilon)$, from which it follows that

$$\boldsymbol{x} + \epsilon\boldsymbol{v} \succeq_{\mathrm{w}} \boldsymbol{x}.$$

Likewise, we can deduce

$$\boldsymbol{x} - \epsilon\boldsymbol{v} \succeq_{\mathrm{w}} \boldsymbol{x}$$

from taking $\boldsymbol{R}'$ as the ground truth. If each element of $\boldsymbol{x}$ is different, Lemma 6.3 concludes that $\boldsymbol{v} = \boldsymbol{0}$ by taking $\boldsymbol{b}_1 = \boldsymbol{x} + \epsilon\boldsymbol{v}, \boldsymbol{b}_2 = \boldsymbol{x} - \epsilon\boldsymbol{v}, \boldsymbol{a} = \boldsymbol{x}$, and ϵ sufficiently small. This is a contradiction.

Therefore, $\boldsymbol{x}$ must have two entries, say, x_i and x_j, with the same value. As $\boldsymbol{x}$ can be an arbitrary point in the interior of any smooth surface of the boundary between S and S', this shows that this surface must be part of a pairwise-comparison hyperplane.

To finish the proof, we show that, by choosing an appropriate unit-norm vector $\boldsymbol{v}$, we will have $\widehat{\boldsymbol{r}} = \boldsymbol{x} + o(\epsilon)$ for sufficiently small ϵ. Note that

$$\sum_{i=1}^{n} D_\phi(y_i, r_i) = \frac{1}{2}(\boldsymbol{y} - \boldsymbol{r})^\top \boldsymbol{H}_\phi(\boldsymbol{r})(\boldsymbol{y} - \boldsymbol{r}) + o(\|\boldsymbol{y} - \boldsymbol{r}\|^2),$$

where $\boldsymbol{H}_\phi(\boldsymbol{r})$ is a diagonal matrix consisting of $\phi''(r_i)$ on its diagonal for $i = 1, \ldots, n$. Owing to the twice continuous differentiability of ϕ, this diagonal Hessian $\boldsymbol{H}_\phi(\boldsymbol{r}) = \boldsymbol{H}_\phi(\boldsymbol{x}) + o(1)$ when $\boldsymbol{r}$ is close to $\boldsymbol{x}$. Recognizing that $\boldsymbol{y} = \boldsymbol{R} = \boldsymbol{x} + \epsilon\boldsymbol{v}$ is close to the boundary when ϵ is sufficiently small, $\widehat{\boldsymbol{r}}$ is the projection of $\boldsymbol{y}$ onto the tangent plane at $\boldsymbol{x}$ under the $\boldsymbol{H}_\phi(\boldsymbol{x})^{-1}$-Mahalanobis distance, up to low-order terms. As such, it suffices to let $\boldsymbol{v}$ be a normal vector to the tangent plane at $\boldsymbol{x}$ under this Mahalanobis distance.

□

Remark 6.2. The proof proceeds by taking the zero noise level. An interesting question for future investigation is to derive a possibly different necessary condition for honesty under the assumption of a nonzero noise level.

We conclude this subsection by proving Lemma 6.3.

Proof of Lemma 6.3. Write $\boldsymbol{\nu} = \boldsymbol{b}_1 - \boldsymbol{a}$, which satisfies $\|\boldsymbol{\nu}\| < \delta/2$. Since $\boldsymbol{a}$ has no ties, both $\boldsymbol{b}_1$ and $\boldsymbol{b}_2$ would have the same ranking as $\boldsymbol{a}$ for sufficiently small δ. Without loss of generality, letting $a_1 \ge a_2 \ge \cdots \ge a_n$, we have $a_1 + \nu_1 \ge a_2 + \nu_2 \ge \cdots \ge a_n + \nu_n$ as well as $a_1 - \nu_1 \ge a_2 - \nu_2 \ge \cdots \ge a_n - \nu_n$.

Assume that both $\boldsymbol{b}_1 \succeq_{\mathrm{w}} \boldsymbol{a}$ and $\boldsymbol{b}_2 \succeq_{\mathrm{w}} \boldsymbol{a}$. By the definition of weak majorization, this immediately gives

$$\nu_1 \ge 0, \nu_1 + \nu_2 \ge 0, \ldots, \nu_1 + \cdots + \nu_n \ge 0$$

and

$$\nu_1 \le 0, \nu_1 + \nu_2 \le 0, \ldots, \nu_1 + \cdots + \nu_n \le 0.$$

Taken together, these two displays show that $\nu_1 = \nu_2 = \cdots = \nu_n = 0$. As such, the only possibility is that $\boldsymbol{b}_1 = \boldsymbol{b}_2 = \boldsymbol{a}$.

□

6.2 Proof of Theorem 2

The following definition and lemma will be used in the proof of this theorem.

Definition 6.4. We say that a vector $\boldsymbol{a} \in \mathbb{R}^n$ *majorizes* $\boldsymbol{b} \in \mathbb{R}^n$ in the *natural order*, denoted $\boldsymbol{a} \succeq_{\mathrm{no}} b$, if

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

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

A departure of this definition from weak majorization or majorization is that majorization in the natural order is not invariant under permutations.

In the lemma below, we write $\boldsymbol{a}^+$ as a shorthand for the projection of $\boldsymbol{a}$ onto the standard isotonic cone $\{\boldsymbol{x} : x_1 \ge x_2 \ge \cdots \ge x_n\}$. A proof of this lemma is given in Section 6.2.1.

Lemma 6.5. *If $\boldsymbol{a} \succeq_{\mathrm{no}} \boldsymbol{b}$, then we have $\boldsymbol{a}^+ \succeq_{\mathrm{no}} \boldsymbol{b}^+$.*

Remark 6.3. Both $\boldsymbol{a}^+$ and $\boldsymbol{b}^+$ have already been ordered from the largest to the smallest. Hence, majorization in the natural order ($\succeq_{\mathrm{no}}$) is identical to majorization ($\succeq$) in the case of $\boldsymbol{a}^+$ and $\boldsymbol{b}^+$.

Now we are in a position to prove Theorem 2. Write $\pi \circ \boldsymbol{a} := (a_{\pi(1)}, a_{\pi(2)}, \ldots, a_{\pi(n)})$ for $\boldsymbol{a}$ under permutation π.

Proof of Theorem 2. Assume without loss of generality that $R_1 \geq R_2 \geq \cdots \geq R_n$. In this case, the ground-truth ranking $\pi^\star$ is the identity, that is, $\pi^\star(i) = i$ for all i, and the optimization program (3.2) for the Isotonic Mechanism is

$$\begin{aligned} \min \ & \|\boldsymbol{y} - \boldsymbol{r}\|^2 \\ \text{s.t.} \ & r_1 \geq r_2 \geq \cdots \geq r_n. \end{aligned}$$

Its solution is the projection of $\boldsymbol{y}$ onto the isotonic cone $\{\boldsymbol{x} : x_1 \geq x_2 \geq \cdots \geq x_n\}$, that is, $\boldsymbol{y}^+ = (\boldsymbol{R} + \boldsymbol{z})^+$.

Consider the optimization program with a different ranking π,

$$\begin{aligned} \min \ & \|\boldsymbol{y} - \boldsymbol{r}\|^2 \\ \text{s.t.} \ & r_{\pi(1)} \geq r_{\pi(2)} \geq \cdots \geq r_{\pi(n)}. \end{aligned} \tag{6.2}$$

This is equivalent to

$$\begin{aligned} \min \ & \|\pi \circ \boldsymbol{y} - \tilde{\boldsymbol{r}}\|^2 \\ \text{s.t.} \ & \tilde{r}_1 \geq \tilde{r}_2 \geq \cdots \geq \tilde{r}_n, \end{aligned}$$

with the relationship $\tilde{\boldsymbol{r}} = \pi \circ \boldsymbol{r}$. From this equivalence it is easy to see that the solution to (6.2) can be written as $\pi^{-1} \circ (\pi \circ \boldsymbol{y})^+ = \pi^{-1} \circ (\pi \circ \boldsymbol{R} + \pi \circ \boldsymbol{z})^+$. It suffices to show that the overall utility obeys

$$\mathbb{E}\, U\left((\boldsymbol{R} + \boldsymbol{z})^+\right) \geq \mathbb{E}\, U\left(\pi^{-1} \circ (\pi \circ \boldsymbol{R} + \pi \circ \boldsymbol{z})^+\right) = \mathbb{E}\, U\left((\pi \circ \boldsymbol{R} + \pi \circ \boldsymbol{z})^+\right),$$

where the equality follows because the overall utility is invariant under permutations. Under Assumption 2.2, the entries $z_1, \ldots, z_n$ of $\boldsymbol{z}$ are exchangeable random variables. This gives

$$\mathbb{E}\, U\left((\pi \circ \boldsymbol{R} + \pi \circ \boldsymbol{z})^+\right) = \mathbb{E}\, U\left((\pi \circ \boldsymbol{R} + \boldsymbol{z})^+\right).$$

Thus, the proof is complete if we prove

$$\mathbb{E}\, U\left((\boldsymbol{R} + \boldsymbol{z})^+\right) \geq \mathbb{E}\, U\left((\pi \circ \boldsymbol{R} + \boldsymbol{z})^+\right). \tag{6.3}$$

To prove (6.3), we utilize the following crucial fact

$$\boldsymbol{R} + \boldsymbol{z} \succeq_{\text{no}} \pi \circ \boldsymbol{R} + \boldsymbol{z}.$$

This holds because $R_1, \ldots, R_n$ are already in descending order. Therefore, it merely follows from Lemma 6.5 that

$$(\boldsymbol{R} + \boldsymbol{z})^+ \succeq_{\text{no}} (\pi \circ \boldsymbol{R} + \boldsymbol{z})^+$$

or, equivalently,

$$(\boldsymbol{R} + \boldsymbol{z})^+ \succeq (\pi \circ \boldsymbol{R} + \boldsymbol{z})^+.$$

By Lemma 6.2, we get

$$\sum_{i=1}^n U\left((\boldsymbol{R} + \boldsymbol{z})_i^+\right) \geq \sum_{i=1}^n U\left((\pi \circ \boldsymbol{R} + \boldsymbol{z})_i^+\right)$$

for any convex function U, which implies (6.3). This completes the proof.

$\square$

6.2.1 Proof of Lemma 6.5

Definition 6.6. We say that $\boldsymbol{c}^1$ is an upward transport of $\boldsymbol{c}^2$ if there exists $1 \le i < j \le n$ such that $c_k^1 = c_k^2$ for all $k \ne i, j$, $c_i^1 + c_j^1 = c_i^2 + c_j^2$, and $c_i^1 \ge c_i^2$.

Equivalently, $\boldsymbol{c}^1$ is an upward transport of $\boldsymbol{c}^2$ if $\boldsymbol{c}^1$ can be obtained by moving some "mass" from an entry of $\boldsymbol{c}^2$ to an earlier entry. As is evident, we have $\boldsymbol{c}^1 \succeq_{\text{no}} \boldsymbol{c}^2$ if $\boldsymbol{c}^1$ is an upward transport of $\boldsymbol{c}^2$.

The following lemmas state two useful properties of this relationship between two vectors.

Lemma 6.7. *Let $\boldsymbol{a} \succeq_{\text{no}} \boldsymbol{b}$. Then there exists a positive integer L and $\boldsymbol{c}^1, \ldots, \boldsymbol{c}^L$ such that $\boldsymbol{c}^1 = \boldsymbol{a}, \boldsymbol{c}^L = \boldsymbol{b}$, and $\boldsymbol{c}^l$ is an upward transport of $\boldsymbol{c}^{l+1}$ for $1 \le l \le L-1$.*

Next, recall that $\boldsymbol{a}^+$ denotes the projection of $\boldsymbol{a}$ onto the standard isotonic cone $\{\boldsymbol{x} : x_1 \ge x_2 \ge \cdots \ge x_n\}$.

Lemma 6.8. *Let $\boldsymbol{a}$ be an upward transport of $\boldsymbol{b}$. Then, we have $\boldsymbol{a}^+ \succeq_{\text{no}} \boldsymbol{b}^+$.*

Lemma 6.5 readily follows from Lemmas 6.7 and 6.8. To see this point, for any $\boldsymbol{a}, \boldsymbol{b}$ satisfying $\boldsymbol{a} \succeq_{\text{no}} \boldsymbol{b}$, note that by Lemma 6.7 we can find $\boldsymbol{c}^1 = \boldsymbol{a}, \boldsymbol{c}^2, \ldots, \boldsymbol{c}^{L-1}, \boldsymbol{c}^L = \boldsymbol{b}$ such that $\boldsymbol{c}^l$ is an upward transport of $\boldsymbol{c}^{l+1}$ for $l = 1, \ldots, L-1$. Next, Lemma 6.8 asserts that

$$(\boldsymbol{c}^l)^+ \succeq_{\text{no}} (\boldsymbol{c}^{l+1})^+$$

for all $l = 1, \ldots, L-1$. Owing to the transitivity of majorization in the natural order, we conclude that $\boldsymbol{a} \succeq_{\text{no}} \boldsymbol{b}$, thereby proving Lemma 6.5.

Below, we prove Lemmas 6.7 and 6.8.

Proof of Lemma 6.7. We prove by induction. The base case $n = 1$ is clearly true. Suppose this lemma is true for n.

Now we prove the lemma for the case $n+1$. Let $\boldsymbol{c}^1 = \boldsymbol{a} = (a_1, a_2, \ldots, a_{n+1})$ and $\boldsymbol{c}^2 := (b_1, a_1 + a_2 - b_1, a_3, a_4, \ldots, a_{n+1})$. As is evident, $\boldsymbol{c}^1$ is an upward transport of $\boldsymbol{c}^2$.

Now we consider operations on the components except for the first one. Let $\boldsymbol{a}' := (a_1 + a_2 - b_1, a_3, a_4, \ldots, a_{n+1})$ and $\boldsymbol{b}' := (b_2, \ldots, b_{n+1})$ be derived by removing the first component of $\boldsymbol{c}^2$ and $\boldsymbol{b}$, respectively. These two vectors obey $\boldsymbol{a}' \succeq_{\text{no}} \boldsymbol{b}'$. To see this, note that $a_1' = a_1 + a_2 - b_1 \ge b_1 + b_2 - b_1 = b_2 = b_1'$, and

$$a_1' + \cdots + a_k' = (a_1 + a_2 - b_1) + a_3 + \cdots + a_{k+1} = \sum_{i=1}^{k+1} a_i - b_1 \ge \sum_{i=1}^{k+1} b_i - b_1 = \sum_{i=2}^{k+1} b_i = b_1' + \cdots + b_k'$$

for $2 \le k \le n-1$. Moreover, it also holds that

$$a_1' + \cdots + a_n' = (a_1 + a_2 - b_1) + a_3 + \cdots + a_{n+1} = \sum_{i=1}^{n+1} a_i - b_1 = \sum_{i=1}^{n+1} b_i - b_1 = b_1' + \cdots + b_n'.$$

Thus, by induction, there must exist $\boldsymbol{c}'^1, \ldots, \boldsymbol{c}'^L$ such that $\boldsymbol{c}'^1 = \boldsymbol{a}', \boldsymbol{c}'^L = \boldsymbol{b}'$, and $\boldsymbol{c}'^l$ is an upward transport of $\boldsymbol{c}'^{l+1}$ for $l = 1, \ldots, L-1$. We finish the proof for $n+1$ by recognizing that $\boldsymbol{c}^1 \equiv \boldsymbol{a}, (b_1, \boldsymbol{c}'^1), (b_1, \boldsymbol{c}'^2), \ldots, (b_1, \boldsymbol{c}'^L) \equiv \boldsymbol{b}$ satisfy the conclusion of this lemma.

□

To prove Lemma 6.8, we need the following two lemmas. We relegate the proofs of these two lemmas to the Appendix. Denote by $\mathbf{e}_i$ the i^{th} canonical-basis vector in $\mathbb{R}^n$.

Lemma 6.9. *For any $\delta > 0$ and $i = 1, \ldots, n$, we have $(\boldsymbol{a} + \delta\mathbf{e}_i)^+ \geq \boldsymbol{a}^+$ in the component-wise sense.*

Remark 6.4. Likewise, the proof of Lemma 6.9 reveals that $(\boldsymbol{a}-\delta\mathbf{e}_i)^+ \leq \boldsymbol{a}^+$. As an aside, recognizing the mean-preserving constraint of isotonic regression, we have $\mathbf{1}^\top(\boldsymbol{a} + \delta\mathbf{e}_i)^+ = \mathbf{1}^\top\boldsymbol{a}^+ + \delta$, where $\mathbf{1} \in \mathbb{R}^n$ denotes the ones vector.

Lemma 6.10. *Denote by $\bar{a}$ the sample mean of $\boldsymbol{a}$. Then $\boldsymbol{a}^+$ has constant entries—that is, $a_1^+ = \cdots = a_n^+$—if and only if*

$$\frac{a_1 + \cdots + a_k}{k} \leq \bar{a}$$

for all $k = 1, \ldots, n$.

Proof of Lemma 6.8. Let $1 \leq i < j \leq n$ be the indices such that $a_i + a_j = b_i + b_j$ and $a_i \geq b_i$. Write $\delta := a_i - b_i \geq 0$. Then, $\boldsymbol{b} = \boldsymbol{a} - \delta\mathbf{e}_i + \delta\mathbf{e}_j$. If $\delta = 0$, then $\boldsymbol{a}^+ = \boldsymbol{b}^+$ because $\boldsymbol{a} = \boldsymbol{b}$, in which case the lemma holds trivially. In the remainder of the proof, we focus on the nontrivial case $\delta > 0$.

The lemma amounts to saying that $\boldsymbol{a}^+ \succeq_{\text{no}} (\boldsymbol{a} - \delta\mathbf{e}_i + \delta\mathbf{e}_j)^+$ for all $\delta > 0$. Owing to the continuity of the projection, it is sufficient to prove the following statement: there exists $\delta_0 > 0$ (depending on $\boldsymbol{a}$) such that $\boldsymbol{a}^+ \succeq_{\text{no}} (\boldsymbol{a} - \delta\mathbf{e}_i + \delta\mathbf{e}_j)^+$.

Let I be the set of indices where the entries of $\boldsymbol{a}^+$ has the same value as i: $I = \{k : a_k^+ = a_i^+\}$. Likewise, define $J = \{k : a_k^+ = a_j^+\}$. There are exactly two cases, namely, $I = J$ and $I \cap J = \emptyset$, which we discuss in the sequel.

Case 1: $I = J$. For convenience, write $I = \{i_1, i_1 + 1, \ldots, i_2\}$. By Lemma 6.10, we have

$$\frac{a_{i_1} + a_{i_1+1} + \ldots + a_{i_1+l-1}}{l} \leq \bar{a}_I := \frac{a_{i_1} + x_{i_1+1} + \ldots + x_{i_2}}{i_2 - i_1 + 1}$$

for $l = 1, \ldots, i_2 - i_1 + 1$.

Now we consider $\boldsymbol{b} = \boldsymbol{a} - \delta\mathbf{e}_i + \delta\mathbf{e}_j$ restricted to I. Assume that δ is sufficiently small so that the constant pieces of $\boldsymbol{b}^+$ before and after I are the same as those of $\boldsymbol{a}^+$. Since $i_1 \leq i < j \leq i_2$, we have

$$b_{i_1} + b_{i_1+1} + \ldots + b_{i_2} = a_{i_1} + a_{i_1+1} + \ldots + a_{i_2}.$$

On the other hand, we have

$$b_{i_1} + b_{i_1+1} + \ldots + b_{i_1+l-1} \leq a_{i_1} + a_{i_1+1} + \ldots + a_{i_1+l-1}$$

since the index i comes earlier than j. Taken together, these observations give

$$\frac{b_{i_1} + b_{i_1+1} + \ldots + b_{i_1+l-1}}{l} \leq \frac{b_{i_1} + b_{i_1+1} + \ldots + b_{i_2}}{i_2 - i_1 + 1}$$

for all $l = 1, \ldots, i_2 - i_1 + 1$. It follows from Lemma 6.10 that the projection $\boldsymbol{b}^+ = (\boldsymbol{a} - \delta\mathbf{e}_i + \delta\mathbf{e}_j)^+$ remains constant on the set I and this value is the same as $\boldsymbol{a}^+$ on I since $b_{i_1} + b_{i_1+1} + \ldots + b_{i_2} = a_{i_1} + a_{i_1+1} + \ldots + a_{i_2}$. That is, we have $\boldsymbol{b}^+ = \boldsymbol{a}^+$ in this case.

Case 2: $I \cap J = \emptyset$. As earlier, let δ be sufficiently small. Write $I = \{i_1, i_1 + 1, \ldots, i_2\}$ and $J = \{j_1, j_1 + 1, \ldots, j_2\}$, where $i_2 < j_1$. Since the isotonic constraint is inactive between the $(i_1 - 1)^{\text{th}}$ and i_1^{th} components, the projection $\boldsymbol{a}_I^+$ restricted to I is the same as projecting $\boldsymbol{a}_I$ onto the $|I| = (i_2 - i_1 + 1)$-dimensional standard isotonic cone. As δ is sufficiently small, the projection $(\boldsymbol{a} - \delta\mathbf{e}_i + \delta\mathbf{e}_j)_I^+$ restricted to I is also the same as projecting $(\boldsymbol{a} - \delta\mathbf{e}_i + \delta\mathbf{e}_j)_I$ onto the $|I| = (i_2 - i_1 + 1)$-dimensional standard isotonic cone.

However, since $i \in I$ but $j \notin J$, we see that $(\boldsymbol{a} - \delta\mathbf{e}_i + \delta\mathbf{e}_j)_I = \boldsymbol{a}_I - \delta\mathbf{e}_i$, where $\mathbf{e}_i$ now should be regarded as the $(i - i_1 + 1)^{\text{th}}$ canonical-basis vector in the reduced $(i_2 - i_1 + 1)$-dimensional space. Then, by Lemma 6.9 and Remark 6.4, we see that

$$\boldsymbol{b}_I^+ = (\boldsymbol{a}_I - \delta\mathbf{e}_i)^+ \le \boldsymbol{a}_I^+$$

in the component-wise sense, which, together with the fact that $b_l^+ = a_l^+$ for $l \in \{1, \ldots, i_1 - 1\} \cup \{i_2 + 1, \ldots, j_1 - 1\} \cup \{j_2 + 1, \ldots, n\}$, gives

$$b_1^+ + \cdots + b_l^+ \le a_1^+ + \cdots + a_l^+$$

for all $l = 1, \ldots, j_1 - 1$. Moreover,

$$\begin{aligned} b_1^+ + \cdots + b_l^+ - (a_1^+ + \cdots + a_l^+) &= b_{i_1}^+ + \cdots + b_{i_2}^+ - (a_{i_1}^+ + \cdots + a_{i_2}^+) \\ &= b_{i_1} + \cdots + b_{i_2} - (a_{i_1} + \cdots + a_{i_2}) \\ &= -\delta \end{aligned} \tag{6.4}$$

when $i_2 + 1 \le l \le j_1 - 1$.

Now we turn to the case $j_1 \le l \le j_2$. As earlier, for sufficiently small δ, the projection $(\boldsymbol{a} - \delta\mathbf{e}_i + \delta\mathbf{e}_j)_J^+$ restricted to J is the same as projecting $(\boldsymbol{a} - \delta\mathbf{e}_i + \delta\mathbf{e}_j)_J$ onto the $|J| = (j_2 - j_1 + 1)$-dimensional standard isotonic cone. Then, since $\boldsymbol{b}_J = (\boldsymbol{a} - \delta\mathbf{e}_i + \delta\mathbf{e}_j)_J = \boldsymbol{a}_J + \delta\mathbf{e}_j$, it follows from Lemma 6.9 that

$$\boldsymbol{b}_J^+ \ge \boldsymbol{a}_J^+, \tag{6.5}$$

and meanwhile, we have

$$b_{j_1}^+ + \cdots + b_{j_2}^+ - (a_{j_1}^+ + \cdots + a_{j_2}^+) = b_{j_1} + \cdots + b_{j_2} - (a_{j_1} + \cdots + a_{j_2}) = \delta. \tag{6.6}$$

Thus, for any $j_1 \le l \le j_2$, (6.5) and (6.6) give

$$b_{j_1}^+ + \cdots + b_l^+ - (a_{j_1}^+ + \cdots + a_l^+) \le b_{j_1}^+ + \cdots + b_{j_2}^+ - (a_{j_1}^+ + \cdots + a_{j_2}^+) = \delta.$$

Therefore, we get

$$\begin{aligned} &b_1^+ + \cdots + b_l^+ - (a_1^+ + \cdots + a_l^+) \\ &= b_1^+ + \cdots + b_{j_1-1}^+ - (a_1^+ + \cdots + a_{j_1-1}^+) + b_{j_1}^+ + \cdots + b_l^+ - (a_{j_1}^+ + \cdots + a_l^+) \\ &= -\delta + b_{j_1}^+ + \cdots + b_l^+ - (a_{j_1}^+ + \cdots + a_l^+) \\ &\le -\delta + \delta \\ &= 0, \end{aligned}$$

where the second equality follows from (6.4).

Taken together, the results above show that

$$b_1^+ + \cdots + b_l^+ \le a_1^+ + \cdots + a_l^+$$

for $1 \le l \le j_2$, with equality when $l \le i_1 - 1$ or $l = j_2$. In addition, this inequality remains true—in fact, reduced to equality—when $l > j_2$. This completes the proof.

□

6.3 Proof of Theorem 3

Define

$$\widetilde{U}(\boldsymbol{x}) = \sum_{i=1}^{n} U(x_i; R_{\rho(i)}), \tag{6.7}$$

where ρ is a permutation such that $\boldsymbol{x}$ and $\boldsymbol{R}_\rho$ have the same descending order. For example, if x_l is the largest element of $\boldsymbol{x}$, so is $R_{\rho(l)}$ the largest element of $\boldsymbol{R}_\rho$. By construction, $\widetilde{U}$ is symmetric. Moreover, this function satisfies the following two lemmas. The proofs are given later in this subsection.

Lemma 6.11. *Under Assumption 3.2, the overall utility satisfies*

$$\sum_{i=1}^{n} U(x_i; R_i) \le \widetilde{U}(\boldsymbol{x}).$$

Lemma 6.12. *Under Assumption 3.2, the function $\widetilde{U}$ is Schur-convex in the sense that $\widetilde{U}(\boldsymbol{a}) \ge \widetilde{U}(\boldsymbol{b})$ is implied by $\boldsymbol{a} \succeq \boldsymbol{b}$.*

Now we are ready to prove Theorem 3.

Proof of Theorem 3. Assume without loss of generality that $R_1 \ge R_2 \ge \cdots \ge R_n$. Denote by $\widehat{\boldsymbol{R}}_\pi$ the output of the Isotonic Mechanism provided ranking π. For simplicity, write $\widehat{\boldsymbol{R}} = \widehat{\boldsymbol{R}}_{\pi^\star}$ when the ranking is the true ranking $\pi^\star$. Note that $\pi^\star(i) = i$ for all i, and $\widehat{\boldsymbol{R}}$ and $\boldsymbol{R}$ have the same descending order. As such, we get

$$\widetilde{U}(\widehat{\boldsymbol{R}}) = \sum_{i=1}^{n} U(\widehat{R}_i; R_i).$$

To prove

$$\mathbb{E}\left[\sum_{i=1}^{n} U(\widehat{R}_i; R_i)\right] \ge \mathbb{E}\left[\sum_{i=1}^{n} U(\widehat{R}_{\pi,i}; R_i)\right],$$

we start by observing that

$$\widetilde{U}(\widehat{\boldsymbol{R}}_\pi) \ge \sum_{i=1}^{n} U(\widehat{R}_{\pi,i}; R_i)$$

is an immediate consequence of Lemma 6.11. Hence, it is sufficient to prove

$$\mathbb{E}\,\widetilde{U}(\widehat{\boldsymbol{R}}) \ge \mathbb{E}\,\widetilde{U}(\widehat{\boldsymbol{R}}_\pi). \tag{6.8}$$

As in the proof of Theorem 2, it follows from Lemma 6.5 that

$$\widehat{\boldsymbol{R}} = (\boldsymbol{R} + \boldsymbol{z})^+ \succeq (\pi \circ \boldsymbol{R} + \boldsymbol{z})^+.$$

As Lemma 6.12 ensures that $\widetilde{U}$ is Schur-convex, the majorization relation above gives

$$\widetilde{U}(\widehat{\boldsymbol{R}}) \geq \widetilde{U}((\pi \circ \boldsymbol{R} + \boldsymbol{z})^+). \tag{6.9}$$

Moreover, the coupling argument in the proof of Theorem 2 implies that $(\pi \circ \boldsymbol{R} + \boldsymbol{z})^+$ has the same probability distribution as $\widehat{\boldsymbol{R}}_\pi$, which gives

$$\mathbb{E}\, \widetilde{U}((\pi \circ \boldsymbol{R} + \boldsymbol{z})^+) = \mathbb{E}\, \widetilde{U}(\widehat{\boldsymbol{R}}_\pi).$$

Together with (6.9), this equality implies (6.8).

□

Next, we turn to the proof of Lemma 6.11.

Proof of Lemma 6.11. Given two permutations π_1 and π_2, if there exist two indices i, j such that $\pi_1(k) = \pi_2(k)$ for all $k \neq i, j$ and $R_{\pi_1(i)} - R_{\pi_1(j)} = -(R_{\pi_2(i)} - R_{\pi_2(j)})$ has the same sign as $x_i - x_j$, we say that π_1 is an upward swap of π_2 with respect to $\boldsymbol{x}$. As is evident, the permutation ρ in (6.7) can be obtained by sequentially swapping the identity permutation in an upward manner with respect to $\boldsymbol{x}$. Therefore, it suffices to prove the lemma in the case $n = 2$. Specifically, we only need to prove that

$$U(x_1; R_1) + U(x_2; R_2) \leq U(x_1; R_2) + U(x_2; R_1) \tag{6.10}$$

if $x_1 \geq x_2$ and $R_1 \leq R_2$.

Define

$$g(x) = U(x; R_2) - U(x_2; R_2) - U(x; R_1) + U(x_2; R_1).$$

Then, (6.10) is equivalent to $g(x) \geq 0$ for $x \geq x_2$. To prove this, observe that

$$g'(x) = \frac{\mathrm{d}U(x; R_2)}{\mathrm{d}x} - \frac{\mathrm{d}U(x; R_1)}{\mathrm{d}x} \geq 0$$

by Assumption 3.2. This establishes (6.10), thereby completing the proof.

□

Next, we turn to the proof of Lemma 6.12, for which we need the following lemma. For a proof of this lemma, see [18].

Lemma 6.13 (Schur–Ostrowski criterion)**.** *If a function $f : \mathbb{R}^n \to \mathbb{R}$ is differentiable. Then f is Schur-convex if and only if it is symmetric and satisfies*

$$(x_i - x_j)\left(\frac{\partial f}{\partial x_i} - \frac{\partial f}{\partial x_j}\right) \geq 0$$

for all $1 \leq i \neq j \leq n$.

Remark 6.5. The condition on the overall utility in Proposition 5.3 is precisely Schur-convexity. Thus, Proposition 5.3 follows from the proof of Theorem 3.

Proof of Lemma 6.12. First, consider the case where all elements of $\boldsymbol{x}$ are different from each other. Without loss of generality, assume $x_i > x_j$. It suffices to prove that

$$\frac{\partial \widetilde{U}(\boldsymbol{x})}{\partial x_i} - \frac{\partial \widetilde{U}(\boldsymbol{x})}{\partial x_j} = \frac{\mathrm{d}U(x; R_{\rho(i)})}{\mathrm{d}x}\Big|_{x=x_i} - \frac{\mathrm{d}U(x; R_{\rho(j)})}{\mathrm{d}x}\Big|_{x=x_j} \geq 0. \tag{6.11}$$

Since $U(x; R_{\rho(i)})$ is a convex function in x, we have

$$\frac{\mathrm{d}U(x; R_{\rho(i)})}{\mathrm{d}x}\Big|_{x=x_i} - \frac{\mathrm{d}U(x; R_{\rho(i)})}{\mathrm{d}x}\Big|_{x=x_j} \geq 0$$

as the derivative of a convex function is a nondecreasing function. Next, recognizing that $R_{\rho(i)} \geq R_{\rho(j)}$ is implied by the construction of the permutation ρ, it follows from Assumption 3.2 that

$$\frac{\mathrm{d}U(x; R_{\rho(i)})}{\mathrm{d}x}\Big|_{x=x_j} - \frac{\mathrm{d}U(x; R_{\rho(j)})}{\mathrm{d}x}\Big|_{x=x_j} \geq 0.$$

Adding the last two inequalities, we arrive at (6.11).

If $\boldsymbol{x}$ has ties—for example, $x_i = x_{i'}$ for some $i' \neq i$—then $\widetilde{U}$ is one-sided differentiable with respect to x_i at $\boldsymbol{x}$. Indeed, the right derivative

$$\frac{\partial_+ \widetilde{U}(\boldsymbol{x})}{\partial x_i} = \frac{\mathrm{d}U(x; \max\{R_{\rho(i)}, R_{\rho(i')}\})}{\mathrm{d}x}\Big|_{x=x_i},$$

while the left derivative

$$\frac{\partial_- \widetilde{U}(\boldsymbol{x})}{\partial x_i} = \frac{\mathrm{d}U(x; \min\{R_{\rho(i)}, R_{\rho(i')}\})}{\mathrm{d}x}\Big|_{x=x_i}.$$

Other than this difference, the remainder resembles the proof in the earlier case. For example, we still have $R_{\rho(j)} \leq \min\{R_{\rho(i)}, R_{\rho(i')}\}$ and $R_{\rho(j')} \leq \min\{R_{\rho(i)}, R_{\rho(i')}\}$ for any j' such that $x_j = x_{j'}$. Thus, details are omitted.

□

6.4 Proof of Theorem 4

Write $\boldsymbol{I} := (I_1, \ldots, I_p)$ for a coarse ranking of sizes $n_1, \ldots, n_p$. Let $\pi_{\boldsymbol{I},\boldsymbol{y}}$ be the permutation that sorts the entries of $\boldsymbol{y}$ in each subset I_i in descending order and subsequently concatenates the p subsets in order. For the first subset I_1, for example, it satisfies $\{\pi_{\boldsymbol{I},\boldsymbol{y}}(1), \ldots, \pi_{\boldsymbol{I},\boldsymbol{y}}(n_1)\} = I_1$ and $\boldsymbol{y}_{\pi_{\boldsymbol{I},\boldsymbol{y}}(1)} \geq \boldsymbol{y}_{\pi_{\boldsymbol{I},\boldsymbol{y}}(2)} \geq \cdots \geq \boldsymbol{y}_{\pi_{\boldsymbol{I},\boldsymbol{y}}(n_1)}$. If $\boldsymbol{y} = (3.5, 7.5, 5, -1), I_1 = \{1, 3\}$, and $I_2 = \{2, 4\}$, this permutation gives

$$(\pi_{\boldsymbol{I},\boldsymbol{y}}(1), \pi_{\boldsymbol{I},\boldsymbol{y}}(2), \pi_{\boldsymbol{I},\boldsymbol{y}}(3), \pi_{\boldsymbol{I},\boldsymbol{y}}(4)) = (3, 1, 2, 4), \quad \pi_{\boldsymbol{I},\boldsymbol{y}} \circ \boldsymbol{y} = (5, 3.5, 7.5, -1).$$

When clear from the context, for simplicity, we often omit the dependence on $\boldsymbol{y}$ by writing $\pi_{\boldsymbol{I}}$ for $\pi_{\boldsymbol{I},\boldsymbol{y}}$.

The proof of Theorem 4 relies heavily on the following two lemmas. In particular, Lemma 6.14 reveals the importance of the permutation constructed above.

Lemma 6.14. *The solution to the coarse Isotonic Mechanism* (4.2) *is given by the Isotonic Mechanism* (3.2) *with* $\pi = \pi_{\boldsymbol{I}}$.

Remark 6.6. Thus, the solution to (4.2) can be expressed as $\pi_{\boldsymbol{I}}^{-1} \circ (\pi_{\boldsymbol{I}} \circ \boldsymbol{y})^+$.

Next, let $\boldsymbol{I}^\star := (I_1^\star, \ldots, I_p^\star)$ be the ground-truth coarse ranking that satisfies (4.1), while $\boldsymbol{I}$ is an arbitrary coarse ranking of the same sizes $n_1, \ldots, n_p$.

Lemma 6.15. *There exists a permutation ρ of the indices $1, \ldots, n$ depending only on $\boldsymbol{I}^\star$ and $\boldsymbol{I}$ such that*

$$\pi_{\boldsymbol{I}^\star} \circ (\boldsymbol{R} + \boldsymbol{a}) \succeq_{\mathrm{no}} \pi_{\boldsymbol{I}} \circ (\boldsymbol{R} + \rho \circ \boldsymbol{a})$$

for any $\boldsymbol{a} \in \mathbb{R}^n$.

To clear up any confusion, note that $\pi_{\boldsymbol{I}^\star} \circ (\boldsymbol{R} + \boldsymbol{a}) = \pi_{\boldsymbol{I}^\star, \boldsymbol{R}+\boldsymbol{a}} \circ (\boldsymbol{R} + \boldsymbol{a})$ and $\pi_{\boldsymbol{I}} \circ (\boldsymbol{R} + \rho \circ \boldsymbol{a}) = \pi_{\boldsymbol{I}, \boldsymbol{R}+\rho\circ\boldsymbol{a}} \circ (\boldsymbol{R} + \rho \circ \boldsymbol{a})$.

The proofs of these two lemmas will be presented once we prove Theorem 4 as follows.

Proof of Theorem 4. Denote by $\widehat{\boldsymbol{R}}_{\boldsymbol{I}}$ the solution to the coarse Isotonic Mechanism (4.2). The overall utility can be written as

$$\begin{aligned} U(\widehat{\boldsymbol{R}}_{\boldsymbol{I}}) &= U(\pi_{\boldsymbol{I}}^{-1} \circ (\pi_{\boldsymbol{I}} \circ \boldsymbol{y})^+) \\ &= U((\pi_{\boldsymbol{I}} \circ \boldsymbol{y})^+) \\ &= \sum_{i=1}^n U((\pi_{\boldsymbol{I}} \circ (\boldsymbol{R} + \boldsymbol{z}))_i^+). \end{aligned}$$

Since the permutation ρ in Lemma 6.15 is deterministic, it follows from Assumption 2.2 that $\boldsymbol{z}$ has the same distribution as $\rho \circ \boldsymbol{z}$. This gives

$$\mathbb{E}\, U(\widehat{\boldsymbol{R}}_{\boldsymbol{I}}) = \mathbb{E}\left[\sum_{i=1}^n U((\pi_{\boldsymbol{I}} \circ (\boldsymbol{R} + \boldsymbol{z}))_i^+)\right] = \mathbb{E}\left[\sum_{i=1}^n U((\pi_{\boldsymbol{I}} \circ (\boldsymbol{R} + \rho \circ \boldsymbol{z}))_i^+)\right].$$

By Lemma 6.15, $\pi_{\boldsymbol{I}^\star} \circ (\boldsymbol{R} + \boldsymbol{z}) \succeq_{\mathrm{no}} \pi_{\boldsymbol{I}} \circ (\boldsymbol{R} + \rho \circ \boldsymbol{z})$. By making use of Lemma 6.5, immediately, we get

$$(\pi_{\boldsymbol{I}^\star} \circ (\boldsymbol{R} + \boldsymbol{z}))^+ \succeq_{\mathrm{no}} (\pi_{\boldsymbol{I}} \circ (\boldsymbol{R} + \rho \circ \boldsymbol{z}))^+$$

or, equivalently,

$$(\pi_{\boldsymbol{I}^\star} \circ (\boldsymbol{R} + \boldsymbol{z}))^+ \succeq (\pi_{\boldsymbol{I}} \circ (\boldsymbol{R} + \rho \circ \boldsymbol{z}))^+. \tag{6.12}$$

Next, applying Lemma 6.2 to (6.12) yields

$$\mathbb{E}\, U(\widehat{\boldsymbol{R}}_{\boldsymbol{I}}) = \mathbb{E}\left[\sum_{i=1}^n U((\pi_{\boldsymbol{I}} \circ (\boldsymbol{R} + \rho \circ \boldsymbol{z}))_i^+)\right] \le \mathbb{E}\left[\sum_{i=1}^n U((\pi_{\boldsymbol{I}^\star} \circ (\boldsymbol{R} + \boldsymbol{z}))_i^+)\right].$$

Recognizing that the right-hand size is just the expected overall utility when the owner reports the ground-truth coarse ranking, we get

$$\mathbb{E}\, U(\widehat{\boldsymbol{R}}_{\boldsymbol{I}}) \le \mathbb{E}\, U(\widehat{\boldsymbol{R}}_{\boldsymbol{I}^\star}).$$

This finishes the proof.

$\square$

Proof of Lemma 6.14. Recognizing that the constraint in (4.2) is less restrictive than that of (3.2) with $\pi = \pi_{\boldsymbol{I}}$, it is sufficient to show that the minimum of (4.2) also satisfies the constraint of (3.2). For notational simplicity, denote by $\widehat{\boldsymbol{R}}$ the optimal solution to (4.2). To prove that $\pi_{\boldsymbol{I}} \circ \widehat{\boldsymbol{R}}$ is in descending order, it is sufficient to show that for each $i = 1, \dots, p$, the subset I_i satisfies the property that $\widehat{\boldsymbol{R}}_{I_i}$ has the same order in magnitude as $\boldsymbol{y}_{I_i}$.

Suppose that on the contrary $\widehat{\boldsymbol{R}}_{I_i}$ does not have the same order in magnitude as $\boldsymbol{y}_{I_i}$ for some $1 \le i \le p$. Now let $\widehat{\boldsymbol{R}}'$ be identical to $\widehat{\boldsymbol{R}}$ except on the subset I_i and on this subset, $\widehat{\boldsymbol{R}}'_{I_i}$ is permuted from $\widehat{\boldsymbol{R}}_{I_i}$ to have the same order in magnitude as $\boldsymbol{y}_{I_i}$. Note that $\widehat{\boldsymbol{R}}'$ continues to satisfy the constraint of (4.2). However, we observe that

$$\begin{aligned}
\|\boldsymbol{y} - \widehat{\boldsymbol{R}}'\|^2 - \|\boldsymbol{y} - \widehat{\boldsymbol{R}}\|^2 &= \|\boldsymbol{y}_{I_i} - \widehat{\boldsymbol{R}}'_{I_i}\|^2 - \|\boldsymbol{y}_{I_i} - \widehat{\boldsymbol{R}}_{I_i}\|^2 \\
&= 2\sum_{j \in I_i} y_j \widehat{R}_j - 2\sum_{j \in I_i} y_j \widehat{R}'_j.
\end{aligned}$$

By the rearrangement inequality, we have

$$\sum_{j \in I_i} y_j \widehat{R}_j \le \sum_{j \in I_i} y_j \widehat{R}'_j,$$

which concludes $\|\boldsymbol{y} - \widehat{\boldsymbol{R}}'\|^2 \le \|\boldsymbol{y} - \widehat{\boldsymbol{R}}\|^2$. This is contrary to the assumption that $\widehat{\boldsymbol{R}}$ is the (unique) optimal solution to (4.2).

□

Proof of Lemma 6.15. We prove this lemma by explicitly constructing such a permutation ρ. Let ρ satisfy the following property: ρ restricted to each subset I_i is identical to $I_i^\star$ for each $i = 1, \dots, p$ in the sense that

$$\{\rho(j) : j \in I_i\} = I_i^\star$$

for each i. Moreover, for any $j \in I_i \cap I_i^\star$, we let $\rho(j) = j$, and for any other $j \in I_i \setminus I_i^\star$, we define ρ to be the (unique) mapping from $I_i \setminus I_i^\star$ to $I_i^\star \setminus I_i$ such that the induced correspondence between $\boldsymbol{R}_{I_i \setminus I_i^\star}$ and $\boldsymbol{R}_{I_i^\star \setminus I_i}$ is nondecreasing. For example, ρ maps the largest entry of $\boldsymbol{R}_{I_i \setminus I_i^\star}$ to the largest entry of $\boldsymbol{R}_{I_i^\star \setminus I_i}$, maps the second largest entry of $\boldsymbol{R}_{I_i \setminus I_i^\star}$ to the second largest entry of $\boldsymbol{R}_{I_i^\star \setminus I_i}$, and so on and so forth.

With the construction of ρ in place, we proceed to prove $\pi_{\boldsymbol{I}^\star} \circ (\boldsymbol{R} + \boldsymbol{a}) \succeq_{\text{no}} \pi_{\boldsymbol{I}} \circ (\boldsymbol{R} + \rho \circ \boldsymbol{a})$. For any $1 \le l \le n$, let i satisfy $n_1 + \cdots + n_{i-1} < l \le n_1 + \cdots + n_{i-1} + n_i$ (if $l \le n_1$, then $i = 1$). Now we aim to prove

$$\sum_{j=1}^{l} (\boldsymbol{R} + \boldsymbol{a})_{\pi_{\boldsymbol{I}^\star}(j)} \ge \sum_{j=1}^{l} (\boldsymbol{R} + \rho \circ \boldsymbol{a})_{\pi_{\boldsymbol{I}}(j)},$$

which is equivalent to

$$\sum_{j=1}^{l} R_{\pi_{\boldsymbol{I}^\star}(j)} + \sum_{j=1}^{l} a_{\pi_{\boldsymbol{I}^\star}(j)} \ge \sum_{j=1}^{l} R_{\pi_{\boldsymbol{I}}(j)} + \sum_{j=1}^{l} a_{\rho \circ \pi_{\boldsymbol{I}}(j)}. \tag{6.13}$$

By the construction of ρ, we have

$$\sum_{j=1}^{n_1 + \cdots + n_{i-1}} a_{\pi_{\boldsymbol{I}^\star}(j)} = \sum_{j=1}^{n_1 + \cdots + n_{i-1}} a_{\rho \circ \pi_{\boldsymbol{I}}(j)}.$$

In addition, the left-hand side of (6.13) sums over the $n_1 + \cdots + n_{i-1}$ largest entries of the true values, that is,

$$\sum_{j=1}^{n_1+\cdots+n_{i-1}} R_{\pi_{\boldsymbol{I}^\star}(j)} = \sum_{j=1}^{n_1+\cdots+n_{i-1}} R_{\pi^\star(j)},$$

where $\pi^\star$ is the true ranking of $\boldsymbol{R}$. Thus, it is sufficient to prove

$$\sum_{j=1}^{n_1+\cdots+n_{i-1}} R_{\pi^\star(j)} + \sum_{j=n_1+\cdots+n_{i-1}+1}^{l} (\boldsymbol{R}+\boldsymbol{a})_{\pi_{\boldsymbol{I}^\star}(j)} \geq \sum_{j=1}^{l} R_{\pi_{\boldsymbol{I}}(j)} + \sum_{j=n_1+\cdots+n_{i-1}+1}^{l} a_{\rho\circ\pi_{\boldsymbol{I}}(j)}. \tag{6.14}$$

Note that $J^l = \{\rho \circ \pi_{\boldsymbol{I}}(j) : n_1 + \cdots + n_{i-1} + 1 \leq j \leq l\}$ is a subset of

$$\{\rho \circ \pi_{\boldsymbol{I}}(j) : n_1 + \cdots + n_{i-1} + 1 \leq j \leq n_1 + \cdots + n_i\} = \{\rho(j') : j' \in I_i\} = I_i^\star.$$

Then, by the definition of $\pi_{\boldsymbol{I}} = \pi_{\boldsymbol{I},\boldsymbol{R}+\boldsymbol{a}}$, we have

$$\begin{aligned}\sum_{j=n_1+\cdots+n_{i-1}+1}^{l} (\boldsymbol{R}+\boldsymbol{a})_{\pi_{\boldsymbol{I}^\star}(j)} &= \sum_{j=n_1+\cdots+n_{i-1}+1}^{l} (\boldsymbol{R}+\boldsymbol{a})_{\rho\circ\pi_{\boldsymbol{I}}(j)} \\ &= \sum_{j=n_1+\cdots+n_{i-1}+1}^{l} R_{\rho\circ\pi_{\boldsymbol{I}}(j)} + \sum_{j=n_1+\cdots+n_{i-1}+1}^{l} a_{\rho\circ\pi_{\boldsymbol{I}}(j)},\end{aligned}$$

which, together with (6.14), shows that we would finish the proof of this lemma once verifying

$$\sum_{j=1}^{n_1+\cdots+n_{i-1}} R_{\pi^\star(j)} + \sum_{j=n_1+\cdots+n_{i-1}+1}^{l} R_{\rho\circ\pi_{\boldsymbol{I}}(j)} \geq \sum_{j=1}^{l} R_{\pi_{\boldsymbol{I}}(j)}. \tag{6.15}$$

Now we prove (6.15) as follows. By the construction of ρ, we have $\rho \circ \pi_{\boldsymbol{I}}(j) = \pi_{\boldsymbol{I}}(j)$ whenever $\pi_{\boldsymbol{I}}(j) \in I_i \cap I_i^\star$. Since any such $\pi_{\boldsymbol{I}}(j)$ with $n_1 + \cdots + n_{i-1} + 1 \leq j \leq l$ contributes equally to both sides of (6.15), without loss of generality, we can assume that $I_i \cap I_i^\star = \emptyset$. To see why (6.15) holds, note that if

$$\sum_{j=n_1+\cdots+n_{i-1}+1}^{l} R_{\pi_{\boldsymbol{I}}(j)} \tag{6.16}$$

is summed over the $l - n_1 - \cdots - n_{i-1}$ largest entries of $\boldsymbol{R}_{I_i}$, then by the construction of ρ,

$$\sum_{j=n_1+\cdots+n_{i-1}+1}^{l} R_{\rho\circ\pi_{\boldsymbol{I}}(j)} \tag{6.17}$$

is summed over the $l - n_1 - \cdots - n_{i-1}$ largest entries of $\boldsymbol{R}_{I_i^\star}$. Thus, (6.15) follows since its right-hand side is the sum of the l largest entries of $\boldsymbol{R}$. The sum (6.16) may skip some large entries, and (6.17) would skip correspondingly. Here, (6.15) remains true since summation and skipping are applied to $\boldsymbol{R}$ that has already been ordered from the largest to the smallest.

$\square$

7 Discussion

In this paper, we have studied how an appraiser can better estimate an unknown vector by eliciting information from an owner who has knowledge of the vector. Assuming convex utility for the owner, among other things, we prove that if the owner truthfully provides the appraiser with information about the ground-truth vector, then the knowledge partition must be generated from pairwise comparisons between some entries of the unknown vector. Next, we show that the owner would indeed be truthful when she is asked to provide a ranking of the entries of the vector or, equivalently, to delineate most precisely the vector using pairwise comparisons. This gives a computationally efficient method that we refer to as the Isotonic Mechanism, which gives the most fine-grained information among all truthful mechanisms to the appraiser for better estimation of the unknown vector. The Isotonic Mechanism yields a more accurate estimate of the ground truth than the raw observation, regardless of the noise distribution. The accuracy gain is markedly more pronounced when the dimension of the vector is high and the noise in the observation is significant. We also have obtained several relaxations of this mechanism; for example, in the case of incomplete knowledge of the ground truth, the owner remains truthful and the appraiser can continue to improve estimation accuracy.

Our work opens a host of avenues for future research. Most immediately, a technical question is to find all pairwise-comparison-based knowledge partitions that are truthful. Another related question is to prove the conjecture that Proposition 2.9 holds for any fixed noise distribution.

Empirical studies. Coincidentally, shortly after a preliminary version of the present paper was submitted, NeurIPS 2021 required all authors to "rank their papers in terms of their own perception of the papers' scientific contributions to the NeurIPS community." Using this dataset, an empirical study may be used to analyze what would be the outcome if this mechanism were used in NeurIPS 2021, or at least were used as a reference for the decision making. On the flip side, peer review involves much sophistication that has not been considered in the development of the Isotonic Mechanism, and therefore more efforts are needed toward the employment of this mechanism (see Section 4 of [28]). Perhaps a more realistic starting point is to integrate the Isotonic Mechanism or its relaxations with the long line of research that aims to incentivize reviewers to provide more accurate review scores [17, 9, 30].

Extensions of statistical models. The appraiser can incorporate the information provided by the owner in a way different from (2.1). For example, an alternative to (2.1) is to solve

$$\min_{\boldsymbol{r}} \ \frac{1}{2}\|\boldsymbol{y} - \boldsymbol{r}\|^2 + \mathrm{Pen}(\boldsymbol{r}),$$

with some penalty term $\mathrm{Pen}(\boldsymbol{r})$ satisfying $\mathrm{Pen}(\boldsymbol{r}) = 0$ if $\boldsymbol{r} \in S$ and otherwise $\mathrm{Pen}(\boldsymbol{r}) > 0$. It is valuable to find truthful knowledge partitions in this case. It is equally important to consider when the observation is generated by complex statistical models such as generalized linear models. However, even the existence of a nontrivial truthful knowledge partition is open at the moment. Moreover, the present work has not explored the setting where the ground truth takes a nonparametric form or is restricted by some constraints such as being nonnegative. This enables a connection with the literature of shape-restricted regression [24].

Relaxation of assumptions. To broaden its applicability, the robustness of the Isotonic Mechanism can be analyzed in a realistic setting where the owner might inadvertently report an incorrect knowledge element. It is also important, but apparently challenging, to design truthful knowledge partitions when grades provided by the agents might not have the same noise distribution, which violates Assumption 2.2. From the same angle, moreover, a crucial direction is to incorporate peer prediction [19, 32, 16, 12, 11] in the design of truthful mechanisms, which assumes strategic agents who receive rewards depending on their forecasts.

Perhaps the most pressing extension is concerned with the utility of the owner. Moving away from the assumption that the utility function is an arbitrary or unknown convex function, interestingly, it may be possible to obtain a more fine-grained truthful mechanism may be obtained regarding a specific convex utility function. For example, we have the following result, with its proof deferred to the Appendix.

Proposition 7.1. *Assume that the noise terms $z_1, \ldots, z_n$ are i.i.d. random variables with mean 0 and the overall utility is $U(\widehat{\boldsymbol{R}}) = \|\widehat{\boldsymbol{R}}\|^2$. Then the collection of all lines passing through the origin in $\mathbb{R}^n$,*

$$\mathcal{S} = \{\{a\boldsymbol{u} : a \in \mathbb{R}\} : \|\boldsymbol{u}\| = 1\},$$

is a truthful knowledge partition.[9]

In light of this result, a meaningful question is to design better truthful knowledge partitions that correspond with given utility functions. Furthermore, a more challenging avenue for future research is to tackle nonconvex utility functions. In the nonconvex regime, however, the isotonic knowledge partition is no longer truthful (see Proposition A.1 in the Appendix). To cope with nonconvex utility, one possible approach is to include additional rewards (by, for example, the conference organizers) to the utility function so that the modified utility function becomes convex or approximately convex.

Outlook. At a high level, our work points to a flexible way of incorporating elements from game theory and mechanism design into statistical learning. More precisely, the owner is concerned with her utility as a function of the appraisal and her move in the game is to pick an element from the knowledge partition, while the appraiser focuses on the estimation accuracy and he can choose any template of estimators that leverages information provided by the owner.

$$\begin{tikzcd}[column sep=huge]
\text{Alice} \arrow[r, shift left, "\text{knowledge of ground truth}"] & \text{Insider-Assisted Mechanism} \arrow[l, shift left, "\text{payoff: utility}"] \arrow[r, shift right, "\text{payoff: estimation accuracy}"'] & \text{Bob} \arrow[l, shift right, "\text{template of estimators}"']
\end{tikzcd}$$

Figure 2: An illustration of the game between the two parties under an insider-assisted mechanism.

Assuming a Stackelberg competition where the appraiser moves first followed by the owner, how can we design a knowledge partition that leads to a good equilibrium such that both parties have a good outcome? More generally, the appraiser may change his strategy given the owner's move. Taking the example of the Isotonic Mechanism, does the appraiser have an incentive to deviate from using isotonic regression? An important question is therefore to find knowledge partitions that lead to an efficient Nash equilibrium.

[9] An knowledge element of this partition is invariant with respect to the sign change: $\{a\boldsymbol{u} : a \in \mathbb{R}\} = \{a(-\boldsymbol{u}) : a \in \mathbb{R}\}$. This shows that this knowledge partition is topologically equivalent to an $(n-1)$-sphere with antipodes identified, which is better known as the real projective space of dimension $n-1$. Note that this knowledge partition is not separated by piecewise smooth surfaces.

Acknowledgments

We would like to thank Patrick Chao, Qiyang Han, Nihar Shah, Haifeng Xu, Rakesh Vohra, and Yuhao Wang for very insightful comments and fruitful discussions. This research was supported in part by NSF grants CCF-1934876 and CAREER DMS-1847415, and an Alfred Sloan Research Fellowship.

References

[1] J. Aitchison and S. Silvey. Maximum-likelihood estimation of parameters subject to restraints. *The Annals of Mathematical Statistics*, 29(3):813–828, 1958.

[2] B. C. Arnold. *Majorization and the Lorenz order: A brief introduction.* Springer-Verlag Lecture Notes in Statistics, vol. 43, 1987.

[3] R. E. Barlow, D. J. Bartholomew, J.-M. Bremner, and H. Brunk. *Statistical inference under order restrictions: The theory and application of isotonic regression.* Wiley New York, 1972.

[4] R. E. Barlow and H. D. Brunk. The isotonic regression problem and its dual. *Journal of the American Statistical Association*, 67(337):140–147, 1972.

[5] M. Bogdan, E. Van Den Berg, C. Sabatti, W. Su, and E. J. Candès. Slope—Adaptive variable selection via convex optimization. *The Annals of Applied Statistics*, 9(3):1103–1140, 2015.

[6] S. Chatterjee, A. Guntuboyina, and B. Sen. On risk bounds in isotonic and other shape restricted regression problems. *The Annals of Statistics*, 43(4):1774–1800, 2015.

[7] Y. Chen and C. Suh. Spectral MLE: Top-k rank aggregation from pairwise comparisons. In *International Conference on Machine Learning*, pages 371–380. PMLR, 2015.

[8] O. François. Arbitrariness of peer review: A bayesian analysis of the nips experiment. *arXiv preprint arXiv:1507.06411*, 2015.

[9] T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

[10] K. G. Jamieson and R. Nowak. Active ranking using pairwise comparisons. *Advances in Neural Information Processing Systems*, 24, 2011.

[11] Y. Kong. Dominantly truthful multi-task peer prediction with a constant number of tasks. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2398–2411. SIAM, 2020.

[12] Y. Kong and G. Schoenebeck. An information theoretic framework for designing information elicitation mechanisms that reward truth-telling. *ACM Transactions on Economics and Computation (TEAC)*, 7(1):1–33, 2019.

[13] D. M. Kreps. *A course in microeconomic theory.* Princeton University Press, 1990.

[14] V. Krishna and E. Maenner. Convex potentials with an application to mechanism design. *Econometrica*, 69(4):1113–1119, 2001.

[15] J. B. Kruskal. Nonmetric multidimensional scaling: a numerical method. *Psychometrika*, 29(2):115–129, 1964.

[16] Y. Liu and Y. Chen. Machine-learning aided peer prediction. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 63–80, 2017.

[17] Y. Liu, J. Wang, and Y. Chen. Surrogate scoring rules. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 853–871, 2020.

[18] A. W. Marshall, I. Olkin, and B. C. Arnold. *Inequalities: theory of majorization and its applications*, volume 143. Springer, 1979.

[19] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359–1373, 2005.

[20] S. Negahban, S. Oh, and D. Shah. Rank centrality: Ranking from pairwise comparisons. *Operations Research*, 65(1):266–287, 2017.

[21] G. Qiao, W. Su, and L. Zhang. Oneshot differentially private top-k selection. In *International Conference on Machine Learning*, pages 8672–8681. PMLR, 2021.

[22] T. Robertson, F. T. Wright, and R. L. Dykstra. *Order restricted statistical inference*. John Wiley & Sons, 1988.

[23] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 2015.

[24] R. J. Samworth and B. Sen. Special issue on "nonparametric inference under shape constraints". *Statistical Science*, 33(4):469–472, 2018.

[25] N. B. Shah. An overview of challenges, experiments, and computational solutions in peer review (extended version). *Communications of the ACM*, 2022.

[26] N. B. Shah and M. J. Wainwright. Simple, robust and optimal ranking from pairwise comparisons. *The Journal of Machine Learning Research*, 18(1):7246–7283, 2017.

[27] I. Stelmakh, N. B. Shah, A. Singh, and H. Daumé III. A novice-reviewer experiment to address scarcity of qualified reviewers in large conferences. *arXiv preprint arXiv:2011.15050*, 2020.

[28] W. J. Su. You are the best reviewer of your own papers: An owner-assisted scoring mechanism. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

[29] S.-H. Sun. ICLR2020-openreviewdata. `https://github.com/shaohua0116/ICLR2020-OpenReviewData`, 2020. [Online].

[30] J. Wang, I. Stelmakh, Y. Wei, and N. B. Shah. Debiasing evaluations that are biased by evaluations. *arXiv preprint arXiv:2012.00714*, 2020.

[31] F. Wauthier, M. Jordan, and N. Jojic. Efficient ranking from pairwise comparisons. In *International Conference on Machine Learning*, pages 109–117. PMLR, 2013.

[32] J. Witkowski and D. C. Parkes. A robust bayesian truth serum for small populations. In *Twenty-Sixth AAAI Conference on Artificial Intelligence*, 2012.

[33] C.-H. Zhang. Risk bounds in isotonic regression. *The Annals of Statistics*, 30(2):528–555, 2002.

A Appendix

Proof of Proposition 2.7. Assume that $\mathcal{S}$ is a nontrivial knowledge partition. Pick any knowledge element $S \in \mathcal{S}$ and let x be an interior point of S. Consider the noiseless setting with utility function $U(x) = x$, which is a nondecreasing convex function. Write $a = \sup_y \{y : \text{interval } (x, y) \subset S\}$. If $a = \infty$, then S contains all sufficiently large numbers. In this case, we instead pick any different knowledge element in order to ensure $a < \infty$.

Therefore, we can assume $a < \infty$. Let S' be the knowledge element that contains a (small) right neighborhood of a. Taking ground truth $R = \frac{x+2a}{3} < a$, if the owner reports S', then the solution would be a. Since $U(a) > U(R)$, the owner would be better off reporting S' instead of S. This contradiction demonstrates that $\mathcal{S}$ must be trivial.

□

Proof of Proposition 2.9. Recall that $\widehat{\boldsymbol{R}}_S$ is the solution to

$$\begin{aligned} \min_{\boldsymbol{r}} \ & \|\boldsymbol{y} - \boldsymbol{r}\|^2 \\ \text{s.t.} \ & \boldsymbol{r} \in S. \end{aligned}$$

When the noise level in $\boldsymbol{y} = \boldsymbol{R} + \boldsymbol{z}$ tends to zero and the ground truth $\boldsymbol{R} \in S$, the projection of $\boldsymbol{y}$ onto S is asymptotically equivalent to the projection of $\boldsymbol{y}$ onto the tangent cone of S at $\boldsymbol{R}$ (see [23]). More precisely, letting $T_S(\boldsymbol{R})$ be the tangent cone of S at $\boldsymbol{R}$ and writing $\widehat{\boldsymbol{R}}_{T_S(\boldsymbol{R})}$ for the projection of $\boldsymbol{y}$ onto $T_S(\boldsymbol{R})$, we have $\widehat{\boldsymbol{R}}_S = \widehat{\boldsymbol{R}}_{T_S(\boldsymbol{R})} + o(\|\widehat{\boldsymbol{R}}_S - \boldsymbol{R}\|)$. This fact implies that, with probability tending to one,

$$\limsup_{\sigma \to 0} \frac{\|\widehat{\boldsymbol{R}}_{T_S(\boldsymbol{R})} - \boldsymbol{R}\|^2}{\|\widehat{\boldsymbol{R}}_S - \boldsymbol{R}\|^2} = \limsup_{\sigma \to 0} \frac{\mathbb{E}\,\|\widehat{\boldsymbol{R}}_{T_S(\boldsymbol{R})} - \boldsymbol{R}\|^2}{\mathbb{E}\,\|\widehat{\boldsymbol{R}}_S - \boldsymbol{R}\|^2} = 1.$$

To prove the first part of Proposition 2.9, therefore, it suffices to show that

$$\|\widehat{\boldsymbol{R}}_{T_{S_2}(\boldsymbol{R})} - \boldsymbol{R}\|^2 \le \|\widehat{\boldsymbol{R}}_{T_{S_1}(\boldsymbol{R})} - \boldsymbol{R}\|^2$$

with probability one. This inequality follows from the fact that $T_{S_2}(\boldsymbol{R}) \subset T_{S_1}(\boldsymbol{R})$ and both cones have apex at $\boldsymbol{R}$.

Next, we prove the second part. Because both S_1 and S_2 are cones, it follows from Moreau's decomposition theorem that

$$\|\widehat{\boldsymbol{R}}_{S_1} - \boldsymbol{y}\|^2 + \|\widehat{\boldsymbol{R}}_{S_1}\|^2 = \|\boldsymbol{y}\|^2$$

and

$$\|\widehat{\boldsymbol{R}}_{S_2} - \boldsymbol{y}\|^2 + \|\widehat{\boldsymbol{R}}_{S_2}\|^2 = \|\boldsymbol{y}\|^2.$$

Since $S_2 \subset S_1$, we get $\|\widehat{\boldsymbol{R}}_{S_1} - \boldsymbol{y}\|^2 \le \|\widehat{\boldsymbol{R}}_{S_2} - \boldsymbol{y}\|^2$, which in conjunction with the two identities above gives

$$\|\widehat{\boldsymbol{R}}_{S_1}\|^2 \ge \|\widehat{\boldsymbol{R}}_{S_2}\|^2. \tag{A.1}$$

In the limit $\sigma \to \infty$, we have $\|\widehat{\boldsymbol{R}}_{S_1} - \boldsymbol{R}\|^2 = (1 + o(1))\|\widehat{\boldsymbol{R}}_{S_1}\|^2$ and $\|\widehat{\boldsymbol{R}}_{S_2} - \boldsymbol{R}\|^2 = (1 + o(1))\|\widehat{\boldsymbol{R}}_{S_2}\|^2$ with probability tending to one. Together with (A.1), this concludes

$$\limsup_{\sigma \to \infty} \frac{\mathbb{E}\,\|\widehat{\boldsymbol{R}}_{S_2} - \boldsymbol{R}\|^2}{\mathbb{E}\,\|\widehat{\boldsymbol{R}}_{S_1} - \boldsymbol{R}\|^2} \le \limsup_{\sigma \to \infty} (1 + o(1)) = 1.$$

□

Proof for the example in Section 4. First, consider the case $U(x) = x^2$. For simplicity, we start by assuming $\boldsymbol{R} = \boldsymbol{0}$. Due to symmetry, the expected overall utility of reporting S_1 is the same as that of reporting an arbitrary isotonic cone S_π. In particular, taking any $S_\pi \neq S_1$, we have $S_\pi \subset S_2$. The proof of Proposition 2.9 above shows that

$$\|\widehat{\boldsymbol{R}}_{S_\pi}\|^2 \le \|\widehat{\boldsymbol{R}}_{S_2}\|^2 \tag{A.2}$$

with probability one. By taking the expectation, we get

$$\mathbb{E}\, U(\widehat{\boldsymbol{R}}_{S_\pi}) = \mathbb{E}\,\|\widehat{\boldsymbol{R}}_{S_\pi}\|^2 < \mathbb{E}\,\|\widehat{\boldsymbol{R}}_{S_2}\|^2 = \mathbb{E}\, U(\widehat{\boldsymbol{R}}_{S_2})$$

since (A.2) is a strict inequality with positive probability. Equivalently, we get

$$\mathbb{E}\, U(\widehat{\boldsymbol{R}}_{S_1}) < \mathbb{E}\, U(\widehat{\boldsymbol{R}}_{S_2}). \tag{A.3}$$

Moving back to the case $\boldsymbol{R} = (n\epsilon, (n-1)\epsilon, \ldots, 2\epsilon, \epsilon) \in S_1$, (A.3) remains valid for sufficiently small ϵ.

Next, we consider $U(x) = \max\{0, x\}^2$. As earlier, we first assume $\boldsymbol{R} = \boldsymbol{0}$. For any isotonic cone S_π, let us take as given for the moment that the empirical distribution of the entries of $\widehat{\boldsymbol{R}}_{S_\pi}$ is symmetric with respect to the origin over the randomness of the Gaussian noise. This symmetry is also true for S_2. Therefore, we get

$$\mathbb{E}\, U(\widehat{\boldsymbol{R}}_{S_1}) = \frac{1}{2}\,\mathbb{E}\,\|\widehat{\boldsymbol{R}}_{S_1}\|^2 < \frac{1}{2}\,\mathbb{E}\,\|\widehat{\boldsymbol{R}}_{S_2}\|^2 = \mathbb{E}\, U(\widehat{\boldsymbol{R}}_{S_2}).$$

This inequality continues to hold for sufficiently small ϵ.

To finish the proof, we explain why the above-mentioned symmetric property of $\widehat{\boldsymbol{R}}_{S_\pi}$ in distribution is true. Let π^- be the reverse ranking of π, that is, $\pi^-(i) = \pi(n+1-i)$ for all $1 \le i \le n$. For any Gaussian noise vector $\boldsymbol{z} = (z_1, \ldots, z_n)$, it is easy to see that the entries (as a set) of the projection of $\boldsymbol{z}$ onto S_π are negative to the entries (as a set) of the projection of $-\pi^- \circ \boldsymbol{z}$ onto S_π. Last, note that $-\pi^- \circ \boldsymbol{z}$ has the same probability distribution as $\boldsymbol{z}$. This completes the proof. □

Proof of Proposition 5.1. Recognizing that the solution to the Isotonic Mechanism takes the form $\pi^{-1} \circ (\pi \circ (\boldsymbol{R} + \boldsymbol{z}))^+$, the expected overall utility is

$$\begin{aligned}
\mathbb{E}\, U(\pi^{-1} \circ (\pi \circ (\boldsymbol{R} + \boldsymbol{z}))^+) &= \mathbb{E}\, U((\pi \circ (\boldsymbol{R} + \boldsymbol{z}))^+) \\
&= \mathbb{E}\, U((\pi \circ \boldsymbol{R} + \pi \circ \boldsymbol{z})^+) \\
&= \mathbb{E}\, U((\pi \circ \boldsymbol{R} + \boldsymbol{z})^+) \\
&= \mathbb{E}\left[\sum_{i=1}^n U((\pi \circ \boldsymbol{R} + \boldsymbol{z})_i^+)\right],
\end{aligned}$$

where we use the exchangeability of the distribution of the noise vector $\boldsymbol{z}$. Next, the assumption that π_1 is more consistent than π_2 with respect to the ground truth $\boldsymbol{R}$ implies

$$\pi_1 \circ \boldsymbol{R} \succeq_{\text{no}} \pi_2 \circ \boldsymbol{R},$$

from which it follows that $\pi_1 \circ \boldsymbol{R} + \boldsymbol{z} \succeq_{\text{no}} \pi_2 \circ \boldsymbol{R} + \boldsymbol{z}$. By the Hardy–Littlewood–Pólya inequality, therefore, we get

$$\sum_{i=1}^{n} U((\pi_1 \circ \boldsymbol{R} + \boldsymbol{z})_i^+) \geq \sum_{i=1}^{n} U((\pi_2 \circ \boldsymbol{R} + \boldsymbol{z})_i^+)$$

for all $\boldsymbol{z}$. This concludes

$$\mathbb{E}\, U(\pi_1^{-1} \circ (\pi_1 \circ (\boldsymbol{R} + \boldsymbol{z}))^+) \geq \mathbb{E}\, U(\pi_2^{-1} \circ (\pi_2 \circ (\boldsymbol{R} + \boldsymbol{z}))^+).$$

□

Proof of Lemma 6.9. The proof relies on the min-max formula of isotonic regression (see Chapter 1 of [22]):

$$a_k^+ = \max_{v \geq k} \min_{u \leq k} \frac{a_u + a_{u+1} + \cdots + a_v}{v - u + 1}$$

for any $k = 1, \ldots, n$. For simplicity, write $\boldsymbol{b} = \boldsymbol{a} + \delta \mathbf{e}_i$. Then it is clear that $\boldsymbol{b}$ is larger than or equal to $\boldsymbol{a}$ in the component-wise sense. Therefore, we get

$$b_k^+ = \max_{v \geq k} \min_{u \leq k} \frac{b_u + b_{u+1} + \cdots + b_v}{v - u + 1} \geq \max_{v \geq k} \min_{u \leq k} \frac{a_u + a_{u+1} + \cdots + a_v}{v - u + 1} = a_k^+$$

for all $k = 1, \ldots, n$. This concludes the proof.

□

Proof of Lemma 6.10. To begin with, assume that $\boldsymbol{a}^+$ has constant entries. Suppose the contrary that

$$\bar{a}_k := \frac{a_1 + \cdots + a_k}{k} > \bar{a}$$

for some k, which implies

$$\frac{a_1 + \cdots + a_k}{k} > \bar{a} > \bar{a}_{-k} := \frac{a_{k+1} + \cdots + a_n}{n - k}.$$

This inequality allows us to get

$$\begin{aligned}
&\sum_{i=1}^{k}(a_i - \bar{a}_k)^2 + \sum_{i=k+1}^{n}(a_i - \bar{a}_{-k})^2 \\
&= \sum_{i=1}^{k}(a_i - \bar{a})^2 - k(\bar{a}_k - \bar{a})^2 + \sum_{i=k+1}^{n}(a_i - \bar{a})^2 - (n-k)(\bar{a}_{-k} - \bar{a})^2 \\
&< \sum_{i=1}^{n}(a_i - \bar{a})^2 \\
&= \|\boldsymbol{a} - \boldsymbol{a}^+\|^2.
\end{aligned}$$

As such, the vector formed by concatenating k copies of $\bar{a}_k$ followed by $n-k$ copies of $\bar{a}_{-k}$, which lies in the standard isotonic cone since $\bar{a}_k > \bar{a}_{-k}$, leads to a smaller squared error than $\boldsymbol{a}^+$. This contradicts the definition of $\boldsymbol{a}^+$.

Next, we assume that

$$\frac{a_1 + \cdots + a_k}{k} \leq \bar{a} \tag{A.4}$$

for all $k = 1, \ldots, n$. To seek a contradiction, suppose that $\boldsymbol{a}^+$ has more than one constant piece. That is, a partition of $\{1, 2, \ldots, n\} = I_1 \cup I_2 \cup \cdots \cup I_L$ with $L \geq 2$ from the left to the right satisfies the following: the entries of $\boldsymbol{a}^+$ on I_l are constant and, denoting by $a^+_{I_l}$ the value on I_l, the isotonic constraint requires

$$a^+_{I_1} > a^+_{I_2} > \cdots > a^+_{I_L}. \tag{A.5}$$

Making use of a basic property of isotonic regression, we have

$$a^+_{I_l} = \frac{\sum_{i \in I_l} a_i}{|I_l|}$$

for $l = 1, \ldots, L$, where $|I_l|$ denotes the set cardinality. This display, together with (A.5), shows

$$\begin{aligned}
\bar{a} &= \frac{a_1 + \cdots + a_n}{n} \\
&= \frac{|I_1|a^+_{I_1} + |I_2|a^+_{I_2} + \cdots + |I_L|a^+_{I_L}}{n} \\
&< \frac{|I_1|a^+_{I_1} + |I_2|a^+_{I_1} + \cdots + |I_L|a^+_{I_1}}{n} \\
&= a^+_{I_1} \\
&= \frac{a_1 + \cdots + a_{|I_1|}}{|I_1|},
\end{aligned}$$

a contradiction to (A.4).

□

Proof of Proposition 7.1. Denote by $\boldsymbol{u}$ the unit-norm vector reported by the owner (equivalently, $-\boldsymbol{u}$). The output of the mechanism is

$$\widehat{\boldsymbol{R}}_{\boldsymbol{u}} = (\boldsymbol{u} \cdot \boldsymbol{y})\boldsymbol{u} = (\boldsymbol{u} \cdot (\boldsymbol{R} + \boldsymbol{z}))\boldsymbol{u} = (\boldsymbol{u} \cdot \boldsymbol{R})\boldsymbol{u} + (\boldsymbol{u} \cdot \boldsymbol{z})\boldsymbol{u}.$$

The overall utility is

$$\|\widehat{\boldsymbol{R}}_{\boldsymbol{u}}\|^2 = \|(\boldsymbol{u} \cdot \boldsymbol{R})\boldsymbol{u} + (\boldsymbol{u} \cdot \boldsymbol{z})\boldsymbol{u}\|^2.$$

Its expectation is

$$\mathbb{E}\,\|(\boldsymbol{u} \cdot \boldsymbol{R})\boldsymbol{u} + (\boldsymbol{u} \cdot \boldsymbol{z})\boldsymbol{u}\|^2 = \|(\boldsymbol{u} \cdot \boldsymbol{R})\boldsymbol{u}\|^2 + \mathbb{E}\,\|(\boldsymbol{u} \cdot \boldsymbol{z})\boldsymbol{u}\|^2 + 2\,\mathbb{E}\left[((\boldsymbol{u} \cdot \boldsymbol{R})\boldsymbol{u}) \cdot ((\boldsymbol{u} \cdot \boldsymbol{z})\boldsymbol{u})\right].$$

We have $\|(\boldsymbol{u} \cdot \boldsymbol{R})\boldsymbol{u}\|^2 = (\boldsymbol{u} \cdot \boldsymbol{R})^2\|\boldsymbol{u}\|^2 = (\boldsymbol{u} \cdot \boldsymbol{R})^2$ and

$$\mathbb{E}\,\|(\boldsymbol{u} \cdot \boldsymbol{z})\boldsymbol{u}\|^2 = \mathbb{E}(\boldsymbol{u} \cdot \boldsymbol{z})^2\|\boldsymbol{u}\|^2 = \mathbb{E}(\boldsymbol{u} \cdot \boldsymbol{z})^2 = \|\boldsymbol{u}\|^2\,\mathbb{E}\,z_1^2 = \mathbb{E}\,z_1^2,$$

where the third equality makes use of the fact that $z_1, \ldots, z_n$ are i.i.d. centered random variables. Besides, we have

$$2\,\mathbb{E}\left[((\boldsymbol{u} \cdot \boldsymbol{R})\boldsymbol{u}) \cdot ((\boldsymbol{u} \cdot \boldsymbol{z})\boldsymbol{u})\right] = 2\left((\boldsymbol{u} \cdot \boldsymbol{R})\boldsymbol{u}\right) \cdot \left((\boldsymbol{u} \cdot \mathbb{E}\,\boldsymbol{z})\boldsymbol{u}\right) = 0.$$

Thus, we get

$$\mathbb{E}\,\|\widehat{\boldsymbol{R}}_{\boldsymbol{u}}\|^2 = (\boldsymbol{u}\cdot\boldsymbol{R})^2 + \mathbb{E}\,z_1^2 \le \|\boldsymbol{R}\|^2 + \mathbb{E}\,z_1^2,$$

with equality if and only if $\boldsymbol{u}$ has the same direction as $\boldsymbol{R}$, that is, $\boldsymbol{R} \in \{a\boldsymbol{u} : a \in \mathbb{R}\}$. In words, the owner would maximize her expected overall utility if and only if she reports the line that truly contains the ground truth.

$\square$

Proposition A.1. *Under Assumptions 2.1 and 2.2, if the utility function U in* (2.2) *is nonconvex, then there exists a certain ground truth $\boldsymbol{R}$ and a noise distribution such that the owner is not truthful under the Isotonic Mechanism.*

Proof of Proposition A.1. Let the noise vector $\boldsymbol{z} = \boldsymbol{0}$. Since U is not convex, there must exist $r_1 > r_2$ such that

$$U(r_1) + U(r_2) < 2U\left(\frac{r_1 + r_2}{2}\right). \tag{A.6}$$

Let the ground truth $\boldsymbol{R}$ satisfy $\boldsymbol{R}_1 = r_1, \boldsymbol{R}_2 = r_2$, and $\boldsymbol{R}_i = r_2 - i$ for $i = 3, \ldots, n$. Note that $\boldsymbol{R}$ is in descending order. If the owner reports the true ranking, the solution to the Isotonic Mechanism is $\boldsymbol{R}$ itself and her overall utility is

$$U(r_1) + U(r_2) + \sum_{i=3}^{n} U(r_2 - i). \tag{A.7}$$

However, if the owner reports π such that $\pi(1) = 2, \pi(2) = 1$, and $\pi(i) = i$ for $i \ge 3$, then the solution with this ranking is

$$\widehat{\boldsymbol{R}}_\pi = \left(\frac{r_1 + r_2}{2}, \frac{r_1 + r_2}{2}, r_2 - 3, r_2 - 4, \ldots, r_2 - n\right).$$

The corresponding overall utility is

$$U\left(\frac{r_1 + r_2}{2}\right) + U\left(\frac{r_1 + r_2}{2}\right) + \sum_{i=3}^{n} U(r_2 - i) = 2U\left(\frac{r_1 + r_2}{2}\right) + \sum_{i=3}^{n} U(r_2 - i). \tag{A.8}$$

It follows from (A.6) that (A.8)>(A.7), thereby implying that the owner would be better off reporting the incorrect ranking π instead of the true ranking.

$\square$