

---

# Carrot and Stick: Eliciting Comparison Data and Beyond

---

**Yiling Chen**

Harvard University

yiling@seas.harvard.edu

**Shi Feng**

Harvard University

shifeng@fas.harvard.edu

**Fang-Yi Yu**

George Mason University

fangyiyu@gmu.edu

## Abstract

Comparison data elicited from people are fundamental to many machine learning tasks, including reinforcement learning from human feedback for large language models and estimating ranking models. They are typically subjective and not directly verifiable. How to truthfully elicit such comparison data from rational individuals? We design peer prediction mechanisms for eliciting comparison data using a bonus-penalty payment [11]. Our design leverages on the strong stochastic transitivity for comparison data [60, 13] to create symmetrically strongly truthful mechanisms such that truth-telling 1) forms a strict Bayesian Nash equilibrium, and 2) yields the highest payment among all symmetric equilibria. Each individual only needs to evaluate one pair of items and report her comparison in our mechanism.

We further extend the bonus-penalty payment concept to eliciting networked data, designing a symmetrically strongly truthful mechanism when agents' private signals are sampled according to the Ising models. We provide the necessary and sufficient conditions for our bonus-penalty payment to have truth-telling as a strict Bayesian Nash equilibrium. Experiments on two real-world datasets further support our theoretical discoveries.

## 1 Introduction

In the past two decades, researchers have been embracing the challenge of eliciting private information from individuals when there is no ground truth available to evaluate the quality of elicited contributions, and have made amazing progress. Many mechanisms, collectively called *peer prediction* [42], have been developed to incentivize individuals to strictly truthfully report their information at a Bayesian Nash equilibrium (BNE), by artful design of payment functions that only depend on reports from individuals. Moreover, in multi-task peer prediction mechanisms, the truthful BNE gives each individual the highest expected payoff among all BNEs (i.e. it's a strongly truthful BNE). [11, 57, 30]

However, all prior multi-task peer prediction mechanisms require tasks being ex-ante identical, and hence individuals' private information is independently and identically distributed (iid) for each task. Multi-task peer prediction leverages this structure of information to succeed at truthful elicitation. But what if such structure of information doesn't hold for an information elicitation problem?

One notable application is to elicit pair-wise comparisons of multiple alternatives, such as preferences for consumer products [53], translation [34], peer grading [55], and relevance of language model outputs [9, 10]. Such pair-wise comparison data are crucial for estimating a ranking of the alternatives and for devising reward functions for reinforcement learning. Comparison tasks for different pairs are clearly not ex-ante identical — answers to the tasks demonstrate a certain degree of transitivity (e.g. if  $a$  is preferred to  $a'$  and  $a'$  is preferred to  $a''$ , then it's more likely that  $a$  is preferred to  $a''$ ), rendering existing peer prediction mechanisms not applicable.

In this paper, we design a peer prediction mechanism for eliciting comparison data. We model individuals’ private information of pair-wise comparisons as Bayesian strongly stochastically transitive (Bayesian SST), which takes many widely used models (e.g. Thurstone [58], Bradley-Terry-Luce [4, 38], and Mallows [39]) as special cases. Our mechanism uses a simple bonus-penalty payment [11] (hence carrot and stick) that takes three reports as inputs and admits a strongly truthful symmetric BNE. The key insight that we develop is a condition of information structure that we call *uniform dominance*. When uniform dominance is satisfied, the bonus-penalty payment is the only type of payment that induces a strictly truthful BNE. Information of individuals,  $i$ ,  $j$ , and  $k$ , satisfies uniform dominance if, conditioned on any realization of agent  $i$ ’s information, the probability for  $j$ ’s information to agree with  $i$ ’s is higher than the probability for  $k$ ’s information to agree with  $i$ ’s. Bayesian SST allows us to group three pairwise comparisons,  $(a, a')$ ,  $(a'', a')$  and  $(a'', a)$ , together such that private information about these pairs satisfies uniform dominance. After identifying uniform dominance as a central structure for incentivizing truthful elicitation, we further generalize the bonus-penalty payment to truthfully elicit private information over social networks that demonstrate homophily (i.e. friends tend to have similar opinions than non-friends) [40], and our mechanism can be integrated with common survey techniques such as snowball sampling [24].

**Our contributions.** Our work is a leap forward for designing mechanisms for complex information elicitation settings where ground truth verification is not available.

- We are the first to design mechanisms to truthfully elicit pairwise comparison data under Bayesian SST and networked data under Ising models. In our mechanisms, truthful reporting forms a BNE and yields a strictly higher payoff than any symmetric non-permutation equilibrium.
- We identify a key structure of information, uniform dominance, as a lever such that the simple bonus-penalty payment is the unique payment inducing a strictly truthful BNE. This identification may offer a path for developing truthful elicitation mechanisms for other settings in the future.
- We use Griffiths’ inequality and Weitz’s self-avoiding walk [65] to prove the uniform dominance property in the Ising model. The resulting correlation bounds may be of independent interest.
- We test our mechanisms on real-world data (sushi preference dataset [26, 27] and Last.fm dataset [8]). Even though these datasets do not perfectly satisfy our theoretical assumptions, our mechanisms still provide a stronger incentive for truthful reporting compared to misreporting.

**Related work.** Information elicitation has two settings according to whether verification is possible. Our paper focuses on elicitation without verification.

For information elicitation without verification, Miller et al. [41] introduce the first mechanism for single task signal elicitation that has truth-telling as a strict Bayesian Nash equilibrium but requires full knowledge of the common prior. Bayesian truth serum (BTS) [47] is the first strongly truthful peer prediction mechanism, but requires complicated reports from agents (their private signal and predictions on other’s reports). A series of works [48, 49, 67, 66, 68, 3, 51, 31] relax certain assumptions of BTS but still require complicated reports from agents. Dasgupta and Ghosh [11] introduces the multi-task setting where agents are assigned batch iid tasks and only report their signals. Several works extend this to multiple-choice questions [29, 33, 56, 11], predictions [37], or even continuous value [50], and investigate the limitation and robustness [52, 6, 70, 17, 70]. Another related line of work is co-training and federated learning, which wants to elicit models [32, 36], or samples [64] when multiple iid data or feature of data are available. For more related works, see Faltings [16].

One popular line of work considers information elicitation when verification is possible. Spot-checking requires direct verification of the agent’s report [21]. Recent work on comparison data elicitation [19] utilizes spot-checking concepts and focuses on incentivizing effort. Another form of verification involves using additional samples to evaluate how the agent’s reports improve model performance [1, 28]. Additionally, the verification may have a general relation to the agent’s signal, e.g., proper scoring rules [23, 46, 35, 20].

## 2 Problem Formulation

We discuss our model for eliciting comparison data in this section and defer the extensions to Section 5. Given a collection of items  $\mathcal{A}$  and a set of strategic agents  $\mathcal{N}$ . Agents privately observe

noisy comparisons between pairs of items. Our goal is to design mechanisms to truthfully elicit agents' private information. We will first introduce the information structure of agents' private information of pairwise comparisons in Section 2.1 and then define the information elicitation problem in Section 2.2.

## 2.1 Bayesian SST Models for Comparison Data

We introduce Bayesian Strong Stochastic Transitivity (Bayesian SST) models to capture the structure of agents' private information for comparison data.

Given the set of items  $\mathcal{A}$ , the underlying unknown state about the items is  $\theta \in \Theta$ .  $\theta$  can be the vector of quality scores for the items (Example 2.2) or a reference ranking (Example 2.4).  $\theta$  is drawn according to a common prior  $P_\Theta$ :  $\theta \sim P_\Theta$ . Any realized  $\theta$  has an associated *stochastic comparison function*  $T_\theta : \mathcal{A}^2 \rightarrow \{-1, 1\}$ . For comparisons of two items  $a$  and  $a'$ ,  $T_\theta(a, a')$  and  $T_\theta(a', a)$  stochastically take value 1 or  $-1$ , with  $\Pr[T_\theta(a, a') = 1] = 1 - \Pr[T_\theta(a', a) = -1]$ . For any  $\theta$ ,  $T_\theta$  is strongly stochastically transitive as defined below.

**Definition 2.1** ([60, 13]). A stochastic comparison function,  $T : \mathcal{A}^2 \rightarrow \{-1, 1\}$ , is *strongly stochastically transitive (SST)* if for all  $a, a', a'' \in \mathcal{A}$  with  $\Pr[T(a, a') = 1] > 1/2$  and  $\Pr[T(a', a'') = 1] > 1/2$ , we have

$$\Pr[T(a, a'') = 1] > \max\{\Pr[T(a, a') = 1], \Pr[T(a', a'') = 1]\}.$$

Intuitively, a comparison function is SST when for any three items  $a, a', a''$ , if  $a$  is more favorable than  $a'$  and  $a'$  is more favorable than  $a''$ , then  $a$  is even more favorable than  $a''$ . The concept of SST is a well-established property of comparisons in social science and psychology [18].

Each agent  $i \in \mathcal{N}$  has the knowledge of  $(T_\theta)_{\theta \in \Theta}$  and  $P_\Theta$ . When asked to compare a pair of items  $(a, a')$ , the agent observes an independent draw according to the stochastic comparison function:  $s_i = T_\theta(a, a')$ , where realization  $s_i = 1$  represent item  $a$  is preferred over item  $a'$  by agent  $i$ . We assume items are a priori similar but ex-post distinct so that for all  $a, a' \in \mathcal{A}$ ,  $\mathbb{E}[T_\theta(a, a')] = \mathbb{E}[\mathbb{E}[T_\theta(a, a') | \theta]] = 0$  and  $\mathbb{E}[T_\theta(a, a') | \theta] \neq 0$  for all  $\theta$ .

**Examples of Bayesian SST models.** Bayesian SST models are a general family of models that take many classical parametric ranking models, including Bradley-Terry Luce [4, 38], Thurstone (Case V) [58], and Mallows  $\eta$ -model [39], as special cases.

**Example 2.2** (Bradley-Terry-Luce, Thurstone model, and more [59]). Let  $\theta \in \mathbb{R}^{\mathcal{A}} = \Theta$  where each coordinate is independently and identically sampled from a fixed non-atomic distribution  $\nu$  on  $\mathbb{R}$ , and each item  $a$  have a scalar quality  $\theta_a \in \mathbb{R}$ . Let  $F : \mathbb{R} \rightarrow [0, 1]$  be any strictly increasing function such that  $F(t) = 1 - F(-t)$  for all  $t \in \mathbb{R}$ . Conditional on a fixed  $\theta$ ,

$$\Pr[T_\theta(a, a') = 1] = F(\theta_a - \theta_{a'}) \text{ for all } a, a' \in \mathcal{A}.$$

This model recovers the Thurstone model [58] by setting  $F(t) = \Phi(t)$  where  $\Phi$  is the Gaussian CDF, and the Bradley-Terry-Luce model [4] by setting  $F(t) = \frac{e^t}{1+e^t}$ , the sigmoid function. Moreover, this model also contains any additive random utility model [2] where  $T(a, a') = 1$  if  $\theta_a + Z > \theta_{a'} + Z'$  with iid noise  $Z$  and  $Z'$ , because we can set  $F$  to be the CDF of the difference of two iid noise.

**Proposition 2.3.** *For any strictly increasing  $F$  and non-atomic  $\nu$  on  $\mathbb{R}$ , the parametric model in Example 2.2 is a Bayesian SST model.*

**Example 2.4** (Mallows  $\eta$ -model [39]). Let  $\Theta$  be the set of rankings on  $\mathcal{A}$  and  $\eta > 0$  be a dispersion parameter. Given a reference ranking  $\theta \in \Theta$ , the *Mallows  $\eta$ -model* generate a ranking  $\phi \in \Theta$  with probability  $\Pr(\phi) \propto \exp(-\eta d(\theta, \phi))$  where  $d(\theta, \phi) = |\{(a, a') \in \mathcal{A}^2 : \theta(a) < \theta(a') \text{ and } \phi(a) > \phi(a')\}|$  is Kendall's tau distance, and  $\theta(a)$  is the rank of item  $a$ . Therefore, to generate comparisons, we first sample a uniform  $\theta$  and

$$\Pr[T_\theta(a, a') = 1] = \sum_{\phi: \phi(a) > \phi(a')} \Pr(\phi), \text{ for all } a, a' \in \mathcal{A}.$$

**Proposition 2.5.** *For any  $\eta > 0$ , Mallows  $\eta$ -model in Example 2.4 with uniform distribution on reference ranking is an Bayesian SST model.*

The proofs for propositions 2.3 and 2.5 are closely related to strong stochastic transitivity [54, 7], but are provided in the appendix for completeness.

## 2.2 Peer Prediction Mechanism Design

To truthfully elicit comparison data from agents, a peer prediction mechanism creates a game between the agents outlined below: First, we choose an *assignment*  $\mathcal{E} = \{e_i = (a_{u_i}, a_{v_i}) : i \in \mathcal{N}\}$  where agent  $i \in \mathcal{N}$  gets a pair of items  $e_i = (a_{u_i}, a_{v_i}) \in \mathcal{A}^2$  to compare. Then each agent  $i \in \mathcal{N}$  privately observes the realization of the comparison (signal)  $s_i \in \{-1, 1\}$ , which is an independent realization of  $T_\theta(a_{u_i}, a_{v_i})$ , and reports  $\hat{s}_i \in \{-1, 1\}$  potentially different from her signal. We use  $S_i = S(a_{u_i}, a_{v_i})$  to denote the random variable of agent  $i$ 's signal, where the randomness of  $S(\cdot, \cdot)$  comes from both  $\theta$  and  $T_\theta$ . Let  $\mathbf{S}$  represent the random vector of all agents' signals,  $\mathbf{s} = (s_i)_{i \in \mathcal{N}}$  be all agents' realized private signals and  $\hat{\mathbf{s}} = (\hat{s}_i)_{i \in \mathcal{N}}$  be all agents' reports. Finally, a *peer prediction mechanism*  $(M_i)_{i \in \mathcal{N}}$  takes all agents' reports  $\hat{\mathbf{s}}$  and pays agent  $i$  with  $M_i(\hat{\mathbf{s}}) \in \mathbb{R}$ .

Each agent  $i$ 's strategy is a random function from her signal to a report  $\sigma_i : s_i \mapsto \hat{s}_i$ , and the randomness of their strategies is independent of each other's and all signals. With slight abuse of notation, we write  $\sigma_i(s_i, \hat{s}_i) = \Pr[\hat{S}_i = \hat{s}_i | S_i = s_i]$  as the conditional probability of reporting  $\hat{s}_i$  given private signal  $s_i$ . A strategy profile  $\sigma$  is a collection of all agent's strategies. All agents are rational and risk-neutral, so they want to maximize their expected payments. Thus, given prior  $P_\Theta$ , randomness of  $T_\theta$  and a strategy profile  $\sigma$ , agent  $i$  wants to maximize her ex-ante payment denoted as  $\mathbb{E}_{\sigma, \theta, T_\theta}[M_i(\hat{\mathbf{s}})]$  where  $\hat{\mathbf{s}}$  is the random vector of all agents' report that depends on the signals  $\mathbf{S}$  and strategy profile  $\sigma$ .

We introduce three families of strategies, truth-telling, permutation, and uninformed strategy profiles, which are central to understanding effective peer prediction mechanisms.

- A strategy  $\sigma_i$  is *truthful* (or truth-telling) if it is a deterministic identity map,  $\sigma_i(s_i) = s_i$ . A strategy profile is truthful if all agents' strategies are truthful.
- A *permutation strategy profile* is where agents simultaneously relabel their signals and then report the relabeled ones. A permutation strategy is indistinguishable from truth-telling unless the peer prediction mechanism has additional knowledge about the prior signal distribution.[33]
- Finally, a strategy is *uninformed* if it has the same report distribution across all signals, and it is informed otherwise. Common examples include consistently reporting all signals as a constant value, such as 1 or -1, or using a random report regardless of the signal. Uninformed strategies are undesirable as the reports bear no relationship to the private signals.

A strategy profile is *symmetric* if all agents use the same strategy. For example, both truth-telling and permutation strategy profiles are symmetric.

We now introduce goals for a peer prediction mechanism that favors truth-telling more than other strategies. First, we want the truth-telling (strategy profile) to be a strict *Bayesian Nash equilibrium* (BNE) so that any agent's payment would strictly decrease if she unilaterally changes to any non-truthful strategy. Moreover, there may be multiple equilibria, and a desirable mechanism should ensure that truth-telling is better than all other equilibria. In this paper, we aim for symmetrically strongly truthful mechanisms defined below.

**Definition 2.6.** A peer prediction mechanism is *symmetrically strongly truthful* if truth-telling is a BNE, and each agent's expected payment in truth-telling is no less than the payment in any other symmetric equilibrium with equality for the equilibrium with a permutation strategy profile.<sup>1</sup>

## 3 Bonus-penalty Payment Mechanism for Comparison Data

We now propose a *bonus-penalty payment mechanism* for eliciting comparison data. The mechanism makes use of a *bonus-penalty payment* function, which can be seen as an agreement payment and introduced by Dasgupta and Ghosh [11] in a different context (see discussion in appendix A). Formally, for any  $\hat{s}_i, \hat{s}_j, \hat{s}_k \in \{-1, 1\}$ , the bonus-penalty payment function is

$$U^{BPP}(\hat{s}_i, \hat{s}_j, \hat{s}_k) = \hat{s}_i \hat{s}_j - \hat{s}_i \hat{s}_k = 2(\mathbf{1}[\hat{s}_i = \hat{s}_j] - \mathbf{1}[\hat{s}_i = \hat{s}_k]), \quad (1)$$

which rewards when the first input agrees with the second but punishes when it agrees with the third.

<sup>1</sup>Kong and Schoenebeck [30] shows that it is impossible to pay the truth-telling strategy profile strictly better than other permutation strategy profiles without additional knowledge of the prior signal distribution.

Mechanism 1 uses the bonus-penalty payment eq. (1) for each agent  $i$  by carefully choosing agent  $j$  and  $k$  such that agent  $j$ 's signal is more likely to agree with agent  $i$ 's than agent  $k$ 's signal is. The crux of finding such pair of agents is to show that if agent  $i$  prefers item  $a$  over  $a'$ , she would expect that others will prefer any third item  $a''$  over  $a'$ , and prefer  $a$  over  $a''$ . Thus, if agent  $j$  has pair  $(a'', a')$  and agent  $k$  has pair  $(a'', a)$ , then agent  $j$ 's signal is more likely to take the same value as  $i$ 's than agent  $k$ 's signal is. This is the main idea behind the proof of Theorem 3.1, where we establish the symmetrically strongly truthfulness of Mechanism 1. To ensure the existence of such pairs are assigned, we require the assignment  $\mathcal{E}$  to be **admissible** where for all  $(a, a') \in \mathcal{E}$ , there exists  $a'' \in \mathcal{A}$  so that  $(a'', a')$  and  $(a'', a) \in \mathcal{E}$ .

---

**Mechanism 1:** BPP mechanism for comparison data

---

**Input:** Let  $\mathcal{A}$  be a collection of items,  $\mathcal{E}$  be an admissible assignment, and  $\mathbf{s}$  be agents' reports.

**for agent  $i \in \mathcal{N}$  with pair  $e_i = (a_{u_i}, a_{v_i}) = (a, a')$  do**

    Find  $a'' \in \mathcal{A}$  and two agents  $j$  and  $k$  so that  $e_j = (a'', a')$  and  $e_k = (a'', a)$ , and pay agent  $i$

$$M_i(\mathbf{s}) = U^{BPP}(\hat{s}_i, \hat{s}_j, \hat{s}_k) = \hat{s}_i \hat{s}_j - \hat{s}_i \hat{s}_k. \quad (2)$$


---

**Theorem 3.1.** *Given a collection of items  $\mathcal{A}$  and a set of agents  $\mathcal{N}$  with  $|\mathcal{A}|, |\mathcal{N}| \geq 3$ , for any admissible assignment matrix  $\mathcal{E}$  and Bayesian SST model with  $(T_\theta)_{\theta \in \Theta}$  and  $P_\Theta$ , the BPP mechanism for comparison (Mechanism 1) is symmetrically strongly truthful.*

We defer the proof of theorem 3.1 to section 4. The admissible condition imposes little overhead on downstream learning problems, including rank recovery [25] and identification of the top  $k$  items [15]. Specifically, the size of assignment  $\mathcal{E}$  is the number of comparisons and corresponds to the sample complexity for these learning problems. If a learning algorithm requires a set of pairs to compare  $\mathcal{E}^{ML}$ , we can construct an admissible superset  $\mathcal{E}$  that introduces a constant factor overhead and can recover  $\mathcal{E}^{ML} \subseteq \mathcal{E}$ .<sup>2</sup>

We remark that the bonus-penalty payment function eq. (2) can be seen as a boolean function for *transitivity* [45]; see remark 3.2 for a formal statement. Hence, theorem 3.1 implies that agents' manipulations can only decrease the probability of transitivity among their reports.

**Remark 3.2.** Note that a deterministic comparison function  $t : \mathcal{A} \times \mathcal{A} \rightarrow \{-1, 1\}$  satisfies transitivity on three items  $a, a', a'' \in \mathcal{A}$  if and only if  $t(a, a'), t(a', a''), t(a'', a)$  are not all equal, that is  $NAE(t(a, a'), t(a', a''), t(a'', a)) = 1$  where

$$NAE(w_1, w_2, w_3) = \frac{3}{4} - \frac{1}{4}w_1w_2 - \frac{1}{4}w_1w_3 - \frac{1}{4}w_2w_3.$$

The agent's random noisy comparisons may or may not satisfy transitivity. The probability of transitivity is the probability that they do.

We can show that the bonus-penalty payment in eq. (2) is equivalent to the above transitivity test when agents are truth-telling. Formally,

$$\begin{aligned} & NAE(S(a, a'), S(a', a''), S(a'', a)) \\ &= \frac{3}{4} - \frac{1}{4}(S(a, a')S(a', a'') + S(a, a')S(a'', a) + S(a', a'')S(a'', a)) \\ &= \frac{1}{4}(S(a, a')S(a'', a') - S(a, a')S(a'', a)) + \frac{3}{4} + \frac{1}{4}S(a'', a')S(a'', a) \quad (S(a', a'') = -S(a'', a')) \\ &= \frac{1}{4}(s_i s_j - s_i s_k) + \frac{3}{4} + \frac{1}{4}s_j s_k = \frac{1}{4}M_i(\mathbf{s}) + \frac{3}{4} + \frac{1}{4}s_j s_k \quad (\text{truth-telling}) \end{aligned}$$

Therefore,  $\arg \max_{\hat{s}_i} \mathbb{E}[NAE(\hat{s}_i, -S_j, S_k) | S_i = s_i] = \arg \max_{\hat{s}_i} \mathbb{E} \left[ \frac{1}{4}M_i(\hat{s}_i, S_j, S_k) + \frac{3}{4} + \frac{1}{4}s_j s_k | S_i = s_i \right] = \arg \max_{\hat{s}_i} \mathbb{E}[M_i(\hat{s}_i, S_j, S_k) | S_i = s_i]$ .

<sup>2</sup>Specifically, given any assignment  $\mathcal{E}^0$ , we can construct a superset  $\mathcal{E}$  so that for any  $(a, a') \in \mathcal{E}^0$  find an arbitrary  $a'' \neq a, a'$  and add  $(a, a'), (a', a), (a, a''), (a'', a), (a', a''), (a'', a')$  into  $\mathcal{E}$ . Thus,  $\mathcal{E}$  is admissible and at most six times larger than  $\mathcal{E}^0$ .

## 4 Proof of Theorem 3.1: from Bayesian SST Model to Uniform Dominance

To prove theorem 3.1, we formalize the idea that agent  $j$ 's signal is more likely to agree with agent  $i$ 's than agent  $k$ 's is as what we call uniform dominance in definition 4.1. We'll show that any Bayesian SST model satisfies this property. Then, we'll prove that BPP mechanism is symmetrically strongly truthful when agents' private signals satisfy uniform dominance.

**Definition 4.1.** Given a random vector  $(S_i, S_j, S_k) \in \{-1, 1\}^3$  with joint distribution  $P$ ,  $S_j$  **uniformly dominates**  $S_k$  for  $S_i$  if  $\Pr[S_j = 1 | S_i = 1] > \Pr[S_k = 1 | S_i = 1]$  and  $\Pr[S_j = -1 | S_i = -1] > \Pr[S_k = -1 | S_i = -1]$ . We call such an ordered tuple  $\langle S_i, S_j, S_k \rangle$  a uniformly dominant tuple.<sup>3</sup>

Lemma 4.2 shows how to identify uniformly dominant tuples under Bayesian SST models.

**Lemma 4.2.** Under any Bayesian SST model, for any agent  $i$  and items  $a, a'$  and  $a''$ , agent  $j$ 's signal  $S_j = S(a'', a')$  uniformly dominates agent  $k$ 's signal  $S_k = S(a'', a)$  for signal  $S_i = S(a, a')$ .

In other words, under any Bayesian SST model, the distribution of  $S(a, a'), S(a'', a'), S(a'', a)$  satisfies uniform dominance for any  $a, a', a''$ . In the rest of this section, we can view  $(S_i, S_j, S_k)$  as an abstract random vector with some joint distribution  $P$ .

We now establish some implications of uniform dominance on the bonus-penalty payment. Lemma 4.3 shows that truth-telling is the best response if other signals are reported truthfully. Lemma 4.4 states that the expected payment is zero if everyone uses uninformed strategies (random functions independent of input). Lemma 4.5 characterizes the best response under symmetric strategy profiles (the same random function on each coordinate).

**Lemma 4.3 (Truthfulness).** Given a uniformly dominant tuple  $\langle S_i, S_j, S_k \rangle$  with distribution  $P$ , for all  $s_i \in \{-1, 1\}$ ,  $s_i = \arg \max_{\hat{s}_i \in \{-1, 1\}} \mathbb{E}_P [U^{BPP}(\hat{s}_i, S_j, S_k) | S_i = s_i]$  and  $\mathbb{E}_P [U^{BPP}(S_i, S_j, S_k)] > 0$ .

**Lemma 4.4.** Given a uniformly dominant tuple  $\langle S_i, S_j, S_k \rangle$ , with joint distribution  $P$  if agent  $j$  and  $k$  both use an uninformed strategy  $\sigma$  so that  $\hat{S}_j = \sigma(S_j)$  and  $\hat{S}_k = \sigma(S_k)$ , for all  $s_i$  and  $\hat{s}_i$  in  $\{-1, 1\}$ ,  $\mathbb{E}_{\sigma, P} [U^{BPP}(\hat{s}_i, \hat{S}_j, \hat{S}_k) | S_i = s_i] = 0$ .

**Lemma 4.5.** Given a uniformly dominant tuple  $\langle S_i, S_j, S_k \rangle$  with distribution  $P$ , for any strategy  $\sigma$  and  $s_i \in \{-1, 1\}$  when agent  $j$  and  $k$  both use  $\sigma$ ,  $\arg \max_{\hat{s}_i \in \{-1, 1\}} \mathbb{E}_{\sigma, P} [U^{BPP}(\hat{s}_i, \hat{S}_j, \hat{S}_k) | S_i = s_i] = \arg \max_{\hat{s}_i \in \{-1, 1\}} \{\sigma(s_i, \hat{s}_i) - \sigma(-s_i, \hat{s}_i)\}$ .

We'd like to highlight that lemmas 4.3 to 4.5 as well as the proof of theorem 3.1 below hold for any uniformly dominant tuple  $\langle S_i, S_j, S_k \rangle$ , not necessarily derived from the Bayesian SST model. This offers a path to generalize our mechanism for comparison data to other settings.

*Proof of theorem 3.1.* By lemma 4.2, for any agent  $i$ , the associated agent  $j$ 's signal  $S_j = S(a'', a')$  uniformly dominates the associated  $k$ 's signal  $S_k = S(a'', a)$  for signal  $S_i = S(a, a')$ . By lemma 4.3, if agent  $j$  and  $k$  are truthful, agent  $i$ 's best response is truthful reporting, so truth-telling is a BNE.

Now we show that all other symmetric equilibria are permutation or uninformed equilibria. For any symmetric equilibrium  $\sigma = (\sigma_\iota)_{\iota \in \mathcal{N}}$  so that everyone uses the same strategy  $\sigma_\iota = \sigma$  for all  $\iota \in \mathcal{N}$ . If  $\sigma$  is not deterministic so that  $\sigma(s, s), \sigma(s, -s) > 0$  for some  $s \in \{-1, 1\}$ , agent  $i$  must be indifferent between reporting  $s$  and  $-s$  when getting  $S_i = s$ .  $\sigma(s, s) - \sigma(-s, s) = \sigma(s, -s) - \sigma(-s, -s)$  by lemma 4.5. This means  $\sigma(s, s) = \sigma(-s, s)$  and  $\sigma(s, -s) = \sigma(-s, -s)$ , and  $\sigma$  is an uninformed strategy. If the strategy is deterministic, there are two cases. If  $\sigma(s) = \sigma(-s)$ , the strategy is also uninformed. If  $\sigma(s) \neq \sigma(-s)$ ,  $\sigma$  is either truth-telling  $s \mapsto s$  or flipping  $s \mapsto -s$  for all  $s$ .

Finally, by lemma 4.4, any uninformed equilibrium's expectation is zero. Additionally, because eq. (1) is invariant when all inputs are flipped, the truth-telling and flipping/permuation equilibria has the same expected payment which is positive by lemma 4.3.  $\square$

<sup>3</sup>If we view  $S_j$  and  $S_k$  as two statistical tests for a binary event  $S_i$ , the two inequalities in definition 4.1 say that  $S_j$  has a better type II and type I error than  $S_k$  respectively.

## 5 Generalization of Bonus-penalty Payment Mechanisms

We now leverage the key idea of uniform dominance to design peer prediction mechanisms for networked data in section 5.1. In section 5.2, we summarize our design approach as a general scheme that first identifies uniform dominance structures and then engages the bonus-penalty payment. We prove the uniqueness of bonus-penalty payment: it is the only payment function, up to some positive affine transformation, that induces truth-telling as a strict BNE for all uniform dominant tuples.

### 5.1 Bonus-penalty Payment Mechanisms for Networked Data

Uniform dominance implies agent  $i$ 's signal is more likely to agree with agent  $j$ 's than with agent  $k$ 's. Social networks are another natural domain exhibiting this property, as homophily [40] suggests that agents' opinions or signals in a social network are more likely to agree with their friends than with non-friends. Leveraging this insight, we use a bonus-penalty payment scheme to elicit binary networked data.

---

#### Mechanism 2: BPP mechanism for networked data

---

**Input:** Let  $(V, E)$  be a graph of agents in  $V$ ,  $\hat{\mathbf{s}} \in \{-1, 1\}^V$  from all agent's reports.

**for** agent  $i \in V$  **do**

Find agents  $j$  (friend) and  $k$  (non-friend) so that  $(i, j) \in E$  but  $(i, k) \notin E$ , and pay agent  $i$

$$M_i(\hat{\mathbf{s}}) = U^{BPP}(\hat{s}_i, \hat{s}_j, \hat{s}_k) = \hat{s}_i \hat{s}_j - \hat{s}_i \hat{s}_k. \quad (3)$$


---

Below, we provide a theoretical guarantee for our mechanism under a popular graphical model for social network data, *Ising model* [14, 43], which captures the correlation between agents and their friends. Formally, an Ising model consists of an undirected graph  $(V, E)$  and correlation parameter  $\beta_{i,j} \geq 0$  for each edge  $(i, j) \in E$ . Each agent is a node in the graph,  $\mathcal{N} = V$ , and has a binary private signal (1 or  $-1$ ) jointly distributed as the following: For all  $\mathbf{s} = (s_i)_{i \in V} \in \{-1, 1\}^V$ ,  $\Pr_{\beta}[\mathbf{S} = \mathbf{s}] \propto \exp(H(\mathbf{s}))$  where the energy function is  $H(\mathbf{s}) = \sum_{(i,j) \in E} \beta_{i,j} s_i s_j$ .

**Theorem 5.1.** *If agents' signals are sampled from an Ising model on undirected graph  $(V, E)$  with correlation parameters  $\beta$ , Mechanism 2 is symmetrically strongly truthful, when  $\frac{2\beta}{d} > \ln \frac{e^{2(d+1)\bar{\beta}} + 1}{e^{2\bar{\beta}} + e^{2d\bar{\beta}}}$  where  $\underline{\beta} = \min_{(i,j) \in E} \beta_{i,j}$ ,  $\bar{\beta} = \max_{(i,j) \in E} \beta_{i,j}$ , and  $d$  is the maximal degree of graph  $(V, E)$ .*

Mechanism 2 does not require knowledge about parameters of the Ising model, but only the connection of the network  $(V, E)$ . Social network platforms, which already possess this knowledge, can easily integrate our mechanism when conducting surveys. Additionally, snowball sampling [24], which relies on participants referring their friends, is also naturally compatible with our mechanism.

The complete proof of theorem 5.1 is quite technical and is deferred to the appendix, where we also explain why the bound between  $\beta$  and  $d$  is necessary. Below, we provide a sketch of the proof.

*Proof sketch for theorem 5.1.* As discussed in section 4, we only need to show that for any agent  $i$ , for all agent  $j$  with  $(i, j) \in E$  and  $k$  with  $(i, k) \notin E$ ,  $j$ 's signal uniformly dominates  $k$ 's signal for  $i$ 's signal. Because the energy function  $H(\mathbf{s})$  above remains invariant when the signs are flipped,  $\Pr[S_i = 1] = \Pr[S_j = 1] = \Pr[S_k = 1] = 1/2$ , it is sufficient to prove that

$$\Pr[S_i = 1 \mid S_j = 1] > \Pr[S_i = 1 \mid S_k = 1]. \quad (4)$$

We then prove a lower bound for the left-hand side and an upper bound for the right-hand side separately. For the left-hand side, we use the Griffiths' inequality [44] to show that the minimum value of  $\Pr[S_i = 1 \mid S_j = 1]$  happens when  $j$  is the only friend of  $i$ . For the right-hand side, we use Weitz's self-avoiding walk [65] and reduce any graph with maximum degree  $d$  into a  $d$ -ary tree.  $\square$

### 5.2 General Design Scheme and Uniqueness

The design of BPP mechanisms for comparison data and networked data has suggested a general design scheme for other elicitation settings. That is, if one can *identify a uniformly dominant tuple for*

---

**Mechanism 3:** General design scheme using BPP

---

**Input:** Let  $\hat{s}$  be reports from agents in  $\mathcal{N}$ .

**for** agent  $i \in \mathcal{N}$  **do**

    Find two agents  $j$  and  $k$  so that  $j$ 's signal uniformly dominates  $k$ 's for  $i$ 's, and pay agent  $i$

$$M_i(\hat{s}) = U^{BPP}(\hat{s}_i, \hat{s}_j, \hat{s}_k) = \hat{s}_i \hat{s}_j - \hat{s}_i \hat{s}_k.$$

---

each agent, adopting the bonus-penalty payment gives a symmetrically strongly truthful mechanism. We further show that the bonus-penalty payment is in some sense unique.

**Theorem 5.2.** *If for each agent  $i$  the associated agent  $j$ 's signal uniformly dominates  $k$ 's signal for  $i$ 's signal, the above scheme is symmetrically strongly truthful.*

When an agent  $i$  has multiple pairs of  $(j_1, k_1), \dots, (j_\ell, k_\ell)$  so that  $j_l$ 's signal uniformly dominates  $k_l$ 's for  $i$ 's for each  $l = 1, \dots, \ell$ , we may pay agent  $i$  the average of bonus-penalty payment on all pairs  $M_i(\hat{s}) = \frac{1}{\ell} \sum_{l=1}^{\ell} U^{BPP}(\hat{s}_i, \hat{s}_{j_l}, \hat{s}_{k_l})$ . This average maintains our symmetrically strongly truthful guarantee while potentially reducing the variance in payments.

Theorem 5.2 shows that bonus-penalty payment is a sufficient condition for designing good elicitation mechanisms for information structures with uniform dominance. We now prove it is also a necessary condition: any payment that induces truth-telling as a strict BNE under all uniformly dominant tuples must be an affine transformation of the bonus-penalty payment.

**Theorem 5.3** (Uniqueness). *A payment  $U : \{-1, 1\}^3 \rightarrow \mathbb{R}$  satisfies that, for all uniformly dominant tuples  $\langle S_i, S_j, S_k \rangle$ ,  $s_i = \arg \max_{\hat{s}_i \in \{-1, 1\}} \mathbb{E} [U(\hat{s}_i, S_j, S_k) \mid S_i = s_i]$ , if and only if there exist  $\lambda > 0$  and  $\mu : \{-1, 1\}^2 \rightarrow \mathbb{R}$  so that*

$$U(s_i, s_j, s_k) = \lambda U^{BPP}(s_i, s_j, s_k) + \mu(s_j, s_k), \text{ for all } s_i, s_j, s_k \in \{-1, 1\}$$

where choice of  $\mu$  does not affect the set of equilibria.

## 6 Experiments

We present experiments on real-world data to evaluate our models and mechanisms. We hope to cast insights on two questions empirically. Does our mechanism provide better rewards when all agents report truthfully than when all agents report randomly? Does our mechanism incentivize truth-telling for each agent if all other agents are truthful? We evaluate Mechanism 1 and 2 by comparing three settings, truth-telling, uninformed, and unilateral deviation, using *empirical cumulative distribution functions* (ECDF) on agents' payments. Each point on ECDF denotes the fraction of agents who get paid less than a particular value.

For both comparison and networked datasets (figs. 1 and 2), we find our mechanisms provide better payments to agents under truthful settings than the other two settings. The ECDF under truth-telling lies below the other two ECDFs, which is known as first-order stochastic dominance. This implies that the truth-telling strategy results in higher average, quantiles (e.g., first quartile, median, and third quartile), and a greater expectation of any monotone function on the empirical distribution than the other two settings. We provide additional

### 6.1 SUSHI Preference Dataset for Comparison Data

We consider preference data for a collection of 10 sushi items (item set A) [26, 27], and focus on a group of 249 agents. Each agent provides a complete ranking of all 10 types of sushi in the dataset. These agents are female, aged thirty to forty-nine, who took more than three hundred seconds to rank the items and mostly lived in Kanto and Shizuoka until age fifteen. We restrict the set of agents to avoid significant violations of transitivity across different agents and to better align with our model assumptions. In the appendix, we will present the experimental results for other groups of users and further test whether the dataset satisfies transitivity.

For the first question, we use Mechanism 1 to compute each agent's payment under the truth-telling or uninformed strategy profile. For each agent  $i$ , we 1) randomly sample three items  $a, a', a''$  and

two agents  $j, k$ , 2) derive agent  $i$ 's comparison on the first two items  $(a, a')$  from her ranking, (and similarly for agent  $j$ 's comparison on  $(a', a'')$ , and agent  $k$ 's comparison on  $(a, a'')$ ), 3) compute bonus-penalty payment on these three comparisons, 4) repeat the above procedure 100 times and pay agent  $i$  with the average of those 100 trials. For the uninformed strategy setting, we replace every agent's comparisons with uniform random bits and compute the payment. The left column of fig. 1 presents the ECDF of payments for the agents in both settings. The figure shows that in the uninformed random strategy setting only about 50% of the agents receive positive payments, while in the original dataset (truthful strategy setting) over 75% of the users receive positive payments. The right column of fig. 1 tests the second question if the agent has the incentive to deviate when every other agent is truthful. The truth-telling curve is identical to the left column of fig. 1. For unilateral deviation, each agent gets the above bonus-penalty payment when her comparisons are replaced by uniform random bits. We plot the ECDFs of payments for both settings in the right column of fig. 1. The figure shows that the ECDF of the unilateral deviation payments is above the ECDF of human users' payments, indicating that our mechanism pays more to the truth-telling agents.

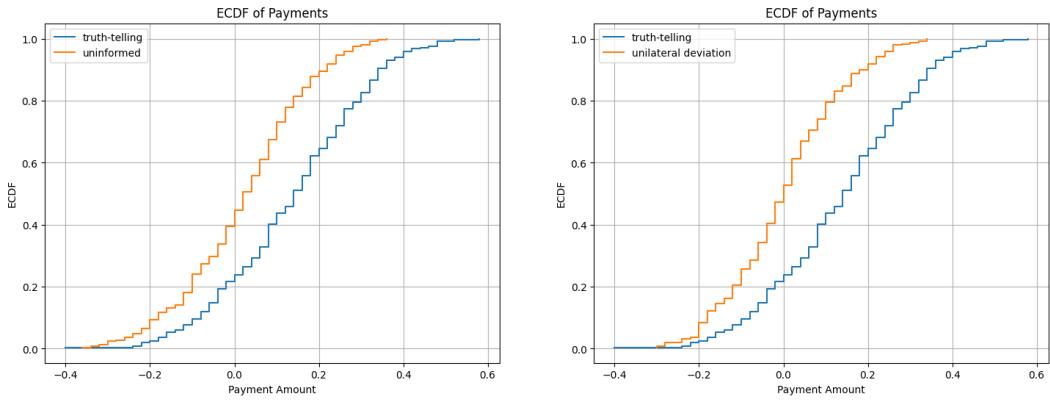


Figure 1: SUSHI preference dataset

## 6.2 Last.fm Dataset for Networked Data

We test our BPP mechanism on the Last.fm dataset from Cantador et al. [8]. This dataset consists of 1892 agents on Last.fm, forming a social network with 12704 edges and an average degree of 6.71. It records agents' top fifty favorite artists whom they have listened to the most. We note that, in the dataset, listener fractions for all artists are much smaller than non-listener fractions. This bias differs from our Ising model in section 5.1 where every agent has the same chance to get both signals. Thus, the result can be seen as a stress test for our mechanism even when the data deviate from the assumption of our theoretical results.

Figure 2 focuses on the most popular artist in the dataset, Lady Gaga, who has a listener fraction of 32.3%. The results for additional artists are presented in the supplementary material. The left column of fig. 2 tests the first question. Each agent has a binary signal about whether or not she listens to a particular artist (Lady Gaga in this section). For the truth-telling setting, everyone reports her signal truthfully and gets payment by the bonus-penalty payment (formally defined in section 5.1). For the uninformed setting, everyone gets the bonus-penalty payment when all reports are iid according to the prior (0.322 for Lady Gaga). When everyone is truthful, more than 76% of agents get positive payments and have an average payment of 0.37 for Lady Gaga, while when agents report randomly, only half get positive payments, and have a near zero average payment. These results suggest that agents got more incentive to choose the truth-telling equilibrium than the uninformed equilibrium. The right column of fig. 2 tests the second question. The truth-telling curve is identical to the left column of fig. 2. For the unilateral deviation setting, each agent gets the bonus-penalty payment when she reports listener/non-listener uniformly at random. The unilateral deviation's payment is worse than the payments for truth-telling, decreasing from 0.37 to near zero for Lady Gaga.

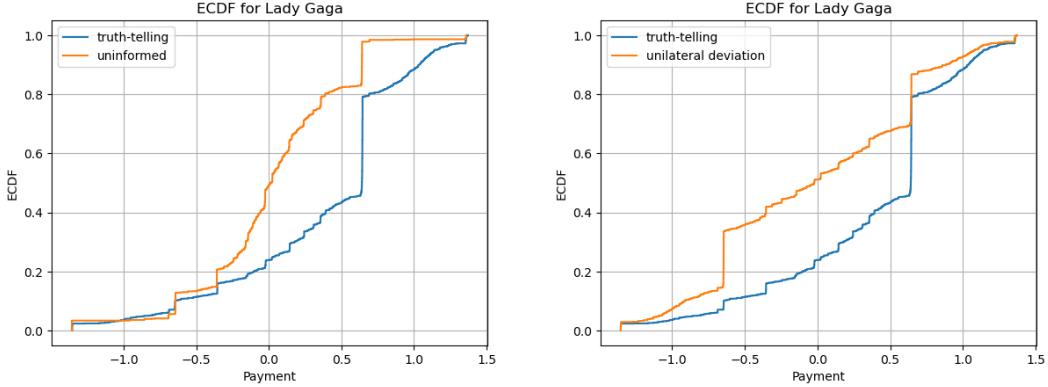


Figure 2: Last.fm dataset for Lady Gaga

## 7 Conclusion and Discussion

We introduce a symmetrically strongly truthful peer prediction mechanism for eliciting comparison data without verification and extend it to eliciting networked data under Ising models. Our mechanisms are evaluated using real-world data. A key insight from our work is the identification of a structure we term “uniform dominance,” which suggests a path for designing mechanisms in more complex elicitation settings. For example, in time-series data, adjacent points tend to be more related than distant ones, and in contextual settings, feedback from similar contexts is typically more related than from different contexts.

A central assumption in this study is that agents are *a priori similar*. Hence, noisy comparisons of item pairs are independent of the assigned agent’s identity. This assumption is reasonable for items with widely agreed-upon rankings, such as quality assessments of large language model (LLM) outputs. However, it may break down in settings where preferences are highly polarized, such as political opinions or social choice problems<sup>4</sup>. Despite this, our additional experiments in appendix F, which relax the selection rule used in obtaining fig. 1, show that the mechanism remains robust even when some dissimilarities among agents exist.

Agents in our model are assumed to focus solely on maximizing their payments, without accounting for efforts or external incentives such as minimizing others’ rewards or intentionally distorting rankings. While our mechanism may be extended to handle binary effort as suggested in previous work [11, 57], accommodating more than two effort levels would require additional assumptions [69]. Moreover, one may hope to incorporate the designer’s utility, by factoring in downstream learning problems along with elicitation payments. This would necessitate a significant overhaul of the existing learning framework.

Our mechanisms achieve a symmetric, strongly truthful equilibrium. This does not rule out the existence of non-symmetric equilibria with potentially higher utility. However, such equilibria would require complex coordination among agents, making them less likely to arise naturally.

From a technical standpoint, our approach involves several assumptions that can be generalized or relaxed. Our Bayesian SST model, which relies on strong stochastic transitivity, serves as a non-parametric extension of several widely used parametric ranking models. In appendix C.2, we present both positive and negative results regarding weaker notions of transitivity (e.g., [5]). While we assume admissible assignments, this can be relaxed to random assignments with full support. Additionally, limited liability can be ensured in our mechanism. For example, adding a constant of 1 to the payment function in eq. (2) ensures that the payment is either 2 or 0.

---

<sup>4</sup>For example, when ranking phone features (e.g., innovation, performance, brand reputation, price, ease of use), consumers often fall into two groups: early adopters, who prioritize cutting-edge technology, and late adopters, who favor stability, affordability, and ease of use. Their opposing preferences violate the *a priori* similarity assumption. Imagine an early adopter whose payment in eq. (2) depends on two late adopters. Since their preferences may differ significantly, the early adopter might have an incentive to misreport her preferences.

## Acknowledgments

This research was partially supported by the National Science Foundation under grant no. IIS-2147187.

## References

- [1] Jacob D Abernethy and Rafael Frongillo. A collaborative mechanism for crowdsourcing prediction problems. *Advances in neural information processing systems*, 24, 2011.
- [2] Hossein Azari, David Parks, and Lirong Xia. Random utility theory for social choice. *Advances in Neural Information Processing Systems*, 25, 2012.
- [3] Aurélien Baillon. Bayesian markets to elicit private information. *Proceedings of the National Academy of Sciences*, 114(30):7958–7962, 2017.
- [4] Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- [5] Mark Braverman and Elchanan Mossel. Noisy sorting without resampling. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’08, page 268–276, USA, 2008. Society for Industrial and Applied Mathematics.
- [6] Noah Burrell and Grant Schoenebeck. Measurement integrity in peer prediction: A peer assessment case study. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pages 369–389, 2023.
- [7] Róbert Busa-Fekete, Eyke Hüllermeier, and Balázs Szörényi. Preference-based rank elicitation using statistical models: The case of mallows. In *International conference on machine learning*, pages 1071–1079. PMLR, 2014.
- [8] Iván Cantador, Peter Brusilovsky, and Tsvi Kuflik. Second workshop on information heterogeneity and fusion in recommender systems. In *Proceedings of the fifth ACM conference on Recommender systems*, pages 387–388, 2011.
- [9] Mingyue Cheng, Hao Zhang, Jiqian Yang, Qi Liu, Li Li, Xin Huang, Liwei Song, Zhi Li, Zhenya Huang, and Enhong Chen. Towards personalized evaluation of large language models with an anonymous crowd-sourcing platform. In *Companion Proceedings of the ACM on Web Conference 2024*, pages 1035–1038, 2024.
- [10] Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E Gonzalez, et al. Chatbot arena: An open platform for evaluating llms by human preference. *arXiv preprint arXiv:2403.04132*, 2024.
- [11] Anirban Dasgupta and Arpita Ghosh. Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd international conference on World Wide Web*, pages 319–330, 2013.
- [12] Constantinos Daskalakis, Nishanth Dikkala, and Gautam Kamath. Concentration of multilinear functions of the ising model with applications to network data. *Advances in Neural Information Processing Systems*, 30, 2017.
- [13] Donald Davidson and Jacob Marschak. Experimental tests of a stochastic decision theory. *Measurement: Definitions and theories*, 17(2), 1959.
- [14] Glenn Ellison. Learning, local interaction, and coordination. *Econometrica: Journal of the Econometric Society*, pages 1047–1071, 1993.
- [15] Brian Eriksson. Learning to top-k search using pairwise comparisons. In *Artificial Intelligence and Statistics*, pages 265–273. PMLR, 2013.
- [16] Boi Faltings. Game-theoretic mechanisms for eliciting accurate information. In *IJCAI*, 2022.

[17] Shi Feng, Fang-Yi Yu, and Yiling Chen. Peer prediction for learning agents. *Advances in Neural Information Processing Systems*, 35:17276–17286, 2022.

[18] Peter C Fishburn. Binary choice probabilities: on the varieties of stochastic transitivity. *Journal of Mathematical psychology*, 10(4):327–352, 1973.

[19] Kiriaki Frangias, Andrew Lin, Ellen Vitercik, and Manolis Zampetakis. Algorithmic contract design for crowdsourced ranking. *arXiv preprint arXiv:2310.09974*, 2023.

[20] Rafael Frongillo and Ian A Kash. Vector-valued property elicitation. In *Conference on Learning Theory*, pages 710–727. PMLR, 2015.

[21] Xi Alice Gao, James R Wright, and Kevin Leyton-Brown. Incentivizing evaluation with peer prediction and limited access to ground truth. *Artificial Intelligence*, 275:618–638, 2019.

[22] Hans-Otto Georgii. *Gibbs measures and phase transitions*. Walter de Gruyter GmbH & Co. KG, Berlin, 2011.

[23] Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.

[24] Leo A Goodman. Snowball sampling. *The annals of mathematical statistics*, pages 148–170, 1961.

[25] David R Hunter. Mm algorithms for generalized bradley-terry models. *The annals of statistics*, 32(1):384–406, 2004.

[26] Toshihiro Kamishima. Nantonac collaborative filtering: recommendation based on order responses. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 583–588, 2003.

[27] Toshihiro Kamishima and Shotaro Akaho. Efficient clustering for orders. In *Sixth IEEE International Conference on Data Mining-Workshops (ICDMW’06)*, pages 274–278. IEEE, 2006.

[28] Sai Praneeth Karimireddy, Wenshuo Guo, and Michael I Jordan. Mechanisms that incentivize data sharing in federated learning. *arXiv preprint arXiv:2207.04557*, 2022.

[29] Yuqing 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.

[30] Yuqing Kong and Grant Schoenebeck. A framework for designing information elicitation mechanisms that reward truth-telling. *CoRR*, abs/1605.01021, 2016. URL <http://arxiv.org/abs/1605.01021>.

[31] Yuqing Kong and Grant Schoenebeck. Equilibrium selection in information elicitation without verification via information monotonicity. In *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2018.

[32] Yuqing Kong and Grant Schoenebeck. Water from two rocks: Maximizing the mutual information. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 177–194, 2018.

[33] Yuqing Kong and Grant 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.

[34] Julia Kreutzer, Joshua Uyheng, and Stefan Riezler. Reliability and learnability of human bandit feedback for sequence-to-sequence reinforcement learning. *arXiv preprint arXiv:1805.10627*, 2018.

[35] Nicolas S Lambert, David M Pennock, and Yoav Shoham. Eliciting properties of probability distributions. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 129–138, 2008.

[36] Yang Liu, Rixing Lou, and Jiaheng Wei. Auditing for federated learning: A model elicitation approach. In *Proceedings of the Fifth International Conference on Distributed Artificial Intelligence*, DAI '23, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400708480. doi: 10.1145/3627676.3627683. URL <https://doi.org/10.1145/3627676.3627683>.

[37] Yang Liu, Juntao Wang, and Yiling Chen. Surrogate scoring rules. *ACM Transactions on Economics and Computation*, 10(3):1–36, 2023.

[38] R Duncan Luce. *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2005.

[39] C. L. Mallows. Non-null ranking models. i. *Biometrika*, 44(1/2):114–130, 1957. ISSN 00063444. URL <http://www.jstor.org/stable/2333244>.

[40] Miller McPherson, Lynn Smith-Lovin, and James M Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444, 2001.

[41] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, pages 1359–1373, 2005.

[42] Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359–1373, 2005.

[43] Andrea Montanari and Amin Saberi. The spread of innovations in social networks. *Proceedings of the National Academy of Sciences*, 107(47):20196–20201, 2010.

[44] Marc Mézard and Andrea Montanari. *Information, Physics, and Computation*. Oxford University Press, 01 2009. ISBN 9780198570837. doi: 10.1093/acprof:oso/9780198570837.001.0001. URL <https://doi.org/10.1093/acprof:oso/9780198570837.001.0001>.

[45] Ryan O’Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.

[46] Kent Harold Osband. *Providing Incentives for Better Cost Forecasting (Prediction, Uncertainty Elicitation)*. University of California, Berkeley, 1985.

[47] Drazen Prelec. A bayesian truth serum for subjective data. *science*, 306(5695):462–466, 2004.

[48] Goran Radanovic and Boi Faltings. A robust bayesian truth serum for non-binary signals. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 27, pages 833–839, 2013.

[49] Goran Radanovic and Boi Faltings. Incentives for truthful information elicitation of continuous signals. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, pages 770–776, 2014.

[50] Grant Schoenebeck and Fang-Yi Yu. Learning and strongly truthful multi-task peer prediction: A variational approach. In *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.

[51] Grant Schoenebeck and Fang-Yi Yu. Two strongly truthful mechanisms for three heterogeneous agents answering one question. *ACM Transactions on Economics and Computation*, 10(4):1–26, 2023.

[52] Grant Schoenebeck, Fang-Yi Yu, and Yichi Zhang. Information elicitation from rowdy crowds. In *Proceedings of the Web Conference 2021*, pages 3974–3986, 2021.

[53] Sören W Scholz, Martin Meissner, and Reinhold Decker. Measuring consumer preferences for complex products: A compositional approach basedonpaired comparisons. *Journal of Marketing Research*, 47(4):685–698, 2010.

[54] Nihar Shah, Sivaraman Balakrishnan, Aditya Guntuboyina, and Martin Wainwright. Stochastically transitive models for pairwise comparisons: Statistical and computational issues. In *International Conference on Machine Learning*, pages 11–20. PMLR, 2016.

[55] Nihar B Shah, Joseph K Bradley, Abhay Parekh, Martin Wainwright, and Kannan Ramchandran. A case for ordinal peer-evaluation in moocs. In *NIPS workshop on data driven education*, volume 15, page 67, 2013.

[56] Victor Shnayder, Arpit Agarwal, Rafael Frongillo, and David C. Parkes. Informed truthfulness in multi-task peer prediction. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 179–196, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-3936-0.

[57] Victor Shnayder, Arpit Agarwal, Rafael M. Frongillo, and David C. Parkes. Informed truthfulness in multi-task peer prediction. *CoRR*, abs/1603.03151, 2016. URL <http://arxiv.org/abs/1603.03151>.

[58] Louis L Thurstone. A law of comparative judgment. In *Scaling*, pages 81–92. Routledge, 2017.

[59] Amos Tversky and J Edward Russo. Substitutability and similarity in binary choices. *Journal of Mathematical psychology*, 6(1):1–12, 1969.

[60] Stephan Vail. A stochastic model for utilities. *Unpublished manuscript*, 1953.

[61] Luis Von Ahn and Laura Dabbish. Labeling images with a computer game. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 319–326, 2004.

[62] Luis von Ahn and Laura Dabbish. Designing games with a purpose. *Commun. ACM*, 51(8):58–67, aug 2008. ISSN 0001-0782. doi: 10.1145/1378704.1378719. URL <https://doi.org/10.1145/1378704.1378719>.

[63] Bo Waggoner and Yiling Chen. Output agreement mechanisms and common knowledge. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 2, pages 220–226, 2014.

[64] Jiaheng Wei, Zuyue Fu, Yang Liu, Xingyu Li, Zhuoran Yang, and Zhaoran Wang. Sample elicitation. In *International Conference on Artificial Intelligence and Statistics*, pages 2692–2700. PMLR, 2021.

[65] Dror Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 140–149, 2006.

[66] Jens Witkowski and David C. Parkes. A Robust Bayesian Truth Serum for Small Populations. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI 2012)*, 2011.

[67] Jens Witkowski and David C. Parkes. Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC 2012, Valencia, Spain, June 4-8, 2012, pages 964–981. ACM, 2012.

[68] Peter Zhang and Yiling Chen. Elicitability and knowledge-free elicitation with peer prediction. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 245–252. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

[69] Yichi Zhang and Grant Schoenebeck. High-effort crowds: Limited liability via tournaments. In *Proceedings of the ACM Web Conference 2023*, WWW '23, page 3467–3477, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450394161. doi: 10.1145/3543507.3583334. URL <https://doi.org/10.1145/3543507.3583334>.

[70] Shuran Zheng, Fang-Yi Yu, and Yiling Chen. The limits of multi-task peer prediction. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 907–926, 2021.

## A Further discussion on BPP payment

In this section, we discuss the connection of bonus-penalty payment and existing peer prediction mechanisms. First, if we substitute the third input with a uniformly random bit, denoted as  $\hat{s}_k = Z \sim_u \{-1, 1\}$ , the bonus-penalty payment simplifies to the *agreement mechanism* [62, 61, 63], one of the most basic peer prediction mechanisms,

$$\mathbb{E}[U^{BPP}(\hat{s}_i, \hat{s}_j, Z)] = \hat{s}_i \hat{s}_j = \mathbf{1}[\hat{s}_i = \hat{s}_j] - 1.$$

However, the agreement mechanism is not symmetrically strongly truthful, as all agents always reporting 1 and  $-1$  can result in higher payments than truth-telling.

The bonus-penalty payment eq. (1) is originally proposed by [11, 57] for the multi-task setting. Our BPP mechanism in Mechanism 3 can be seen as a generalization of multi-task setting. In the multi-task setting, agents works on multiple tasks and for each task the private signals are jointly identically and independently (iid) sampled from a fixed distribution and the each agent's strategy also are iid. Take two agents (Isabel and Julia) and two tasks as an example: Isabel has a private signal  $(s_i^1, s_i^2)$  and reports  $(\hat{s}_i^1, \hat{s}_i^2)$  and Julia has  $(s_j^1, s_j^2)$  and reports  $(\hat{s}_j^1, \hat{s}_j^2)$  where  $(s_i^l, s_j^l)$  are iid from random vector  $(S_i, S_j)$ . Isabel and Julia decide their reports on each task using random function  $\sigma_i, \sigma_j : \{-1, 1\} \mapsto \{-1, 1\}$  respectively. Dasgupta and Ghosh [11] use the following payments for Isabel

$$\mathbf{1}[\hat{s}_i^1 = \hat{s}_j^1] - \mathbf{1}[\hat{s}_i^1 = \hat{s}_j^2] = \frac{1}{2} U^{BPP}(\hat{s}_i^1, \hat{s}_j^1, \hat{s}_j^2).$$

The payment is a special case of Mechanism 3 by taking the second input as  $\hat{s}_j^1$  and the third input as  $\hat{s}_j^2$ . Additionally,  $S_j^1$  uniform dominates  $S_j^2$  for  $S_i^1$  if and only if

$$\Pr[S_j = 1 \mid S_i = 1] > \Pr[S_j = 1], \text{ and } \Pr[S_j = -1 \mid S_i = -1] > \Pr[S_j = -1]$$

which is called *categorical signal distributions* [57].

Finally, similar to Shnayder et al. [57], we may extend to non-binary signal setting by extending the payment to

$$U^{BPP}(\hat{s}_i, \hat{s}_j, \hat{s}_k) = 2(\mathbf{1}[\hat{s}_i = \hat{s}_j] - \mathbf{1}[\hat{s}_i = \hat{s}_k])$$

and the definition of uniform dominance to the following.

**Definition A.1.** Given a random vector  $(S_i, S_j, S_k) \in \Omega^3$  on a discrete domain, we say  $S_j$  uniformly dominates  $S_k$  for  $S_i$  if

$$\begin{aligned} \Pr[S_j = s \mid S_i = s] - \Pr[S_k = s \mid S_i = s] &> 0 \text{ and} \\ \Pr[S_j = s' \mid S_i = s] - \Pr[S_k = s' \mid S_i = s] &< 0 \end{aligned}$$

for all  $s, s' \in \Omega$  with  $s \neq s'$ .

However, the guarantee for truth-telling (*informed truthfulness*) is weaker than the binary setting.

**Theorem A.2.** Given any discrete domain  $\Omega$ , if for each agent  $i$  the associated agent  $j$ 's signal uniformly dominates  $k$ 's signal for  $i$ 's signal (definition A.1), Mechanism 3's scheme is symmetrically informed truthful so that

1. truth-telling is a strict equilibrium, and
2. each agent's expected payment in truth-telling is no less than the payment in any other symmetric equilibria and strictly better than any uninformed equilibrium's.

*Proof.* First truth-telling is a strict equilibrium, because if  $S_i = s$ ,

$$\begin{aligned} &\arg \max_{\hat{s}} \mathbb{E}[U^{BPP}(\hat{s}, S_j, S_k) \mid S_i = s] \\ &= \arg \max_{\hat{s}} \Pr[S_j = \hat{s} \mid S_i = s] - \Pr[S_k = \hat{s} \mid S_i = s] \\ &= s \end{aligned} \tag{by definition A.1}$$

Additionally, because  $\Pr[S_j = s \mid S_i = s] - \Pr[S_k = s \mid S_i = s] > \Pr[S_j = s' \mid S_i = s] - \Pr[S_k = s' \mid S_i = s]$  for all  $s' \neq s$ , summing over all possible  $s' \in \Omega$  on both sides gets  $\Pr[S_j = s \mid S_i = s] - \Pr[S_k = s \mid S_i = s] > 0$  and

$$\mathbb{E}[U^{BPP}(S_i, S_j, S_k)] > 0.$$

For any informed equilibrium, by a direct computation  $\mathbb{E}[U^{BPP}(\hat{S}_i, \hat{S}_j, \hat{S}_k)] = 0$ .

Finally, we show that the truth-telling has the maximum expected payment for each agents. When all agent use a strategy  $\sigma : \Omega \rightarrow \Omega$ , agent  $i$ 's expected payment is

$$\begin{aligned} & \sum_{s_i, \hat{s}_i \in \Omega} \Pr[S_i = s_i] \sigma(s_i, \hat{s}_i) \mathbb{E}[U^{BPP}(\hat{s}_i, \hat{S}_j, \hat{S}_k) \mid S_i = s_i] \\ &= 2 \sum_{s_i, \hat{s}_i \in \Omega} \Pr[S_i = s_i] \sigma(s_i, \hat{s}_i) \sum_{s \in \Omega} (\Pr[S_j = s \mid S_i = s_i] - \Pr[S_k = s \mid S_i = s_i]) \sigma(s, \hat{s}_i) \\ &= 2 \sum_{s_i \in \Omega} \Pr[S_i = s_i] \sum_{\hat{s}_i, s \in \Omega} \sigma(s_i, \hat{s}_i) \sigma(s, \hat{s}_i) (\Pr[S_j = s \mid S_i = s_i] - \Pr[S_k = s \mid S_i = s_i]) \end{aligned}$$

Let  $f_{s_i}(s) := \sum_{\hat{s}_i \in \Omega} \sigma(s_i, \hat{s}_i) \sigma(s, \hat{s}_i)$  which is between 0 and 1, because  $f_{s_i}(s) \leq \sum_{\hat{s}_i \in \Omega} \sigma(s_i, \hat{s}_i) \sum_{\hat{s}_i \in \Omega} \sigma(s, \hat{s}_i) = 1$ . Then the expectation becomes

$$\begin{aligned} & \sum_{s_i, \hat{s}_i \in \Omega} \Pr[S_i = s_i] \sigma(s_i, \hat{s}_i) \mathbb{E}[U^{BPP}(\hat{s}_i, \hat{S}_j, \hat{S}_k) \mid S_i = s_i] \\ &= 2 \sum_{s_i \in \Omega} \Pr[S_i = s_i] \sum_{s \in \Omega} (\Pr[S_j = s \mid S_i = s_i] - \Pr[S_k = s \mid S_i = s_i]) f_{s_i}(s) \\ &\leq 2 \sum_{s_i \in \Omega} \Pr[S_i = s_i] (\Pr[S_j = s_i \mid S_i = s_i] - \Pr[S_k = s_i \mid S_i = s_i]) \\ &= \mathbb{E}[U^{BPP}(S_i, S_j, S_k)] \end{aligned}$$

The inequality holds because  $f_{s_i} \in [0, 1]$  and definition A.1. Therefore, we complete the proof.  $\square$

## B Proofs in Section 2: Bayesian SST model and other models

The proofs of propositions 2.3 and 2.5 are standard, and variations can be found in related literature. We include proofs here for completeness.

*Proof of proposition 2.3.* First given  $\theta \in \mathbb{R}^{\mathcal{A}}$ , for all distinct  $a, a', a'' \in \mathcal{A}$ ,  $\Pr[T_\theta(a, a') = 1], \Pr[T_\theta(a', a'') = 1] > 1/2$  implies that  $\theta_a - \theta_{a'} > 0$  and  $\theta_{a'} - \theta_{a''} > 0$  because  $F$  is strictly increasing and  $F(0) = 1/2$ . Because  $\theta_a - \theta_{a''} = \theta_a - \theta_{a'} + \theta_{a'} - \theta_{a''} > \max(\theta_a - \theta_{a'}, \theta_{a'} - \theta_{a''})$ , we have

$$\begin{aligned} \Pr[T_\theta(a, a'') = 1] &= F(\theta_a - \theta_{a''}) \\ &> \max F(\theta_a - \theta_{a'}), F(\theta_{a'} - \theta_{a''}) \\ &= \max \Pr[T_\theta(a, a') = 1], \Pr[T_\theta(a', a'') = 1] \end{aligned}$$

and thus  $T_\theta$  is strongly stochastically transitive for all  $\theta$  with distinct coordinates which happens surely as  $\nu$  is non-atomic. Finally, since the distribution on  $\theta$  is exchangeable on each coordinate,  $\mathbb{E}[\mathbb{E}[T_\theta(a, a')]] = 0$  for all  $a, a'$ .  $\square$

*Proof of proposition 2.5.* First given  $\theta \in \Theta$ , for all distinct  $a, a' \in \mathcal{A}$ , if the rank of  $a$  is higher than  $a'$ ,

$$\Pr[T_\theta(a, a') = 1] = h_\eta(\theta(a') - \theta(a) + 1) - h_\eta(\theta(a') - \theta(a))$$

where  $h_\eta(x) = \frac{x}{1 - \exp(-\eta x)}$  by Busa-Fekete et al. [7].

**Claim B.1.** For any  $\eta > 0$  and  $x \in \mathbb{Z}_{>0}$ , the difference  $h_\eta(x+1) - h_\eta(x)$  is increasing and larger than  $1/2$  where  $h_\eta(x) = \frac{x}{1 - \exp(-\eta x)}$ .

By claim B.1,  $\Pr[T_\theta(a, a') = 1], \Pr[T_\theta(a', a'') = 1] > 1/2$  implies that  $\theta(a') - \theta(a) > 0$  and  $\theta(a'') - \theta(a') > 0$ . Thus,  $\theta(a'') - \theta(a) > \max(\theta(a'') - \theta(a'), \theta(a'') - \theta(a'))$ , and

$$\begin{aligned} \Pr[T_\theta(a, a'') = 1] &= h(\theta(a'') - \theta(a) + 1) - h(\theta(a'') - \theta(a)) \\ &> \max h(\theta(a'') - \theta(a') + 1) - h(\theta(a'') - \theta(a')), h(\theta(a') - \theta(a) + 1) - h(\theta(a') - \theta(a)) \\ &= \max \Pr[T_\theta(a, a') = 1], \Pr[T_\theta(a', a'') = 1] \end{aligned}$$

where the second inequality is due to claim B.1. Therefore,  $T_\theta$  is strongly stochastically transitive for all  $\theta$ . Finally,  $\mathbb{E}[\mathbb{E}[T_\theta(a, a')]] = 0$  for all  $a, a'$  since  $\theta$  is an uniform distribution on rankings.  $\square$

*Proof of claim B.1.* We first prove that the function  $h_\eta(x) = \frac{x}{1-\exp(-\eta x)}$  is increasing and strictly convex on  $x \geq 0$ . Because  $h_\eta(x) = \frac{1}{\eta}h_1(\eta x)$ , for all  $\eta, x$ , it is sufficient to consider  $\eta = 1$ . First,  $h'_1(x) = \frac{1-(x+1)e^{-x}}{(1-e^{-x})^2} > 0$ , so  $h_1$  is increasing. Second, as  $h''_1(x) = \frac{e^{-x}((x-2)+(x+2)e^{-x})}{(1-e^{-x})^3}$ , to show  $h''_1(x) > 0$  for all  $x > 0$ , it is sufficient to show that  $g(x) = (x-2) + (x+2)e^{-x} > 0$ . Because  $g(0) = 0$  and  $g'(x) = 1 - (x+1)e^{-x} > 0$ ,  $g(x) > 0$  for all  $x > 0$ . Therefore,  $h_1$  is strictly convex.

On the other hand,  $h_\eta(x+2) - h_\eta(x+1) > h_\eta(x+1) - h_\eta(x)$  for all  $x$  by convexity, and  $h_\eta(2) - h_\eta(1) = \frac{1}{1+e^{-\eta}} > \frac{1}{2}$  which completes the proof.  $\square$

## C Proofs in Section 3 and 4

### C.1 Uniform dominance from Bayesian SST

*Proof of lemma 4.2.* With a prior similar assumption for Bayesian SST model, we only need to show

$$\Pr[S(a'', a') = 1 \mid S(a, a') = 1] > \Pr[S(a'', a) = 1 \mid S(a, a') = 1], \quad (5)$$

and the other case  $\Pr[S(a'', a') = -1 \mid S(a, a') = -1] > \Pr[S(a'', a) = -1 \mid S(a, a') = -1]$  follows by symmetry. To prove eq. (5), we can rewrite the conditional probability in expectations of  $T_\theta$ .

$$\begin{aligned} &\Pr[S(a'', a') = 1 \mid S(a, a') = 1] \\ &= \frac{\int \Pr[T_\theta(a'', a') = 1, T_\theta(a, a') = 1 \mid \theta] dP_\Theta}{\int \Pr[T_\theta(a, a') = 1 \mid \theta] dP_\Theta} \\ &= \frac{\int \Pr[T_\theta(a'', a') = 1 \mid \theta] \Pr[T_\theta(a, a') = 1 \mid \theta] dP_\Theta}{\int \Pr[T_\theta(a, a') = 1 \mid \theta] dP_\Theta} \quad (\text{conditional independent}) \\ &= 2 \int \Pr[T_\theta(a'', a') = 1 \mid \theta] \Pr[T_\theta(a, a') = 1 \mid \theta] dP_\Theta \quad (\text{a prior similar}) \\ &= 2 \int \frac{\mathbb{E}[T_\theta(a'', a') \mid \theta] + 1}{2} \frac{\mathbb{E}[T_\theta(a, a') \mid \theta] + 1}{2} dP_\Theta \quad (\text{binary value}) \\ &= \frac{1}{2} \int \mathbb{E}[T_\theta(a'', a') \mid \theta] \mathbb{E}[T_\theta(a, a') \mid \theta] + \mathbb{E}[T_\theta(a'', a') \mid \theta] + \mathbb{E}[T_\theta(a, a') \mid \theta] + 1 dP_\Theta \\ &= \frac{1}{2} \int \mathbb{E}[T_\theta(a'', a') \mid \theta] \mathbb{E}[T_\theta(a, a') \mid \theta] + 1 dP_\Theta. \quad (\text{a prior similar}) \end{aligned}$$

**Claim C.1.** For any strongly stochastically transitive  $T_\theta$  on  $\mathcal{A}$ , and distinct  $a, a', a'' \in \mathcal{A}$

$$\mathbb{E}[T_\theta(a, a') \mid \theta] \mathbb{E}[T_\theta(a'', a') \mid \theta] > \mathbb{E}[T_\theta(a, a') \mid \theta] \mathbb{E}[T_\theta(a'', a) \mid \theta].$$

With claim C.1, we have

$$\begin{aligned} \Pr[S(a'', a') = 1 \mid S(a, a') = 1] &= \frac{1}{2} \int \mathbb{E}[T_\theta(a'', a') \mid \theta] \mathbb{E}[T_\theta(a, a') \mid \theta] + 1 dP_\Theta \\ &> \frac{1}{2} \int \mathbb{E}[T_\theta(a'', a) \mid \theta] \mathbb{E}[T_\theta(a, a') \mid \theta] + 1 dP_\Theta = \Pr[S(a'', a) = 1 \mid S(a, a') = 1]. \end{aligned}$$

This completes the proof of eq. (5), and thus the uniform dominance.  $\square$

*Proof of claim C.1.* We let  $Q(\alpha, \alpha') := \mathbb{E}[T_\theta(\alpha, \alpha') \mid \theta] = 2\Pr[T_\theta(\alpha, \alpha') = 1 \mid \theta] - 1$  for all  $\alpha, \alpha'$ . Note that  $Q(\alpha, \alpha') > 0$  if and only if  $\Pr[T_\theta(\alpha, \alpha') = 1 \mid \theta] > 1/2$  and  $Q(\alpha, \alpha') = -Q(\alpha', \alpha)$ .

By symmetry, let  $Q(a, a') > 0$ . It is sufficient to show that

$$Q(a'', a') > Q(a'', a).$$

If  $Q(a', a'') > 0$ , by definition 2.1  $Q(a, a'') > Q(a', a'') > 0$  so  $Q(a'', a') > Q(a'', a)$ . Now consider  $Q(a', a'') < 0$ . If  $Q(a'', a) < 0$ ,  $Q(a'', a') > 0 > Q(a'', a)$ . If  $Q(a'', a) > 0$ , we have  $Q(a'', a) > 0, Q(a, a') > 0$ , and thus  $Q(a'', a') > Q(a'', a)$  by definition 2.1  $\square$

## C.2 Uniform dominance and weak notions of stochastic transitivity

There are weaker forms of stochastic transitivity, raising the question of whether they are sufficient for uniform dominance as in lemma 4.2. We show that general weak stochastic transitivity is not sufficient. Additionally, we show that although the noisy sorting model from [5] is only weakly stochastically transitive but does not satisfy definition 2.1, it exhibits uniform dominance.

**Definition C.2** ([13]). A stochastic comparison function,  $T : \mathcal{A}^2 \rightarrow \{-1, 1\}$ , is *weakly stochastically transitive* if for all  $a, a', a'' \in \mathcal{A}$  with  $\Pr[T(a, a') = 1] > 1/2$  and  $\Pr[T(a', a'') = 1] > 1/2$ ,

$$\Pr[T(a, a'') = 1] > 1/2.$$

Compared to definition 2.1, the weak stochastic transitivity only require the item  $a$  is favorable than  $a''$ . Below we provide a simple weakly stochastically transitive example with a prior similar property that does not satisfy the uniform dominance in eq. (5).

**Example C.3.** Consider the set of three items and  $\Theta$  consists of all ranking on  $\mathcal{A}$  with uniform prior where  $\theta$  maps each items to its value. Given  $\theta \in \Theta$  so that if  $\theta(a) > \theta(a') > \theta(a'')$ ,

$$\Pr[T_\theta(a, a') = 1] = \Pr[T_\theta(a', a'') = 1] = 0.9 \text{ and } \Pr[T_\theta(a, a'') = 1] = 0.6.$$

Note that the model is weakly stochastically transitive, because an item with a larger value is more favorable and the weak stochastic transitivity is reduced to transitivity on the values. However, the model is not strongly stochastically transitive, because  $\Pr[T_\theta(a, a'') = 1] = 0.6 < \max\{\Pr[(T(a, a') = 1], \Pr[(T(a', a'') = 1)]\} = 0.9$ . Finally, as the rank  $\theta$  has a uniform prior, the model satisfies a prior similar assumption.

To conclude the example, we show that eq. (5) does not hold for the above model. By direct computation over all six possible ranking  $\theta$ , we have

$$\begin{aligned} & \Pr[S(a'', a') = 1 \mid S(a, a') = 1] \\ &= \frac{1}{2} \int \mathbb{E}[T_\theta(a'', a') \mid \theta] \mathbb{E}[T_\theta(a, a') \mid \theta] + 1 dP_\Theta \\ &= \frac{1}{2} \left(1 - \frac{64}{6}\right), \end{aligned}$$

but  $\Pr[S(a'', a) = 1 \mid S(a, a') = 1] = \frac{1}{2} \left(1 + \frac{64}{6}\right)$ . Therefore, we have  $\Pr[S(a'', a') = 1 \mid S(a, a') = 1] < \Pr[S(a'', a) = 1 \mid S(a, a') = 1]$ , and show that eq. (5) does not hold.

Though the above example shows that weak stochastic transitivity is not sufficient.<sup>5</sup> Below we show a popular weakly stochastically transitive model in Braverman and Mossel [5] has uniform dominance as in lemma 4.2.

**Example C.4.** Let  $\Theta$  be the set of rankings on  $\mathcal{A}$  and  $\eta > 0$  be a parameter. Given a uniformly distributed reference ranking  $\theta \in \Theta$ , the noise ranking model [5] ensures that for all  $\theta(a) > \theta(a')$

$$\Pr[T_\theta(a, a') = 1] = \frac{1}{2} + \eta$$

Note that the above model does not satisfy the strict inequality in definition 2.1, but by direct computation,  $\Pr[S(a'', a') = 1 \mid S(a, a') = 1] = \frac{1}{2} \left(1 + \frac{4\eta^2}{3}\right)$  and  $\Pr[S(a'', a) = 1 \mid S(a, a') = 1] = \frac{1}{2} \left(1 - \frac{4\eta^2}{3}\right)$ , which satisfies lemma 4.2.

<sup>5</sup>In the above example, we can also decrease 0.9 to a smaller number that satisfies both uniform dominance and weak stochastic transitivity.

### C.3 Symmetrically strongly truthful from uniform dominance

*Proof of lemma 4.3.* Suppose  $S_i = 1$ . Because  $\Pr[S_j = 1 | S_i = 1] > \Pr[S_k = 1 | S_i = 1]$ ,  $\Pr[S_j = -1 | S_i = 1] < \Pr[S_k = -1 | S_i = 1]$ . Therefore,  $\arg \max_{\hat{s} \in \{-1, 1\}} \Pr[S_j = \hat{s} | S_i = 1] - \Pr[S_k = \hat{s} | S_i = 1] = 1$ . Identical argument holds for the case of  $S_i = -1$  which completes the proof.

Additionally, the expected payment of truth-telling is

$$\begin{aligned} \mathbb{E}[U^{BPP}(S_i, S_j, S_k)] &= \sum_a \Pr[S_i = s_i] \sum_{s_j, s_k} \Pr[S_j = s_j, S_k = s_k | S_i = s_i] U^{BPP}(s_i, s_j, s_k) \\ &= 2 \sum_a \Pr[S_i = s_i] \sum_{s_j, s_k} \Pr[S_j = s_j, S_k = s_k | S_i = s_i] (\mathbf{1}[s_i = s_k] - \mathbf{1}[s_i = s_k]) \\ &= 2 \sum_a \Pr[S_i = s_i] (\Pr[S_j = s_i | S_i = s_i] - \Pr[S_k = s_i | S_i = s_i]) \\ &> 0 \end{aligned}$$

The last inequality holds due to definition 4.1.  $\square$

*Proof of lemma 4.4.* As  $\sigma$  is uninformed, let  $\mu(s) = \sigma(s, s) = \sigma(-s, s)$  and  $\mu(-s) = \sigma(s, -s) = \sigma(-s, -s)$  for all  $s$ .

$$\mathbb{E}[U^{BPP}(\hat{s}_i, \hat{S}_j, \hat{S}_k) | S_i = s_i] = \sum_{\hat{s}_j, \hat{s}_k} \mu(\hat{s}_j) \mu(\hat{s}_k) U^{BPP}(\hat{s}_i, \hat{S}_j, \hat{S}_k) = \sum_{\hat{s}_j, \hat{s}_k} \mu(\hat{s}_j) \mu(\hat{s}_k) (\hat{s}_i \hat{s}_j - \hat{s}_i \hat{s}_k) = 0$$

The first equality holds as the reports are independent of signals.  $\square$

*Proof of lemma 4.5.*

$$\begin{aligned} &\mathbb{E}_{P, \sigma}[U^{BPP}(\hat{s}_i, \hat{S}_j, \hat{S}_k) | S_i = s_i] \\ &= \sum_{s_j, s_k, \hat{s}_j, \hat{s}_k} \Pr[S_j = s_j, S_k = s_k | S_i = s_i] \sigma(s_j, \hat{s}_j) \sigma(s_k, \hat{s}_k) U^{BPP}(\hat{s}_i, \hat{s}_j, \hat{s}_k) \\ &= 2 \sum_{s_j, s_k, \hat{s}_j, \hat{s}_k} \Pr[S_j = s_j, S_k = s_k | S_i = s_i] \sigma(s_j, \hat{s}_j) \sigma(s_k, \hat{s}_k) (\mathbf{1}[\hat{s}_i = \hat{s}_j] - \mathbf{1}[\hat{s}_i = \hat{s}_k]) \\ &\quad \text{(by eq. (1))} \\ &= 2 \sum_{s_j, \hat{s}_j} \Pr[S_j = s_j | S_i = s_i] \sigma(s_j, \hat{s}_j) \mathbf{1}[\hat{s}_i = \hat{s}_j] - 2 \sum_{s_k, \hat{s}_k} \Pr[S_k = s_k | S_i = s_i] \sigma(s_k, \hat{s}_k) \mathbf{1}[\hat{s}_i = \hat{s}_k] \\ &= 2 \sum_{s, \hat{s}} (\Pr[S_j = s | S_i = s_i] - \Pr[S_k = s | S_i = s_i]) \sigma(s, \hat{s}) \mathbf{1}[\hat{s}_i = \hat{s}] \quad \text{(renaming dummy variables)} \\ &= 2 \sum_s (\Pr[S_j = s | S_i = s_i] - \Pr[S_k = s | S_i = s_i]) \sigma(s, \hat{s}_i) \end{aligned}$$

Let  $\delta = \Pr[S_j = s_i | S_i = s_i] - \Pr[S_k = s_i | S_i = s_i] > 0$ , because  $S_j$  uniformly dominates  $S_k$  for  $S_i$ . Additionally,  $\Pr[S_j = -s_i | S_i = s_i] - \Pr[S_k = -s_i | S_i = s_i] = 1 - \Pr[S_j = s_i | S_i = s_i] - 1 + \Pr[S_k = s_i | S_i = s_i] = -\delta$ . We have

$$\begin{aligned} &\mathbb{E}_{P, \sigma}[U^{BPP}(\hat{s}_i, \hat{S}_j, \hat{S}_k) | S_i = s_i] \\ &= 2 \sum_s (\Pr[S_j = s | S_i = s_i] - \Pr[S_k = s | S_i = s_i]) \sigma(s, \hat{s}_i) \\ &= 2\delta (\sigma(s_i, \hat{s}_i) - \sigma(-s_i, \hat{s}_i)), \end{aligned}$$

so  $\arg \max_{\hat{s}_i \in \{-1, 1\}} \mathbb{E}_{P, \sigma}[U^{BPP}(\hat{s}_i, \hat{S}_j, \hat{S}_k) | S_i = s_i] = \arg \max_{\hat{s}_i \in \{-1, 1\}} \{\sigma(s_i, \hat{s}_i) - \sigma(-s_i, \hat{s}_i)\}$  which completes the proof.  $\square$

## D Proofs in Section 5.1

Before diving into the proof, we introduce some notations. We further introduce Ising models with bias parameter  $\alpha \in \mathbb{R}_{\geq 0}^V$  in addition to  $\beta$  where

$$H(\mathbf{s}) = \sum_{i, j \in V} \beta_{i, j} s_i s_j + \sum_{i \in V} \alpha_i s_i$$

and  $\Pr_{\alpha, \beta}[\mathbf{S} = \mathbf{s}] \propto \exp(H(\mathbf{s}))$ , for all configuration  $\mathbf{s}$ . Given  $i \in V$ , let the expectation and ratio be

$$\nu_i(\alpha, \beta) = \mathbb{E}_{\alpha, \beta}[S_i] = \Pr_{\alpha, \beta}[S_i = 1] - \Pr_{\alpha, \beta}[S_i = -1] \text{ and } \rho_i(\alpha, \beta) = \frac{\Pr_{\alpha, \beta}[S_i = 1]}{\Pr_{\alpha, \beta}[S_i = -1]}$$

respectively which are monotone to each other. We will omit  $\alpha, \beta$  when clear. Given a subset  $U \subseteq V$ ,  $\mathbf{s}_U \in \{-1, 1\}^U$  is a configuration over the nodes in  $U$ , and  $\mathbf{s}_U = 1$  if  $x_\iota = 1$  for all  $\iota \in U$ . We write  $\Pr[\cdot | \mathbf{S}_U = \mathbf{s}_U]$ ,  $\nu_i | \mathbf{S}_U = \mathbf{s}_U$ , and  $\rho_i | \mathbf{S}_U = \mathbf{s}_U$  for the conditional probability, expectation and ratio when the configuration in  $U$  is fixed as specified by  $\mathbf{s}_U$ .

**A lower bound for LHS** Informally, we want to lower bound the correlation between adjacent  $i$  and  $j$  (friends). Note that as we remove edges (setting coordinates of  $\beta$  to zeros), the correlation should decrease, and the smallest correlation between neighboring nodes  $i$  and  $j$  happens when  $E = \{(i, j)\}$ . Lemma D.2 formalizes this idea using the following monotone inequality [44, Theorem 17.2].

**Theorem D.1** (Griffiths' inequality). *For any  $i \in V$ ,  $\nu_i(\alpha, \beta) = \mathbb{E}_{\alpha, \beta}[S_i]$  is non-negative and non-decreasing in all  $\beta_{j,k} \geq 0$  and  $\alpha_j \geq 0$  with  $j, k \in V$ .*

**Lemma D.2.** *Given  $V$  and  $i, j \in V$ , for all  $\alpha, \beta$ , and  $\beta'$ , if  $\beta'_e = \beta_e$  when  $e = (i, j)$  and  $\beta'_e = 0$  otherwise, we have*

$$\nu_{i|S_j=1}(\alpha, \beta) \geq \nu_{i|S_j=1}(\alpha, \beta') \text{ and } \rho_{i|S_j=1}(\alpha, \beta) \geq \rho_{i|S_j=1}(\alpha, \beta').$$

*Proof.* First, note that we can write the conditional expectation  $\mathbb{E}_{\alpha, \beta}[S_i | S_j = 1]$  as marginal expectation. Formally, consider  $\alpha^\eta$  so that  $\alpha_\iota^\eta = \alpha_\iota$  if  $\iota \neq j$  and  $\alpha_j^\eta = \alpha_j + \eta$ . Because  $\eta \rightarrow \alpha^\eta$  is non-decreasing,  $\eta \rightarrow \nu_i(\alpha^\eta, \beta)$  is non-decreasing by theorem D.1. In addition,  $\Pr_{\alpha^\eta, \beta}[S_i | S_j = 1] = \Pr_{\alpha, \beta}[S_i | S_j = 1]$  for all  $s$ , and  $\Pr_{\alpha^\eta, \beta}[S_j = 1] = O(e^{-2\eta})$ , so

$$\nu_{i|S_j=1}(\alpha, \beta) = \mathbb{E}_{\alpha, \beta}[S_i | S_j = 1] = \lim_{\eta \rightarrow +\infty} \nu_i(\alpha^\eta, \beta).$$

Similarly,

$$\nu_{i|S_j=1}(\alpha, \beta') = \mathbb{E}_{\alpha, \beta'}[S_i | S_j = 1] = \lim_{\eta \rightarrow +\infty} \nu_i(\alpha^\eta, \beta').$$

On the other hand, consider  $\beta^\lambda$  so that  $\beta_e^\lambda = \beta_e$  if  $e \neq (i, j)$  and  $\beta_{i,j}^\lambda = \beta_{i,j} + \lambda$ . By theorem D.1,  $\nu_i(\alpha^\eta, \beta^\lambda)$  is non-decreasing in  $\lambda$  for all  $\eta$ . Because  $\beta^0 = \beta'$  and  $\beta^1 = \beta$ , we have

$$\nu_{i|S_j=1}(\alpha, \beta') = \lim_{\eta \rightarrow +\infty} \nu_i(\alpha^\eta, \beta') \leq \lim_{\eta \rightarrow +\infty} \nu_i(\alpha^\eta, \beta) = \nu_{i|S_j=1}(\alpha, \beta)$$

Because  $\rho_i = \frac{1+\nu_i}{1-\nu_i}$  is monotone in  $\nu_i$ , the second part follows.  $\square$

Given  $\beta'$  defined in lemma D.2, by some direct computation with  $\alpha = 0$

$$\rho_{i|S_j=1}(\alpha, \beta) \geq \rho_{i|S_j=1}(\alpha, \beta') = e^{2\alpha_i + 2\beta_{i,j}} = e^{2\beta}. \quad (6)$$

**An upper bound for RHS** Now, we need to upper bound the correlation between non-adjacent  $i$  and  $k$  (non-friends). We will use Weitz's self-avoiding walks reduction [65] to upper bound the correlation on general graph  $G$  by the correlation on trees.

Given a general graph  $G$ , and an arbitrary node  $i$ , we can construct the Self Avoiding Walk Tree of  $G$  rooted at  $i$ , denoted  $T_{SAW}(G, i)$ , so that  $\Pr[S_i = 1 | \mathbf{S}_U = \mathbf{s}_U]$  is the same in  $G$  as in the tree. We outline the construction.  $T_{SAW}(G, i)$  enumerates all self-avoiding walks in  $G$  starting at  $i$  which terminates when it revisits a previous node (closes a cycle). Then,  $T_{SAW}(G, i)$  introduces a leaf with a certain boundary condition. The self-avoiding walk never revisits a node immediately, so there all the leaves with fixed boundary conditions are at least three hops away from node  $i$ . Note that if  $G$  has maximum degree  $d$ ,  $T_{SAW}$  is a  $d$ -ary tree.

**Theorem D.3** ([65]). *For any  $\alpha, \beta$ , node  $i \in V$ , and configuration  $\mathbf{s}_U$  on  $U \subset V$ ,*

$$\Pr_{\alpha, \beta}[S_i = 1 | \mathbf{S}_U = \mathbf{s}_U] = \Pr_{T_{SAW}(G, i)}[S_i = 1 | \mathbf{S}_U = \mathbf{s}_U].$$

First, with the above theorem, we have  $\nu_{i|S_k=1}(\alpha, \beta) = \mathbb{E}_{\alpha, \beta}[S_i | S_k = 1] = \mathbb{E}_{T_{SAW}(G, i)}[S_i | S_k = 1]$ . By the monotone property in theorem D.1, setting all two-hop neighbors  $U$  in  $T_{SAW}(G, i)$  to 1 (recalled that any boundary conditions for  $T_{SAW}(G, i)$  being at least three hops away) increases the conditional expectation,

$$\mathbb{E}_{T_{SAW}(G, i)}[S_i | S_k = 1] \leq \mathbb{E}_{T_{SAW}(G, i)}[S_i | \mathbf{S}_U = 1, S_k = 1].$$

Let  $T$  be the tree by truncating  $T_{SAW}(G, i)$  at level 2. By the Markov property of Ising models, the expectation is equal to the expectation on  $T$ .

$$\mathbb{E}_{\alpha, \beta}[S_i | S_k = 1] \leq \mathbb{E}_{T_{SAW}(G, i)}[S_i | \mathbf{S}_U = 1] = \mathbb{E}_T[S_i | \mathbf{S}_U = 1]. \quad (7)$$

Finally, we can recursively compute the probability ratio  $\rho_i$  (and thus expectation  $\nu_i$ ) on trees. Specifically, given a rooted tree  $T'$ , we define  $\rho_{T'}$  as the ratio of probabilities for the root to be +1 and -1 respectively, and  $\rho_{T'|S_U=s_U}$  for the ratio of conditional probabilities. As stated in the following lemma, it is well known (see, for example, [22]) that the ratio of each node can be computed recursively over the children's ratio.

**Lemma D.4.** *Given a tree  $T$  rooted at  $i$  with parameter  $(\alpha, \beta)$  and boundary condition  $\mathbf{s}_U$ ,*

$$\rho_{T|\mathbf{S}_U=\mathbf{s}_U} = e^{2\alpha_i} \prod_{l=1}^d \frac{\rho_{T_l|\mathbf{S}_U=\mathbf{s}_U} e^{2\beta_{i,j_l}} + 1}{e^{2\beta_{i,j_l}} + \rho_{T_l|\mathbf{S}_U=\mathbf{s}_U}}$$

where  $j_1, \dots, j_d$  are children of  $i$  and  $T_l$  is the subtree rooted at  $j_l$  for all  $l$ .

By the monotone property in theorem D.1, the maximum of right-hand side of eq. (7) happens when  $T$  is a complete  $d$ -ary tree with  $\beta = \bar{\beta}$ . Therefore,

$$\rho_{i|S_k=1}(\alpha, \beta) \leq \left( \frac{e^{2(d+1)\bar{\beta}} + 1}{e^{2\bar{\beta}} + e^{2d\bar{\beta}}} \right)^d. \quad (8)$$

Finally, with eqs. (6) and (8), we have  $\rho_{i|S_j=1}(\alpha, \beta) \geq e^{2\beta} \geq \left( \frac{e^{2(d+1)\bar{\beta}} + 1}{e^{2\bar{\beta}} + e^{2d\bar{\beta}}} \right)^d \geq \rho_{i|S_k=1}(\alpha, \beta)$  which implies eq. (4).

**Remark D.5.** Note that for any graph  $G$  there exists small enough  $\bar{\beta}, \underline{\beta}$  so that the condition in theorem 5.1 is satisfied, because the inequality become equality when  $\bar{\beta} = \underline{\beta} = 0$ , and  $\frac{\partial}{\partial \beta} \frac{2\beta}{d} > 0 = \frac{\partial}{\partial \beta} \ln \frac{e^{2(d+1)\beta} + 1}{e^{2\beta} + e^{2d\beta}}$ .

The bound between  $\beta$  and  $d$  is necessary as shown in fig. 3. On the other hand, by the Markov property of the Ising model, the majority of all neighbor's signals is a sufficient statistic, and we can show the majority of all neighbor's signals are uniformly dominant to a non-neighbor's signal. Therefore, we can get a symmetrically strongly truthful mechanism by replacing  $j$ 's reports with the majority of reports from  $i$ 's neighbors.

## E Proof of Theorem 5.3

The sufficient condition is done by lemma 4.3, because

$$\begin{aligned} & \arg \max_{\hat{s}_i \in \{-1, 1\}} \mathbb{E} [\lambda U^{BPP}(\hat{s}_i, S_j, S_k) + \mu(S_j, S_k) | S_i = s_i] \\ &= \arg \max_{\hat{s}_i \in \{-1, 1\}} \mathbb{E} [\lambda U^{BPP}(\hat{s}_i, S_j, S_k) | S_i = s_i] \\ &= \arg \max_{\hat{s}_i \in \{-1, 1\}} \mathbb{E} [U^{BPP}(\hat{s}_i, S_j, S_k) | S_i = s_i] \quad (\lambda > 0) \\ &= s_i \quad (\text{by lemma 4.3}) \end{aligned}$$

For the necessary, given  $U$ , define  $D(s_j, s_k) = \frac{1}{2}(U(1, s_j, s_k) - U(-1, s_j, s_k))$  and  $\mu(s_j, s_k) = \frac{1}{2}(U(1, s_j, s_k) + U(-1, s_j, s_k))$  for all  $s_j$  and  $s_k$  in  $\{-1, 1\}$ . Hence

$$U(s_i, s_j, s_k) = s_i \cdot D(s_j, s_k) + \mu(s_j, s_k), \forall s_i, s_j, s_k \in \{-1, 1\} \quad (9)$$

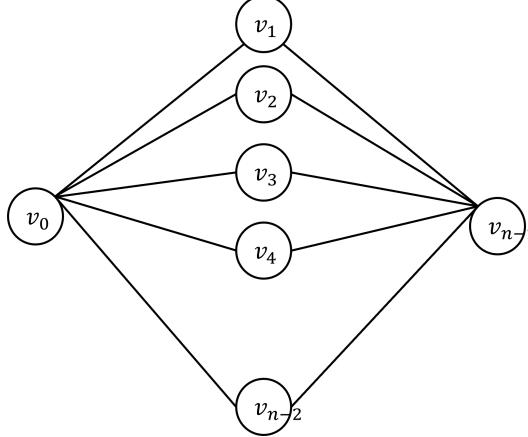


Figure 3: As fixing any  $\beta, \bar{\beta}$ , we can construct a simple graph with  $V = \{v_0, \dots, v_{n-1}\}$  and  $E = \{(v_0, v_l), (v_l, v_{n-1}) : l = 1, \dots, n-2\}$  where agent  $v_0$  and  $v_{n-1}$  are not connected but share  $n-2$  common friends. We can show that the correlation between  $S_0$  and  $S_{n-1}$  converge to 1 as the number of common friends  $d$  increases, while the correlation between  $S_0$  and  $S_1$  is bounded away from 1.

Given a joint distribution satisfying definition 4.1, we let  $p^{s_i}(s_j, s_k) = \Pr[S_j = s_j, S_k = s_k \mid S_i = s_i]$  and additionally write  $p^{s_i} = \begin{bmatrix} p^{s_i}(1, 1) & p^{s_i}(1, -1) \\ p^{s_i}(-1, 1) & p^{s_i}(-1, -1) \end{bmatrix}$ . Then definition 4.1 ensures that

$$p^1(1, -1) > p^1(-1, 1) \text{ and } p^{-1}(1, -1) < p^{-1}(-1, 1).$$

Because  $U$  is truthful for all uniformly dominant tuples, we have

$$\begin{aligned} 0 &< \mathbb{E}[U(1, S_j, S_k) \mid S_i = 1] - \mathbb{E}[U(-1, S_j, S_k) \mid S_i = 1] = 2 \sum_{s_j, s_k} D(s_j, s_k) p^1(s_i, s_j) \\ 0 &> \mathbb{E}[U(1, S_j, S_k) \mid S_i = -1] - \mathbb{E}[U(-1, S_j, S_k) \mid S_i = -1] = 2 \sum_{s_j, s_k} D(s_j, s_k) p^{-1}(s_i, s_j). \end{aligned} \quad (10)$$

Suppose the following are true

$$D(1, -1) = -D(-1, 1) > 0 \quad (11)$$

$$D(1, 1) = D(-1, -1) = 0 \quad (12)$$

Let  $\lambda = D(1, -1) > 0$ . By eqs. (11) and (12), we have

$$\begin{aligned} U(s_i, s_j, s_k) &= s_i \cdot D(s_j, s_k) + \mu(s_j, s_k) && \text{(by eq. (9))} \\ &= \lambda \cdot s_i (s_j - s_k) + \mu(s_j, s_k) && \text{(by eqs. (9) and (11))} \end{aligned}$$

which completes the proof. Thus, we will construct three joint distributions satisfying definition 4.1 to prove eqs. (11) and (12).

The first joint distribution  $p_1^{s_i}(s_j, s_k)$  with  $0 < \delta \leq 1/2$

$$p^1 = \begin{bmatrix} 0 & 1/2 + \delta \\ 1/2 - \delta & 0 \end{bmatrix} \text{ and } p^{-1} = \begin{bmatrix} 0 & 1/2 - \delta \\ 1/2 + \delta & 0 \end{bmatrix}.$$

Then eq. (10) on the first distribution reduces to

$$0 < D(1, -1)p_1^1(1, -1) + D(-1, 1)p_1^1(-1, 1) = \frac{1}{2}(D(1, -1) + D(-1, 1)) + \delta(D(1, -1) - D(-1, 1))$$

$$0 > D(1, -1)p_1^{-1}(1, -1) + D(-1, 1)p_1^{-1}(-1, 1) = \frac{1}{2}(D(1, -1) + D(-1, 1)) - \delta(D(1, -1) - D(-1, 1)).$$

As we take  $\delta$  to zero, we prove  $D(1, -1) = -D(-1, 1)$ . Then plugging in with nonzero  $\delta$ , we have  $D(1, -1) > 0$  and complete the proof of eq. (11).

The second joint distribution  $p_2^{s_i}(s_j, s_k)$  with  $0 \leq \epsilon \leq 1$  is

$$p^1 = \begin{bmatrix} 1 - \epsilon & \frac{3\epsilon}{4} \\ \frac{\epsilon}{4} & 0 \end{bmatrix} \text{ and } p^{-1} = \begin{bmatrix} 1 - \epsilon & \frac{\epsilon}{4} \\ \frac{3\epsilon}{4} & 0 \end{bmatrix}.$$

With eq. (11), eq. (10) reduces to

$$\begin{aligned} 0 &< (1 - \epsilon)D(1, 1) + \frac{\epsilon}{4}(D(1, -1) - D(-1, 1)) \\ 0 &> (1 - \epsilon)D(1, 1) - \frac{\epsilon}{4}(D(1, -1) - D(-1, 1)). \end{aligned}$$

By taking  $\epsilon$  to zero, we prove  $D(1, 1) = 0$ . We can prove  $D(-1, -1) = 0$  using the similar trick and complete the proof of eq. (12).

## F Additional empirical results

### F.1 Comparison data

Here we test if the dataset satisfy transitivity property. We denote the proportion of rankings such that item  $a$  is higher than item  $a'$  in the dataset by  $p_{a>a'}$ . If  $p_{a>a'} > 1/2$ ,  $p_{a'>a''} > 1/2$ , and  $p_{a>a''} > 1/2$ , we say the triple of items  $\{a, a', a''\}$  empirically satisfies transitivity. If  $p_{a>a'} > 1/2$ ,  $p_{a'>a''} > 1/2$ , and  $p_{a>a''} > \max\{p_{a>a'}, p_{a'>a''}\}$ , we say the triple of items  $\{a, a', a''\}$  empirically satisfies strong transitivity. We first test the transitivity of the SUSHI subdataset selected in section 6.1. We find that 100% of the item triples empirically satisfy transitivity, and 69.17% of the item triples empirically satisfy strong transitivity. This suggests that our transitivity assumption for the comparison data is mostly aligned.

Moreover, we conducted an experiment on the entire SUSHI dataset without any selection criteria and demonstrated the results in fig. 4. Observe that the ECDF of payments from original human users also dominates the payments under the uninformed strategy and the unilateral deviating strategy. This is consistent with our experimental results in section 6.1. However, there are two minor difference. First the separation of truth-telling from the other two is slightly less prominent than fig. 1 with the selection criteria. This may be due to a slightly lower degree of transitivity across agents with different backgrounds. In particular, we found the average value of  $p_{a>a''} - \max\{p_{a>a'}, p_{a'>a''}\}$  is 0.0559 without the selection criteria which is less than 0.0604 with the selection criteria in fig. 1. Second, the fraction of agents receiving positive payments is slightly higher than in fig. 1 (0.785 and 0.763 respectively). This aligned with or empirical (strong) transitively which are 1 and 0.7667 compared to the above 1 and 0.69117. Furthermore, we also conducted experiments on other groups of users by changing the selection criteria. Those interested can refer to fig. 5, fig. 6 and table 1 for the results, which further verify the effectiveness of our mechanism.

Selection criteria	Number of users	Average utility	Fraction of positive utility
All (No selection)	5000	0.138	78.5%
Female, 30-49, Kanto/Shizuoka	249	0.137	76.3%
Male, 30-49, Kanto/Shizuoka	185	0.167	82.2%
Female, 5-29, Kanto/Shizuoka	146	0.175	84.2%
Female, 50+, Kanto/Shizuoka	26	0.13	80.8%
Female, 30-49, Tohoku	30	0.174	83.3%
Female, 30-49, Hokuriku	23	0.105	69.6%

Table 1: Summary of truth-telling utility in appendix F.1.

### F.2 Networked data

Alongside fig. 2, Figure 7 and table 2 present empirical results for the top five popular artists in the dataset, excluding Lady Gaga, who are Britney Spears, Rihanna, The Beatles, and Katy Perry. All these settings show similar results. However, the Beatles' data is less conclusive as the payment distribution under the uninformed strategy profile is close to the truth-telling. This observation is

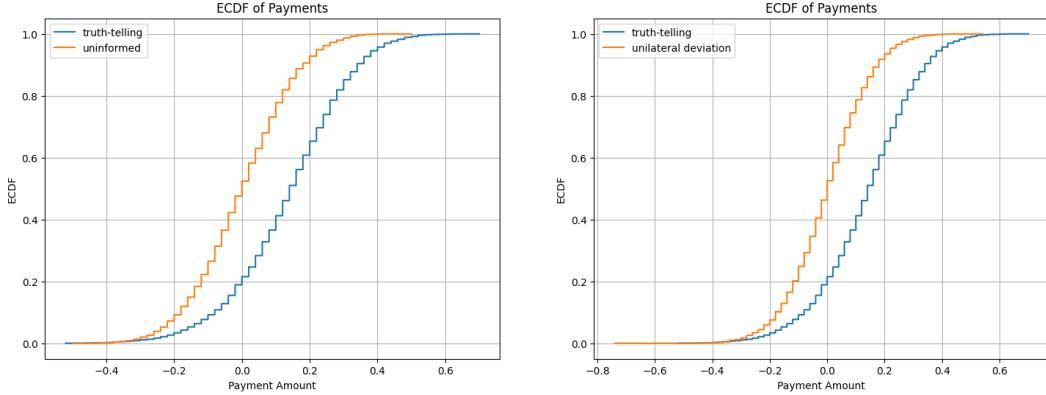


Figure 4: ECDF comparisons on all users without any selection.

also documented in Daskalakis et al. [12] which notes that the Ising model performs much better for rock artists than for pop artists. The authors conjecture that this may be due to the highly divisive popularity of pop artists like Lady Gaga and Britney Spears, whose listeners may form dense cliques within the graph.

Note that there is a buck of agent with a payment of around 0.5 under the truth-telling. This is because many non-listeners have no listener friends, and payment is  $1 - [(1 - p) - p] = 2p$  is twice the popularity  $p \approx 0.25$ . Moreover, the jump is most minor for the Beatles, and indicates less agreement between non-listeners. Additionally, by the definition of bonus-penalty payment, we can see the payment of deviation is the minus of the truthful payment, so that the ECDF is symmetric around  $(0, 0.5)$ .

Artists	Fraction of listener	Average utility	Fraction of positive utility
Lady Gaga	32.2%	0.37	76%
Britney Spears	27.6%	0.420	82.6%
Rihanna	25.6%	0.422	83.4%
The Beatles	25.4%	0.137	68.5%
Katy Perry	25.0%	0.361	79.9%

Table 2: Summary of truth-telling utility in appendix F.2.

Figure 8 further shows the scatter plot of average payment and fraction of agents with positive payments across the top fifty popular artists where all settings have more than 60% percent of agents get positive payment. However, for less popular artists, the performance of our mechanism declines. This is expected, as we cannot provide effective incentives when only one agent listens to an artist.

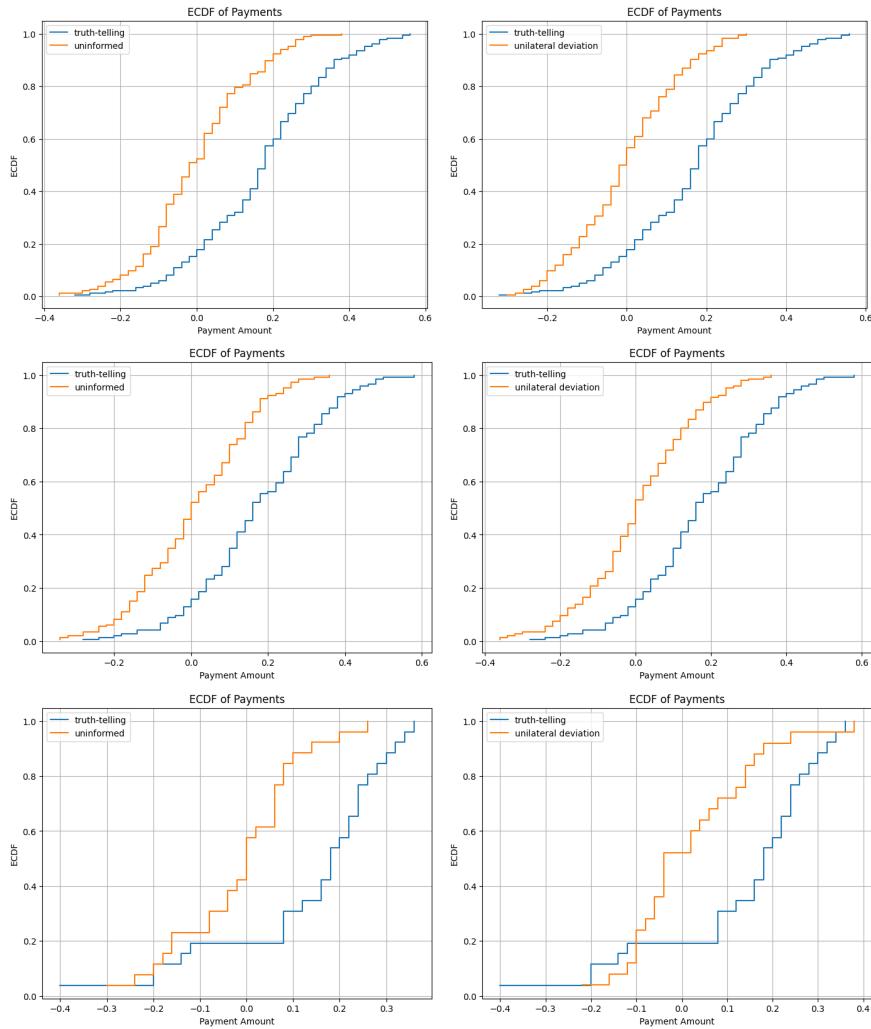


Figure 5: In each of the rows, we present the ECDF comparisons after changing the selection criteria for the user group as follows: from female to male, from ages 30–49 to ages 5–29, from ages 30–49 to ages 50+, respectively.

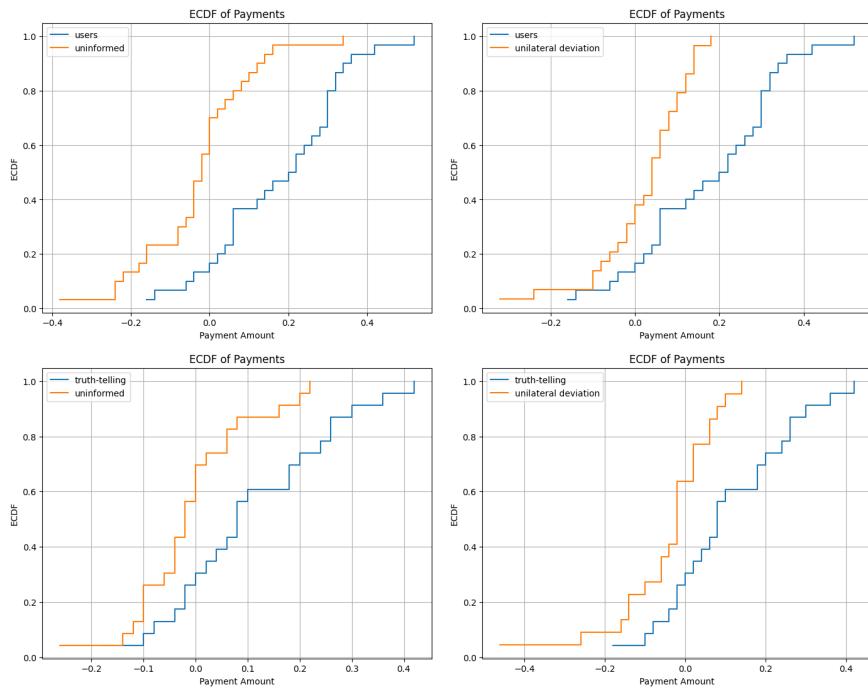


Figure 6: In each of the rows, we present the ECDF comparisons after changing the location criteria for the user group as follows: from mostly living in Kanto or Shizuoka to Tohoku until age 15, and from mostly living in Kanto or Shizuoka to Hokuriku until age 15, respectively.

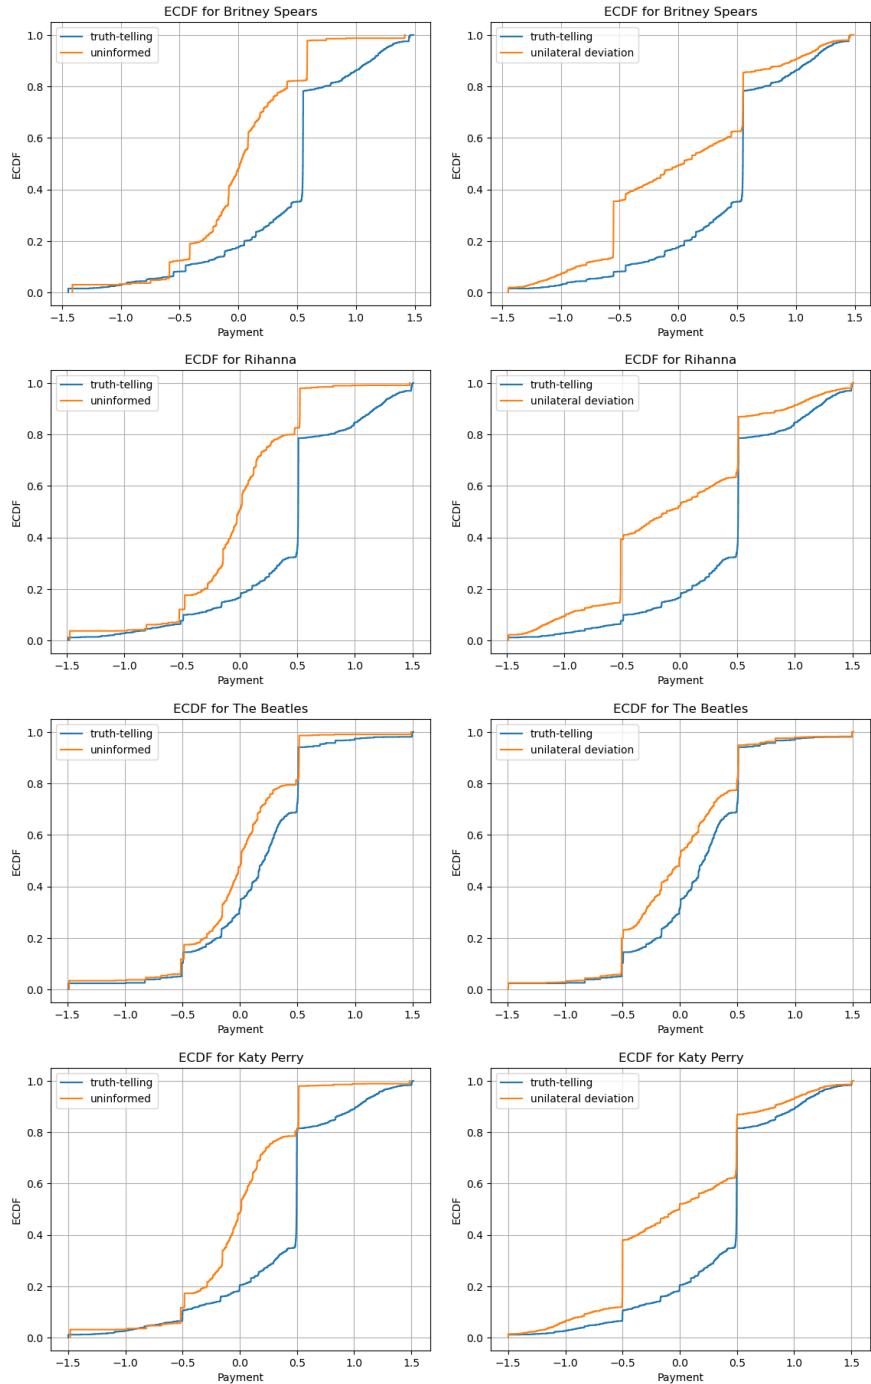


Figure 7: Last.fm dataset for other top five popular artists excluding Lady Gaga.

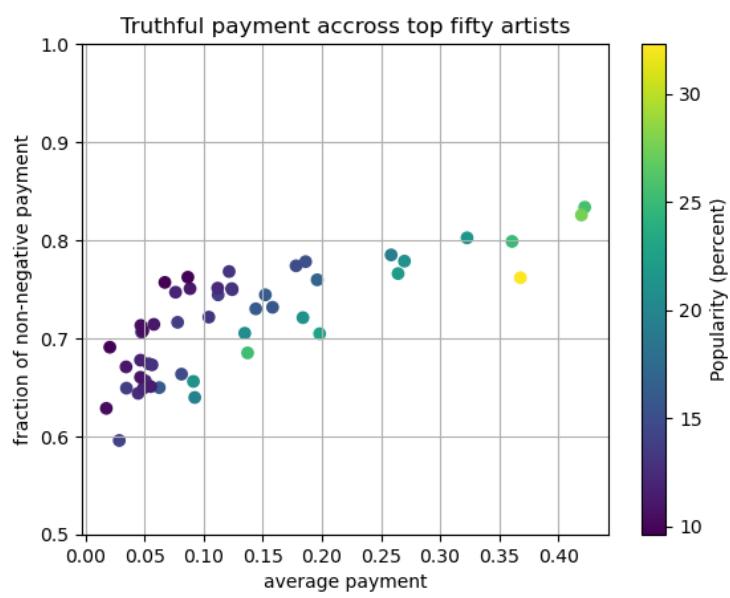


Figure 8: Average payment and fraction of positive payment under the truth-telling across top fifty popular artists.

## NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer **[Yes]** , **[No]** , or **[NA]** .
- **[NA]** means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

**The checklist answers are an integral part of your paper submission.** They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "**[Yes]**" is generally preferable to "**[No]**", it is perfectly acceptable to answer "**[No]**" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "**[No]**" or "**[NA]**" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer **[Yes]** to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading “NeurIPS paper checklist”.**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

### 1. Claims

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

Answer: **[Yes]**

Justification: Our claims accurately reflect the contributions and scope of the paper.

Guidelines:

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

### 2. Limitations

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

Answer: **[Yes]**

Justification: In section 7, we discuss potential future research directions, which are the limitations of our current work.

Guidelines:

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

### 3. Theory Assumptions and Proofs

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

Answer: [\[Yes\]](#)

Justification: We present all the assumptions. The complete proofs are provided in the appendix.

Guidelines:

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

### 4. Experimental Result Reproducibility

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

Answer: [\[Yes\]](#)

Justification: The code is uploaded in the supplementary material. All the information required to reproduce the experimental results is provided.

Guidelines:

- The answer NA means that the paper does not include experiments.

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

## 5. Open access to data and code

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

Answer: **[Yes]**

Justification: The code is uploaded in the supplementary material.

Guidelines:

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

- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

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

Answer: **[Yes]**

Justification: Most of these details are explained in section 6 and in the provided code.

Guidelines:

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

## 7. Experiment Statistical Significance

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

Answer: **[No]**

Justification: We believe the error bars are not relevant to our empirical metrics.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

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

Answer: **[Yes]**

Justification: We believe the computer resources are not relevant to our main contributions.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.

- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

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

Answer: [\[Yes\]](#)

Justification: The research conducted in our paper conforms with the NeurIPS Code of Ethics.

Guidelines:

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

## 10. Broader Impacts

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

Answer: [\[Yes\]](#)

Justification: Our work contributes to the theory of information elicitation. We discussed the applicability and limitations for elicitation settings.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

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

Answer: [\[NA\]](#)

Justification: We believe this paper poses no such risks.

Guidelines:

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

## 12. Licenses for existing assets

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

Answer: **[Yes]**

Justification: The datasets used in this paper are mentioned with URLs and the licenses and terms of use are properly respected.

Guidelines:

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

## 13. New Assets

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

Answer: **[NA]**

Justification: The paper does not release new assets.

Guidelines:

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

## 14. Crowdsourcing and Research with Human Subjects

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

Answer: [NA]

Justification: The paper does not involve crowdsourcing or research with human subjects.

Guidelines:

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

**15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

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

Answer: [NA]

Justification: The paper does not involve crowdsourcing or research with human subjects.

Guidelines:

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