Two Strongly Truthful Mechanisms for Three Heterogeneous Agents Answering One Question

GRANT SCHOENEBECK, University of Michigan
FANG-YI YU, George Mason University

Peer prediction mechanisms incentivize self-interested agents to truthfully report their signals even in the absence of verification by comparing agents' reports with their peers. We propose two new mechanisms, Source and Target Differential Peer Prediction, and prove very strong guarantees for a very general setting.

Our Differential Peer Prediction mechanisms are *strongly truthful*: Truth-telling is a strict Bayesian Nash equilibrium. Also, truth-telling pays strictly higher than any other equilibria, excluding permutation equilibria, which pays the same amount as truth-telling. The guarantees hold for *asymmetric priors* among agents, which the mechanisms need not know (*detail-free*) in the *single question setting*. Moreover, they only require *three agents*, each of which submits a *single item report*: two report their signals (answers), and the other reports her forecast (prediction of one of the other agent's reports). Our proof technique is straightforward, conceptually motivated, and turns on the logarithmic scoring rule's special properties.

Moreover, we can recast the Bayesian Truth Serum mechanism [20] into our framework. We can also extend our results to the setting of *continuous signals* with a slightly weaker guarantee on the optimality of the truthful equilibrium.

CCS Concepts: • **Information systems** → **Incentive schemes**; • **Theory of computation** → **Quality of equilibria**; • **Mathematics of computing** → *Information theory;*

Additional Key Words and Phrases: Peer prediction, log scoring rule, prediction market

ACM Reference format:
Grant Schoenebeck and Fang-Yi Yu. 2023. Two Strongly Truthful Mechanisms for Three Heterogeneous Agents Answering One Question. *ACM Trans. Econ. Comput.* 10, 4, Article 14 (February 2023), 25 pages.
https://doi.org/10.1145/3565560

1 INTRODUCTION

Three friends, Alice, Bob, and Chloe, watch a political debate on television. We want to ask their opinions on who won the debate. We are afraid they may be less than truthful unless we can pay them for truthful answers. Thus, we seek to design mechanisms that reward the agents for truth-telling. Their opinions may systematically differ, but are nonetheless related. For example, it turns out Alice values description and argumentation, Bob values argumentation and presentation, and Chloe values description and presentation.

Grant Schoenebeck and Fang-Yi Yu contributed equally to this research.

Grant Schoenebeck and Fang-Yi Yu are pleased to acknowledge the support of the National Science Foundation grants 1618187 and 2007256. Fang-Yi Yu is pleased to acknowledge the support of the National Science Foundation grant 2007887.

Authors' addresses: G. Schoenebeck, University of Michigan; email: schoeneb@umich.edu; F.-Y. Yu, George Mason University; email: fyy2412@gmu.edu.

2167-8375/2023/02-ART14
https://doi.org/10.1145/3565560

In this article, we design two peer prediction mechanisms by asking Alice, Bob, and Chloe to play three characters: the *expert,* who makes predictions; the *target* who, is being predicted; and the *source,* who helps predictions. The source and target are asked for their opinions. In the most straightforward setting, the expert makes two predictions (e.g., 70% chance of "yes" and 30% chance of "no") of the target's opinion: An initial prediction before the source's opinion is revealed and an improved prediction afterwards.

Our political debate motivation might be quixotic (at the very least, we need to ensure that they do not communicate during the debate so we can elicit the target's prediction with and without the source's information). However, peer-grading can easily fit into this paradigm: We might ask Bob and Carol to grade a paper, while Alice tries to predict Carol's mark for the paper before and after seeing Bob's mark. Similarly, Alice, Bob, and Carol might be peer-reviewing a paper, filling out a survey,[1] or doing any crowdsourcing task (e.g., labeling data for machine learning applications). In such cases, it is natural to reward agents for doing a good job, and also to have them update a prediction with additional information.

For simplicity (and to collect a grade from all three agents), Alice, Bob, and Chloe, might play all three characters: Alice could predict Chloe's signal before and after seeing Bob's; Bob could predict Alice's signal before and after seeing Chloe's; and Chloe could predict Bob's signal before and after seeing Alice's.

This problem is known in the literature as peer prediction or information elicitation without verification. In the *single-question setting* agents are only asked one question. Incentivizing agents is important so they not only participate, but provide thoughtful and accurate information. Our goal is to elicit truthful information from agents with minimal requirements.

Drawing from previous peer prediction literature, we would like our mechanisms to have the following desirable properties:

Strongly Truthful [12]. Providing truthful answers is a **Bayesian Nash equilibrium (BNE)** and also guarantees the maximum agents' welfare among any equilibrium. This maximum is "strict" with the exception of a few unnatural permutation equilibria where agents report according to a relabeling of the signals (defined more formally in Section 2).[2] This will incentivize the agents to tell the truth–even if they believe the other agents will disagree with them. Moreover, they have no incentive to coordinate on an equilibrium where they do not report truthfully. In particular, note that playing a permutation equilibrium still requires as much effort from the agents as playing truth-telling.

General Signals. The mechanism should work for *heterogeneous* agents who may even have *continuous* signals (with a weaker truthfulness guarantee). In our above example, the friends may not have the same political leanings, and the mechanism should be robust to that. Furthermore, instead of a single winner, we may want to elicit the magnitude of their (perceived) victory.

Detail-Free. The mechanism is not required to know the specifics about the different agents (e.g., the aforementioned joint prior). In the above example, the mechanism should not be required to know the *a priori* political leanings of the different agents.

On Few Agents. We would like our mechanisms to work using as few agents as possible—in our case, three.

Single-item Reports. We would like to make it easy for agents so they provide very little information: only one item, either their signal or a prediction. In our case, two agents will

[1]Here, anonymity may be required to preserve privacy.

[2]Kong and Schoenebeck [12] show that it is not possible for truth-telling to pay strictly more than permutation equilibrium in detail-free mechanisms.

need to provide their signals (e.g., whom they believe won the debate). The remaining agent will need to provide a prediction on one outcome—a single real value. (e.g., their forecast for how likely a particular other agent was to choose a particular candidate as the victor).

1.1 Our Contributions

- We define two Differential Peer Prediction mechanisms (Mechanisms 1 and 2) that are strongly truthful and detail-free for the single question setting and only require a single item report from three agents. Moreover, the agents need not be homogeneous and their signals may be continuous.
- Mechanism 1 rewards the source for the improvement of the experts prediction. We can use any strictly proper scoring rule (see Definition 2.2) to measure the improvement, and truth-telling is an equilibrium. Moreover, if we use the log scoring rule, then truth-telling has the highest total payment among all equilibria.
- Mechanism 2, which rewards the target for the improvement of the experts prediction, exploits special properties of the log scoring rule (see Techniques below for details), which may be of independent interest. Here, the mechanism can be generalized by replacing the expert with a suitable predictor that predicts the target's report given information from a source (which could be the collection of many agents). We show how to recast the Bayesian Truth Serum mechanism into the framework of this Mechanism (Section 4). This gives added intuition for its guarantees.
- We provide a simple, conceptually motivated proof for the guarantees of Differential Peer Prediction mechanisms. Especially in contrast to the most closely related work [10] our proof is very simple.

1.2 Summary of Our Techniques

Target Incentive Mechanisms. Many of the mechanisms for the single question use what we call *source* incentives: They pay agents for reporting a signal that improves the prediction of another agent's signal. The original peer prediction mechanism [15] does exactly this. To apply this idea to the detail-free setting [31, 33], mechanisms take a two-step approach: They first elicit an agent's prediction of some target agent's report and then measure how much that prediction improves given a report from a source agent.

In Section 3.2, we explicitly develop a technique, which we call *target* incentives, for rewarding certain agents for signal reports that agree with a prediction about them. We show that log scoring rules can elicit signals as well as forecasts by paying the difference of log scoring rule on the signal between an initial prediction and an improved prediction. This may be of independent interest and is also the foundation for the results in Sections 3.2 and 4.

Information Monotonicity. We use information monotonicity, a tool from information theory, to obtain strong truthfulness. Like the present article, the core of the argument that the Disagreement Mechanism [10] is strongly truthful (for symmetric equilibrium) is based on information monotonicity. However, because it is hard to characterize the equilibria in the Disagreement Mechanism, the analysis ends up being quite complex. A framework for deriving strongly truthful mechanisms from information monotonicity, which we implicitly employ, is distilled in Kong and Schoenebeck [12].

In Section 3, we use the above techniques to develop strongly truthful mechanisms, source-Differential Peer Prediction and target-Differential Peer Prediction, for the single question setting. Source-Differential Peer Prediction is quite similar to the Knowledge-Free Peer Prediction Mechanism [33], however, it is strongly truthful, which we show using information monotonicity of log

scoring rule. Target-Differential Peer Prediction additionally uses the target incentive techniques above to show it is strongly truthful.

1.3 Related Work

Single Task Setting. In this setting, each agent receives a single signal from a common prior. Miller et al. [15] introduce the first mechanism for single task signal elicitation that has truth-telling as a strict Bayesian Nash equilibrium and does not need verification. However, their mechanism requires full knowledge of the common prior and there exist some equilibria where agents get paid more than truth-telling. At a high level, the agents can all simply submit the reports with the highest expected payment and this will typically yield a payment much higher than that of truth-telling. Note that this is both natural to coordinate on (in fact, Gao et al. [6] found that in an online experiment, agents did exactly this) and does not require any effort toward the task from the agents. Kong et al. [9] modify the above mechanism such that truth-telling pays strictly better than any other equilibrium but still requires the full knowledge of the common prior.

Prelec [20] designs the first detail-free peer prediction mechanism—**Bayesian truth serum (BTS)**. Moreover, BTS is strongly truthful and can easily be made to have one-item reports. However, BTS requires an infinite number of participants, does not work for heterogeneous agents, and requires the signal space to be finite. The analysis, while rather short, is equally opaque. A key insight of this work is to ask agents not only about their own signals, but forecasts (prediction) of the other agents' reports.

A series of works [1, 22, 23, 31–33] relax the large population requirement of BTS but lose the strongly truthful property. Zhang and Chen [33] is unique among prior work in the single question setting in that it works for heterogeneous agents, whereas other previous detail-free mechanisms require homogeneous agents with conditionally independent signals.

To obtain the strongly truthful property, Kong and Schoenebeck [10] introduce the Disagreement Mechanism, which is detail-free, strongly truthful (for symmetric equilibrium), and works for six agents. Thus, it generalizes BTS to the finite agent setting while retaining strong truthfulness. However, it requires homogeneous agents, cannot handle continuous signals, and fundamentally requires that each agent reports both a signal and a prediction. Moreover, its analysis is quite involved. However, it is within the BTS framework, in that it only asks for agents' signals and predictions, whereas our mechanism typically asks at least one agent for a prediction after seeing the signal of another agent.

Finally, most of these works either have multiple rounds [32, 33] or work only if the common prior is symmetric [1, 13, 20, 22, 31], though sometimes this can be relaxed to a restriction more like positive correlation [32]. Our mechanisms also have multiple rounds; however, we can simplify them to single round, but this requires asking questions that may be slightly more complex than the BTS framework.

Prelec [21], posted subsequently to the conference publication of this work [25] but developed independently, uses very similar techniques to this work combined with the setting explored in Reference [32] where agents are asked questions before and after seeing their signal. Similar to our target DPP mechanism, the mechanisms in Prelec [21] are target incentive mechanisms and pay the target by log scoring rule on different pairs of initial and improved predictions (e.g., one agent's predictions before and after getting her signal that requires additional temporal coordination). However, with the above additional temporal coordination, the mechanisms can work on two agents, and our mechanism requires at least three agents for the setting we consider.

Surprisingly, and, with the exception of a footnote in Miller et al. [15], unmentioned by any of the above works, the idea of target incentive mechanisms with the log scoring rule can be dated back over 20 years to a (so far unpublished) working paper [19], which studies information pump games

that also use improvement of predictions on the log scoring rule to encourage truthful reports. In particular, the paper presents a special case of our main technical lemma (Lemma 3.4) that requires a slightly stronger assumption than our second-order stochastic relevant (Definition 2.1). Besides a weaker assumption, our connection to information theory enables us to design strongly truthful mechanisms instead of truthful mechanisms.

Continuous Single Task Setting. Kong et al. [13] show how to generalize both BTS and the Disagreement Mechanism (with similar properties including homogeneous agents) into a restricted continuous setting where signals are Gaussians related in a simple manner. The generalization of the Disagreement Mechanism requires the number of agents to increase with the dimension of the continuous space.

The aforementioned Radanovic and Faltings [23] considers continuous signals. However, it uses a discretization approach that yields exceedingly complex reports. Additionally, it requires homogeneous agents.

In a slightly different setting, Kong and Schoenebeck [11] study eliciting agents' forecasts for some (possibly unverifiable) event, which are continuous values between 0 and 1. However, here, we are concerned with eliciting signals that can be from a much richer space.

Multi-task Setting. In the multi-task setting, introduced in Dasgupta and Ghosh [5], agents are assigned a batch of *a priori* similar tasks that require each agents' private information to be a binary signal. Several works extend this to multiple-choice questions [5, 8, 12, 24, 27]. Recently, a sequence of works study the robustness and limitation of the multi-task setting [3, 26, 34].

The multi-task mechanisms and our single-task mechanism each offer advantages. The key advantage of the multi-task mechanisms is that agents are only asked for their signal, and not a prediction. Multi-task mechanisms accomplish this by, implicitly or explicitly, learning some relation between the reports of different agents. However, because of this, multi-task mechanisms strongly depend on an assumption that both the joint distribution of signals on different questions are i.i.d. and that the agents apply the same (possibly random) strategy to each task in an i.i.d. manner. This assumption is not unreasonable in certain crowd-sourcing, peer review, and peer grading settings, but is likely violated in a survey setting. In the setting of the present article, no such assumption is needed, as the mechanism can be applied individually to each question or task.

Even in settings where the i.i.d. assumption holds, it may be the case that (in practice) agents receive information in addition to the elicited signal so the above learning approach fails. For example, an agent may like a paper, but believe it to be on a generally unpopular topic, and therefore conclude that the mechanism will incorrectly predict her rating. This is because the relation between agents' reports are learned on all topics and so may be incorrect when applied to the subset of papers on unpopular topics. In such a case the strategic guarantees of the multi-task mechanisms may fail. Our mechanism mitigates this problem by having agents themselves doing the prediction, who also have access to the contextual information that will naturally be incorporated into their prediction.

Another drawback of the multi-task setting, as its name suggests, is the number of questions required for each agent. Mechanisms tend to either make assumptions about the correlation between signals (e.g., Reference [5]) or the structure must be learned (e.g., References [24, 27]). In the latter case, the strategic guarantees are parameterized by an ϵ that only decreases asymptotically in the number of agents [24]. An exception to this is the DMI mechanism [8], but this still often requires a fairly large number of tasks to work at all and has additional restrictions. However, recent work [3] shows that the pairing mechanism [24] combined with proper machine learning can work in settings with as few as four tasks per agent. In contrast, our mechanism only requires a single task.

2 PRELIMINARIES

2.1 Peer Prediction Mechanism

There are three characters, Alice, Bob, and Chloe, in our mechanisms. Alice (and, respectively, Bob, Chloe) has a privately observed signal a (respectively, b, c) from a set $\mathcal{A}$ (respectively, $\mathcal{B}$, C). They all share a common belief that their signals (a, b, c) are generated from a random variable (A, B, C) that takes values from $\mathcal{A} \times \mathcal{B} \times C$ with a probability measure P called the *common prior*. P describes how agents' private signals relate to each other.

Agents are Bayesian. For instance, after Alice receives $A = a$, she updates her belief to the *posterior* $P((B, C) = (\cdot, \cdot) \mid A = a)$, which is a distribution over the remaining signals. We will use $P_{B,C|A}(\cdot \mid a)$ instead to simplify the notion. Similarly Alice's posterior of Bob's signal is denoted by $P_{B|A}(\cdot \mid a)$, which is a distribution on $\mathcal{B}$.

A peer prediction mechanism on Alice, Bob, and Chloe has three payment functions (U_A, U_B, U_C). The mechanism first collects reports $\mathbf{r} := (r_A, r_B, r_C)$ from agents. It pays Alice with $U_A(\mathbf{r})$ (and Bob and Chloe analogously). Alice's strategy θ_A is a (random) function from her signal to a report. All agents are rational and risk-neutral, so they are only interested in maximizing their (expected) payment. Thus, given a strategy profile $\boldsymbol{\theta} := (\theta_A, \theta_B, \theta_C)$, Alice, for example, wants to maximize her *ex-ante payment* under common prior P, which is $u_A(\boldsymbol{\theta}; P) := \mathbb{E}_{P,\boldsymbol{\theta}}[U_A(\mathbf{r})]$. Let ex-ante *agents' welfare* denote the sum of ex-ante payment to all agents, $u_A(\boldsymbol{\theta}; P) + u_B(\boldsymbol{\theta}; P) + u_C(\boldsymbol{\theta}; P)$. A strategy profile $\boldsymbol{\theta}$ is a *Bayesian Nash equilibrium* under common prior P if by changing the strategy unilaterally, an agent's payment can only weakly decrease. It is a *strict Bayesian Nash equilibrium* if an agent's payment strictly decreases as her strategy changes.

We want to design peer prediction mechanisms to "elicit" all agents to report their information truthfully without verification. We say Alice's strategy τ_A is *truthful* for a mechanism $\mathcal{M}$ if Alice truthfully reports the information requested by the mechanism.[3] We call the strategy profile $\boldsymbol{\tau}$ truth-telling if each agent reports truthfully. Moreover, we want to design *detail-free* mechanisms that have no knowledge about the common prior P except agents' (possible non-truthful) reports. However, agents can always relabel their signals, and detail-free mechanisms cannot distinguish such a strategy profile from the truth-telling strategy profile. We call these strategy profiles *permutation strategy profiles*. They can be translated back to truth-telling reports by some permutations applied to each component of $\mathcal{A} \times \mathcal{B} \times C$—that is, the agents report according to a relabeling of the signals.

We now define some goals for our mechanism that differ in how unique the high payoff of truth-telling is. We call a mechanism ***truthful*** if the truth-telling strategy profile $\boldsymbol{\tau}$ is a strict Bayesian Nash equilibrium. However, in a truthful mechanism, non-truth-telling equilibria may yield a higher ex-ante payment for each agent. In this article, we aim for ***strongly truthful mechanisms*** [12] that are not only truthful but also ensure the ex-ante agents' welfare in the truth-telling strategy profile $\boldsymbol{\tau}$ is strictly better than all non-permutation equilibria. Note that in a symmetric game, this ensures that each agent's individual expected ex-ante payment is maximized by truth-telling compared to any other symmetric equilibrium.

Now, we define the set of common priors that our detail-free mechanisms can work on. Note that peer reports are not useful when the agents' signals are independent of each other. Thus, a peer prediction mechanism needs to exploit some interdependence between agents' signals.

Definition 2.1 (Zhang and Chen [33]). A common prior P is $\langle A, B, C\rangle$*-second-order stochastic relevant* if for any distinct signals $b, b' \in \mathcal{B}$, there is $a \in \mathcal{A}$, such that $P_{C|A,B}(\cdot \mid a, b) \neq P_{C|A,B}(\cdot \mid a, b')$.

[3]Here, we do not define the notion of truthful reports formally, because it is intuitive in our mechanisms. For general setting, we can use query models to formalize it [29].

Thus, when Alice with a is making a prediction of Chloe's signal, Bob's signal is relevant, so his signal induces different predictions when $B = b$ or $B = b'$.

We call P ***second-order stochastic relevant*** if the above statement holds for any permutation of $\{A, B, C\}$.[4]

To avoid measure theoretic concerns, we initially require that P has full support, and the joint signal space $\mathcal{A} \times \mathcal{B} \times C$ to be finite. In Appendix G, we will show how to extend our results to general measurable spaces.

2.2 Proper Scoring Rules

Scoring rules are powerful tools to design mechanisms for eliciting predictions. Consider a finite set of possible outcomes Ω, e.g., $\Omega = \{\text{sunny}, \text{rainy}\}$. An expert, Alice, first reports a distribution $\hat{P} \in \mathcal{P}(\Omega)$ as her prediction of the outcome, where $\mathcal{P}(\Omega)$ denotes the set of all probability measures on Ω. Then, the mechanism and Alice observe the outcome ω. The mechanism gives Alice a score $\text{PS}[\omega, \hat{P}]$. Alice maximizes her expected score by reporting her true belief for the outcome, P (the probability of each possibly outcome of ω):

Definition 2.2. A ***scoring rule*** $\text{PS} : \Omega \times \mathcal{P}(\Omega) \mapsto \mathbb{R}$ is ***proper*** if for any distributions $P, \hat{P} \in \mathcal{P}(\Omega)$ we have $\mathbb{E}_{\omega \sim P}[\text{PS}[\omega, P]] \geq \mathbb{E}_{\omega \sim P}[\text{PS}[\omega, \hat{P}]]$. A scoring rule PS is ***strictly proper*** when the equality holds only if $\hat{P} = P$.

Given any convex function f, one can define a new proper scoring rule PS^f [12]. In this article, we consider a special scoring rule called the ***logarithmic scoring rule*** [30], defined as

$$\text{LSR}[\omega, P] := \log(p(\omega)), \tag{1}$$

where $p : \Omega \to \mathbb{R}$ is the probability density function of P. Another popular scoring rule is *Brier scoring rule* (quadratic scoring rule) [2], defined as

$$\text{QSR}[\omega, P] := 2p(\omega) - \sum_{\omega' \in \Omega} p(\omega')^2. \tag{2}$$

2.3 Information Theory

Peer prediction mechanisms and prediction markets incentivize agents to truthfully report their signals. One key idea these mechanisms use is that agents' signals are interdependent and strategic manipulation can only dismantle this structure. Here, we introduce several basic notions from information theory [4].

The KL-divergence is a measure of the dissimilarity of two distributions: Let P and Q be probability measures on a finite set Ω with density functions p and q, respectively. The ***KL divergence*** (also called relative entropy) from Q to P is $D_{KL}(P\|Q) := \sum_{\omega \in \Omega} -p(\omega) \log(q(\omega)/p(\omega))$.

We now introduce mutual information, which measures the amount of information between two random variables: Given a random variable (X, Y) on a finite set $\mathcal{X} \times \mathcal{Y}$, let $p_{X,Y}(x, y)$ be the probability density of the random variable (X, Y), and let $p_X(x)$ and $p_Y(y)$ be the marginal probability density of X and Y, respectively. The ***mutual information*** $I(X; Y)$ is the KL-divergence from the joint distribution to the product of marginals:

$$I(X; Y) := \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p_{X,Y}(x, y) \log \frac{p_{X,Y}(x, y)}{p_X(x) p_Y(y)} = D_{KL}(P_{X,Y} \| P_X \otimes P_Y),$$

[4]Our definition has some minor differences from Zhang and Chen [33]'s, for ease of exposition. For instance, they only require the statement holds for one permutation of $\{A, B, C\}$ instead of all the permutations.

where $\otimes$ denotes the tensor product between distributions. Moreover, if (X, Y, Z) is a random variable, then the *mutual information between X and Y conditional on Z* is

$$I(X;Y \mid Z) := \mathbb{E}_Z[D_{KL}(P_{(X,Y)|Z} \| P_{X|Z} \otimes P_{Y|Z})].$$

The data-processing inequality shows no manipulation of the signals can improve mutual information between two random variables, and the inequality is of fundamental importance in information theory.

Theorem 2.3 (Data Processing Inequality). *If $X \to Y \to Z$ forms a Markov chain,*[5]

$$I(X;Y) \geq I(X;Z).$$

Because the mutual information is symmetric, neither can manipulating X increase the mutual information between X and Y. Thus, we say mutual information is information-monotone in both coordinates.

By basic algebraic manipulations, Kong and Schoenebeck [12] relate proper scoring rules to mutual information. For two random variables X and Y,

$$\mathbb{E}_{x,y}\left[\mathrm{LSR}[y, P(Y \mid x)] - \mathrm{LSR}[y, P(Y)]\right] = I(X;Y). \tag{3}$$

We can generalize the mutual information in two ways [12]. The first is to define $f-MI$ using the f-divergence, where f is a convex function, to measure the distance between the joint distribution and the product of the marginal distributions. The KL-divergence is just a special case of the f-divergence. This retains the symmetry between the inputs.

The second way is to use a different proper scoring rule. As mentioned, any convex function f gives rise to a proper scoring rule PS^f. Then the Bregman mutual information can be defined as in Equation (3): $BMI^f(X,Y) := \mathbb{E}_{x,y}[\mathrm{PS}^f[y, P_{Y|X}(\cdot \mid x)] - \mathrm{PS}^f[y, P_Y(\cdot)]]$. Note that by the properties of proper scoring rules BMI is information-monotone in the first coordinate; however, in general, it is not information-monotone in the second.

Thus, by Equation (3), mutual information is the unique measure that is both a Bregman mutual information and an f-MI. This observation is one key for designing our strongly truthful mechanisms.

3 EXPERTS, TARGETS, AND SOURCES: STRONGLY TRUTHFUL PEER PREDICTION MECHANISMS

In this section, we show how to design strongly truthful mechanisms to elicit agents' *signals* by implicitly running a prediction market.

Our mechanisms have three characters, Alice, Bob, and Chloe, and there are three roles: expert, target, and source:

- An expert makes predictions on a target's report,
- a target is asked to report his signal, and
- a source provides her information to an expert to improve the expert's prediction.

By asking agents to play these three roles, we design two strongly truthful mechanisms based on two different ideas.

The first mechanism is ***source differential peer prediction* (S-DPP)**. This mechanism is based on the *knowledge-free peer prediction* mechanism by Zhang and Chen [33], which rewards a *source* by how useful her signal is for an expert to predict a target's report. Their mechanism is only

[5]Random variables X, Y, and Z form a Markov chain if the conditional distribution of Z depends only on Y and is conditionally independent of X.

truthful but not strongly truthful. We carefully shift the payment functions and employ Equation (3) and the data-processing inequality on log scoring rule to achieve the strongly truthful guarantee.

We further propose a second mechanism, ***target differential peer prediction* (T-DPP)**. Instead of rewarding a source, the T-DPP mechanism rewards a *target* by the difference of the logarithmic scoring rule on her signal between an initial prediction and an improved prediction. Later, in Section 4, we show Bayesian truth serum can be seen as a special case of our T-DPP mechanism.

Then, we discuss how to remove the temporal separation between agents making reports in Section 3.3 where agents only need to report once, and their reports do not depend on other agents' reports.

3.1 The Source Differential Peer Prediction Mechanism

The main idea of the S-DPP mechanism is that it rewards a source by the usefulness of her signal for predictions. Specifically, suppose Alice acts as an expert, Bob as the target, and Chloe as the source. Our mechanism first asks Alice to make an ***initial prediction*** $\hat{Q}$ on Bob's report. Then, after Chloe reports her signal, we collect Alice's ***improved prediction*** $\hat{Q}^+$ after seeing Chloe's additional information. In each case, Alice maximizes her utility by reporting her Bayesian posterior conditioned on her information.

The payments for Alice and Bob are simple. S-DPP pays Alice the sum of the logarithmic scoring rule on those two predictions. S-DPP pays Bob zero. Chloe's payment consists of two parts. First, we pay her the prediction score of the improved prediction $\hat{Q}^+$. By the definition of a proper scoring rule (Definition 2.2), Chloe will report truthfully to maximize it. For the second part, we subtract from Chloe's payment three times the score of the initial prediction $\hat{Q}$. This ensures the ex-ante agent welfare equals the mutual information, which is maximized at the truth-telling strategy profile. To ensure Bob also reports his signal truthfully, we permute Bob and Chloe's roles in the mechanism uniformly at random.

Mechanism 1: Two-round Source Differential Peer Prediction

Require: Alice, Bob, and Chloe have private signals $a \in \mathcal{A}$, $b \in \mathcal{B}$, and $c \in C$ drawn from second-order stochastic relevant common prior P known to all three agents. LSR is the logarithmic scoring rule (1).

1: Bob and Chloe report their signals, $\hat{b}$ and $\hat{c}$.
2: Set Alice as the expert. Set Bob or Chloe as the *target* and the other as the *source* uniformly at random. We use t to denote the target's report and use s to denote the source's report.
3: Alice is informed who is the target and predicts the target's report t with $\hat{Q}$.
4: Given the source's report s, the expert makes another prediction $\hat{Q}^+$.
5: The payment to the expert is $\mathrm{LSR}[t, \hat{Q}] + \mathrm{LSR}[t, \hat{Q}^+]$.
6: The payment to the target is 0.
7: The payment to the source is $\mathrm{LSR}[t, \hat{Q}^+] - 3\,\mathrm{LSR}[t, \hat{Q}]$.

Theorem 3.1. *If the common prior P is second-order stochastic relevant on a finite set with full support, then Mechanism 1 is strongly truthful:*

(1) *The truth-telling strategy profile $\boldsymbol{\tau}$ is a strict Bayesian Nash equilibrium.*
(2) *The ex-ante agents' welfare in the truth-telling strategy profile $\boldsymbol{\tau}$ is strictly better than all non-permutation strategy profiles.*

We defer the proof to Appendix C. Intuitively, because the logarithmic scoring rule is proper, Alice (the expert) will make truthful predictions when Bob and Chloe report their signals truthfully. Similarly, the source is willing to report her signal truthfully to maximize the improved prediction score. This shows Mechanism 1 is truthful.

To show the source is willing to report truthfully, we show Lemma 3.2, which is a data processing inequality for second-order stochastic relevant distributions, and present the proof in Appendix C.

Lemma 3.2. *Let random variable (X, Y, Z) be $\langle X, Y, Z\rangle$-stochastic relevant on a finite space $\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$ with full support. Given a deterministic function $\theta : \mathcal{Y} \to \mathcal{Y}$,*

$$\mathbb{E}_{x,y,z}\left[\log\left(\frac{P_{Z|XY}(z \mid x, y)}{P_{Z|X}(z \mid x)}\right)\right] - \mathbb{E}_{x,y,z}\left[\log\left(\frac{P_{Z|XY}(z \mid x, \theta(y))}{P_{Z|X}(z \mid x)}\right)\right] \geq 0.$$

Moreover, equality occurs only if θ is an identity function, $\theta(y) = y$.

Though Lemma 3.2 only considers the log scoring rule, it is straightforward to show the source is willing to report truthfully when we use any strictly proper scoring rule. Consequentially, the S-DPP mechanism will still have truth-telling as an equilibrium. However, the total payment at the truth-telling strategy profile may not be maximum.

Note that we can ask Alice, Bob, and Chloe to play all three characters and have the identical guarantee as Theorem 3.1. We illustrate this modification on n agents in Section 5. Furthermore, if the agents' common prior P is symmetric, then the above modification creates a symmetric game where each agent's expected payment at the truth-telling strategy profile is both non-negative and maximized among all symmetric equilibria.

3.2 Target Differential Peer Prediction Mechanism

The **target differential peer prediction mechanism (T-DPP)** is identical to the S-DPP except for the payment functions. In contrast to the S-DPP mechanism, T-DPP rewards a target. We show that paying the difference between initial prediction and an improved prediction on a target's signal can incentivize the target to report truthfully (Lemma 3.4).

Our mechanism pays Alice the sum of the log scoring rule on those two predictions. The mechanism pays Bob the improvement from the initial prediction $\hat{Q}$ to the improved prediction $\hat{Q}^+$. Finally, Chloe's payment depends on Alice's first initial prediction $\hat{Q}$, which is independent of Chloe's action. To ensure Chloe also reports her signal truthfully, we permute the roles of Bob and Chloe uniformly at random in the mechanism as well.

Mechanism 2: Two-round Target Differential Peer Prediction

Require: Alice, Bob, and Chloe have private signals $a \in \mathcal{A}$, $b \in \mathcal{B}$, and $c \in C$ drawn from second-order stochastic relevant common prior P known to all three agents. LSR is the logarithmic scoring rule (1).

1: Bob and Chloe report their signals, $\hat{b}$ and $\hat{c}$.
2: Set Alice as the expert. Set Bob or Chloe as the *target* and the other as the *source* uniformly at random. We use t to denote the target's report and use s to denote the source's report.
3: Alice is informed who is the target and predicts the target's report t with $\hat{Q}$.
4: Given the source's report s, the expert makes another prediction $\hat{Q}^+$.
5: The payment to the expert is $\mathrm{LSR}[t, \hat{Q}] + \mathrm{LSR}[t, \hat{Q}^+]$.
6: The payment to the target is $\mathrm{LSR}[t, \hat{Q}^+] - \mathrm{LSR}[t, \hat{Q}]$.
7: The payment to the source is $-2\,\mathrm{LSR}[t, \hat{Q}]$.

Theorem 3.3. *If the common prior P is second-order stochastic relevant on a finite set with full support, Mechanism 2 is strongly truthful.*

Although the theoretical guarantee in Theorems 3.1 and 3.3 are identical, in Section 5, we discuss that target DPP may be more robust if we want to replace the expert as an machine learning algorithm.

We defer the proof to Appendix D and provide a sketch here. We first show Mechanism 2 is truthful. Because the log scoring rule is proper, Alice (the expert) will make the truthful predictions when Bob and Chloe report their signals truthfully. Thus, the difficult part is to show the target is willing to report his signal truthfully if the expert and the source are truthful. Because the roles of Bob and Chloe are symmetric in the mechanism, we can assume Bob is the target and Chloe is the source from now on.

Lemma 3.4 (Logarithmic Proper Scoring Rule Reversed). *Suppose Alice and Chloe are truthful, and the common prior is $\langle A, B, C\rangle$-second-order stochastic relevant. As the target, Bob's best response is to report his signal truthfully.*

This is a generalization of a lemma in Prelec [20] and Kong and Schoenebeck [12] and extends to non-symmetric prior and finite agent setting. The main idea to prove Lemma 3.4 is to show that *maximizing Bob's expected payment is equivalent to maximizing the reward of a proper scoring rule applied to predicting Chloe's report with prediction $P(C \mid \theta(b))$*. Therefore, by the property of proper scoring rules, Bob is incentivized to tell the truth.[6] With Lemma 3.4, the rest of the proof of Theorem 3.3 is identical to the proof of Theorem 3.1, which is included in Appendix D.

Proof of Lemma 3.4. Given Alice and Chloe are truthful, let $\theta : \mathcal{B} \to \mathcal{B}$ be Bob's (deterministic) best response. Let Alice, Bob, and Chloe's signals be a, b, and c, respectively. When Alice and Chloe both report truthfully, Chloe's report is $\hat{c} = c$. Alice's initial prediction is $Q = P_{B|A}(\cdot \mid a)$, and her improved prediction is $Q^+ = P_{B|A,C}(\cdot \mid a, c)$. Hence, Bob with strategy θ gets payment

$$\mathrm{LSR}[\theta(b), P_{B|AC}(\cdot \mid a, c)] - \mathrm{LSR}[\theta(b), P_{B|A}(\cdot \mid a)].$$

Because θ is a best response, for all $b \in \mathcal{B}$, reporting $\theta(b)$ maximizes Bob's expected payment conditional on $B = b$,

$$\mathbb{E}_{(a,c)\sim A,C|B=b}\left[\mathrm{LSR}[\theta(b), P_{B|A,C}(\cdot \mid a, c)] - \mathrm{LSR}[\theta(b), P_{B|A}(\cdot \mid a)]\right]. \tag{4}$$

The ex-ante payment of Bob is computed by summing over Equation (4) with weight P_B, as:

$$u(\theta) := \mathbb{E}_{(a,b,c)\sim P}\left[\mathrm{LSR}[\theta(b), P_{B|A,C}(\cdot \mid a, c)] - \mathrm{LSR}[\theta(b), P_{B|A}(\cdot \mid a)]\right],$$

which is maximized over θ. Now, we can swap the role of B and C.

$$\begin{aligned}
u(\theta) &= \mathbb{E}_{(a,b,c)\sim P}\left[\mathrm{LSR}[\theta(b), P_{B|A,C}(\cdot \mid a, c)] - \mathrm{LSR}[\theta(b), P_{B|A}(\cdot \mid a)]\right] \\
&= \mathbb{E}_{a,b,c}\left[\log(P_{B|A,C}(\theta(b) \mid a, c)) - \log(P_{B|A}(\theta(b) \mid a)\right] \qquad \text{(by the definition (1))} \\
&= \mathbb{E}_{a,b,c}\left[\log\left(\frac{P_{B|A,C}(\theta(b) \mid a, c)}{P_{B|A}(\theta(b) \mid a)}\right)\right] \\
&= \mathbb{E}_{a,b,c}\left[\log\left(\frac{P_{B,C|A}(\theta(b), c \mid a)}{P_{B|A}(\theta(b) \mid a)P_{C|A}(c \mid a)}\right)\right] \\
&= \mathbb{E}_{a,b,c}\left[\log\left(\frac{P_{C|A,B}(c \mid a, \theta(b))}{P_{C|A}(c \mid a)}\right)\right].
\end{aligned}$$

[6] Prelec [19] also shows a weaker version of the above lemma. However, his proof requires a stronger assumption than second-order stochastic relevant: For any distinct signals $b, b' \in \mathcal{B}$ and signals $a \in \mathcal{A}$ $c \in C$, $P_{C|A,B}(c \mid a, b) \neq P_{C|A,B}(c \mid a, b')$.

The above value can be seen as the ex-ante prediction score of Bob who reports prediction $P_{C|A,C}(\cdot \mid a, \theta(b))$ for Chloe's signal. Similarly, the ex-ante payment of Bob when his strategy is truth-telling τ is

$$u(\tau) = \mathbb{E}_{a,b,c}\left[\log\left(\frac{P_{C|A,B}(c \mid a, b)}{P_{C|A}(c \mid a)}\right)\right].$$

The difference between $u(\tau)$ and $u(\theta)$ is

$$u(\tau) - u(\theta) = \mathbb{E}_{a,b,c}\left[\log\left(\frac{P_{C|A,B}(c \mid a, b)}{P_{C|A}(c \mid a)}\right)\right] - \mathbb{E}_{a,b,c}\left[\log\left(\frac{P_{C|A,B}(c \mid a, \theta(b))}{P_{C|A}(c \mid a)}\right)\right].$$

First, by Lemma 3.2, we know $u(\tau) \geq u(\theta)$. However, because θ is a best response, the inequality is in fact equality $u(\tau) = u(\theta)$. By the second part of Lemma 3.2, this shows θ is an identity and $\theta = \tau$. □

Note that the proof uses (A) the log scoring rule is a Bregman mutual information, which can be written as the difference between two proper scoring rules, and (B) the log scoring rule is also an f-mutual information, which is symmetric between the inputs. Furthermore, though both mechanisms work with the log scoring rule, the S-DPP can work with general proper scoring rule, but the T-DPP cannot. Proposition 3.5 provides a counter-example where the Brier scoring rule (2) applied in the reverse way does not elicit the target to report truthfully, which shows a distinction between the log scoring rule and other scoring rules.

Proposition 3.5. *If we replace the log scoring rule with the Brier scoring rule (2), then there exists an $\langle A, B, C\rangle$-second-order stochastic relevant prior P such that reporting his signal truthfully is not a best response for Bob.*

Proof. Let $\mathcal{A} = \{1\}$, and $\mathcal{B} = C = \{1, 2, 3\}$. We define an $\langle A, B, C\rangle$-second-order stochastic relevant prior

$$(P(1, b, c))_{b,c} = \begin{pmatrix} 0.12 & 0.11 & 0.16 \\ 0.04 & 0.05 & 0.18 \\ 0.15 & 0.18 & 0.01 \end{pmatrix}.$$

By direct computations, Bob's payment is 0.0878 under truth-telling strategy, but he can get 0.0990 if he misreports 1 as 2. □

3.3 Single-round DPP Mechanism for Finite Signal Spaces

When the signal spaces are finite, the above two-round mechanisms (Mechanisms 1 and 2) can be reduced to single-round mechanisms by using a virtual signal w. That is, for Alice's improved prediction, we provide Alice with a random virtual signal w instead of the actual report from the source and pay her the prediction score when the source's report is equal to the virtual signal $s = w$. Here, we state only the single-round target-DPP; the single-round source-DPP can be defined analogously.

Mechanism 3 has the same truthfulness guarantees as Mechanism 2. The proof is the same and is presented in Appendix E.

Theorem 3.6. *If agents' common beliefs are stochastic relevant and the set $\mathcal{B}$ and C are finite, then Mechanism 3 is strongly truthful.*

Remark 3.7. Mechanism 3 uses the virtual signal trick to decouple the dependency between the expert's (Alice's) prediction and the source's (Chloe's) signal, $w \in \mathcal{X}_s$. Furthermore, the logarithmic scoring rule is a local proper scoring rule [17] such that the score $\mathrm{LSR}[w, P] = \log p(w)$ only depends on the probability at w. Hence, we can further simplify Alice's report by asking her to

Mechanism 3: Single Round T-DPP

Require: Alice, Bob, and Chloe have private signals $a \in \mathcal{A}$, $b \in \mathcal{B}$, and $c \in C$ drawn from second-order stochastic relevant common prior P known to all three agents. The empty set $\emptyset$ is neither in $\mathcal{B}$ nor C.

1: Bob and Chloe report their signals, $\hat{b}$ and $\hat{c}$.
2: Set Alice as the expert. Set Bob or Chloe as the *target* and the other as the *source* uniformly at random. We use t to denote the target's report and use s to denote the source's report.
3: Sample w uniformly from $\mathcal{X}_s \cup \{\emptyset\}$ where $\mathcal{X}_s$ is the signal space of the source, and tell the expert w and who is the target.
4: **if** $w = \emptyset$ **then** ▷ initial prediction
5: The expert makes a prediction $\hat{Q}$ of t.
6: **else** ▷ improved prediction
7: The expert makes prediction $\hat{Q}$ of t pretending the source's report $s = w$.
8: **end if**
9: The payment to the expert is $\mathbf{1}[w = s] \cdot \mathrm{LSR}[t, \hat{Q}] + \mathbf{1}[w = \emptyset] \cdot \mathrm{LSR}[t, \hat{Q}]$.
10: The target's payment has three cases: $\mathbf{1}[w = s] \cdot \mathrm{LSR}[t, \hat{Q}] - \mathbf{1}[w = \emptyset] \cdot \mathrm{LSR}[t, \hat{Q}]$.
11: The payment to the source is $-2 \cdot \mathbf{1}[w = \emptyset]\, \mathrm{LSR}[t, \hat{Q}]$.

predict the probability density $\in [0, 1]$ of a single virtual signal $z \in \mathcal{X}_t$ in the target's (e.g., Bob's) signal space.

This trick can be extended to settings with a countably infinite set of signals. For example, for signals in $\mathbb{N}$, we can generate the virtual signal from a Poisson distribution (which dominates the counting measure) and normalize the payments correspondingly. However, this trick does not work on general measurable spaces, e.g., real numbers, because the probability of the virtual signal matching the source's report can be zero.

4 BAYESIAN TRUTH SERUM AS A PREDICTION MARKET

In this section, we revisit the original **Bayesian Truth Serum (BTS)** by Prelec [20] from the perspective of prediction markets. We first define the setting, which is a special case of ours (Mechanism 2), and use the idea of prediction markets (Appendix A) to understand BTS.

4.1 Setting of BTS

There are n agents. They all share a common prior P. We call P *admissible* if it consists of two main elements: states and signals. The *state* T is a random variable in $\{1, \ldots, m\}$, $m \geq 2$, which represents the true state of the world. Each agent i observes a *signal* X_i from a finite set Ω. The agents have a common prior consisting of $P_T(t)$ and $P_{X|T}(\cdot \mid t)$ such that the prior joint distribution of $x_1, \ldots, x_n$ is $\Pr(X_1 = x_1, \ldots, X_n = x_n) = \prod_{t \in [m]} P_T(t) \prod_{i \in [n]} P_{X|T}(x_i \mid t)$.

Now, we restate the main theorem concerning Bayesian Truth Serum:

Theorem 4.1 ([20]). *For all $\alpha > 1$, if the common prior P is admissible and $n \to \infty$, Mechanism 4 is strongly truthful.*

4.2 Information Score as Prediction Market

Prelec [20] uses a clever algebraic calculation to prove this main result. Kong and Schoenebeck [12] use information theory to show that for BTS the ex-ante agents' welfare for the truth-telling strategy profile is strictly better than for all other non-permutation equilibria. Here, we use the

Mechanism 4: The original BTS

Require: The common prior is admissible, and $\alpha > 1$.

1: Agent i reports $\hat{x}_i \in \Omega$ and $\hat{Q}_i \in \mathcal{P}(\Omega)$.

2: For each agent i, choose a reference agent $j \neq i$ uniformly at random. Compute $Q^{(n)}_{-ij} \in \mathcal{P}(\Omega)$ such that for all $x \in \Omega$

$$Q^{(n)}_{-ij}(x) = \frac{1}{n-2} \sum_{k \neq i,j} \mathbf{1}[\hat{x}_k = x], \tag{5}$$

which is the empirical distribution of the other $n-2$ agents' reports.

3: The prediction score and information score of i are

$$S_{\mathrm{Pre}} = \mathrm{LSR}\left[\hat{x}_j, \hat{Q}_i\right] - \mathrm{LSR}\left[\hat{x}_j, Q^{(n)}_{-ij}\right] \text{ and } S_{\mathrm{Im}} = \mathrm{LSR}\left[\hat{x}_i, Q^{(n)}_{-ij}\right] - \mathrm{LSR}\left[\hat{x}_i, \hat{Q}_j\right].$$

And the payment to i is $S_{\mathrm{Pre}} + \alpha\, S_{\mathrm{Im}}$.

idea of prediction markets to show BTS is a truthful mechanism and use Mechanism 2 to reproduce BTS when the common prior is admissible and $n \to \infty$.

The payment from BTS consists of two parts: the *information score*, S_{Im} and the *prediction score*, S_{Pre}. The prediction score is exactly the log scoring rule that is well-studied in the previous literature. However, the role of the information score is more complicated. Here, we provide an interpretation based on Mechanism 2. Informally, the information score is the improvement from one agent's prediction to the aggregating prediction from all agents on one agent's signal, which is formalized in Proposition 4.2. Thus, by Lemma 3.4, reporting signal truthfully maximizes the agent's information score.

Now, we formalize this idea. Consider $i = 1$ and $j = 2$ in BTS and call them Bob and Alice, respectively. We let Chloe be the collection of other agents $\{3, 4, \ldots, n\}$. Let us run Mechanism 2 on this information structure. Bob is the target. Alice's initial prediction is $Q = P_{X_1|X_2}(\cdot \mid x_2)$. When Chloe's signal is $x_3, x_4, \ldots, x_n$, Alice's improved prediction is $Q^+ = P_{X_1|X_{-1}}(\cdot \mid x_{-1})$, where $x_{-1} := (x_2, x_3, \ldots, x_n)$ is the collection of all agents' reports expect Bob's. By Lemma 3.4, Bob's payment, $\mathrm{LSR}[\hat{x}_1, Q^+] - \mathrm{LSR}[\hat{x}_1, Q]$, which equals

$$\mathrm{LSR}[\hat{x}_1, P_{X_1|X_{-1}}(\cdot \mid x_{-1})] - \mathrm{LSR}[\hat{x}_1, P_{X_1|X_2}(\cdot \mid x_2)], \tag{6}$$

is maximized in expectation when Bob reports his private signal x_1.

Note that Bob's payment here (Equation (6)) is nearly identical to Bob's information score in the BTS (Mechanism 4) at the truth-telling strategy profile: $\mathrm{LSR}[\hat{x}_1, Q^{(n)}_{-\{1,2\}}] - \mathrm{LSR}[\hat{x}_1, \hat{Q}_2]$, which equals

$$\mathrm{LSR}\left[\hat{x}_1, Q^{(n)}_{-\{1,2\}}\right] - \mathrm{LSR}\left[\hat{x}_1, P_{X_1|X_2}(\cdot \mid x_2)\right]. \tag{7}$$

The only difference between Equations (6) and (7) is that the former predicts $\hat{x}$ using $P_{X_1|X_{-1}}(\cdot \mid x_{-1})$ in the first term, while the latter uses $Q^{(n)}_{-\{1,2\}}$. Therefore, the original BTS reduces to a special case of Mechanism 2 as $n \to \infty$, if we can show $\lim_{n\to\infty} P(X_1 \mid x_{-1}) = \lim_{n\to\infty} Q^{(n)}_{-\{1,2\}}$. Formally,

Proposition 4.2. *For all $t = 1, \ldots, m$ and $w \in \Omega$,*

$$Q^{(n)}_{-\{1,2\}}(w) - P_{X_1|X_{-1}}(w \mid x_{-1}) \xrightarrow{P_{X|T}(\cdot|t)} 0 \text{ as } n \to \infty.$$

That is, the difference between these estimators converges to zero in probability as n goes to infinity.

The proposition follows by seeing that, fixing the state of the world t, both $Q^{(n)}_{-\{1,2\}}(\cdot)$ and $P_{X_1|X_{-1}}(\cdot \mid x_{-1})$ converge to $P_{X|T}(w \mid t)$, which is the posterior distribution of Bob's signal given

the state of the world. However, Proposition 4.2 requires agents' signal are symmetric and conditionally independent. In Section 5, we discuss that, in practice, we may replace the simple average $Q^{(n)}_{-\{1,2\}}(\cdot)$ with another learning algorithm to relax these assumptions.

5 DISCUSSION AND APPLICATIONS

We define two Differential Peer Prediction mechanisms, S-DPP and T-DPP, which are strongly truthful, detail-free, and only require a single item report from three agents. In addition to the nice theoretical guarantees, our core observation is that paying an agent the difference between an initial and improved prediction of his signal is a powerful peer prediction tool.

We believe that our mechanism can be applied in several domains, including peer grading, peer review, surveys, and crowd-sourcing. In fact, Srinivasan and Morgenstern [28] use our DPP mechanisms in their proposed market for peer-review process. Moreover, in peer grading, multi-round review processes are already used in practice [18].

As discussed in the related work, existing single-task mechanisms either make strong assumptions on the signal distribution, e.g., symmetric, or also use multiple rounds. We believe our multi-round mechanism are more practical than mechanisms requiring strong assumptions on the signal distribution, which may not hold in these domains.

While multi-task peer-prediction mechanism can also sometimes be deployed in this area, the present mechanism has three advantages: (1) it only requires one question or task, while the multi-task peer prediction mechanisms often require many. This is of great importance in peer grading and peer review, where each agent may only grade a small handful of items. (2) Other mechanisms require learning the relationship between signals; however, in the proposed applications (e.g., peer grading and peer review) the agents typically see much more information than the mere score, and the relationship of signals may depend on particular traits of different items. Our mechanisms mitigate this problem, because the expert can also see the item and use its traits to inform her prediction. (3) Unlike the multi-task setting, our mechanism does not require questions and responses to be i.i.d.

Our presentation involves each agent only giving a one item response. As highlighted in the introduction, our mechanism can easily be adapted so each agent plays all three roles, and thus provides a signal and a prediction. Specifically, given n agents, we could assign each agent an index in $[n]$. In the first round, each agent reports her signal and an initial prediction on $i+1$'s reported signal. In the second round, agent i receives $i-1$'s reported signal and makes an improved prediction on $i+1$'s reported signal.[7] This variant mechanism treats agents symmetrically and can collect more signals, which is often the goal. Furthermore, this symmetric design may be more fair.

Our mechanism does require some coordination between agents, but in general it is quite minimal. First, we assume that the identities of agents are established. Because we allow heterogeneous agents, the expert must know who the target is to respond. However, in practice, this could be relaxed to knowing the "type" of each of the agents as long as knowing the type is sufficient to specify the joint prior. Additionally, if agents are homogeneous, then agents' identities are irrelevant. Second, agents cannot be paid until all the reports are in, because some payments rely on all reports. However, in the single-round mechanisms, no additional coordination is required: Agents can interact with the mechanism in any order. Even in the two-round mechanisms, the only requirement is that the expert must participate after the source. In the case where roles can safely be correlated with arrival times, the first arrivals can be assigned to source/target and the final to be the expert, and then no further coordination is required.

[7] We use modular arithmetic here.

Machine Learning-aided Peer Prediction. Given the ubiquity of learning algorithms, in our S-DPP and T-DPP, we may use a learning algorithm to replace the character of the expert that makes predictions on the target's report. With this modification, agents only need to report their signals without making complicated predictions. Therefore, using learning algorithms as a surrogate may greatly simplify the complication of our mechanisms.

Now, we discuss possible conditions for those learning algorithms to ensure truthfulness. S-DPP requires that the learning algorithms can improve their prediction based on one agent's signal, namely, the source's. However, by Lemma 3.4, T-DPP really only requires that the learning algorithms can make two predictions on the target's report such that the improved prediction is better than the initial one. The condition in T-DPP is weaker than the condition in S-DPP, because a learning algorithm may not have discernible improvement based on one agent's (source's) signal, but can still make an improved prediction with enough information. For instance, the initial predictions in the BTS (Mechanism 4) is one agent's prediction $\hat{Q}_j$, and the improved prediction is the empirical average $Q^{(n)}_{-i,j}$. We can replace the empirical average with any learning algorithm that uses all other agents' signals to make improved predictions.

As mentioned before, if the agents are privy to additional information that systematically changes the relationship between agents' signals, then the machine learning algorithms applied to the entire data, but not given access to the instances themselves, may not work; for example, if two agents agree in their assessments of dramatic movies but always disagree in their assessments of comedy movies. The issue is that the relationship cannot be properly learned without information about the movie itself. To combat this issue, the machine learner could take as input the instances themselves [14].

One future direction is to use this machinery to analyze when BTS retains its strongly truthful guarantee, e.g., for what parameters of finite and/or heterogeneous agents.

APPENDICES

A INTRODUCTION TO PREDICTION MARKETS

Now, we want to get the collective prediction from a large group of experts. If we ask them all to report the prediction simultaneously and pay each of them the log scoring rule on their predictions, then we only receive many different predictions and it is not clear how to aggregate those predictions into a single prediction.

Hanson's [7] idea is to approach the experts *sequentially*. The mechanism asks experts to predict, *given predictions that previous experts have made*, and pays the experts the difference of score between their prediction minus the score of the previous one. Formally,

(1) The designer chooses an initial prediction $\hat{y}_0$, e.g., the uniform distribution on Ω.
(2) The experts $i = 1, 2, \ldots, n$ arrive in order. Each expert i changes the prediction from $\hat{y}_{i-1}$ to $\hat{y}_i$
(3) The market ends and the event's outcome $w \in \Omega$ is observed.
(4) Expert i receives a payoff $\mathrm{PS}[w, \hat{y}_i] - \mathrm{PS}[w, \hat{y}_{i-1}]$.

Therefore, each expert (strictly) maximizes his expected score by reporting his truth belief given his own knowledge and the prediction of the previous experts.

Suppose, instead of multiple experts arriving in order, we have one expert (Alice) but multiple signals arrive in order. For example, Alice is asked to predict the champion of a tennis tournament where $w \in \Omega$ is the set of players. As the tournament proceeds, Alice collects additional signals $(x_i)_{i=1,\ldots,n}$ that inform the outcome. Formally,

(1) The designer chooses an initial prediction $\hat{y}_0$.
(2) In round $i = 1, 2, \ldots, n$, a signal x_i arrives, and Alice changes the prediction from $\hat{y}_{i-1}$ to $\hat{y}_i$

(3) At the end, the outcome $w \in \Omega$ is observed.
(4) Alice receives a payoff $\sum_{i=1}^{n} (\mathrm{PS}[w, \hat{y}_i] - \mathrm{PS}[w, \hat{y}_{i-1}])$.

With belief P, if Alice reports truthfully in each round, then she will report $P(W \mid y_1, y_2, \ldots, y_i)$ at round i. If we use log scoring rule, then her payment at round i will be $I(Y_i; W | Y_1, \ldots, Y_{i-1})$. Her overall payment will be $I(Y_1, \ldots, Y_n; W)$, which maximizes her payment. This is an illustration of the chain rule for Mutual Information: $I(X, Y; Z) = I(Y; ZY|X) + I(X; Z)$.

B DATA PROCESSING INEQUALITY

There are several proofs for the data processing inequality (Theorem 2.3). However, for information elicitation, we often aim for a strict data processing inequality such that given a pair of random variables (X, Y) if a random function $\theta : \mathcal{Y} \to \mathcal{Y}$ is not a invertible function, $I(X; Y) > I(X; \theta(Y))$. In this section, we will show such guarantee holds if X and Y are stochastic relevant (defined later).

We say a pair of random variable X, Y on a finite space $\mathcal{X} \times \mathcal{Y}$ is ***stochastic relevant*** if for any distinct x and x' in $\mathcal{X}$, $P_{Y|X}(\cdot \mid x) \neq P_{Y|X}(\cdot \mid x')$. And the above condition also holds when we exchange X and Y.

Theorem B.1. *If (X, Y) on a finite space $\mathcal{X} \times \mathcal{Y}$ is stochastic relevant and has full support. For all random function θ from $\mathcal{Y}$ to $\mathcal{Y}$ where the randomness of θ is independent of (X, Y),*

$$I(X; Y) = I(X; \theta(Y))$$

if and only if θ is a deterministic invertible function. Otherwise, $I(X; Y) > I(X; \theta(Y))$.

Moreover, we can extend this to conditional mutual information when the random variable is second-order stochastic relevant (Definition 2.1).

Proposition B.2. *If (W, X, Y) on a finite space $\mathcal{W} \times \mathcal{X} \times \mathcal{Y}$ is second-order stochastic relevant and has full support. For any random function θ from $\mathcal{Y}$ to $\mathcal{Y}$, if the randomness of θ is independent of random variable (W, X, Y),*

$$I(X; Y \mid W) = I(X; \theta(Y) \mid W)$$

if and only if θ is an one-to-one function. Otherwise, $I(X; Y \mid W) > I(X; \theta(Y) \mid W)$.

B.1 Proof of Theorem B.1

Theorem B.3 (Jensen's Inequality). *Let X be a random variable on a probability space $(\mathcal{X}, \mathcal{F}, \mu)$ and let $f : \mathbb{R} \to \mathbb{R}$ be a convex function. Then $f(\mathbb{E}[X]) \leq \mathbb{E}[f(X)]$. The equality holds if and only if f agree almost everywhere on the range of X with a linear function.*

Given a random function $\theta : \mathcal{Y} \to \mathcal{Y}$, we use $q : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$ to denote its transition matrix where $q(y, \hat{y}) = \Pr[\theta(y) = \hat{y}]$ for all $y, \hat{y} \in \mathcal{Y}$. Let $\hat{Y}$ be the random variable $\theta(Y)$.

Variational representation. By the variational representation of mutual information [16, 24], let $\Phi(a) = a \log a$, $\Phi^*(b) = \exp(b - 1)$ and $\Phi'(a) = 1 + \log a$ the mutual information between X and Y is

$$I(X; Y) = \sup_{k: \mathcal{X} \times \mathcal{Y} \to \mathbb{R}} \left\{ \mathbb{E}_{P_{X,Y}}[k(X, Y)] - \mathbb{E}_{P_X \otimes P_Y}\left[\Phi^*\left(k(X, Y)\right)\right] \right\}$$

and the maximum happens when

$$K(x, y) := \Phi'\left(\frac{P_{X,Y}(x, y)}{P_X(x) P_Y(y)}\right). \tag{8}$$

We define $\hat{K}$ for X and $\hat{Y}$ similarly. With these notions, the mutual information between X and $\hat{Y}$ is

$$\begin{aligned} I(X;\hat{Y}) &= \mathbb{E}_{P_{X,\hat{Y}}}[\hat{K}(X,\hat{Y})] - \mathbb{E}_{P_X \otimes P_{\hat{Y}}}\left[\Phi^*\left(\hat{K}(X,\hat{Y})\right)\right] \\ &= \mathbb{E}_{P_{X,Y}}\left[\int \hat{K}(x,\hat{y})q(y,\hat{y})d\hat{y}\right] - \mathbb{E}_{P_X \otimes P_Y}\left[\int \Phi^*\left(\hat{K}(x,\hat{y})\right) q(y,\hat{y})d\hat{y}\right] \\ &\le \mathbb{E}_{P_{X,Y}}\left[\int \hat{K}(x,\hat{y})q(y,\hat{y})d\hat{y}\right] - \mathbb{E}_{P_X \otimes P_Y}\left[\Phi^*\left(\int \hat{K}(x,\hat{y})q(y,\hat{y})d\hat{y}\right)\right]. \end{aligned}$$

The last inequality holds due to convexity of Φ^* and Jensen's inequality. Let $L(x,y) := \int \hat{K}(x,\hat{y})q(y,\hat{y})d\hat{y}$ for all x, y. We have

$$\begin{aligned} I(X;\hat{Y}) &\le \mathbb{E}_{P_{X,Y}}\left[L(x,y)\right] - \mathbb{E}_{P_X \otimes P_Y}\left[\Phi^*\left(L(x,y)\right)\right] && (9) \\ &\le \sup_{k:\mathcal{X}\times\mathcal{Y}\to\mathbb{R}} \left\{\mathbb{E}_{P_{X,Y}}[k(X,Y)] - \mathbb{E}_{P_X \otimes P_Y}\left[\Phi^*\left(k(X,Y)\right)\right]\right\} && (10) \\ &= I(X;Y). \end{aligned}$$

Sufficient condition. We first show the equality holds if θ is an invertible function. Hence, we need to show (9) and (10) are equalities. Because θ is an invertible function, q is a permutation matrix. Thus, for all x, y $\int \Phi^*(\hat{K}(x,\hat{y}))q(y,\hat{y})d\hat{y} = \Phi^*(\int \hat{K}(x,\hat{y})q(y,\hat{y})d\hat{y})$, and (9) is equality. For (10), for all x and y,

$$\begin{aligned} L(x,y) &= \int \hat{K}(x,\hat{y})q(y,\hat{y})d\hat{y} \\ &= \hat{K}(x,\theta(y)) && \text{(deterministic function)} \\ &= \Phi'\left(\frac{P_{X,\hat{Y}}(x,\theta(y))}{P_X(x)P_{\hat{Y}}(\theta(y))}\right) && \text{(by (8))} \\ &= \Phi'\left(\frac{P_{X,Y}(x,y)}{P_X(x)P_Y(y)}\right) && \text{(invertible)} \\ &= K(x,y). \end{aligned}$$

Therefore, (10) is an equality. This completes the proof.

Necessary condition. Now, we show the equality holds only if θ is an invertible function, i.e., q is a permutation matrix. We first show a weaker statement, q is injective. Formally, let $R_q(y) := \{\hat{y} : q(y,\hat{y}) > 0\}$ is the support of q on y. We say q is injective if for all distinct y, y' the support of $q(y,\cdot)$ and $q(y',\cdot)$ are disjoint, $R_q(y) \cap R_q(y') = \emptyset$.

We prove this by contradiction: if q is not injective and $I(X;Y) = I(X;\hat{Y})$, (X,Y) is not stochastic relevant. $I(X;Y) = I(X;\hat{Y})$ implies (9) and (10) are equalities. Because (9) is an equality, given x and y for all $\hat{y} \in R_q(y)$,

$$L(x,y) = \hat{K}(x,\hat{y}). \tag{11}$$

Because (10) is an equality, for all x and y,

$$L(x,y) = K(x,y). \tag{12}$$

Suppose q is not injective. There exist y_1, y_2, and y^* in $\mathcal{Y}$ such that $y_1 \neq y_2$ and $y^* \in R_q(y_1) \cap R_q(y_2)$. For all x,

$$\begin{aligned} K(x, y_1) &= L(x, y_1) && \text{(by (12))} \\ &= \hat{K}(x, y^*) && \text{(by (11) and } \hat{y}^* \in R_q(y_1)) \\ &= L(x, y_2) && \text{(by (11) and } \hat{y}^* \in R_q(y_2)) \\ &= K(x, y_2). && \text{(by (12))} \end{aligned}$$

Since Φ' is invertible, for all x

$$\frac{P_{X,Y}(x, y_1)}{P_X(x)P_Y(y_1)} = \frac{P_{X,Y}(x, y_2)}{P_X(x)P_Y(y_2)}.$$

Therefore, $P_{X|Y}(\cdot \mid y_1) = P_{X|Y}(\cdot \mid y_2)$, and (X, Y) is not stochastic relevant. This shows the Markov kernel q is injective and has a deterministic inverse function.

Now, we show if q is injective, then q is a permutation when $\mathcal{Y}$ is a finite space. Because q is a Markov kernel $|R_q(y)| \geq 1$ for all y. Moreover, because q is injective, $|\cup_y R_q(y)| = \sum_y |R_q(y)| \geq |\mathcal{Y}|$. However, $\cup_y R_q(y) = \{\hat{y} : \exists y, (y, \hat{y}) \in R_q\} \subseteq \mathcal{Y}$, $|\cup_y R_q(y)| \leq |\mathcal{Y}|$. Therefore, by pigeonhole principle, $|R_q(y)| = 1$ for all y, which is one-to-one.[8]

B.2 Proof of Proposition B.2

Proposition B.2. Given random variable (W, X, Y), define pointwise conditional mutual information between X and Y given $W = w$ as

$$I(X; Y \mid W = w) := D_{KL}\left(P_{X|W}(\cdot \mid w) \otimes P_{Y|W}(\cdot \mid w) \parallel P_{(X,Y)|W}(\cdot \mid w)\right),$$

which is the mutual information between $X|W = w$ and $Y|W = w$.

First observe that conditional mutual information $I(X; Y \mid W)$ is the average pointwise conditional mutual information between X and Y across different W,

$$I(X; Y \mid W) = \int I(X; Y \mid W = w)\, p_W(w) dw.$$

Thus, we can apply Theorem B.1 to each pointwise conditional mutual information.

The sufficient condition is straightforward. For the necessary condition, we can reuse the argument in the proof of Theorem B.1. Let $\Phi(a) = a \log a$ and

$$K(x, y \mid w) := \Phi'\left(\frac{P_{X,Y|W}(x, y \mid w)}{P_{X|W}(x \mid w)P_{Y|W}(y \mid w)}\right).$$

We define $\hat{K}(x, y \mid w)$ for X, $\hat{Y}$, and W similarly, and we let $L(x, y \mid w) := \int \hat{K}(x, \hat{y} \mid w) q(y, \hat{y}) d\hat{y}$. By similar derivation, we have analogy of Equations (11) and (12): For all x, y, w and $\hat{y} \in R_q(y)$

$$L(x, y \mid w) = \hat{K}(x, \hat{y} \mid w) \tag{13}$$

and

$$L(x, y \mid w) = K(x, y \mid w). \tag{14}$$

[8] Note that the proof implicitly uses the property that the distribution of $(X, Y, \hat{Y})$ has a full support. In particular, Equations (11) and (12) only hold on the support of the distribution.

Suppose q is not injective. There exists y_1, y_2 and y^* such that $y_1 \neq y_2$ and $y^* \in R_q(y_1) \cap R_q(y_2)$. For all x and w

$$\begin{aligned} K(x, y_1 \mid w) &= L(x, y_1 \mid w) && \text{(by (14))} \\ &= \hat{K}(x, y^* \mid w) && \text{(by (13) and } y^* \in R_q(y_1)) \\ &= L(x, y_2 \mid w) && \text{(by (13) and } y^* \in R_q(y_2)) \\ &= K(x, y_2 \mid w). \end{aligned}$$

Since Φ' is injective, for all x and w

$$\frac{P_{X,Y|W}(x, y_1 \mid w)}{P_{X|W}(x \mid w) P_{Y|W}(y_1 \mid w)} = \frac{P_{X,Y|W}(x, y_2 \mid w)}{P_{X|W}(x \mid w) P_{Y|W}(y_2 \mid w)}.$$

Therefore, there exists distinct y_1 and y_2 such that for all w

$$P_{X|Y,W}(\cdot \mid y_1, w) = P_{X|Y,W}(\cdot \mid y_2, w).$$

This contradicts the condition that (X, Y, W) is second-order stochastic relevant. □

C PROOFS IN SECTION 3.1

Proof of Lemma 3.2.

$$\begin{aligned} & \mathbb{E}_{x,y,z}\left[\log\left(\frac{P_{Z|X,Y}(z \mid x, y)}{P_{Z|X}(z \mid x)}\right)\right] - \mathbb{E}_{x,y,z}\left[\log\left(\frac{P_{Z|X,Y}(z \mid x, \theta(y))}{P_{Z|X}(z \mid x)}\right)\right] \\ &= \mathbb{E}_{x,y,z}\left[\log\left(\frac{P_{Z|X,Y}(z \mid x, y)}{P_{Z|X,Y}(z \mid x, \theta(y))}\right)\right] \\ &= \mathbb{E}_{x,y}\left[\mathbb{E}_z\left[\log\left(\frac{P_{Z|X,Y}(z \mid x, y)}{P_{Z|X,Y}(z \mid x, \theta(y))}\right) \mid X = x, Y = y\right]\right] \\ &= \mathbb{E}_{x,y}\left[D_{KL}(P_{Z|X,Y}(\cdot \mid x, \theta(y)) \| P_{Z|X,Y}(\cdot \mid x, y))\right]. \end{aligned}$$

Let $d(x, y, y') := D_{KL}(P_{Z|X,Y}(\cdot \mid x, y') \| P_{Z|X,Y}(\cdot \mid x, y))$, which is the KL-divergence from random variable Z conditional on $X = x$ and $Y = y$ to Z conditional on $X = x$ and $Y = y'$. Thus, we have

$$\mathbb{E}_{x,y}\left[D_{KL}(P_{Z|X,Y}(\cdot \mid x, \theta(y)) \| P_{Z|X,Y}(\cdot \mid x, y))\right] = \mathbb{E}_{x,y}\left[d(x, y, \theta(y))\right]. \tag{15}$$

First, note that by Jensen's inequality (Theorem B.3) $d(x, y, \theta(y)) \geq 0$ for all x and y, so Equation (15) is non-negative. This shows the first part.

Let $E_\theta = \{y : \theta(y) \neq y\} \subseteq \mathcal{Y}$, which is the event such that θ disagrees with the identity mapping. Because P is $\langle X, Y, Z\rangle$-second-order stochastic relevant, for all $y \in E_\theta$ there is x, $P_{Z|X,Y}(\cdot \mid x, y) \neq P_{Z|X,Y}(\cdot \mid x, \theta(y))$, so $d(x, y, \theta(y)) > 0$ by Jensen's inequality (Theorem B.3). Therefore, when equality holds, the probability of event E_θ is zero, and θ is an identity because $\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$ is a finite space. □

Proof of Theorem 3.1. The proof has two parts. Mechanism 1 is truthful and the truth-telling strategy profile maximizes the ex-ante agent welfare.

Truthfulness. We first show Mechanism 1 is truthful. For the expert Alice, suppose Bob and Chloe provide their signals truthfully. Her expected payment consists of two prediction scores $\mathrm{LSR}[b, \hat{Q}]$ and $\mathrm{LSR}[b, \hat{Q}^+]$, where $\hat{Q}$ is her first prediction and $\hat{Q}^+$ is the second. The expected first prediction score (under the randomness of Bob's signal B conditional on Alice's signal being a) is

$$\mathbb{E}_{b \sim P_{B|A}(\cdot|a)}[\mathrm{LSR}[b, \hat{Q}]] \leq \mathbb{E}_{b \sim P_{B|A}(\cdot|a)}[\mathrm{LSR}[b, P_{B|A}(\cdot \mid a)]],$$

which is less than reporting truthful prediction $P_{B|A}(\cdot \mid a)$, since log scoring rule is proper (Definition 2.2). Similarly, her expected payment is maximized when her improved prediction $\hat{Q}^+$ is $P_{B|A,C}(\cdot \mid a, c)$.

If Chloe is the source, then she will tell the truth given Alice and Bob report truthfully by Lemma 3.2. Formally, let Alice's, Bob's, and Chloe's signal be a, b, and c, respectively. Let $\theta : C \to C$ denote Chloe's (deterministic) best response. Alice's initial prediction and Bob's signal is $P_{B|A}(\cdot \mid a)$. Because Chloe unilaterally deviates, Alice's improved prediction is $P_{B|A,C}(\cdot \mid a, \theta(c))$. Therefore, Chloe's payment is $\mathrm{LSR}[b, P_{B|A,C}(\cdot \mid a, \theta(c))] - 3\,\mathrm{LSR}[b, P_{B|A}(\cdot \mid a)]$.

Note that regardless of Chloe's report, the initial prediction is $\hat{Q} = P_{B|A}(\cdot \mid a)$. Hence, equivalently, Chloe's best response also maximizes $\mathrm{LSR}[b, P_{B|A,\hat{C}}(\cdot \mid a, \hat{c})] - \mathrm{LSR}[b, P_{B|A}(\cdot \mid a)]$. Taking expectation over signal A, B, C and strategy θ, we have

$$\begin{aligned} v(\theta) &:= \sum_{a,b,c} P_{A,B,C}(a,b,c)\,(\mathrm{LSR}[b, P_{B|A,C}(\cdot \mid a, \theta(c))] - \mathrm{LSR}[b, P_{B|A}(\cdot \mid a)]) \\ &= \mathbb{E}_{a,b,c}\left[\log(P_{B|A,C}(b \mid a, \theta(c))) - \log(P_{B|A}(b \mid a))\right] \qquad \text{(by (1))} \\ &= \mathbb{E}_{a,b,\hat{c}}\left[\log\left(\frac{P_{B|A,C}(b \mid a, \theta(c))}{P_{B|A}(b \mid a)}\right)\right]. \end{aligned}$$

Similarly, the ex-ante payment of Chloe when her strategy is truth-telling τ is

$$v(\tau) = \mathbb{E}_{a,b,c}\left[\log\left(\frac{P_{B|A,C}(b \mid a, c)}{P_{B|A}(b \mid a)}\right)\right].$$

The difference between $v(\tau)$ and $v(\theta)$ is

$$v(\tau) - v(\theta) = \mathbb{E}_{a,b,c}\left[\log\left(\frac{P_{B|A,C}(b \mid a, c)}{P_{B|A}(b \mid a)}\right)\right] - \mathbb{E}_{a,b,c}\left[\log\left(\frac{P_{B|A,C}(b \mid a, \theta(c))}{P_{B|A}(b \mid a)}\right)\right].$$

First, by Lemma 3.2, we know $v(\tau) \geq v(\theta)$. However, because θ is a best response, the inequality is in fact equality, $v(\tau) = v(\theta)$. By the second part of Lemma 3.2, this shows θ is an identity and $\theta = \tau$.

If Chloe is the target, then her action does not affect her expected payment, so reporting her signal truthfully is a best response strategy. By randomizing the roles of source and target, both Bob and Chloe will report their signals truthfully.

Strongly truthful. Now, we show the truth-telling strategy profile $\boldsymbol{\tau}$ maximizes the ex-ante agent welfare under P. If Bob is the target, then the ex-ante agent welfare (before anyone receives signals) in truth-telling strategy profile $\boldsymbol{\tau}$ is

$$\begin{aligned} \sum_i u_i(\boldsymbol{\tau}; P) &= \mathbb{E}_{(a,b,c)\sim P}\left[2\,(\mathrm{LSR}[b, P_{B|A,C}(\cdot \mid a, c)] - \mathrm{LSR}[b, P_{B|A}(\cdot \mid a)])\right] \\ &= 2\,\mathbb{E}_{(a,b,c)\sim P}\left[\log\left(\frac{P_{B|A,C}(b \mid a, c)}{P_{B|A}(b \mid a)}\right)\right] \\ &= 2I(B; C \mid A), \end{aligned}$$

which is the conditional information between Bob's and Chloe's signals given Alice's signal.

However, let $\boldsymbol{\theta} = (\theta_A, \theta_B, \theta_C)$ be an equilibrium strategy profile, where Bob and Chloe report signals $\theta_B(B)$ and $\theta_C(C)$, respectively. Since $\boldsymbol{\theta}$ is an equilibrium, if Bob is the target, then Alice with signal a will predict truthfully, and report $\hat{Q} = P_{\theta_B(B)|A}(\cdot \mid a)$ and $\hat{Q}^+ = P_{\theta_B(B)|A,\theta_C(C)}(\cdot \mid a, \theta_C(c))$. By a similar computation, the ex-ante agent welfare is

$$\sum_i u_i(\boldsymbol{\theta}; P) = 2I(\theta_B(B); \theta_C(C) \mid A) \leq 2I(B; C \mid A) = \sum_i u_i(P, \boldsymbol{\tau}).$$

The inequality is based on the data processing inequality (Theorem 2.3). Moreover, by Proposition B.2, the equality holds only if θ is a permutation strategy profile. □

D PROOF IN SECTION 3.2

D.1 Proof of Theorem 3.3

The proof is mostly identical to Theorem 3.1 in Appendix C. We include it for completeness.

Proof of Theorem 3.3. The proof also has two parts. Mechanism 2 is truthful and the truth-telling strategy profile maximizes the ex-ante agent welfare.

We first show Mechanism 2 is truthful. For the expert Alice, the proof is identical to the proof of Theorem 3.1. By Lemma 3.4, if Bob is the target, then he will tell the truth given Alice and Chloe report truthfully. If Bob is the source, then his action does not affect his expect payment, so reporting his signal truthfully is a best response strategy. By randomizing the role of source and target, both Bob and Chloe will report their signals truthfully.

The proof for strongly truthful is identical to the proof of Theorem 3.1. □

Note that if we randomize the roles amount Alice, Bob, and Chloe, then each agent has a non-negative expected payment at the truth-telling equilibrium.

E PROOF OF THEOREM 3.6

For the expert Alice, suppose Bob and Chloe provide their signals truthfully. Her payment consists of two prediction scores: When the random variable $w = \emptyset$, the prediction score (under the randomness of Bob's signal B conditional on Alice's signal being a) is

$$\mathbb{E}_{b\sim P_{B|A}(\cdot|a)}[\mathrm{LSR}[b, \hat{Q}]] \le \mathbb{E}_{b\sim P_{B|A}(\cdot|a)}[\mathrm{LSR}[b, P_{B|A}(\cdot \mid a)]].$$

Since log scoring rule is proper (Definition 2.2), reporting truthful prediction $P_{B|A}(\cdot \mid a)$ maximizes it. Similarly, when $w \neq \emptyset$, her (conditional) expected payment is maximized when her improved prediction is $P_{B|A,C}(\cdot \mid a, w)$. For the target Bob, suppose Alice and Chloe report truthfully. We will follow the proof of Lemma 3.4 to show Bob's best response is truth-telling. Let $\theta : \mathcal{B} \to \mathcal{B}$ be Bob's (deterministic) best response. Bob's expected payment depends on four values: signals a, b, c, and virtual signal w:

$$U_B = \mathbf{1}[w = c]\,\mathrm{LSR}[\theta(b), P_{B|A,C}(\cdot \mid a, w)] - \mathbf{1}[w = \emptyset]\,\mathrm{LSR}[\theta(b), P_{B|A}(\cdot \mid a)].$$

And Bob's expected payment is

$$u_B(\theta) = \frac{1}{|C|+1}\,\mathbb{E}_{a,b,c}\left[\mathrm{LSR}[\theta(b), P_{B|A,C}(\cdot \mid a, c)] - \mathrm{LSR}[\theta(b), P_{B|A}(\cdot \mid a)]\right].$$

Thus, by the same argument in Lemma 3.4, Bob's best response is truth-telling. If Bob is the source, then his action does not affect his expect payment, so reporting his signal truthfully is a best response strategy. By randomizing the role of source and target, both Bob and Chloe will report their signals truthfully.

The proof of strongly truthful is identical to the proof of Theorem 3.3.

F SKETCH PROOF FOR PROPOSITION 4.2

A consistent predictor f of a value Y given evidence $X_1, X_2, \ldots$ is one where more information leads to a better prediction such that

$$\lim_{n\to\infty} \Pr[|f(x_1, x_2, \ldots, x_n) - Y| \ge \epsilon] \to 0.$$

The proposition follows by seeing that, fixing t and w, both $Q^{(n)}_{-\{1,2\}}(w)$ and $P_{X_1|X_{-1}}(w \mid x_2, x_3, \ldots, x_n)$ are both consistent estimators for $P_{X_1|T}(w \mid t)$.

$Q^{(n)}_{-\{1,2\}}(w)$ is the empirical distribution of $n-2$ independent samples from $P_{X|T}(\cdot \mid t)$ to estimate $P_{X|T}(w \mid t)$ and is therefore a consistent estimator.

However, because X_1 and $X_2, X_3, \ldots, X_n$ are independent conditional on t, the posterior distribution $P_{T|X_{-1}}(t \mid x_2, x_3, \ldots, x_n)$ is consistent. That is, for all $t \in [m]$, $\Pr[|P(T = t \mid x_2, x_3, \ldots, x_n) - 1| \geq \epsilon \mid T = t] \to 0$. Thus,

$$P_{X_1|X_{-1}}(\cdot \mid x_2, x_3, \ldots, x_n) = \sum_t P_{X_1|T}(\cdot \mid t) P_{T|X_{-1}}(t \mid x_2, x_3, \ldots, x_n)$$

is also a consistent predictor of $P_{X_1|T}(w \mid t)$, which completes the proof.

G GENERAL MEASURE SPACES

G.1 Settings

There are three characters, Alice, Bob, and Chloe. Consider three measure spaces $(\mathcal{A}, \mathcal{S}_A, \mu_A)$, $(\mathcal{B}, \mathcal{S}_B, \mu_B)$, and $(\mathcal{C}, \mathcal{S}_C, \mu_C)$. Let $\mathcal{X} := \mathcal{A} \times \mathcal{B} \times \mathcal{C}$, $\mathcal{S} := \mathcal{S}_A \times \mathcal{S}_B \times \mathcal{S}_C$, and $\mu_X = \mu_A \otimes \mu_B \otimes \mu_C$, where $\otimes$ denotes the product between distributions. Let $\mathcal{P}(\mathcal{X})$ be the set of probability density function on $\mathcal{X}$ with respect to μ_X.[9]

Alice (and, respectively, Bob, Chloe) has a privately observed signal a (respectively, b, c) from set $\mathcal{A}$ (respectively, $\mathcal{B}$ $\mathcal{C}$). They all share a *common prior belief* that their signals (a, b, c) are generated from a random variable $\mathbf{X} := (A, B, C)$ on $(\mathcal{X}, \mathcal{S})$ with a probability measure $P \in \mathcal{P}(\mathcal{X})$, and a positive density function $p > 0$. We consider a *uniform second-order stochastic relevant* for general measure space as follows[10]:

Definition G.1. A random variable (A, B, C) in $\mathcal{A} \times \mathcal{B} \times \mathcal{C}$ with a probability measure P is not $\langle A, B, C\rangle$*-uniform stochastic relevant* if there exist a signal $a \in \mathcal{A}$ and two distinct signals $b, b' \in \mathcal{B}$ such that the posterior on C is identical whether $B = b$ with $A = a$ or $B = b'$ with $A = a$,

$$P_{C|A,B}(\cdot \mid a, b) = P_{C|A,B}(\cdot \mid a, b') \text{ almost surely on } \mu_C.$$

Otherwise, we call P $\langle A, B, C\rangle$***-uniform stochastic relevant***. Thus, when Alice is making a prediction to Chloe's signal, Bob's signal is always relevant and induces different predictions when $B = b$ or $B = b'$.

We call P *uniform second-order stochastic relevant* if it is $\langle X, Y, Z\rangle$-uniform stochastic relevant where $\langle X, Y, Z\rangle$ is any permutation of $\{A, B, C\}$.

G.2 Theorems 3.1 and 3.3 on general measure spaces

Here, we state analogous results to Theorems 3.1 and 3.3. The proofs are mostly identical.

Theorem G.2. *Given a measure space $(\mathcal{X}, \mathcal{S}, \mu_X)$, if the common prior P is uniform second-order stochastic relevant on the measurable space $(\mathcal{X}, \mathcal{S})$, and P is absolutely continuous with respect to μ_X, then Mechanism 1 has the following properties:*

[9]Formally, $\mathcal{P}(\mathcal{X})$ is the set of all distributions on $\mathcal{X}$ that are absolutely continuous with respect to measure μ_X. For $P \in \mathcal{P}(\mathcal{X})$, we denote the density of P with respect to μ by $p(\cdot)$. For example, if $\mathcal{X}$ is a discrete space, then we can set μ as the counting measure. If $\mathcal{X}$ is a Euclidean space $\mathbb{R}^d$, then we can use the Lebesgue measure.

[10]One major difference between $\langle A, B, C\rangle$-stochastic relevant (Definition 2.1) and $\langle A, B, C\rangle$-uniform second-order stochastic relevant (Definition G.1) is the quantifier of A: Given all distinct pair b, b', it is sufficient to have one a^* such that $P_{C|AB}(\cdot|a^*, b) \neq P_{C|AB}(\cdot|a^*, b')$. However, for uniform stochastic relevant, it requires for all a, $P_{C|AB}(\cdot|a, b) \neq P_{C|AB}(\cdot|a, b')$. One issue for second-order stochastic relevant in general measure space is that we can change measure zero point to make such distribution stochastic irrelevant, and the probability to derive a^* such that $P_{C|AB}(\cdot|a^*, b) \neq P_{C|AB}(\cdot|a^*, b')$ may be zero.

(1) *The truth-telling strategy profile $\boldsymbol{\tau}$ is a strict Bayesian Nash Equilibrium.*
(2) *The ex-ante agent welfare in the truth-telling strategy profile $\boldsymbol{\tau}$ is strictly better than all non-invertible strategy profiles.*

Here, the maximum agent welfare happens not only at permutation strategy profiles, but also invertible strategy profile. This limitation is due to the strictness of data processing inequality (Theorem B.1). For example, consider a pair of random variables (X, Y) on $\mathbb{Z}_{>0} \times \mathbb{Z}_{>0}$. Let θ be a Markov operator such that for $x \in \mathbb{Z}_{>0}$, $\theta(x) = x$ with probability $1/2$ and $\theta(x) = -x$ otherwise. Although θ is not a one-to-one function, $I(X;Y) = I(\theta(X);Y)$. However, following the proof of Theorem B.1, we can say the equality holds when θ is injective.

The guarantee of Mechanism 2 is the same.

Theorem G.3. *Given a measure space $(\mathcal{X}, \mathcal{S}, \mu_X)$ if the common prior P is uniform second-order stochastic relevant on the measurable space $(\mathcal{X}, \mathcal{S})$, and P is absolutely continuous with respect to μ_X, then Mechanism 2 has the following properties:*

(1) *The truth-telling strategy profile $\boldsymbol{\tau}$ is a strict Bayesian Nash Equilibrium.*
(2) *The ex-ante agent welfare in the truth-telling strategy profile $\boldsymbol{\tau}$ is strictly better than all non-invertible strategy profiles.*

ACKNOWLEDGMENTS

We would like to thank Drazen Prelec for reference to a related work [19].

REFERENCES

[1] Aurélien Baillon. 2017. Bayesian markets to elicit private information. *Proc. Nat. Acad. Sci.* 114, 30 (2017), 7958–7962.
[2] Glenn W. Brier. 1950. Verification of forecasts expressed in terms of probability. *Month. Weath. Rev.* 78, 1 (1950), 1–3.
[3] Noah Burrell and Grant Schoenebeck. 2021. Measurement integrity in peer prediction: A peer assessment case study. *arXiv preprint arXiv:2108.05521* (2021).
[4] Thomas M. Cover and Joy A. Thomas. 2001. *Elements of Information Theory*. Wiley, USA. DOI: https://doi.org/10.1002/0471200611
[5] Anirban Dasgupta and Arpita Ghosh. 2013. Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd International World Wide Web Conference*. International World Wide Web Conferences Steering Committee/ACM, 319–330. DOI: https://doi.org/10.1145/2488388.2488417
[6] Xi Alice Gao, Andrew Mao, Yiling Chen, and Ryan Prescott Adams. 2014. Trick or treat: Putting peer prediction to the test. In *Proceedings of the 15th ACM Conference on Economics and Computation*. ACM, 507–524.
[7] Robin Hanson. 2003. Combinatorial information market design. *Inf. Syst. Front.* 5, 1 (2003), 107–119.
[8] Yuqing Kong. 2020. Dominantly truthful multi-task peer prediction with a constant number of tasks. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2398–2411.
[9] Yuqing Kong, Katrina Ligett, and Grant Schoenebeck. 2016. Putting peer prediction under the micro (economic) scope and making truth-telling focal. In *Proceedings of the International Conference on Web and Internet Economics*. Springer, 251–264.
[10] Yuqing Kong and Grant Schoenebeck. 2018. Equilibrium selection in information elicitation without verification via information monotonicity. In *Proceedings of the 9th Innovations in Theoretical Computer Science Conference*.
[11] Yuqing Kong and Grant Schoenebeck. 2018. Water from two rocks: Maximizing the mutual information. In *Proceedings of the ACM Conference on Economics and Computation*. ACM, 177–194.
[12] Yuqing Kong and Grant Schoenebeck. 2019. An information theoretic framework for designing information elicitation mechanisms that reward truth-telling. *ACM Trans. Econ. Computat.* 7, 1 (2019), 2.
[13] Yuqing Kong, Grant Schoenebeck, Fang-Yi Yu, and Biaoshuai Tao. 2020. Information elicitation mechanisms for statistical estimation. In *Proceedings of the 34th AAAI Conference on Arificial Intelligence (AAAI'20)*.
[14] Yang Liu and Yiling Chen. 2017. Machine-learning aided peer prediction. In *Proceedings of the ACM Conference on Economics and Computation*. ACM, 63–80.
[15] N. Miller, P. Resnick, and R. Zeckhauser. 2005. Eliciting informative feedback: The peer-prediction method. *Manag. Sci.* (2005), 1359–1373.
[16] XuanLong Nguyen, Martin J. Wainwright, and Michael I. Jordan. 2010. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Trans. Inf. Theor.* 56, 11 (2010), 5847–5861.

[17] Matthew Parry, A. Philip Dawid, Steffen Lauritzen et al. 2012. Proper local scoring rules. *Ann. Statist.* 40, 1 (2012), 561–592.
[18] Chris Piech, Jonathan Huang, Zhenghao Chen, Chuong Do, Andrew Ng, and Daphne Koller. 2013. Tuned models of peer assessment in MOOCs. *arXiv preprint arXiv:1307.2579* (2013).
[19] Drazen Prelec. 2001. A two-person scoring rule for subjective reports. Massachusetts Institute of Technology working paper.
[20] Drazen Prelec. 2004. A Bayesian truth serum for subjective data. *Science* 306, 5695 (2004), 462–466. https://doi.org/10.1126/science.1102081
[21] Drazen Prelec. 2021. Bilateral Bayesian truth serum: The nxm signals case. Retrieved from SSRN 3908446.
[22] Goran Radanovic and Boi Faltings. 2013. A robust Bayesian truth serum for non-binary signals. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI'13)*. 833–839.
[23] Goran Radanovic and Boi Faltings. 2014. Incentives for truthful information elicitation of continuous signals. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI'14)*. 770–776.
[24] Grant Schoenebeck and Fang-Yi Yu. 2020. Learning and strongly truthful multi-task peer prediction: A variational approach. *arXiv preprint arXiv:2009.14730* (2020).
[25] Grant Schoenebeck and Fang-Yi Yu. 2020. Two strongly truthful mechanisms for three heterogeneous agents answering one question. In *Proceedings of the International Conference on Web and Internet Economics*. Springer, 119–132.
[26] Grant Schoenebeck, Fang-Yi Yu, and Yichi Zhang. 2021. Information elicitation from rowdy crowds. In *Proceedings of the Web Conference*. 3974–3986.
[27] Victor Shnayder, Arpit Agarwal, Rafael Frongillo, and David C. Parkes. 2016. Informed truthfulness in multi-task peer prediction. In *Proceedings of the ACM Conference on Economics and Computation (EC'16)*. ACM, New York, NY, 179–196.
[28] Siddarth Srinivasan and Jamie Morgenstern. 2021. Auctions and prediction markets for scientific peer review. *arXiv preprint arXiv:2109.00923* (2021).
[29] Bo Waggoner and Yiling Chen. 2013. Information elicitation sans verification. In *Proceedings of the 3rd Workshop on Social Computing and User Generated Content (SC'13)*.
[30] Robert L. Winkler. 1969. Scoring rules and the evaluation of probability assessors. *J. Amer. Statist. Assoc.* 64, 327 (1969), 1073–1078.
[31] Jens Witkowski and David C. Parkes. 2011. A robust Bayesian truth serum for small populations. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI'12)*.
[32] Jens Witkowski and David C. Parkes. 2012. Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 964–981.
[33] Peter Zhang and Yiling Chen. 2014. Elicitability and knowledge-free elicitation with peer prediction. In *Proceedings of the International Conference on Autonomous Agents and Multi-agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 245–252.
[34] Shuran Zheng, Fang-Yi Yu, and Yiling Chen. 2021. The limits of multi-task peer prediction. *CoRR* abs/2106.03176 (2021).

Received 31 January 2021; revised 25 September 2022; accepted 30 September 2022