



Two Strongly Truthful Mechanisms for Three Heterogeneous Agents Answering One Question

Grant Schoenebeck¹  and Fang-Yi Yu²  

¹ University of Michigan, Ann Arbor, USA
schoeneb@umich.edu

² Harvard University, Cambridge, USA
fayu@umich.edu

Abstract. 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 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* which the mechanisms need not know (*prior-free*) in the *signal question setting*. Moreover, they only require *three agents*, each of which submits a *signal item report*: one reports her forecast and the others their signals.

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 [11] 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.

Keywords: Peer prediction · Log scoring rule · Prediction market

1 Introduction

Crowd-sourcing relies on eliciting truthful information from agents. Peer prediction is the problem of information elicitation without verification. Incentivizing agents is important so that they not only participate, but provide thoughtful and accurate information. This has a multitude of applications including peer-grading, reviews, and labeling data (for machine learning or research). In the

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 grants 2007887.

single-question setting agents are only asked one question. Our goal is to elicit truthful information from agents with minimal requirements.

For example, say three friends watch a political debate on television. We would like to ask each of them who won the debate and pay them to incentivize truthful answers. This situation will be modeled as each agent receiving some information from the debate about which candidate won. Moreover, prior to the debate, there is a joint prior distribution over the signals of the different agents which is common knowledge among the agents. Thus, one friend’s belief on who won yields some insights about the perceived winners of the other friends.

We will design mechanisms to compensate the agents for their information. We would like our mechanisms to have the following desirable properties:

Strongly Truthful [8]. 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 Sect. 2).¹ 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 that they provide very little information: only one item, either their signal or a prediction. In our case, two agents will 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 Related Work

Single Task Setting. In this setting, each agent receives a single signal from a common prior. Miller et al. [10] introduce the first mechanism for single task signal elicitation that has truth-telling as a strict Bayesian Nash equilibrium and

¹ Kong and Schoenebeck [8] show that it is not possible for truth-telling to pay strictly more than permutation equilibrium in detail-free mechanisms.

does not need verification. However, their mechanism requires full knowledge of the common prior and there exist some equilibria that 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. [3] 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. [5] 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 [11] 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 [12, 13, 17–19] relax the large population requirement of BTS but lose the strongly truthful property. Zhang and Chen [19] 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.

Kong and Schoenebeck [6] 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 symmetric 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.

	Truthful	# Agents	Strongly truthful	General signals
BTS [11]	✓	∞	✓	
Robust BTS [17]	✓	3		
Disagreement [6]	✓	6	✓	
Knowledge-free peer prediction [19]	✓	3		✓
Differential peer prediction	✓	3	✓	✓

Continuous Single Task Setting. Kong et al. [9] shows 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 [13] considers continuous singles. However, it uses a discretization approach which yields exceedingly complex reports. Additionally, it requires homogeneous agents.

In a slightly different setting, Kong and Schoenebeck [7] 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 which can be from a much richer space.

Multi-task Setting. In the multitask setting, introduced in Dasgupta and Ghosh [2], agents are assigned a batch of a priori similar tasks which require each agents' private information to be a binary signal. Several works extend this to multiple-choice questions [2, 4, 8, 14]. The multi-task setting is easier to work in than the single-task setting because the mechanism can better deduce the strategy of any particular agent by comparing reports across questions. However, this setting is substantially more restrictive than the single-question setting of the present paper in that it is important the questions are all similar and gives no guarantees when questions have different priors. An example of when this requirement holds is asking agents to label images as "cat" or "no cat".

1.2 Our Contributions

- We define two Differential Peer Prediction mechanisms (Mechanism 1 and 2) which 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.
- We provide a simple, conceptually motivated proof for the guarantees of Differential Peer Prediction mechanisms. Especially in contrast to the most closely related work [6] our proof is very simple.
- We show special properties of the logarithmic scoring rules (see Techniques below for details). This allows the construction of *target* incentives where an agent is rewarded when its signal is predicted well, and we believe will also be of independent interest.
- We recast the Bayesian Truth Serum mechanism into our framework, showing that it is a *target* incentive mechanism (Sect. 4). This gives added intuition for its guarantees.

1.3 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 [10] does exactly this. To apply this idea to the detail-free setting [17, 19], 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 Sect. 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. In particular, we show that log scoring rules can elicit signals as well as forecasts. This may be of independent interest, and is also the foundation for the results in Sects. 3.2 and 4.

Information Monotonicity. We use information monotonicity, a tool from information theory, to obtain strong truthfulness. Like the present paper, the core of the argument that the Disagreement Mechanism [6] 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 [8].

In Sect. 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 [19], however, it is strongly truthful. Target-Differential Peer Prediction also uses the target incentive techniques above.

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} , \mathcal{C}). They all share a common belief that their signals (a, b, c) are generated from a random variable (A, B, C) which takes values from $\mathcal{A} \times \mathcal{B} \times \mathcal{C}$ with a probability measure P called *common prior*. P describes how agents' private signals relate to each other's.

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 $P_{B,C \mid A}(\cdot \mid a)$ instead to simplify the notion. Similarly Alice's posterior of Bob's signal is denoted by $P_{B \mid A}(\cdot \mid a)$, 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. We pay 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 that 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 expected *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.² We call the strategy profile τ truth-telling if each agent reports truthfully. Moreover, we want to design *detail-free* mechanisms which 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 $\mathcal{A} \times \mathcal{B} \times \mathcal{C}$ —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 τ is a strict Bayesian Nash equilibrium. However, in a truthful mechanism, often non-truth-telling equilibria may yield a higher ex-ante payment for each agent. In this paper, we aim for *strongly truthful mechanisms* [8] which are not only truthful but also ensure the ex-ante agents’ welfare in truth-telling strategy profile τ 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 peer’s reports are not useful when every agent’s signal are independent of each other. Thus, a peer prediction mechanism needs to exploit some interdependence between agents’ signals.

Definition 1 (Zhang and Chen [19]). *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 $P_{C|A,B}(\cdot | a, b) \neq P_{C|A,B}(\cdot | a, b')$. Thus, when Alice with a is making a prediction to Chloe’s signal, Bob’s signal is relevant so that 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 $\{A, B, C\}$.*³

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

² 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 [15].

³ Our definition has some minor differences from Zhang and Chen [19]’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.

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., $\Omega = \{\text{sunny, rainy}\}$. 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 $\text{PS}[\omega, \hat{P}]$. To the end, if Alice believes the distribution of ω to be P , she maximizes her expected score by reporting P truthfully. We call such scoring function proper defined as follow:

Definition 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 [8]. In this paper, we consider a special scoring rule called the **logarithmic scoring rule** [16], defined as

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

where $p : \Omega \rightarrow \mathbb{R}$ is the probability density function of P .

2.3 Information Theory

Peer prediction mechanisms and prediction markets incentivize agents to truthfully report their signals even in the absence of verification. 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 [1].

The *KL-divergence* is a measure of the dissimilarity between 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:

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

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

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

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 1 (Data processing inequality). *If $X \rightarrow Y \rightarrow Z$ form a Markov chain,⁴*

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

By basic algebraic manipulations, Kong and Schoenebeck [8] relate proper scoring rules to mutual information as follows: For two random variables X and Y ,

$$\mathbb{E}_{x,y} [\text{LSR}[y, P(Y | x)] - \text{LSR}[y, P(Y)]] = I(X; Y). \quad (2)$$

We can generalize the mutual information in two ways [8]. 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 Eq. (2): $BMI^f(X, Y) := \mathbb{E}_{x,y} [\text{PS}^f(y, P_{Y|X}(\cdot | x)) - \text{PS}^f(y, P_Y(\cdot))]$. 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 Eq. (2), 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

⁴ 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 .

and Chen [19], which rewards a *source* by how useful her signal is for an expert to predict a target’s report. Their mechanism is only truthful but not strongly truthful. We carefully shift the payment functions and employ Eq. (2) 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 Sect. 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 Sect. 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’s reporting 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 by the sum of the logarithmic scoring rule on those two predictions. And S-DPP pays Bob 0. Chloe’s payment consists of two parts: First, we pay her the prediction score of the improved prediction \hat{Q}^+ . By the definition of proper scoring rule (Definition 2), Chloe will report truthfully to maximize it. For the second part, we subtract Chloe’s payment by 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 randomly permute Bob and Chloe’s roles in the mechanism.

Theorem 2. *If the common prior P is second order stochastic relevant on a finite set with full support, Mechanism 1 is strongly truthful:*

1. *The truth-telling strategy profile τ is a strict Bayesian Nash equilibrium.*
2. *The ex-ante agents’ welfare in the truth-telling strategy profile τ is strictly better than all non-permutation strategy profiles.*

We defer the proof to the full version. Intuitively, because the logarithmic scoring rule is proper, Alice (the expert) will make the 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.

Note that if the agents’ common prior P is symmetric, we can randomize the roles among Alice, Bob, and Chloe to create a symmetric game where

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 $c \in \mathcal{C}$ 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. Randomly set Bob or Chloe as the *target* and the other as the *source*. 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 $\text{LSR}[t, \hat{Q}] + \text{LSR}[t, \hat{Q}^+]$.
 - 6: The payment to the target is 0.
 - 7: The payment to the source is $\text{LSR}[t, \hat{Q}^+] - 3\text{LSR}[t, \hat{Q}]$.
-

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 1).

Our mechanism pays Alice by the sum of log scoring on those two predictions. And the mechanism pays Bob by 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 randomly 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 $c \in \mathcal{C}$ 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. Randomly set Bob or Chloe as the *target* and the other as the *source*. 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 $\text{LSR}[t, \hat{Q}] + \text{LSR}[t, \hat{Q}^+]$.
 - 6: The payment to the target is $\text{LSR}[t, \hat{Q}^+] - \text{LSR}[t, \hat{Q}]$.
 - 7: The payment to the source is $-2\text{LSR}[t, \hat{Q}]$.
-

Theorem 3. *If the common prior P is second order stochastic relevant on a finite set with full support, Mechanism 2 is strongly truthful*

We defer the proof to the full version, 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 1 (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 [11] and Kong and Schoenebeck [8], and extends to non-symmetric prior and finite agent setting. The main idea is that to maximize Bob's expected payment, we show that equivalently Bob wants to maximize a proper scoring rule with prediction $P(C \mid \theta(b))$ on predicting Chloe's report. Therefore, by the property of proper scoring rules, Bob is incentivized to tell the truth. We defer the proof to the full version. With Lemma 1, the rest of the proof is very similar to the proof of Theorem 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 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 $s = w$. We defer the formal mechanism to the full version.

4 Bayesian Truth Serum as a Prediction Market

In this section, we revisit the original Bayesian Truth Serum (BTS) by Prelec [11] 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 to understand BTS.

4.1 Setting of BTS

There are n agents. They all share a common prior P . We call P is *admissible* if it consists of two main elements: states and signals. The *state* T is a random variable in $\{1, \dots, m\}$, $m \geq 2$ 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, \dots, x_n is $\Pr(X_1 = x_1, \dots, X_n = x_n) = \prod_{t \in [m]} P_T(t) \prod_{i \in [n]} P_{X|T}(x_i \mid t)$.

Now we restate the main theorem concerning Bayesian Truth Serum:

Mechanism 3. The original BTS

Require: $\alpha > 1$ **Ensure:** The common prior is admissible

- 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 $Q_{-ij}^{(n)} \in \mathcal{P}(\Omega)$ such that for all $x \in \Omega$

$$Q_{-ij}^{(n)}(x) = \frac{1}{n-2} \sum_{k \neq i, j} \mathbf{1}[\hat{x}_k = x] \quad (3)$$

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

- 3: The prediction score and information score of i are

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

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

Theorem 4 [11]. *For all $\alpha > 1$, if the common prior P is admissible and $n \rightarrow \infty$, Mechanism 3 is strongly truthful.*

4.2 Information Score and Prediction Market

Prelec [11] uses clever algebraic calculation to prove this main results. Kong and Schoenebeck [8] 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 prediction markets to show BTS is a truthful mechanism, and use Mechanism 2 to reproduce BTS.

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 and is well-studied in the previous literature. However, the role of information score is more complicated. Here we provide an interpretation based on Mechanism 2.

We consider $i = 2$ and $j = 1$ in BTS and call them Bob and Alice respectively. We let Chloe be the collection of other agent $\{3, 4, \dots, n\}$. Let's run Mechanism 2 on this information structure. Bob is the target. Alice's initial prediction is $Q = P_{X_2|X_1}(\cdot | x_1)$. When Chloe's signal is x_3, x_4, \dots, x_n , Alice's improved prediction is $Q^+ = P_{X_2|X_{-2}}(\cdot | x_{-2})$ where $x_{-2} = (x_1, x_3, \dots, x_n)$ is the collection of all agents' reports except Bob's. By Lemma 1, Bob is still incentivized to report his private signal x_2 which maximizes the expectation, $\text{LSR}[\hat{x}_2, Q^+] - \text{LSR}[\hat{x}_2, Q]$ that equals to

$$\text{LSR}[\hat{x}_2, P_{X_2|X_{-2}}(\cdot | x_{-2})] - \text{LSR}[\hat{x}_2, P_{X_2|X_1}(\cdot | x_1)]. \quad (4)$$

For the BTS (Mechanism 3), the information score in BTS at truth-telling strategy profile is $\text{LSR}[\hat{x}_i, Q_{-ij}^{(n)}] - \text{LSR}[\hat{x}_i, \hat{Q}_j]$ which equals to

$$\text{LSR}[\hat{x}_2, Q_{-ij}^{(n)}] - \text{LSR}[\hat{x}_2, P_{X_2|X_1}(\cdot | x_1)]. \quad (5)$$

The only difference between (4) and (5) is the first term: $P_{X_2|X_{-2}}(\cdot | x_1, x_3, \dots, x_n)$ and $Q_{-ij}^{(n)}$. Therefore, the original BTS reduces to a special case of Mechanism 2 as $n \rightarrow \infty$, if we can show $\lim_{n \rightarrow \infty} P(X_2 | x_1, x_3, \dots, x_n) = \lim_{n \rightarrow \infty} Q_{-ij}^{(n)}$. Formally,

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

$$Q_{-ij}^{(n)}(w) - P_{X_2|X_{-2}}(w | x_1, x_3, \dots, x_n) \xrightarrow{P_{\mathbf{X}|T}(\cdot|t)} 0 \text{ as } n \rightarrow \infty.$$

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

5 Conclusion

We define two Differential Peer Prediction mechanisms for the single question setting which are strongly-truthful, detail-free, and only require a single item report from three agents. Moreover, the agents need not to be homogeneous and their signals may be continuous. We also show a new property of the logarithmic scoring rules, apply it to make target incentive mechanisms, and show that BTS can be seen as such a mechanism. One future direction is to use this machinery to analyse when BTS retains its strongly truthful guarantee, e.g. for what parameters of finite and/or heterogeneous agents. We define Differential Peer Prediction, a strongly-truthful, detail-free, mechanism for the single question setting that only requires a single item report from three agents. Moreover, the agents need not be homogeneous and their signals may be continuous. We provide a simple, conceptually motivated proof for the guarantees of Differential Peer Prediction, which ties together several themes in the information elicitation literature.

References

1. Cover, T.M.: Elements of Information Theory. Wiley, New York (1999)
2. Dasgupta, A., Ghosh, A.: Crowdsourced judgement elicitation with endogenous proficiency. In: Proceedings of the 22nd International Conference on World Wide Web, pp. 319–330, International World Wide Web Conferences Steering Committee (2013)
3. Gao, X.A., Mao, A., Chen, Y., Adams, R.P.: Trick or treat: putting peer prediction to the test. In: Proceedings of the Fifteenth ACM Conference on Economics and Computation, pp. 507–524. ACM (2014)
4. Kong, Y.: Dominantly truthful multi-task peer prediction with a constant number of tasks. In: Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 2398–2411. SIAM (2020)
5. Kong, Y., Ligett, K., Schoenebeck, G.: Putting peer prediction under the micro(economic)scope and making truth-telling focal. In: Cai, Y., Vetta, A. (eds.) WINE 2016. LNCS, vol. 10123, pp. 251–264. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-54110-4_18

6. Kong, Y., Schoenebeck, G.: Equilibrium selection in information elicitation without verification via information monotonicity. In: 9th Innovations in Theoretical Computer Science Conference (2018)
7. Kong, Y., Schoenebeck, G.: Water from two rocks: maximizing the mutual information. In: Proceedings of the 2018 ACM Conference on Economics and Computation, pp. 177–194. ACM (2018)
8. Kong, Y., Schoenebeck, G.: An information theoretic framework for designing information elicitation mechanisms that reward truth-telling. *ACM Trans. Econ. Comput. (TEAC)* **7**(1), 2 (2019)
9. Kong, Y., Schoenebeck, G., Yu, F.Y., Tao, B.: Information elicitation mechanisms for statistical estimation. In: Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020), February 2020
10. Miller, N., Resnick, P., Zeckhauser, R.: Eliciting informative feedback: the peer-prediction method. *Manage. Sci.* **51**, 1359–1373 (2005)
11. Prelec, D.: A Bayesian truth serum for subjective data. *Science* **306**(5695), 462–466 (2004)
12. Radanovic, G., Faltings, B.: A robust Bayesian truth serum for non-binary signals. In: Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI 2013), pp. 833–839. EPFL-CONF-197486 (2013)
13. Radanovic, G., Faltings, B.: Incentives for truthful information elicitation of continuous signals. In: Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI 2014), pp. 770–776. EPFL-CONF-215878 (2014)
14. Shnayder, V., Agarwal, A., Frongillo, R., Parkes, D.C.: Informed truthfulness in multi-task peer prediction. In: Proceedings of the 2016 ACM Conference on Economics and Computation, EC 2016, pp. 179–196. ACM, New York (2016). ISBN 978-1-4503-3936-0
15. Waggoner, B., Chen, Y.: Information elicitation sans verification. In: Proceedings of the 3rd Workshop on Social Computing and User Generated Content (SC13) (2013)
16. Winkler, R.L.: Scoring rules and the evaluation of probability assessors. *J. Am. Stat. Assoc.* **64**(327), 1073–1078 (1969)
17. Witkowski, J., Parkes, D.C.: A robust Bayesian truth serum for small populations. In: Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI 2012) (2011)
18. Witkowski, J., Parkes, D.C.: Peer prediction without a common prior. In: Proceedings of the 13th ACM Conference on Electronic Commerce, EC 2012, pp. 964–981. Association for Computing Machinery, New York (2012). ISBN 9781450314152
19. Zhang, P., Chen, Y.: Elicitability and knowledge-free elicitation with peer prediction. In: Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems, pp. 245–252. International Foundation for Autonomous Agents and Multiagent Systems (2014)