

A Robust Bayesian Truth Serum for Non-Binary Signals

Goran Radanovic and Boi Faltings

Ecole Polytechnique Federale de Lausanne
(EPFL)

Artificial Intelligence Laboratory
CH-1015 Lausanne, Switzerland
{goran.radanovic, boi.faltings}@epfl.ch

Abstract

Several mechanisms have been proposed for incentivizing truthful reports of a private signals owned by rational agents, among them the peer prediction method and the Bayesian truth serum. The robust Bayesian truth serum (RBTS) for small populations and binary signals is particularly interesting since it does not require a common prior to be known to the mechanism. We further analyze the problem of the common prior not known to the mechanism and give several results regarding the restrictions that need to be placed in order to have an incentive-compatible mechanism. Moreover, we construct a Bayes-Nash incentive-compatible scheme called multi-valued RBTS that generalizes RBTS to operate on both small populations and non-binary signals.

Introduction

We are interested in techniques for collecting information from different sources, such as in crowdsourcing (Surowiecki 2005), opinion polls (Campbell 1996) or measurements from sensor data (Aberer et al. 2010). Besides techniques for filtering and aggregating the information that has been submitted, an important aspect is also to incentivize the agents providing the information to do so with the best possible accuracy. In this paper, we consider game-theoretic techniques for deriving such incentive schemes.

In situations where a group of agents makes forecasts about an event that will eventually become common knowledge, strictly proper scoring rules can be used as a tool for truth elicitation (Savage 1971)(Gneiting and Raftery 2007). In this setting, truth telling is a dominant strategy. Moreover, as it is indicated in (Lambert and Shoham 2009), proper scoring rules can be also used to truthfully elicit averages, medians and modes of an unknown quantities.

However, it is often the case that one wants to obtain opinions about events that will happen in a distant future, or information about events whose outcomes are hard to verify (Prelec 2004)(Faltings, Jurca, and Li 2012). We consider a scenario that can be described by a group of agents that observe a certain phenomenon. An agent i receives a signal $S_i = s_i$, updates her belief $Pr(S_j|S_i = s_i)$ about what another agent j has observed, and reports her observation s_i to the center through an information report x_i (see Figure

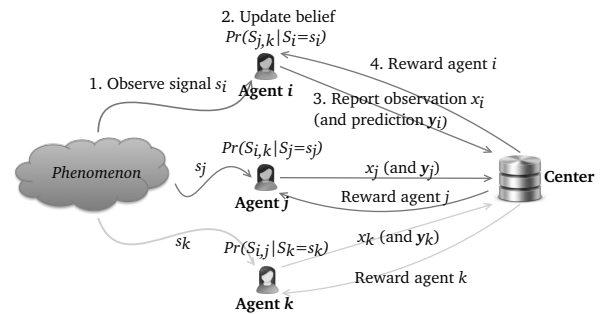


Figure 1: The setting analyzed in this paper¹.

1). Moreover, the center might also ask agent i to submit a prediction y_i about the frequencies of signal values in the population, i.e. her belief $Pr(S_j|S_i = s_i)$. In order to obtain the true observations, the center incentivizes agents to make honest reports. As it cannot verify agents' observations directly, the incentive schemes have to be based on a comparison of reports.

The peer prediction method (Miller, Resnick, and Zeckhauser 2005) is a mechanism based on the comparison of reports that is able to elicit true signal values: the main idea is to extract an agent's posterior belief from her reported value and score it using a proper scoring rule and a report obtained from her peer agent². Due to the usage of proper scoring rules, truth telling is a Bayes-Nash equilibrium. However, the method assumes that:

- the setting has a specific structure, where the state of the word is modeled as a random variable T and observations are modeled as random variables S_i which are conditionally independent given T . The associated probabilities form a prior belief.
- the prior belief is common knowledge, i.e. agents share it and it is known to the mechanism.

Several works investigate modifications of the peer prediction method. Instead of applying proper scoring rules, Jurca and Faltings (Jurca and Faltings 2006)(Jurca and Faltings

¹User icons: David Hopkins (<http://semllabs.co.uk/>). Database icon: Barry Mieny (<http://barrymieny.deviantart.com/>).

²A peer agent is some other agent assigned to the agent that is being scored by the mechanism.

ings 2007) constructed budget minimizing payment schemes using automated mechanism design. They prove that if an agent is scored on the comparison of several reports rather than just one, then the minimum budget required to achieve incentive compatibility decreases. The results also indicate that small deviations of agents' beliefs from the common prior may lead to large increases in payments. If the mechanism can distinguish the time before observations are made from the time after the observations are made, it is possible to exploit the temporal structure to elicit binary signals even when the prior is private and subjective (Witkowski and Parkes 2012a). The key idea is that agents first report their private prior belief about what other agents will observe, then observe a binary signal, and after the observation report their signal values.

The Bayesian truth serum (BTS) (Prelec 2004) assumes a setting similar to the one used in the peer prediction method, but does not require the common prior to be known to the mechanism. Instead, agents are obliged to provide two reports: the information report (their observation) and the prediction report (the prediction of what other agents have observed). BTS is Bayes-Nash incentive-compatible for large populations of agents. Moreover, Prelec and Seung (Prelec and Seung 2006) describe how to extract the true state of the world using the BTS score even when the majority is wrong, and they provide experimental evidences which show that the BTS score performs better than the majority rule in selecting truth. The robust Bayesian truth serum (RBTS) (Witkowski and Parkes 2012b) corrects the main drawback of BTS: its inadequacy to operate on small populations. RBTS is Bayes-Nash incentive-compatible for small populations, but it applies only when binary information is being elicited.

In this paper, we consider incentive schemes that operate both on small populations and non-binary signals for scenarios where agents share a common prior belief not known to the mechanism. We start by showing that a mechanism cannot be incentive-compatible if it elicits only an information report (which corresponds to the agent's observation) and no assumptions are placed on an agent's beliefs³. This motivates the usage of a prediction report, which corresponds to the agent's prediction of frequencies of observed values. In particular, we focus on decomposable payment schemes, meaning that the incentives are the sum of an information score and a prediction score, where the information score is independent of the prediction report and the prediction score is independent of the information report. We show that even with two reports there is no decomposable incentive scheme that does not require additional assumptions about an agent's beliefs. However, as positive results we show two different sufficient conditions on an agent's beliefs that enable incentive-compatible elicitation schemes. For the *self-dominant* condition, we show an incentive-compatible scheme that uses only an information report. For the weaker *self-predicting* condition, we

³Jurca and Faltings (Jurca and Faltings 2011) prove a similar result: in general, without knowing prior beliefs, one cannot construct a mechanism that is Bayes-Nash incentive-compatible.

construct a decomposable payment scheme for non-binary values that is Bayes-Nash incentive-compatible, and generalizes the original RBTS scheme (Witkowski and Parkes 2012b).

We believe that the points made in this paper will contribute to the characterization of incentive-compatible payment schemes, especially when it comes to providing sufficient and necessary conditions for payment schemes to be Bayes-Nash incentive-compatible.

The Setting

Instead of defining a setting using a specific underlying model, as it is done in (Prelec 2004)(Miller, Resnick, and Zeckhauser 2005)(Witkowski and Parkes 2012b), one can describe it in more general terms: with prior beliefs about observations and posterior beliefs about what other agents have observed. We use the following setting:

- There are $n \geq 2$ agents who report their observations to the center. Based on the quality of their observations, the agents receive a score. The agents are risk-neutral and seek to maximize their expected score. The scoring function is denoted by τ and it depends on reports made by all the agents.
- The agents observe a same phenomenon, and their observations are modeled as random variables S_i that take values from $\{0, \dots, m - 1\}$.
- The agents share a common prior belief about the possible observations, i.e. for a possible observation s and two agents i and j we have $Pr(S_i = s) = Pr(S_j = s)$.
- Once an agent observes her value, she updates her posterior belief about what other agents have observed. The updating process is the same for every agent, i.e. for any three agents i, j and k , if agents i and j observe the same value $s_i = s_j = s$, then $Pr(S_k = s_k | S_i = s) = Pr(S_k = s_k | S_j = s)$. Since the updating process is universal, we can simplify our notation by denoting a posterior belief $Pr(S_k = s_k | S_i = s_i)$ by $\mathbf{p}_{s_i}(s_k)$.
- The prior and posterior beliefs are fully mixed, meaning that $0 < Pr(S_i = s_i) < 1 \wedge 0 < Pr(S_j = s_j | S_i = s_i) < 1, \forall s_i, s_j \in \{0, \dots, m - 1\}$. Moreover, the random variables are stochastically relevant: the distribution of S_j conditional on S_i is different for different realizations of S_i (Miller, Resnick, and Zeckhauser 2005), i.e. $\forall s_i \neq \tilde{s}_i, \exists s_j : Pr(S_j = s_j | S_i = s_i) \neq Pr(S_j = s_j | S_i = \tilde{s}_i)$.

Beliefs about Peers

The mechanisms we are investigating are based on comparing reports that different agents make about the same phenomenon. Clearly, there are cases where such comparisons don't make sense. For example, if agents all interpret the phenomenon differently, or use different scales for measurement, their reports cannot be compared directly. Furthermore, what matters is not the true situation, but what agents *believe* about other agents: to provide the right incentives, it is sufficient that they believe other agents to be comparable to themselves, even if in reality they may be very different.

In this paper, we consider two different assumptions about the strength of these beliefs. The strongest assumption we consider is the *self-dominant* assumption, where an agent i believes that the value x she observes is also the most likely value observed by another agent k :

$$Pr(S_k = \tilde{x} | S_i = x) < Pr(S_k = x | S_i = x), \forall x \neq \tilde{x} \quad (1)$$

An example where this assumption holds is when all agents are in the same restaurant and are asked to count the number of customers. A general class of cases where the self-dominant assumption holds is when agents believe that they observe the exact same signal only perturbed by unbiased noise.

As many settings do not satisfy this assumption, we introduce a weaker assumption, the *self-predicting* assumption. Here an agent i believes that another agent k is most likely to observe a certain value x when she herself also observes this value:

$$Pr(S_k = x | S_i = \tilde{x}) < Pr(S_k = x | S_i = x), \forall \tilde{x} \neq x \quad (2)$$

An example where this assumption holds is when agents are customers of a particular restaurant and are asked to report the speed of the restaurant's service based on how quickly they receive their meal. In that case, agent i models the prior probability (her belief) of getting a meal within certain amount of time using a categorical distribution (e.g. the service can be slow, normal or fast). Once her meal is served, she updates her belief using Bayesian inference: the parameters of the categorical distribution are updated using it's conjugate prior (Dirichlet distribution) so the updating process satisfies the expression (2). The posterior distribution represents agent i 's belief about what other agents have reported. A general class of cases where the self-predicting assumption holds is when agents believe that they observe different samples drawn from the same random distribution, provided the agents use Bayesian updating and the phenomenon follows one of many natural probability distributions including categorical or multinomial distributions. In the example above, agents know that serving times are drawn from the same distribution so if they obtain unusually slow service they might still believe that other agents are getting faster service. This makes the self-predicting assumption weaker than the self-dominant assumption, and widely applicable to practical scenarios. Note that, when signals S_i are binary, the self-predicting assumption holds whenever the self-dominant assumption is satisfied. Moreover, the self predicting assumption is satisfied in the original (binary) RBTS setting (Witkowski and Parkes 2012b).

An example where neither of the two assumptions hold is when agents report observations about different or even anti-correlated phenomena, for example the number of customers in the restaurant they are in, which is not necessarily the same for all agents. It may also fail to hold if other agents' measurements are believed to be inaccurate and biased in some way. In this case, none of the methods we discuss in this paper can be applied.

Mechanisms Based on One Report

We begin our analysis with the mechanisms that elicit only the observed values. In other words, agents are asked to

provide only the privately owned signal: the corresponding report is called *information report*. Since the center only knows the reported values, one can expect that there does not exist a mechanism which would elicit honest reports in the general case. Before we formally show this, let us take a closer look at peer payment schemes.

Definition 1. A payment scheme is 1-peer if the score depends only on the reports of the agent that is being scored and her peer agent.

Let us denote reported values by x_i . As it is stated in the setting section, a scoring function depends on all the reported values, i.e. $\tau(x_1, \dots, x_n, x_{agent})$ where we put the report of the agent that is being scored at the last position. On the other hand, 1-peer payment schemes represent restricted version of general scoring function as they have a form $\tau(x_{peer}, x_{agent})$. Nevertheless, their structure allows easier theoretical analysis.

Lemma 1. If it is possible to construct a payment scheme that requires only the information report and is Bayes-Nash incentive-compatible when no restrictions are placed on the updating process, then it is possible to construct a Bayes-Nash incentive-compatible 1-peer payment scheme.

Proof. Let τ be a Bayes-Nash incentive-compatible payment scheme. The expected score of an agent i for reporting x_i , provided that all other agents are honest (i.e. $x_k = s_k$ where s_k is a value observed by agent k), is equal to:

$$\begin{aligned} & \sum_{s_j, s_{k_1}, \dots} Pr(S_j = s_j, S_{k_1} = s_{k_1}, \dots | S_i = s_i) \tau(s_j, s_{k_1}, \dots, x_i) \\ &= \sum_{s_j} (Pr(S_j = s_j | S_i = s_i) \cdot \\ & \cdot \sum_{s_{k_1}, \dots} Pr(S_{k_1} = s_{k_1}, \dots | S_j = s_j, S_i = s_i) \tau(s_j, s_{k_1}, \dots, x_i)) \\ &= \sum_{s_j} Pr(S_j = s_j | S_i = s_i) \tilde{\tau}(s_j, x_i) \end{aligned} \quad (3)$$

Notice that $\tilde{\tau}(s_j, x_i)$ depends on $Pr(S_{k_1} = s_{k_1}, S_{k_2} = s_{k_2}, \dots | S_j = s_j, S_i = s_i)$. However, the original scheme is incentive-compatible in the general case, so it must be incentive-compatible when the updating process keeps $Pr(S_{k_1} = s_{k_1}, S_{k_2} = s_{k_2}, \dots | S_j = s_j, S_i = s_i)$ fixed, but alters $Pr(S_j = s_j | S_i = s_i)$. This implies the existence of a 1-peer payment scheme that is incentive-compatible in general because $\tilde{\tau}$ is incentive-compatible for arbitrary beliefs $Pr(S_j = s_j | S_i = s_i)$. \square

Lemma 1 is useful if one wants to prove impossibility results: if incentive compatibility is required for the general case, it is enough to consider 1-peer payment schemes⁴.

Theorem 1. If only an information report is being elicited and no restrictions are placed on the updating process, there exists no payment scheme that is Bayes-Nash incentive-compatible. Moreover, even if the self-predicting assumption

⁴However, Lemma 1 does not imply that it is enough to observe 1-peer payment schemes in order to achieve incentive compatibility when certain restrictions are put on the updating process.

holds, there exists no payment scheme that is Bayes-Nash incentive-compatible.

Proof. Let us assume that there exists a Bayes-Nash incentive-compatible payment scheme that works on a general set of beliefs. Due to Lemma 1, we restrict our attention to 1-peer payment schemes. Let $\mathbf{p}_0 = Pr(S_j|S_i = 0), \dots, \mathbf{p}_{m-1} = Pr(S_j|S_i = m-1)$ be some arbitrary distribution functions. Since the payment scheme should be incentive-compatible for arbitrary distribution functions, it should also be incentive-compatible when $\mathbf{p}_1(1) > \mathbf{p}_0(1) > \mathbf{p}_0(0) > \mathbf{p}_1(0)$ and $\mathbf{p}_k(k) > \mathbf{p}_l(k), k, l \in \{0, 1, \dots, m-1\}, l \neq k$. Notice that these distributions satisfy the self-predicting assumption. Due to incentive compatibility, we have:

$$\mathbf{p}_0(0)[\tau(0, 0) - \tau(0, 1)] + \mathbf{p}_0(1)[\tau(1, 0) - \tau(1, 1)] + \sum_{i>1} \mathbf{p}_0(i)[\tau(i, 0) - \tau(i, 1)] > 0 \quad (4)$$

$$\mathbf{p}_1(0)[\tau(0, 1) - \tau(0, 0)] + \mathbf{p}_1(1)[\tau(1, 1) - \tau(1, 0)] + \sum_{i>1} \mathbf{p}_1(i)[\tau(i, 1) - \tau(i, 0)] > 0 \quad (5)$$

To simplify the notation, let $\Delta_i = \tau(i, 0) - \tau(i, 1)$. The above expressions are then equal to:

$$\mathbf{p}_0(0)\Delta_0 + \mathbf{p}_0(1)\Delta_1 + \sum_{i>1} \mathbf{p}_0(i)\Delta_i > 0 \quad (6)$$

$$-\mathbf{p}_1(0)\Delta_0 - \mathbf{p}_1(1)\Delta_1 - \sum_{i>1} \mathbf{p}_1(i)\Delta_i > 0 \quad (7)$$

Let us consider new set of beliefs $\mathbf{p}'_0, \mathbf{p}'_1, \mathbf{p}_2, \dots, \mathbf{p}_{m-1}$, where \mathbf{p}'_0 and \mathbf{p}'_1 are equal to: $\mathbf{p}'_0(0) = \mathbf{p}'_1(1), \mathbf{p}'_0(1) = \mathbf{p}_1(0), \mathbf{p}'_1(0) = \mathbf{p}_0(1), \mathbf{p}'_1(1) = \mathbf{p}_0(0), \mathbf{p}'_0(k) = \mathbf{p}_1(k)$ and $\mathbf{p}'_1(k) = \mathbf{p}_0(k)$ for $k \neq 0, 1$. Notice that the new set of beliefs satisfy the self-predicting assumption. Since the incentive compatibility also has to hold for the new posterior beliefs, we have:

$$\mathbf{p}_1(1)\Delta_0 + \mathbf{p}_1(0)\Delta_1 + \sum_{i>1} \mathbf{p}_1(i)\Delta_i > 0 \quad (8)$$

$$-\mathbf{p}_0(1)\Delta_0 - \mathbf{p}_0(0)\Delta_1 - \sum_{i>1} \mathbf{p}_0(i)\Delta_i > 0 \quad (9)$$

The last 4 inequalities give us:

$$(\mathbf{p}_0(0) - \mathbf{p}_0(1))(\Delta_0 - \Delta_1) > 0 \quad (10)$$

$$(\mathbf{p}_1(1) - \mathbf{p}_1(0))(\Delta_0 - \Delta_1) > 0 \quad (11)$$

Because we set $\mathbf{p}_1(1) > \mathbf{p}_0(1) > \mathbf{p}_0(0) > \mathbf{p}_1(0)$, it cannot be that both (10) and (11) are satisfied. That is, we have a contradiction. \square

The significance of Theorem 1 is that it motivates the usage of mechanisms that require an additional report. Namely, as we show in the next section, under the same set of assumptions (i.e. the self-predicting assumption), there exists a mechanism that requires 2 reports and is Bayes-Nash incentive-compatible. To make our analysis complete, we also give a sufficient condition for existence of an incentive-compatible payment scheme that requires only one report.

Proposition 1. *If the self-dominant assumption holds, then there exists an incentive-compatible payment scheme that requires only the information report.*

Proof. Let us analyze the following 1-peer payment scheme:

$$\tau(x_{peer}, x_{agent}) = \mathbb{1}_{x_{peer}=x_{agent}} \quad (12)$$

where $\mathbb{1}_{x_{peer}=x_{agent}}$ is an indicator variable (equal to 1 when $x_{peer} = x_{agent}$, otherwise it is equal to 0). An agent who aims to maximize her reward will choose to report:

$$\begin{aligned} \operatorname{argmax}_{x_{agent}} E(\tau(x_1, \dots, x_n, x_{agent})) &= \\ &= \operatorname{argmax}_{x_{agent}} Pr(S_{peer} = x_{agent} | S_{agent} = s_{agent}) = \\ &= s_{agent} \end{aligned} \quad (13)$$

where s_{agent} is the agent's true observation. \square

Mechanisms Based on Two Reports

Theorem 1 motivates us to introduce an additional report, the *prediction report*. The prediction report represents an agent's prediction of frequencies of observed values, i.e. it is a probability distribution function over possible observations of other agents. Provided that an agent is honest, her prediction report is her belief about what other agents have reported. Let us denote the information report by x_i and the prediction report by \mathbf{y}_i . In general, a scoring function depends on all the reported values, i.e. $\tau(x_1, \mathbf{y}_1, \dots, x_{m-1}, \mathbf{y}_{m-1}, x_{agent}, \mathbf{y}_{agent})$. However, we restrict our attention to *decomposable* payment schemes that reward information reports and prediction reports independently.

Definition 2. *A payment scheme is a decomposable payment scheme if an agent's total score is calculated as the sum of her information score and her prediction score, where the information score does not depend on the agent's prediction report and the prediction score does not depend on the agent's information report. More precisely:*

$$\begin{aligned} \tau_{total}(x_1, \mathbf{y}_1, \dots, x_{m-1}, \mathbf{y}_{m-1}, x_{agent}, \mathbf{y}_{agent}) &= \\ &= \underbrace{\tau(x_1, \mathbf{y}_1, \dots, x_{m-1}, \mathbf{y}_{m-1}, x_{agent})}_{\text{information score}} + \\ &+ \underbrace{\tau(x_1, \mathbf{y}_1, \dots, x_{m-1}, \mathbf{y}_{m-1}, \mathbf{y}_{agent})}_{\text{prediction score}} \end{aligned} \quad (14)$$

Having the decomposable structure where an agent's information score is independent of her prediction report and her prediction score is independent of her information report simplifies the analysis of the incentives (scores) as they do not influence each other. Notice that RBTS is an example of decomposable payment scheme. For the limit case, when number of agents approaches infinity, BTS converges to a decomposable payment scheme since a single agent does not have large impact on the frequencies of signal reports nor on the average of prediction reports. We keep the notion of 1-peer payment schemes from the previous section, i.e. 1-peer payment schemes represent a restricted version of general scoring functions and have a form $\tau(x_{peer}, \mathbf{y}_{peer}, x_{agent}, \mathbf{y}_{agent})$.

The structure of decomposable schemes allows us to analyze information and prediction scores separately. The prediction score elicits an agent's beliefs about what other agents have reported. Since the outcome is known to the mechanism (the mechanism knows what other agents have reported), truthful elicitation of the prediction report can be achieved using proper scoring rules.

Lemma 2. *There exists a decomposable scheme that is Bayes-Nash incentive-compatible if and only if there exists an information score that is Bayes-Nash incentive-compatible.*

Proof. Let us associate every agent with her peer agent, and let R be a proper scoring rule. Consider the following scoring function:

$$\tau(x_1, \mathbf{y}_1, \dots, x_n, \mathbf{y}_n, \mathbf{y}_{agent}) = R(\mathbf{y}_{agent}, x_{peer}) \quad (15)$$

Provided that the peer agent is honest ($x_{peer} = s_{peer}$), the best strategy for the agent that is being scored is to report $\mathbf{y}_{agent} = \mathbf{p}_{agent}$ (her posterior belief about what other agents have observed). Therefore, the payment scheme is incentive-compatible. Since there exists a payment scheme that elicits honest prediction reports regardless of the updating assumptions, the existence of a decomposable scheme that is Bayes-Nash incentive-compatible depends only on the existence of an information score that is Bayes-Nash incentive-compatible. Hence, we proved the statement. \square

Thanks to Lemma 2, we can restrict our analysis to the information score in order to see whether there exists a decomposable payment scheme that is incentive-compatible. Moreover, because we can assume that agents are incentivized to provide honest prediction reports, our notation for the information score can be changed to $\tau(x_1, \mathbf{p}_1, \dots, x_n, \mathbf{p}_n, x_{agent})$. To simplify our analysis, we prove the statement analogous to the one made in Lemma 1.

Lemma 3. *If it is possible to construct a decomposable payment scheme that is Bayes-Nash incentive-compatible when no restrictions are placed on the updating process, then it is possible to construct a Bayes-Nash incentive-compatible 1-peer decomposable payment scheme.*

Proof. Due to Lemma 2, we restrict our analysis to the information score. Let τ be a Bayes-Nash incentive-compatible payment scheme. The expected score of an agent i for reporting x_i , provided that all other agents are honest (i.e. $x_k = s_k$ and $\mathbf{y}_k = \mathbf{p}_{s_k}$, where s_k is a value observed by agent k), is equal to:

$$\begin{aligned} & \sum_{s_j, s_{k_1}, \dots} Pr(S_j = s_j, S_{k_1} = s_{k_1}, \dots | S_i = s_i) \cdot \\ & \cdot \tau(s_j, \mathbf{p}_{s_j}, s_{k_1}, \mathbf{p}_{s_{k_1}}, \dots, x_i) \\ & = \sum_{s_j} (Pr(S_j = s_j | S_i = 0)) \cdot \\ & \cdot \sum_{s_{k_1}, \dots} Pr(S_{k_1} = s_{k_1}, \dots | S_j = s_j, S_i = s_i) \tau(s_j, \mathbf{p}_{s_j}, \dots, x_i) \\ & = \sum_{s_j} Pr(S_j = s_j | S_i = 0) \tilde{\tau}(s_j, \mathbf{p}_{s_j}, x_i) \end{aligned} \quad (16)$$

Notice that $\tilde{\tau}(s_j, \mathbf{p}_{s_j}, x_i)$ depends on $Pr(S_{k_1} = s_{k_1}, S_{k_2} = s_{k_2}, \dots | S_j = s_j, S_i = s_i)$. However, the original scheme is incentive-compatible in the general case, so it must be incentive-compatible when the updating process keeps $Pr(S_{k_1} = s_{k_1}, S_{k_2} = s_{k_2}, \dots | S_j = s_j, S_i = s_i)$ fixed, but alters $Pr(S_j = s_j | S_i = s_i)$. This implies the existence of a 1-peer payment scheme that is incentive-compatible in general because $\tilde{\tau}$ is incentive-compatible for arbitrary beliefs $Pr(S_j = s_j | S_i = s_i)$. \square

As in the previous section, we first provide the impossibility result: if no restrictions are placed on the updating process, even though it is the same for all agents, we cannot guarantee incentive compatibility in general.

Theorem 2. *If no restrictions are placed on the updating process, there exists no decomposable payment scheme that is Bayes-Nash incentive-compatible when observations S_i take more than two values ($m \geq 3$).*

Proof. Let us assume that there exists a Bayes-Nash incentive-compatible payment scheme that works on general set of beliefs. Because of Lemma 2, we restrict our attention to 1-peer payment schemes for information reports. Let $\mathbf{p}_0 = Pr(S_j | S_i = 0)$, ..., $\mathbf{p}_{m-1} = Pr(S_j | S_i = m-1)$ be some arbitrary distribution functions. Due to the incentive compatibility, we have:

$$\begin{aligned} & \mathbf{p}_0(0)(\tau(\mathbf{p}_0, 0, 0) - \tau(\mathbf{p}_0, 0, 1)) + \\ & + \mathbf{p}_0(1)(\tau(\mathbf{p}_1, 1, 0) - \tau(\mathbf{p}_1, 1, 1)) + \\ & + \sum_{i>1} \mathbf{p}_0(i)(\tau(\mathbf{p}_i, i, 0) - \tau(\mathbf{p}_i, i, 1)) > 0 \end{aligned} \quad (17)$$

$$\begin{aligned} & \mathbf{p}_1(0)(\tau(\mathbf{p}_0, 0, 1) - \tau(\mathbf{p}_0, 0, 0)) + \\ & + \mathbf{p}_1(1)(\tau(\mathbf{p}_1, 1, 1) - \tau(\mathbf{p}_1, 1, 0)) + \\ & + \sum_{i>1} \mathbf{p}_1(i)(\tau(\mathbf{p}_i, i, 1) - \tau(\mathbf{p}_i, i, 0)) > 0 \end{aligned} \quad (18)$$

This gives us:

$$\begin{aligned} & (\mathbf{p}_0(0) - \mathbf{p}_1(0))\Delta_0(\mathbf{p}_0) + (\mathbf{p}_0(1) - \mathbf{p}_1(1))\Delta_1(\mathbf{p}_1) + \\ & + \sum_{i>1} (\mathbf{p}_0(i) - \mathbf{p}_1(i))\Delta_i(\mathbf{p}_i) > 0 \end{aligned} \quad (19)$$

where $\Delta_i(\mathbf{p}_i) = (\tau(\mathbf{p}_i, i, 0) - \tau(\mathbf{p}_i, i, 1))$. Since the mechanism should be incentive-compatible for arbitrary distribution functions, it should also be incentive-compatible for the following two cases:

1. When $\mathbf{p}_0(0) = \mathbf{p}_1(0) - \epsilon$, $\mathbf{p}_0(k) = \mathbf{p}_1(k)$ for $k \neq 0, m-1$, and $\mathbf{p}_0(m-1) = \mathbf{p}_1(m-1) + \epsilon$, where $\epsilon > 0$. We denote this distribution by \mathbf{p}_0^- .
2. When $\mathbf{p}_0(0) = \mathbf{p}_1(0) + \epsilon$, $\mathbf{p}_0(k) = \mathbf{p}_1(k)$ for $k \neq 0, m-1$, and $\mathbf{p}_0(m-1) = \mathbf{p}_1(m-1) - \epsilon$, where $\epsilon > 0$. We denote this distribution by \mathbf{p}_0^+ .

Note that due to the *stochastic relevance* assumption, we cannot put $\mathbf{p}_0 = \mathbf{p}_1$. From (19) and the fact that $\mathbf{p}_0(m-1) - \mathbf{p}_1(m-1) = (1 - \sum_{i<m-1} \mathbf{p}_0(i)) - (1 - \sum_{i<m-1} \mathbf{p}_1(i)) = \sum_{i<m-1} (\mathbf{p}_1(i) - \mathbf{p}_0(i))$, we obtain:

$$-\epsilon \cdot (\Delta_0(\mathbf{p}_0^-) - \Delta_{m-1}(\mathbf{p}_{m-1})) > 0 \quad (20)$$

$$\epsilon \cdot (\Delta_0(\mathbf{p}_0^+) - \Delta_{m-1}(\mathbf{p}_{m-1})) > 0 \quad (21)$$

In other words:

$$\Delta_0(\mathbf{p}_0^-) < \Delta_{m-1}(\mathbf{p}_{m-1}) \quad (22)$$

$$\Delta_0(\mathbf{p}_0^+) > \Delta_{m-1}(\mathbf{p}_{m-1}) \quad (23)$$

Let us consider a new \mathbf{p}_1^{++} equal to: $\mathbf{p}_1^{++}(0) = \mathbf{p}_1(0) + 2\epsilon$, $\mathbf{p}_1^{++}(k) = \mathbf{p}_1(k)$ for $k \neq 0, m-1$, and $\mathbf{p}_1^{++}(m-1) = \mathbf{p}_1(m-1) - 2\epsilon$. By applying the previous steps on \mathbf{p}_1^{++} , we obtain:

$$\Delta_0(\mathbf{p}_0^+) = \Delta_0(\mathbf{p}_0^{++}) < \Delta_{m-1}(\mathbf{p}_{m-1}) \quad (24)$$

$$\Delta_0(\mathbf{p}_0^{++}) > \Delta_{m-1}(\mathbf{p}_{m-1}) \quad (25)$$

Hence we have a contradiction (expressions (23) and (24)). Therefore, there exists no decomposable payment scheme that is incentive-compatible for arbitrary updating process. \square

The impossibility result obtained by Theorem 2 forces us to introduce additional assumptions on the updating process so that we can construct incentive-compatible payment scheme. Before giving a sufficient condition for existence of an incentive-compatible payment scheme, we describe a novel payment scheme: the multi-valued RBTS.

Multi-valued RBTS. The multi-valued RBTS has two steps:

1. Each agent i is asked to provide two reports:
 - *Information report* x_i which represents agent i 's reported signal.
 - *Prediction report* \mathbf{y}_i which represents agent i 's prediction about the frequencies of signal values in the overall population.
2. Each agent i is linked with her *peer* agent $j = i + 1 \pmod{n}$ and is rewarded with a score:

$$\underbrace{\frac{1}{\mathbf{y}_j(x_i)} \cdot \mathbb{1}_{x_j=x_i}}_{\text{information score}} + \underbrace{R(\mathbf{y}_i, x_j)}_{\text{prediction score}} \quad (26)$$

where $\mathbb{1}_{x_j=x_i}$ is the indicator variable and R is a strictly proper scoring rule.

The multi-valued RBTS has the following property.

Theorem 3. *If the self-predicting assumption is satisfied, the multi-valued RBTS is Bayes-Nash incentive-compatible.*

Proof. Let us rewrite the self-predicting assumption as $\mathbf{p}_{\tilde{x}_i}(x_i) = Pr(S_k = x_i | S_i = \tilde{x}_i) < Pr(S_k = x_i | S_i = x_i) = \mathbf{p}_{x_i}(x_i), \forall x_i \neq \tilde{x}_i$. By Lemma 2 we know that proper scoring rules can be used as the prediction score. Therefore, it is enough to examine the properties of the information score.

Let us examine the behavior of a rational *agent*. Since we want to show that the truth telling is a Bayes-Nash equilibrium, we can assume that her peer is honest. Thus, the peers's prediction report is equal to her posterior belief about what other agents have reported, which means that $\mathbf{y}_{s_{peer}}(x_{agent}) = \mathbf{p}_{s_{peer}}(x_{agent})$. Also, honesty of the peer

implies that $s_{peer} = x_{peer}$. This means that the expected value of the agent's information score is equal to:

$$\begin{aligned} E(\tau(x_1, \mathbf{y}_1, \dots, x_n, \mathbf{y}_n, x_{agent})) &= \\ &= \frac{\mathbf{p}_{s_{agent}}(x_{agent})}{\mathbf{y}_{s_{peer}}(x_{agent})} \Big|_{s_{peer}=x_{agent}} \\ &= \frac{\mathbf{p}_{s_{agent}}(x_{agent})}{\mathbf{p}_{x_{agent}}(x_{agent})} \end{aligned} \quad (27)$$

By taking into account the assumption $\mathbf{p}_{\tilde{x}_i}(x_i) < \mathbf{p}_{x_i}(x_i), \forall x_i \neq \tilde{x}_i$, we get that the expectation of the information score is maximized for:

$$\begin{aligned} \operatorname{argmax}_{x_{agent}} E(\tau(x_1, \mathbf{y}_1, \dots, x_n, \mathbf{y}_n, x_{agent})) &= \\ = \operatorname{argmax}_{x_{agent}} \frac{\mathbf{p}_{s_{agent}}(x_{agent})}{\mathbf{p}_{x_{agent}}(x_{agent})} &= s_{agent} \end{aligned} \quad (28)$$

Therefore, the maximal value of the information score is achieved when the agent reports her true observation. \square

Note that the direct consequence of Theorem 3 is that the self-predicting assumption represents a sufficient condition for existence of an incentive-compatible decomposable payment scheme.

Example 1. To illustrate how the described protocol elicits honest reports, consider two agents A and B that are observing the same phenomenon and, hence, share a common prior and have the same updating process. Suppose that they measure the event that takes values from $\{0, 1, 2\}$ and that their belief system is the one shown in Table 1. Notice that the self-predicting assumption is satisfied.

Table 1: Agents' prior and posterior beliefs

S_i	0	1	2
$Pr(S_i)$	0.1	0.5	0.4
$Pr(S_j = 0 S_i)$	0.3	0.2	0.2
$Pr(S_j = 1 S_i)$	0.4	0.6	0.3
$Pr(S_j = 2 S_i)$	0.3	0.2	0.5

The center interested in their observations wants to elicit their true measurements so is applying the multi-valued RBTS with the quadratic scoring rule:

$$\begin{aligned} R(\mathbf{y}_{agent}, x_{peer}) &= \frac{1}{2} + \mathbf{y}_{agent}(x_{peer}) - \\ &- \frac{1}{2} \sum_x \mathbf{y}_{agent}^2(x) \end{aligned} \quad (29)$$

Suppose that agent A measures 0; she updates her belief about what her peer has observed. Since agents do not collude, she assumes that agent B is honest and, hence, reports her true observations (i.e. $x_B = s_B$ and $\mathbf{y}_B = Pr(S_A | S_B = s_B)$). Agent A calculates that the expected prediction score for reporting $\mathbf{y}_A = Pr(S_B | S_A = 0)$ is equal to:

$$\begin{aligned} E(\tau_{pred}(\mathbf{y}_A)) &= Pr(S_B = 0 | S_A = 0)R(\mathbf{y}_A, 0) + \\ &+ Pr(S_B = 1 | S_A = 0)R(\mathbf{y}_A, 1) + \\ &+ Pr(S_B = 2 | S_A = 0)R(\mathbf{y}_A, 2) = \\ &= 0.3 \cdot 0.63 + 0.4 \cdot 0.73 + 0.3 \cdot 0.63 = 0.67 \end{aligned} \quad (30)$$

Since R is a proper scoring rule, she knows that the maximum of the expectation is achieved when $y_A = Pr(S_B|S_A = 0)$. For example, if y_A were equal to $y_A(0) = 0.5$, $y_A(1) = 0.2$ and $y_A(2) = 0.3$, the agent would calculate $E(\tau_{pred}(y_A)) = 0.63$, which is less than for the truthful report. Therefore, agent A reports $Pr(S_B|S_A = 0)$ as her prediction report.

In order to choose the best signal value to report, agent A calculates what her expected payoff is for different values of her information report. For example, $E(\tau_{info}(x_A = 0))$ is calculated as:

$$\begin{aligned} E(\tau_{info}(x_A = 0)) &= \frac{Pr(S_B = 0|S_A = 0)}{y_B(0)} \\ &= \frac{Pr(S_B = 0|S_A = 0)}{Pr(S_A = 0|S_B = 0)} = 1 \end{aligned} \quad (31)$$

Similarly, she obtains $E(\tau_{info}(x_A = 1)) = 0.67$ and $E(\tau_{info}(x_A = 2)) = 0.6$. Therefore, the best option for agent A is to report 0, which is her true observation.

Finally, note that:

Corollary 1. *The multi-valued RBTS is Bayes-Nash incentive-compatible in the original RBTS setting.*

Proof. The original RBTS setting is a specific case of the setting introduced in this paper. Moreover, the original RBTS setting satisfies the self-predicting assumption (see Lemma 5 in (Witkowski and Parkes 2012b)). Therefore, the multi-valued RBTS is Bayes-Nash incentive-compatible in the original RBTS setting. \square

Conclusion

This paper provides a new insight into incentive mechanisms for truthful information elicitation based on information and prediction reports as proposed in the (robust) Bayesian truth serum. The focus of the paper is put on decomposable incentive schemes where the information score is independent of the prediction report and the prediction score is independent of the information report. We constructed a new payment scheme for non-binary variables that is Bayes-Nash incentive-compatible provided the self-predicting assumption is satisfied. The introduced payment scheme can be also applied for the original (binary) RBTS setting and hence is more general than the original RBTS mechanism.

In order to allow incentive-compatible schemes, we introduced two different assumptions, the self-dominant and self-predicting assumptions. Future work should include further characterization of settings for which incentive compatibility can be achieved. For example, it is an open question whether the self-predicting assumption is also a necessary condition for the existence of incentive-compatible mechanisms, or whether there are other sufficient conditions that would also admit such mechanisms.

Acknowledgments

The work reported in this paper was supported by Nanotera.ch as part of the Openseense project. We thank the anonymous reviewers for useful comments and feedback.

References

- Aberer, K.; Sathe, S.; Chakraborty, D.; Martinoli, A.; Barrenetxea, G.; Faltings, B.; and Thiele, L. 2010. Openseense: Open community driven sensing of environment. In *ACM SIGSPATIAL International Workshop on GeoStreaming (IWGS)*, 39–42.
- Campbell, J. E. 1996. Polls and votes: The trial-heat presidential election forecasting model, certainty, and political campaigns. *American Politics Quarterly* 24(4):408–433.
- Faltings, B.; Jurca, R.; and Li, J. J. 2012. Eliciting truthful measurements from a community of sensors. In *Proceedings of the Third International Conference on the Internet of Things*, 51–18.
- Gneiting, T., and Raftery, A. E. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102:359378.
- Jurca, R., and Faltings, B. 2006. Minimum payments that reward honest reputation feedback. In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC'06)*, 190–199.
- Jurca, R., and Faltings, B. 2007. Robust incentive-compatible feedback payments. In *Agent-Mediated Electronic Commerce*, volume LNAI 4452, 204–218. Springer-Verlag.
- Jurca, R., and Faltings, B. 2011. Incentives for answering hypothetical questions. In *Workshop on Social Computing and User Generated Content (EC-11)*.
- Lambert, N., and Shoham, Y. 2009. Eliciting truthful answers to multiple-choice questions. In *Proceedings of the tenth ACM conference on Electronic Commerce*, 109–118.
- Miller, N.; Resnick, P.; and Zeckhauser, R. 2005. Eliciting informative feedback: The peer-prediction method. *Management Science* 51:13591373.
- Prelec, D., and Seung, S. 2006. An algorithm that finds truth even if most people are wrong. Working paper.
- Prelec, D. 2004. A bayesian truth serum for subjective data. *Science* 34(5695):462–466.
- Savage, L. J. 1971. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* 66(336):783–801.
- Surowiecki, J. 2005. *The Wisdom of Crowds*. Anchor.
- Witkowski, J., and Parkes, D. C. 2012a. Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC'12)*, 964–981.
- Witkowski, J., and Parkes, D. C. 2012b. A robust bayesian truth serum for small populations. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI'12)*.