

Incentives for Truthful Information Elicitation of Continuous 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

We consider settings where a collective intelligence is formed by aggregating information contributed from many independent agents, such as product reviews, community sensing, or opinion polls. We propose a novel mechanism that elicits both private signals and beliefs. The mechanism extends the previous versions of the Bayesian Truth Serum (the original BTS, the RBTS, and the multi-valued BTS), by allowing small populations and non-binary private signals, while not requiring additional assumptions on the belief updating process. For priors that are sufficiently smooth, such as Gaussians, the mechanism allows signals to be continuous.

Introduction

We consider settings where a collective intelligence is formed by aggregating information contributed from many independent agents, such as product reviews, community sensing or opinion polls. To encourage participation and avoid selection bias, agents should be rewarded for the information they provide. It is important that the rewards provide incentives for relevant and truthful information and discourage random or malicious reports.

Often the collected information cannot be easily verified because it requires a large amount of effort to do so, or the data is entirely private and subjective. This means that scoring techniques based on direct verification, such as strictly proper scoring rules (Savage 1971; Gneiting and Raftery 2007; Lambert and Shoham 2009) or prediction markets (Hanson 2003; Chen and Pennock 2007), cannot be used to elicit effort nor private information. In other words, incentive schemes have to rely solely on the reported data.

The recent impossibility results (Waggoner and Chen 2013; Radanovic and Faltings 2013; Jurca and Faltings 2011) indicate that in order to allow incentive compatibility when direct verification is not applicable, one cannot have arbitrarily structured setting. Our setting is structured according to the one introduced in (Prelec 2004), where agents share a common prior belief, not known to the mechanism. The model is depicted in Figure 1. An agent i receives a signal $S_i = s_i$; updates her belief $Pr(S_j|S_i = s_i)$ regarding what another agent j has observed; and makes the informa-

tion report x_i about her observation s_i and the prediction report y_i about her posterior belief $Pr(S_j|S_i = s_i)$.




Figure 1: The setting analyzed in this paper.

We want to obtain a mechanism for elicitation of agents' private signals that is incentive compatible without requiring additional restrictions on the setting. Moreover, the mechanism should extend to continuous domains. We note that many information aggregation procedures operate on continuous signals; a typical example is community sensing, such as air quality monitoring (Aberer et al. 2010).

Our approach is closely related to the following three techniques: the Bayesian Truth Serum (BTS) (Prelec 2004), the Robust Bayesian Truth Serum (RBTS) (Witkowski and Parkes 2012b) and the Multi-valued Robust Bayesian Truth Serum (Multi-valued RBTS) (Radanovic and Faltings 2013). In contrast to the Peer Prediction Method (Miller, Resnick, and Zeckhauser 2005), BTS introduces the possibility of scoring agents in peer prediction manner without knowing their common prior. However, it assumes a large population of agents. Using the *shadowing* mechanism, RBTS provides bounded incentives for elicitation of binary signals, while requiring no more than three agents. Multi-valued RBTS extends RBTS to non-binary settings at the expense of potentially large payments. Recently, (Witkowski 2014) has provided a generalization of the shadowing method, and hence RBTS, to the non-binary case. Notice that both generalized RBTS and multi-valued RBTS put additional restrictions on the belief updating process when non-binary information is elicited. That is, they require the belief change (absolute and relative respectively) from prior

to posterior to be largest for the observed value. None of the aforementioned BTS mechanisms allows continuous signals.

Contributions. We propose a novel mechanism called the Divergence-based Bayesian Truth Serum. It allows non-binary signals and is incentive-compatible even for small populations, without requiring additional restrictions on the BTS setting. Moreover, the divergence-based BTS is guaranteed to be individually rational with bounded payments, and, for discrete signals, it permits differences in agents’ prior beliefs. Furthermore, it is the first BTS mechanism that can be applied to continuous domains.

Related Work

Peer prediction (Miller, Resnick, and Zeckhauser 2005) elicits signal values, both discrete and continuous, using proper scoring rules. The main drawback is that it requires the mechanism to know agents’ priors. There are several variations of the peer prediction mechanism, such as budget minimizing payment schemes using automated mechanism design (Jurca and Faltings 2006; 2007).

(Witkowski and Parkes 2012a) describe a peer prediction mechanism that does not require agents to have a common prior. However, it requires that there is a temporal structure in the setting and that agents’ private signals are binary. In the mechanism, an agent first reports her private prior belief regarding the possible observations, and then makes an observation and reports it to the mechanism.

Collective revelation (Goel, Reeves, and Pennock 2009) elicits individual predictions and aggregate estimates. It has a setting similar to the peer prediction mechanisms, with the common prior known to the mechanism, and agents that may make multiple observations, generated from a distribution of a particular form (e.g. Bernoulli distribution).

Truthful surveys (Lambert and Shoham 2008) is an elicitation mechanism developed for truthfully sampling opinions. It does not assume that agents have common prior beliefs, but the mechanism provides only weak incentives, so the agents are indifferent between truthful reporting and lying.

The helpful reporting mechanism of (Jurca and Faltings 2011) provides incentives for opinion polls. When a public distribution (announced by the mechanism) is not close to agents’ common prior, the mechanism is not incentive compatible. Instead, agents are incentivized to report values that push the public distribution towards their common prior. Once the public signal is close enough to the prior, the mechanism becomes incentive compatible.

The output agreement mechanism of (Waggoner and Chen 2013) rewards agents based on how close their reports are according to a metric distance. It elicits common-knowledge (e.g. mean), rather than private signals, but it does not require strong assumptions on the structure of agents’ beliefs.

The effort elicitation mechanism of (Dasgupta and Ghosh 2013) is developed for crowdsourcing settings. It applies for elicitation of binary signals and relies on the fact that agents solve multiple tasks that are a priori equivalent. Maximal effort and truthful reporting result in a maximal reward, while

a small number of trusted workers can be used to eliminate the low effort equilibria.

Three interesting results relate to the BTS mechanism mentioned in the introduction: (Prelec and Seung 2006) describe how to use the BTS mechanism in order to obtain the ground truth even when the majority is wrong, while (Shaw, Chen, and Horton 2011; Weaver and Prelec 2013) demonstrate that the BTS mechanism rewards truthful responses and has a positive effect in quality control.

The Setting

We model agents’ reasoning similar to (Prelec 2004; Miller, Resnick, and Zeckhauser 2005; Witkowski and Parkes 2012b; Radanovic and Faltings 2013), where agents have a common prior belief and use the same belief updating procedure. Our setting has the following structure:

- There are $n \geq 2$ risk-neutral agents who make observations of a certain phenomenon, and report their observations to an entity called *center*. In return, the center rewards agents based on the quality of their reports, and the quality is estimated by comparing the reports of different agents. That is, the scoring function τ does not only depend on the report of the agent that is being rewarded, but also on the reports of other agents called *peers*.
- The agents have the common probabilistic belief that consists of the state of the phenomenon T , which takes values from Ω , and agents’ observations S_i (private signals), which take values from Σ . Since in this paper we deal with both continuous and discrete sets Ω and Σ , we use both cumulative probability distributions Pr and probability density distributions p . For simplicity, we describe the setting using only Pr .
- The signals S_i are conditionally independent given T , meaning that their signals are generated according to some distribution function dependent on the state T . In probabilistic terms, this means that for two different agents i and j , $Pr(S_i, S_j|T) = Pr(S_i|T)Pr(S_j|T)$.
- Once an agent i measures the phenomenon, she updates her beliefs $Pr(S_j|S_i)$ regarding what another agent j has observed. The belief updating procedure follows Bayes’ rule. Note that our mechanism, in its general form, applies also for the setting from (Radanovic and Faltings 2013), where agents might have some other belief updating procedure.
- After the observation, an agent i submits two reports:
 - *Information report* x_i , which represents agent i ’s reported signal.
 - *Prediction report* y_i , which represents agent i ’s prediction regarding the frequencies of signal values in the overall population. When agents are honest, this report corresponds to agent i ’s posterior belief $Pr(S_j|S_i)$.
- We assume a fully mixed prior, that is $\forall s_i \in \Sigma, \forall t \in \Omega$:

$$0 < Pr(T = t) < 1 \wedge 0 < Pr(S_i = s_i|T = t) < 1$$

Using Bayes’ rule, $Pr(S_i|T)$ and $Pr(T)$, it follows that the posterior $Pr(S_j|S_i)$ is also fully mixed.

- Finally, the signals S_i 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. $\forall s_i \in \Sigma, \forall \tilde{s}_i \in \Sigma \setminus \{s_i\}, \exists \epsilon > 0$:

$$D(\Pr(S_j|S_i = s_i) || \Pr(S_j|S_i = \tilde{s}_i)) > \epsilon$$

where $D(\cdot)$ is a divergence function (e.g. KL divergence).

Background

Strictly proper scoring rules elicit agents' beliefs regarding the event whose outcome eventually becomes a common knowledge. When the event realizes, the center rewards agent i for her prediction y_i with a strictly proper scoring rule $R(y_i, t)$, where t is the realization of the event (the ground truth). Every strictly proper scoring rule is associated with the divergence function $D(\cdot)$, that measures the difference between the expected scores when a true belief is reported and when some other prediction is reported. Examples of strictly proper scoring rules are the logarithmic scoring rule and the quadratic scoring rule, associated with the KL divergence and the Euclidean distance respectively (see (Gneiting and Raftery 2007) for more details).

When the ground truth is not available to the center, the scoring functions τ have to be based on a comparison of the reports. The primary goal of the center is to elicit private signals S_i , so agents' reports should at least contain their information reports. The information reports alone, however, do not allow incentive compatibility (see Theorem 1 in (Radanovic and Faltings 2013)).

Bayesian Truth Serums thwart the issue by introducing an additional report: the prediction report. Accordingly, the BTS scores are composed of the *information score* and the *prediction score*:

$$\tau_{total} = \underbrace{\tau_{info}}_{\text{information score}} + \underbrace{\tau_{pred}}_{\text{prediction score}}$$

If an agent's information score does not depend on her prediction report and her prediction score does not depend on her information report, we call the payment scheme τ_{total} *decomposable* (Radanovic and Faltings 2013). Both RBTS and Multi-valued RBTS are decomposable, while the original BTS is decomposable in the limit case when the number of agents goes to infinity, which coincides with the requirement for incentive compatibility. It is not surprising that these mechanisms require some additional restrictions of the BTS setting to be incentive compatible - with no specific conditions on the agents' belief updating process, it is not possible to construct an incentive compatible decomposable payment scheme (see Theorem 2 in (Radanovic and Faltings 2013)). This leads us to mechanisms that do not have decomposable structure.

Divergence-based Bayesian Truth Serum

As mentioned in the previous section, decomposable payment schemes cannot achieve incentive compatibility when no restrictions are placed on the belief updating process. Thus, we investigate a broader class of mechanisms where

the information score is not independent of the prediction report. Indeed, this is exactly what makes our mechanism robust: it successfully copes with small populations, non-binary signal values and arbitrary (but common among agents) belief updating procedures.

Divergence-based BTS. Consider observations S_i taking values from a countable discrete set, in particular, $\Sigma \subseteq \mathbb{N}_0 = \{0, 1, \dots\}$. The divergence-based BTS has two steps:

1. Each agent i is asked to provide her information report x_i and her prediction report y_i .
2. Each agent i is linked with a randomly chosen *peer* agent j and is rewarded with a score:

$$\underbrace{-\mathbb{1}_{x_j=x_i \wedge D(y_i||y_j) > \Theta}}_{\text{information score}} + \underbrace{R(y_i, x_j)}_{\text{prediction score}} \quad (1)$$

where $\mathbb{1}_{x_j=x_i}$ is the indicator variable, R is a strictly proper scoring rule, $D(\cdot)$ is the divergence associated to a strictly proper scoring rule, and Θ is a parameter of the mechanism.

The prediction score part of the mechanism rewards agent i if her prediction report y_i matches the distribution of information reports x_j submitted by other agents. Contrary to all earlier versions of BTS, the information score *penalizes* the agent if its information report agrees with its peer while its prediction report does not. Disagreement between prediction reports is characterized by the condition that the divergence between the reports is larger than a threshold Θ . The intuition behind this penalty is that honest agents will not have such an inconsistency with their peers. The following theorem shows the condition on the belief structure and the choice of Θ that make this intuition true.

Theorem 1. *Let $\Sigma \subseteq \mathbb{N}_0$. The divergence-based BTS is strictly Bayes-Nash incentive-compatible when $n \geq 2$ and agents' posteriors satisfy $\forall x \in \Sigma, \forall \tilde{x} \in \Sigma \setminus \{x\}$:*

$$D(\Pr(S_j|S_i = x) || \Pr(S_i|S_j = x)) \leq \Theta < D(\Pr(S_j|S_i = x) || \Pr(S_i|S_j = \tilde{x})) \quad (2)$$

Proof. Consider an agent i who observes s_i and believes that her peer agent is honest, and suppose condition (2) is satisfied.

Due to the properties of the strictly proper scoring rules, agent i 's prediction score is in expectation maximized when she reports $y_i = \Pr(S_j|S_i = s_i)$, and because stochastic relevance holds, this is a strict optimum.

If agent i 's prediction report is $y_i = \Pr(S_j|S_i = s_i)$, then from condition (2) we conclude that the maximum of her information score is achieved when she reports $x_i = s_i$, and is equal to 0.

Since the optimal value of the information score is equal to 0 and the prediction score is maximized when $y_i = \Pr(S_j|S_i = s_i)$, it follows that $x_i = s_i$ and $y_i = \Pr(S_j|S_i = s_i)$ is agent i 's best response.

We still need to prove that this is the strictly optimal response. Since $y_i = \Pr(S_j|S_i = s_i)$ is strictly optimal response for the prediction score, and $x_i = s_i$ achieves the optimal value of the information score, it is enough to show

that agent i 's information score is negative in expectation for $x_i \neq s_i$ and $\mathbf{y}_i = Pr(S_j|S_i = s_i)$. Due to condition (2) and the fully mixed posteriors, the expected score for reporting $x_i \neq s_i$ and $\mathbf{y}_i = Pr(S_j|S_i = s_i)$ is:

$$Pr(S_j = x_i|S_i = s_i)(-\mathbb{1}_{D(Pr(S_j|S_i=s_i)||Pr(S_i|S_j=x_i))>\Theta}) \\ = -Pr(S_j = x_i|S_i = s_i) < 0$$

Putting it all together, truthful reporting is a strict Bayes-Nash equilibrium of the divergence-based BTS. \square

The direct consequence of Theorem 1 is that the divergence-based BTS applies to all former BTS settings, with no specific restrictions.

Corollary 1. *The divergence-based BTS is strictly Bayes-Nash incentive-compatible in the BTS, the RBTS and the Multi-valued RBTS settings for $\Theta = 0$ and $n \geq 2$.*

Proof. These settings are special cases of the setting introduced in this paper that satisfy $0 = D(Pr(S_j|S_i = x)||Pr(S_i|S_j = x)) < D(Pr(S_j|S_i = x)||Pr(S_i|S_j = \tilde{x}))$. So Theorem 1 can be applied with $\Theta = 0$ \square

A convenient feature of the divergence-based BTS is that it allows a population of agents to have different prior beliefs, as long as the agents' posteriors are more similar when they observe the same value than when their observations are different. This is exactly what condition (2) states, and is very realistic if agents indeed have similar observations. Notice that we formalized similarities between posteriors of two different agents using parameter Θ . Although we have assumed that the center knows Θ , it is possible to make the divergence-based BTS a non-parametric method.

Non-parametric divergence-based BTS. To make the divergence-based BTS a non-parametric method, we change its second step. In addition to a peer agent j , the modified method also uses a randomly chosen reference agent k , and the overall score for agent i becomes:

$$\underbrace{-\mathbb{1}_{x_j=x_i \wedge x_k \neq x_i \wedge D(\mathbf{y}_i||\mathbf{y}_j) > D(\mathbf{y}_i||\mathbf{y}_k)}}_{\text{information score}} + \underbrace{R(\mathbf{y}_i, x_j)}_{\text{prediction score}} \quad (3)$$

Theorem 2. *Let $\Sigma \subseteq \mathbb{N}_0$. The non-parametric divergence-based BTS is strictly Bayes-Nash incentive-compatible when $n \geq 3$ and agents' posteriors satisfy condition (2) of Theorem 1.*

Proof (Sketch). For agent i who observes s_i and honest agents j and k , $D(\mathbf{y}_i||\mathbf{y}_j) < D(\mathbf{y}_i||\mathbf{y}_k)$ holds whenever agent i reports $x_i = x_j \neq x_k$ and $\mathbf{y}_i = Pr(S_j|S_i = x_i)$. In that case the information score achieves the optimal value (the optimum is also achieved when $x_i \neq x_j$ or $x_i = x_k$). Because the prediction score is a strictly proper scoring rule, agent i 's best response is to report $x_i = s_i$ and $\mathbf{y}_i = Pr(S_j|S_i = s_i)$. The strictness can be proven the same way as in Theorem 1. \square

Individual Rationality

Appropriate scaling of the divergence-based BTS with a bounded scoring rule R leads to ex-post individual rationality and bounded payments.

Proposition 1. *Consider $\Sigma \subseteq \mathbb{N}_0$ and the divergence-based BTS scheme τ with the quadratic scoring rule defined by the equation:*

$$R(\mathbf{y}_i, x_j) = 2 - (1 - \mathbf{y}_i(x_j))^2 - \sum_{x \in \Sigma \setminus \{x_j\}} \mathbf{y}_i(x)^2$$

Then, $\alpha(\tau + 1)$, where $\alpha > 0$, produces scores in $[0, 3\alpha]$.

Proof. The minimum of the quadratic score is 0 and this can happen when $\mathbf{y}_i(x \neq x_j) = 1$. The maximum is achieved when $\mathbf{y}_i(x_j) = 1$ and is equal to 2. On the other hand, the minimum and the maximum of the information score is -1 and 0, respectively. Therefore $\alpha(\tau + 1)$ produces values in $[0, 3\alpha]$. \square

In addition to the guarantee that participation in the mechanism is individually rational, the proposition tells us that the center can scale α so that the total payoff to all agents does not exceed a fixed budget. For example, when the center has budget β , it can use $\alpha = \frac{\beta}{3n}$ to ensure that the budget is not exceeded. The same holds for the continuous BTS (described in the next section) with the continuous version of the quadratic scoring rule.

Divergence-based BTS for Continuous Signals

All of the BTS mechanisms are based either on matching information reports of an agent and her peer, or the shadowing method, i.e. appropriately shifting a peer's posterior in the direction of the agent's information report. These approaches do not directly extend to continuous domain. One way of dealing with this issue is to discretize the continuous domain and consider all the values from a certain interval to be the same when matching is done.

Continuous BTS. Consider observations S_i taking continuous values, in particular, $\Sigma = \mathbb{R}$. The continuous BTS has the following steps:

1. Each agent i is asked to provide the information report x_i and the prediction report \mathbf{y}_i , as in the divergence based BTS.
2. For each agent i , the mechanism samples a number δ_i from a uniform distribution, i.e. $\delta_i = rand((0, 1))$. The continuous domain $\Sigma = \mathbb{R}$ is then uniformly discretized with the discretization interval of a size δ_i and the constraint that value x_i is in the middle of the interval it belongs to. We denote the interval of a value x_i by Δ_x^i . The constraint can then be written as $x_i = \frac{\max \Delta_x^i - \min \Delta_x^i}{2}$.
3. Finally, an agent i is scored using a modified version of the divergence-based BTS score:

$$\underbrace{-\mathbb{1}_{x_j \in \Delta_x^i \wedge D(\mathbf{y}_i||\mathbf{y}_j) > \delta_i \Theta}}_{\text{information score}} + \underbrace{R(\mathbf{y}_i, x_j)}_{\text{prediction score}} \quad (4)$$

The parameter Θ reflects how close the posteriors of two similar signals are. When agents are fully confident Θ should be big, because posteriors of two similar signals can be significantly different. On the other hand, when agents make mistakes the posteriors of two similar signals are close to each other, making the lower bound on Θ smaller. This

fact can be used to set the appropriate value of Θ . For example, in community sensing, the center can assume that every sensor is worse than some accurate sensor, so the center can adjust Θ according to the specifics of the accurate sensor. For a Θ parameter that never underestimates the divergence $D(\cdot)$ of agents posteriors that observe similar values, the continuous BTS is incentive compatible:

Theorem 3. Consider $\Sigma = \mathbb{R}$, $n \geq 2$ and suppose $\Theta \in (0, \infty)$ is such that $\forall x_i \in \Sigma, \forall \delta_i \in (0, 1), \forall x_j \in \Delta_x^i$:

$$D(p(S_j|S_i = x_i)||p(S_i|S_j = x_j)) \leq \delta_i \Theta \quad (5)$$

Then the continuous BTS is strictly Bayes-Nash incentive-compatible.¹

Proof. Suppose agent i observes s_i and believes that her peer j is honest. Whenever agent i reports x_i and $\mathbf{y}_i = p(S_j|S_i = x_i)$, her information score is equal to 0, because (5) holds. The prediction score is a strictly proper scoring rule, so in expectation the optimal choice for the prediction report is agent i 's posterior $\mathbf{y}_i = p(S_j|S_i = s_i)$ - this is a strict optimum due to stochastic relevance. Therefore, reporting s_i and $p(S_j|S_i = s_i)$ is a Bayes-Nash equilibrium. As it was the case with Theorem 1, in order to show that truthful reporting is a strict equilibrium, we need to prove that, for any information report other than s_i , agent i 's information score is in expectation negative when her prediction report is $\mathbf{y}_i = p(S_j|S_i = s_i)$.

Let $x_i \neq s_i$. Since δ_i can be arbitrarily small, consider δ_{i1} such that $s_i \notin \Delta_x^{i1}$. From stochastic relevance, we know that there exists ϵ such that:

$$\forall x_j \in \Delta_x^{i1} : D(p(S_j|S_i = s_i)||p(S_i|S_j = x_j)) > \epsilon \quad (6)$$

Now, consider $\delta_{i2} = \min(\delta_{i1}, \frac{\epsilon}{\Theta})$. Since $\Delta_x^{i2} \subseteq \Delta_x^{i1}$, it follows from (6) that:

$$\forall x_j \in \Delta_x^{i2} : D(p(S_j|S_i = s_i)||p(S_i|S_j = x_j)) > \epsilon \geq \delta_{i2} \Theta$$

Moreover, $Pr(x_j \in \Delta_x^{i2}|S_i = s_i) = \int_{x_j \in \Delta_x^{i2}} p(S_j = x_j|S_i = s_i) dx_j > 0$ due to the fully mixed posteriors. So, for any $x_i \neq s_i$, the expected information score of agent i who reports x_i and $\mathbf{y}_i = p(S_j|S_i = s_i)$ is strictly negative.

Therefore, reporting s_i and $p(S_j|S_i = s_i)$ is a strict Bayes-Nash equilibrium. \square

It remains to see how to set the parameter Θ . Consider $D(p(S_j|S_i = x_i)||p(S_i|S_j = x_j))$ as a function of x_j for a fixed $x_i = 2$. Condition (5) simply states that one can find a coefficient c such that $c|x_j - 2| \geq D(p(S_j|S_i = 2)||p(S_i|S_j = x_j))$ for $x_j \in (1.5, 2.5)$. As shown in Figure 2, this corresponds to the divergence being bounded by two lines. More formally:

Proposition 2. Consider $\Sigma = \mathbb{R}$. If $\forall x_i \in \mathbb{R}$, $D(p(S_j|S_i = x_i)||p(S_i|S_j = x_j))$ is a continuously differentiable and bounded function of $x_j \in (x_i - 1/2, x_i + 1/2) \setminus \{x_i\}$, then:

$$\Theta \geq \max_{x_i} \max_{x_j \in (x_i - 1/2, x_i + 1/2) \setminus \{x_i\}} \left| \frac{\partial D(p(S_j|S_i = x_i)||p(S_i|S_j = x_j))}{\partial x_j} \right|$$

satisfies condition (5) of Theorem 3.

¹Notice that $\Theta < \infty$.




Figure 2: The divergence of posteriors as a function of a peer's report.

Proof (Sketch). The statement follows from the fact that function $f(x_j) = \Theta|x_j - x_i| - D(p(S_j|S_i = x_i)||p(S_i|S_j = x_j))$ is equal to 0 when $x_j = 0$, and increases as $|x_j - x_i|$ increases. \square

Proposition 2 applies for many divergence functions $D(\cdot)$, like KL divergence or Euclidean distance, as long as the common prior is sufficiently smooth. For example, if $\Omega = \mathbb{R}$, $p(S_i = x_i|T = t)$ and $p(T = t)$ are continuously differentiable functions of t , $p(S_i = x_i|T = t)$ is a continuously differentiable function of x_i , and $D(\cdot)$ is the KL divergence, then one can use Proposition 2 to obtain a lower bound on Θ .

The continuous BTS is a parametric mechanism, so the center needs to set the parameter Θ . Notice that the only restriction for incentive compatibility is that the center sets Θ big enough. However, there is a tradeoff between the value of Θ and the expected value of margin difference of the information score between truthful and non-truthful reporting. That is, the larger Θ is, the smaller the expected punishment is for an agent who deviates from truthful reporting.

If the divergence function $D(p(S_j|S_i = x_i)||p(S_i|S_j = x_j))$ increases as $|x_j - x_i|$ increases, it is possible to make the continuous BTS a non-parametric method.

Non-parametric continuous BTS. To make the continuous BTS a parameter-free method, we introduce a randomly selected reference agent k and change the score to:

$$\underbrace{-\mathbb{1}_{x_j \in \Delta_x^i \wedge x_k \notin \Delta_x^i \wedge D(\mathbf{y}_i||\mathbf{y}_j) > D(\mathbf{y}_i||\mathbf{y}_k)}}_{\text{information score}} + \underbrace{R(\mathbf{y}_i, x_j)}_{\text{prediction score}} \quad (7)$$

Proposition 3. Let $\Sigma = \mathbb{R}$. Suppose that $\forall x_i, x_j, x_k \in \Sigma$:

$$\begin{aligned} |x_j - x_i| < |x_k - x_i| &\implies \\ D(p(S_j|S_i = x_i)||p(S_i|S_j = x_j)) &< D(p(S_j|S_i = x_i)||p(S_i|S_k = x_k)) \end{aligned}$$

Then the non-parametric continuous BTS is strictly Bayes-Nash incentive compatible for $n \geq 3$.

Proof (Sketch). For agent i who observes s_i and honest agents j and k , $D(\mathbf{y}_i||\mathbf{y}_j) < D(\mathbf{y}_i||\mathbf{y}_k)$ holds whenever $x_j \in \Delta_x^i$, $x_k \notin \Delta_x^i$ and agent i reports x_i and $\mathbf{y}_i = p(S_j|S_i = x_i)$. In that case (or if $x_j \notin \Delta_x^i$ or $x_k \in \Delta_x^i$) the information score achieves the optimal value. Because the prediction score is a strictly proper scoring rule, agent i 's best response is to report $x_i = s_i$ and $\mathbf{y}_i = p(S_j|S_i = s_i)$. The strictness can be proven as in Theorem 3. \square

Gaussian Prior

Consider $\Omega = \Sigma = \mathbb{R}$ and agents whose belief system is based on the Gaussian distribution:

- The signal values S_i are generated by a Gaussian $p(S_i) \sim \mathcal{N}(\mu_T, \sigma)$, where σ is fixed, while μ_T defines the state T and is distributed according to the Normal distribution $p(\mu_T) \sim \mathcal{N}(\mu_0, \sigma_0)$.
- An agent i uses Bayesian updating to obtain her belief $p(S_j|S_i)$ regarding what an agent j has observed.

That is, agents' belief system is composed of four parameters $\{\mu_T, \mu_0, \sigma, \sigma_0\}$ (that define the Gaussian prior) and the Bayesian updating rule.

Now, suppose that agent i observes signal $S_i = s_i$. From the Bayesian updating of Gaussian distributions (Bishop 2006), it follows that agent i 's posterior belief regarding agent j 's observations is a Gaussian $p(S_j|S_x = s_i) \sim \mathcal{N}(\mu_{s_i}, \sigma_{s_i})$ with the parameters equal to:

$$\mu_{s_i} = \frac{\frac{\mu_0}{\sigma_0^2} + \frac{s_i}{\sigma^2}}{\frac{1}{\sigma_0^2} + \frac{1}{\sigma^2}}, \sigma_{s_i}^2 = \frac{1}{\frac{1}{\sigma_0^2} + \frac{1}{\sigma^2}} + \sigma^2 \quad (8)$$

The KL divergence of two normal distributions $\mathcal{N}(\mu_1, \sigma_1)$ and $\mathcal{N}(\mu_2, \sigma_2)$ is equal to (Ihara 1993):

$$\log \frac{\sigma_2}{\sigma_1} + \frac{\sigma_1^2}{2\sigma_2^2} + \frac{(\mu_1 - \mu_2)^2}{2\sigma_2^2} - \frac{1}{2} \quad (9)$$

From the expressions (8) and (9) it follows that the KL divergence of agent i 's and agent j 's posteriors is:

$$D(p(S_j|S_i = x_i) || p(S_i|S_j = x_j)) = k(\sigma, \sigma_0) \cdot \frac{(x_j - x_i)^2}{2\sigma^2}$$

where $k(\sigma, \sigma_0) = \frac{\sigma_0^2}{\sigma_0^2 + \sigma^2} \frac{\sigma_0^2}{2\sigma_0^2 + \sigma^2}$. Using Proposition 2, we obtain Θ that satisfies the conditions of Theorem 3:

$$\begin{aligned} \Theta &\geq \max_{x_i} \max_{x_j \in (x_i - 1/2, x_i + 1/2) \setminus \{x_i\}} \left| k(\sigma, \sigma_0) \cdot \frac{2(x_j - x_i)}{2\sigma^2} \right| \\ &\geq k(\sigma, \sigma_0) \cdot \frac{1}{2\sigma^2} \end{aligned} \quad (10)$$

The center does not need to know parameters σ and σ_0^2 : it is sufficient to overestimate (10). We often have that $\sigma_0 \gg \sigma$, and hence $k(\sigma, \sigma_0) \approx \frac{1}{2}$. In that case, the center only needs to underestimate the value of σ . For example, if agents are sensors with accuracy below a certain threshold, the center can infer the minimal value of σ .

The KL divergence of the Gaussian posteriors satisfies the conditions of Proposition 3, which means that one can also use the non-parametric continuous BTS. The Gaussian prior has also another convenient property. Namely, the BTS mechanisms ask agents to report their observations and their posterior beliefs. For an agent i , reporting her posterior belief in the Gaussian model comes down to reporting two parameters μ_{s_i} and σ_{s_i} , where s_i is the agent's observation. So the whole report consists of only three scalar values. One might wonder if it is possible to relax the common prior assumption, considering that the divergence-based BTS can cope with a heterogeneous population. If continuous signals are allowed and natural distributions are used for prior beliefs (e.g. Gaussians), then it is not possible to achieve incentive compatibility unless agents have a common prior.

Proposition 4. *Let $\Omega = \Sigma = \mathbb{R}$ and suppose two agents i and j use the Gaussian priors that might be different. Then no mechanism τ based on the information and prediction reports is strictly Bayes-Nash incentive compatible.*

Proof. Let us assume the opposite, i.e. there exists a mechanism τ that incentivizes agents to reveal their private signals, regardless of their prior. Consider a population consisting of two agents i and j with the priors $\{\mu_T, \mu_0^i, \sigma, \sigma_0\}$ and $\{\mu_T, \mu_0^j, \sigma, \sigma_0\}$. Suppose they observe signals s_i and s_j respectively, and that agent j believes agent i is honest. Let us denote the expected payoff of agent j by τ_j . Now, consider $\mu_0^i = \mu_0^j$. Incentive compatibility of τ implies:

$$\forall s'_j \neq s_j : \tau_j(s_j, \{\mu_{s_j}, \sigma_{s_j}\}) > \tau_j(s'_j, \{\mu_{s_j}, \sigma_{s_j}\}) \quad (11)$$

Consider now $\mu_0^j \neq \mu_0^i$ and $s'_j = \frac{\sigma_0^2}{\sigma_0^2}(\mu_0^i - \mu_0^j) + s_j$. From expression (8), we know that the agent j 's posterior remains the same, i.e. $\mu_{s_j} = \mu_{s'_j}$ and $\sigma_{s_j} = \sigma_{s'_j}$. However, from (11) it follows that the best response of agent j is to report s_j , not s'_j . That is, τ cannot incentivize both agent j , who has the same prior as agent i , and agent j , who has different prior than agent i , to report honestly. \square

Conclusion

This paper explores information elicitation mechanisms where the mechanism designer does not have access to the ground truth nor to the participants' beliefs. We constructed a new payment scheme that operates in the BTS settings and applies to small populations of agents and non-binary signals. When discrete signals are elicited, the scheme permits differences in agents' priors, as long as the agents' posteriors are more similar when they observe the same value than when their observations are different.

As many settings require elicitation of continuous variables, we also derive a mechanism that can be applied for continuous domains. The mechanism depends on a parameter, but, when agents' common prior is smooth, the only condition for incentive compatibility is that the parameter is sufficiently large. We also investigated under what conditions it is possible to obtain a non-parametric version of our method, and showed that the sufficient conditions are satisfied for Gaussian priors.

The non-existence of an incentive compatible mechanism for heterogeneous populations and continuous signals (Proposition 4) shows that not many improvements can be made in the investigated setting. Therefore, the main direction of our future research is to adapt the mechanism for settings where additional information is known to the mechanism, e.g. where the setting has a temporal structure (Witkowski and Parkes 2012a) or the mechanism can extract useful statistics from statistically independent reports (Dasgupta and Ghosh 2013).

Acknowledgments

The work reported in this paper was supported by NanoTera.ch as part of the OpenSense2 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. Opensense: Open community driven sensing of environment. In *ACM SIGSPATIAL International Workshop on GeoStreaming (IWGS)*, 39–42.
- Bishop, C. M. 2006. *Pattern Recognition and Machine Learning*. Springer.
- Chen, Y., and Pennock, D. M. 2007. A utility framework for bounded-loss market makers. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence (UAI2007)*, 49–56.
- Dasgupta, A., and Ghosh, A. 2013. Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd ACM International World Wide Web Conference (WWW13)*.
- Gneiting, T., and Raftery, A. E. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102:359–378.
- Goel, S.; Reeves, D. M.; and Pennock, D. M. 2009. Collective revelation: A mechanism for self-verified, weighted, and truthful predictions. In *Proceedings of the 10th ACM conference on Electronic commerce (EC 2009)*.
- Hanson, R. D. 2003. Combinatorial information market design. *Information Systems Frontiers* 5(1):107–119.
- Ihara, S. 1993. *Information Theory for Continuous Systems*. World Scientific Publisher Co. Pte. Ltd.
- 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. 2008. Truthful surveys. In *Proceedings of the 3rd International Workshop on Internet and Network Economics (WINE 2008)*.
- 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* 51:1359–1373.
- 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.
- Radanovic, G., and Faltings, B. 2013. A robust bayesian truth serum for non-binary signals. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI'13)*.
- Savage, L. J. 1971. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* 66(336):783–801.
- Shaw, A.; Chen, D. L.; and Horton, J. 2011. Designing incentives for inexpert human raters. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work (CSCW 11)*.
- Waggoner, B., and Chen, Y. 2013. Information elicitation sans verification. In *Information Elicitation Sans Verification. In Proceedings of the 3rd Workshop on Social Computing and User Generated Content (SC13)*.
- Weaver, R., and Prelec, D. 2013. Creating truth-telling incentives with the bayesian truth serum. *Journal of Marketing Research* 50:289–302.
- 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)*.
- Witkowski, J. 2014. *Robust Peer Prediction Mechanisms*. Ph.D. Dissertation, Albert-Ludwigs-Universitat Freiburg; Institut fur Informatik.