

Incentive Schemes for Participatory Sensing

Goran Radanovic
Ecole Polytechnique Federale de Lausanne
Artificial Intelligence Laboratory
CH-1015 Lausanne, Switzerland
goran.radanovic@epfl.ch

Boi Faltings
Ecole Polytechnique Federale de Lausanne
Artificial Intelligence Laboratory
CH-1015 Lausanne, Switzerland
boi.faltings@epfl.ch

ABSTRACT

We consider a participatory sensing scenario where a group of private sensors observes the same phenomenon, such as air pollution. Since sensors need to be installed and maintained, owners of sensors are inclined to provide inaccurate or random data. We design a novel payment mechanism that incentivizes honest behavior by scoring sensors based on the quality of their reports. The basic principle follows the standard Bayesian Truth Serum (BTS) paradigm, where highest rewards are obtained for reports that are surprisingly common. The mechanism, however, eliminates the main drawback of the BTS in a sensing scenario since it does not require sensors to report predictions regarding the overall distribution of sensors' measurements. As it is the case with other peer prediction methods, the mechanism admits uninformed equilibria. However, in the novel mechanism these equilibria result in worse payoff than truthful reporting.

Categories and Subject Descriptors

J.0 [Computer Applications]: General

General Terms

Economics; Measurement

Keywords

Game Theory; Mechanisms Design; Participatory Sensing

1. INTRODUCTION

Private mobile devices equipped with sensors represent a great opportunity in acquiring information about spatially distributed phenomena, such as, air pollution or weather. Instead of deploying sensors themselves, parties interested in monitoring these phenomena can ask the crowd (owners of the mobile devices) to report data collected by ubiquitous sensing devices. This approach is called *participatory* [2] or *community* sensing [1].

Although mobile devices are ubiquitous, sensing induces a certain amount of cost due to the fact that sensing modules need to be installed and maintained. Therefore, beside techniques for filtering and aggregating information, one needs

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*
Copyright © 2015, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

to consider mechanisms that incentivize the crowd to participate in the first place. For example, providing monetary rewards in return for sensors' measurements is one form of incentive. In general, a rewarding system should take into account many orthogonal aspects when considering sensors' utility, ranging from privacy concerns to location preferences. We solely focus on the problem of eliciting accurate measurements and assume that sensors' utility is defined through monetary incentives that are appropriately scaled to cover the cost of sensing.

From a game theoretic perspective, a single sensor and its owner can be considered to be a rational agent whose goal is to maximize her profit. On the other hand, a party interested in monitoring a certain phenomenon, often called *aggregator* or *center*, can be regarded as a mechanism designer that wants to elicit honest information.

A peculiar property of a participatory sensing setting is that the center has no control over the sensing devices, nor it has a way of directly verifying the correctness of the obtained data. This means that standard approaches of constructing incentives based on the quality of the provided information, e. g. proper scoring rules [21, 9] or prediction markets [10, 3], are not applicable in our case.

Instead of directly verifying the collected data, the center can use peer evaluation techniques. This is the basic idea of methods based on the peer prediction principle [15, 11], where an agent's score reflects the information her report carries about the reports of other participants.

In general, several aspects are important when it comes to designing a peer prediction mechanism suitable for the participatory sensing setting. First, a mechanism should have a structure that is knowledge-free of sensors' beliefs because sensors may not believe that the center can easily acquire their private beliefs without directly eliciting them. The second desirable property is a low complexity in terms of the amount of information being elicited. In other words, a mechanism should elicit only sensors' measurements, without requiring sensors to provide additional information, e.g. their posterior beliefs as it is the case for the *Bayesian Truth Serum* [16]. Finally, a mechanism should make uninformed equilibria (equilibria where sensors do not perform measurements) less desirable than informed equilibria, thus providing incentives that are resilient to collusion.

One of the peer prediction methods proposed for information elicitation in community sensing setting is the *Peer Truth Serum* (PTS) from [8]. It rewards a sensor for reporting its measurement if the sensor's report matches the report of a peer sensor, with a score inversely proportional

to an estimated prior probability of the sensor’s report. The incentive compatibility of the mechanism, however, is only guaranteed if the estimated prior corresponds to sensors’ beliefs. Furthermore, the mechanism admits uninformed equilibria that can result in significantly higher expected payoffs than truthful reporting. For example, reporting the least likely value according to the estimated prior is an equilibrium strategy with the highest payoff.

While uninformed equilibria are not avoidable in the peer prediction settings, we would like them to be less profitable than the informed ones. We design a mechanism based on the principles of the Bayesian Truth Serum (BTS) paradigm [16], in which the highest reward is obtained for the report that is more common than expected. Unlike the original BTS mechanism, our novel mechanism does not require sensors to have a common prior belief and is minimal in a sense that sensors report only their measurements. To emphasize that the mechanism extends the PTS mechanism using the BTS principle, we refer to it as *Logarithmic Peer Truth Serum* (logarithmic PTS).

The novel mechanism is applicable to a setting where non-binary information is elicited from a large group of agents whose observations are statistically similar. In particular, we focus on a pollution sensing scenario where a dense network of sensors with similar characteristics measure air quality in different locations of a certain urban area.

The logarithmic PTS mechanism has several desirable properties:

- It is incentive compatible and allows sensors to have different and private prior beliefs.
- Any uninformed equilibrium, including random reporting or collusion on one value, results in worse payoff than truthful reporting.
- Collusion strategies that are based on sensors’ measurements do not lead to higher expected payoff than truthful reporting.

Furthermore, using a realistic dataset of pollution measurements, similar to the one from [8], we analyze the properties of our mechanism and verify its practicality.

2. RELATED WORK

Although there is a vast literature about different aspects of participatory and community sensing paradigms, such as an optimal sensor placement [14] or privacy and trustworthiness of sensors [7, 4, 22], we focus in this section on work related to incentives for reporting accurate information.

The original version of the peer prediction method [15] and its extensions [12, 13] are minimal in a sense that they elicit only desired information, but they require knowledge of agents’ prior beliefs. This requirement represents a significant drawback that is not easy to fulfill.

The Bayesian Truth Serum [16], and its variants [26, 18, 19], use additional report to remove the need of agents’ prior beliefs. However, the BTS mechanisms assume that agents share common prior belief, which in our case might not hold since different sensors can have different prior information about the pollution phenomenon. For example, some sensors might have access to prior concentration estimates of pollutants, while others do not. Common prior is also assumed by the mechanisms in [27] and [20].

Peer prediction without a common prior [25] elicits binary information without requiring the knowledge of agents’ priors, nor common prior belief among agents. However, the mechanism assumes a temporal segmentation of the elicitation process that allows the mechanism to elicit agents’ beliefs prior to their observations. This sort of temporal segmentation is hard to achieve in participatory sensing settings.

It is also worth mentioning mechanisms that are not necessarily truthful, but do elicit useful information. Here we can include the mechanism from [8], already mentioned in the introduction, as well as the mechanism of [23, 24] that elicits common knowledge rather than private information.

A setting similar to ours has been analyzed by [6]. The authors construct a mechanism that has strong incentive properties, similar to the properties of our mechanism, but the mechanism can only be applied for elicitation of binary information. Most pollution sensing applications have more than two levels of pollution.

3. PRELIMINARIES

We consider a scenario where a large group of similar sensors measure air quality over a certain urban area. In order to compensate for the cost of sensing, the center provides incentives to sensors in return for their measurements.

We model sensors as rational agents who seek to maximize their rewards. The set of all sensors is denoted by S . After measuring the air pollution phenomenon, a sensor s obtains a private signal X_s that corresponds to the level of pollution at its location. We consider private signal X_s that takes values in finite discrete set $\{x, y, z, \dots\}$; this is a reasonable assumption considering the fact that ubiquitous sensors are prone to measurement errors. For example, there could be three levels of pollution $\{low, medium, high\}$.

Pollution is a localized phenomenon, meaning that its value significantly varies with distance. We model it simply, using two hidden variables Γ and Λ , which intuitively describe how global and local parameters influence a sensor’s measurement at a certain location. More precisely, our pollution model has the following structure.

- Random variable Γ describes a global state of the world and is generated according to a distribution $Pr(\Gamma)$. Random variable Λ_s describes local variations in the vicinity of a sensor s . It is related to Γ through a distribution $Pr(\Lambda_s|\Gamma)$.
- Sensor s ’s measurement X_s is dependent on both Γ and Λ_s , and the dependency is described by a distribution $Pr(X_s|\Lambda_s, \Gamma)$. Since sensors are assumed to be statistically similar, we have that $Pr(X_s = x|\Lambda_s, \Gamma) = Pr(X_{s'} = x|\Lambda_{s'}, \Gamma)$ for $\Lambda_s = \Lambda_{s'}$, which further implies $Pr(X_s = x|\Gamma) = Pr(X_{s'} = x|\Gamma)$.
- Two sensors s and s' that are close to each other have equal local variables $\Lambda_s = \Lambda_{s'}$, while X_s is conditionally independent of $X_{s'}$ given Γ and Λ_s , i.e. $Pr(X_s|X_{s'}, \Lambda_s, \Gamma) = Pr(X_s|\Lambda_s, \Gamma)$. For two sensors s and s' that are significantly away from each other, we assume conditional independence of X_s and $X_{s'}$ given global state Γ , i.e. $Pr(X_s|X_{s'}, \Gamma) = Pr(X_s|\Gamma)$.
- All probabilities have full support (they are greater than 0), and measurements are stochastically relevant

for Λ_s , i.e. $Pr(\Lambda_s|X_s = x, \Gamma)$ and $Pr(\Lambda_s|X_s = y, \Gamma)$ differ for $y \neq x$ (and similarly for sets of measurements, $\{x_1, x_2, \dots\} \neq \{y_1, y_2, \dots\}$ implies inequality between $Pr(\Lambda_s|X_s \in \{x_1, x_2, \dots\}, \Gamma)$ and $Pr(\Lambda_s|X_s \in \{y_1, y_2, \dots\}, \Gamma)$). The latter is a standard assumption in the peer prediction settings [15].

We do not require the center to have access to the pollution model. Sensors, on the other hand, form beliefs about the parameters (distributions) of the pollution model. Different sensors can have different beliefs about the model, and these beliefs are not known to the center. Since our analysis uses only beliefs of sensor s , we denote them in the same way as the parameters of the pollution model. For example, the belief of sensor s about how measurements are generated is denoted by $Pr(X_s|\Lambda_s, \Gamma)$.

From a sensor s 's perspective, there are two groups of sensors: the set of all sensors, denoted by S ; and *peer* sensors P that are placed in the vicinity of sensor s . Since the pollution model can always be applied to smaller scales, for example, by considering separately different parts of an urban area, we assume that the network of sensors is equally dense everywhere so that $1 \ll |P| \ll |S|$ holds.

Sensors might not be honest, so we differentiate sensor s 's measurement X_s , from its reported values Y_s . We classify reporting strategies into three types:

- *Honest reporting*, i.e. $Y_s = X_s$.
- *Heuristic reporting* without making any measurements, described by a distribution π_{heur} , where a sensor s reports $Y_s = x$ with probability $\pi_{heur}(x)$.
- *Misreporting* in which reports are obtained from measurements using function $\rho : \{x, y, z, \dots\} \rightarrow \{x, y, z, \dots\}$, i.e. $Y_s = \rho(X_s)$.

These three types of strategies cover the most interesting cases of sensor behaviour, including the collusion strategy where all agents report the same value. We have not considered strategies that are dependent on sensors' locations. In a dense network of mobile sensors, these strategies are hard to coordinate, while because of privacy concerns, sensors might not be willing to share their locations.

4. LOGARITHMIC SCORE

When making probabilistic estimates of data, one can define a loss function to measure the quality of the estimates. One of the most well known loss functions in machine learning and information theory is the logarithmic loss function $\log \frac{1}{\mathbf{p}(x)}$, where \mathbf{p} is a probabilistic estimate and x is the true value of the data being estimated. For example, if one evaluates that a probability of a binary variable being equal to 1 is 0.9, and the true value is 1, the logarithmic loss is equal to $\log \frac{1}{0.9}$.

The negative value of the logarithmic loss can be used as a scoring function in elicitation of private beliefs, and it is part of a wider class of mechanisms called proper scoring rules [21, 9]. In game theoretic terms, an agent provides her subjective probability \mathbf{p} about a certain event, and upon the realization of the event is scored with:

$$score = a \cdot \log(\mathbf{p}(x)) + b$$

where x is the outcome of the event, and $a > 0$ and b are scaling parameters. It can be shown that the expected logarithmic score is maximized when the agent provides her true beliefs about the event. Although the logarithmic score is not bounded, from a practical point of view this is almost never a problem. Namely, the lower bound on possible values of $\mathbf{p}(x)$ is usually not hard to estimate, so by using scaling parameters a and b , one can easily fit scores to an arbitrary interval. This implies that individual rationality and bounded payments are in practice achievable. Furthermore, we can also scale payments so that they cover the cost of measurements. For simplicity, we set $a = 1$ and $b = 0$ in the remaining part of the paper.

The logarithmic score is associated with Kullback-Leibler divergence (*KL divergence*), which measures the difference between the expected score when the true (optimal) belief \mathbf{p}_{true} is reported and the expected score when an agent reports belief \mathbf{p}_{report} :

$$KL(\mathbf{p}_{true}||\mathbf{p}_{report}) = \sum_x \mathbf{p}_{true}(x) \log \frac{\mathbf{p}_{true}(x)}{\mathbf{p}_{report}(x)}$$

It is a well known fact that KL divergence is positive and is equal to 0 if and only if $\mathbf{p}_{report} = \mathbf{p}_{true}$.

5. LOGARITHMIC PEER TRUTH SERUM

The basic idea behind our mechanism is to score a sensor s based on how statistically significant its report is. To determine the statistical significance, we first sample reports on a global scale and make a normalized histogram \mathbf{x}_{global} of reported values. That is, for each possible measurement value x , we evaluate the fraction of reports in the sample that are equal to x . Second, we calculate the normalized histogram \mathbf{x}_{local} of reports that are in vicinity of sensor s . Finally, the statistical significance of a report equal to x is then defined as $\log \frac{\mathbf{x}_{local}(x)}{\mathbf{x}_{global}(x)}$.

Although the (original) Bayesian Truth Serum [16] has a different and more complex structure of the score, our approach has some similarities to it. Namely, the idea behind our approach is not to assign highest rewards to the most common reports, but to the reports that are more common than expected. To do so, we also use the logarithmic scoring rule. Thus, we call our mechanism *Logarithmic Peer Truth Serum*.

Logarithmic Peer Truth Serum has the following structure:

- Consider a sensor whose report is equal to Y_s . Let us denote by P , sensor s 's *peers*, i.e. sensors that are in vicinity of sensor s .
- Calculate two empirical frequencies:
 - Frequency of reports equal to x among sensor s 's peers:

$$\mathbf{x}_{local}(x) = \frac{1}{|P|} \sum_{p \in P} \mathbb{1}_{Y_p=x}$$

- Frequency of reports equal to x among reference sensors σ ($|\sigma| \gg 1$) that are not each other's peers nor peers of sensor s :

$$\mathbf{x}_{global}(x) = \frac{1}{|\sigma|} \sum_{s' \in \sigma} \mathbb{1}_{Y_{s'}=x}$$

- Finally, reward sensor s with:

$$\text{score} = \log \frac{\mathbf{x}_{\text{local}}(Y_s)}{\mathbf{x}_{\text{global}}(Y_s)}$$

To avoid potential issues with 0 values in $\mathbf{x}_{\text{local}}$ and $\mathbf{x}_{\text{global}}$ histograms, one can apply Laplace (additive) smoothing with small smoothing parameters, or simply include the report of sensor s in both histograms. The latter would, for example, make the score equal to 0 when $\mathbf{x}_{\text{local}}(Y_s) = 0$ and $\mathbf{x}_{\text{global}}(Y_s) = 0$.

The optimal selection of peers is out of the scope of this paper. The simplest criteria for selecting peers would be to choose m closest sensors, where m should be much smaller than $|S|$. Once the peers of each sensor are obtained, it is fairly easy to determine the corresponding reference sensors σ . In practice, however, one can consider σ to be the set of all sensors, without affecting any incentive properties. Namely, it suffices that $\mathbf{x}_{\text{global}}(x)$ converges to $Pr(X_s = x|\Gamma)$, which is for $\sigma \approx S$ naturally satisfied in the considered setting.

As an example, consider a sensor s and a binary measurement space $\{x, y\}$ representing low and high levels of pollution, respectively. Suppose that sensor s observes x and believes that the frequency of reports on a global scale is $\mathbf{x}_{\text{global}}(x) = 0.2$ and $\mathbf{x}_{\text{global}}(y) = 0.8$, while on a local scale it is $\mathbf{x}_{\text{local}}(x) = 0.4$ and $\mathbf{x}_{\text{local}}(y) = 0.6$. These beliefs are consistent with the Bayesian updating. Clearly, a constant score to a sensor that provides a report consistent with a report of its random peer (e.g. see [11]) would incentivize sensor s to report y . However, the logarithmic PTS score incentivizes sensor s to report x since $\log \frac{0.4}{0.2} > \log \frac{0.6}{0.8}$.

In the following subsection we show several important properties of our mechanism. First, the novel mechanism is incentive compatible, i.e. it admits truthful reporting as an equilibrium with expected payoff greater than 0. What distinguishes our mechanism from other peer prediction mechanisms is that it allows agents to have both different prior beliefs and non-binary measurement space. Next, we show that any heuristic reporting equilibrium, in which sensors do not perform measurements, result in expected payoff equal to 0. Moreover, if sensor s reports heuristically, while other sensors are honest, it expects negative payoff. Finally, we analyze collusive strategies and show that they are not profitable, i.e. they do not lead to higher payoff than truthful reporting.

5.1 Truthful Reporting

The first property we show is incentive compatibility of the logarithmic PTS. Before providing a formal proof, we describe how one can see from the properties of the logarithmic scoring rule that truthful reporting is an equilibrium strategy.

Consider a sensor s and assume other sensors are honest. It can be shown that logarithm of ratio $\frac{\mathbf{x}_{\text{local}}(\tilde{x})}{\mathbf{x}_{\text{global}}(\tilde{x})}$ from sensor s 's perspective converges to:

$$\lim_{|P|, |\sigma| \rightarrow \infty} \log \frac{\mathbf{x}_{\text{local}}(\tilde{x})}{\mathbf{x}_{\text{global}}(\tilde{x})} = \log \frac{Pr(X_s = \tilde{x}|\Lambda_s, \Gamma)}{Pr(X_s = \tilde{x}|\Gamma)}$$

where \tilde{x} is sensor s 's report. Using Bayes' rule we obtain:

$$\begin{aligned} \lim_{|P|, |\sigma| \rightarrow \infty} \log \frac{\mathbf{x}_{\text{local}}(\tilde{x})}{\mathbf{x}_{\text{global}}(\tilde{x})} &= \log \frac{Pr(\Lambda_s|X_s = \tilde{x}, \Gamma)}{Pr(\Lambda_s|\Gamma)} \\ &= \log Pr(\Lambda_s|X_s = \tilde{x}, \Gamma) + c \end{aligned}$$

where c does not depend on report \tilde{x} . The score has one indicative feature: sensor s is scored based on how well it predicts local state Λ_s given global state Γ . More precisely, $\log \frac{\mathbf{x}_{\text{local}}(\tilde{x})}{\mathbf{x}_{\text{global}}(\tilde{x})}$ is related to the logarithmic score of the posterior belief of a sensor whose measurement is equal to \tilde{x} . The true belief of sensor s is $Pr(\Lambda_s|X_s = x, \Gamma)$, where x is its measurement, so in order to be scored with its true beliefs, the sensor should report $\tilde{x} = x$. From the previous sections, we know that the logarithmic score incentivizes agents to report their true beliefs, which is in this case equivalent to reporting true measurements.

The logarithmic PTS rewards a sensor based on how well its measurement indicates the local properties of pollution, rather than global ones. This is in accordance with the main idea of the (original) BTS score: a sensor is not rewarded based on how common her report is on the global scale, but based on how surprisingly common her report is on the local scale given the global state.

Finally, notice that c in the expression above is equal to $-\log Pr(\Lambda_s|\Gamma)$. This can be interpreted as if sensor s is being scored based on how much her report improves the knowledge about local state Λ_s . That is, if a sensor s did not make any measurement, her belief regarding Λ_s given global state Γ would be equal to $Pr(\Lambda_s|\Gamma)$. Hence, the difference between $\log Pr(\Lambda_s|X_s = x, \Gamma)$ and $\log Pr(\Lambda_s|\Gamma)$ represents the increase of the sensor's knowledge about Λ_s after performing a measurement. Furthermore, using the same reasoning as above, we obtain that sensor s 's score is in expectation positive: due to the properties of the logarithmic score we know that $-\log Pr(\Lambda_s|\Gamma)$ is in expectation greater than $-\log Pr(\Lambda_s|X_s = x, \Gamma)$ for sensor s whose belief is equal to $Pr(\Lambda_s|X_s = x, \Gamma)$. Putting it all together, we have:

THEOREM 1. *The Logarithmic Peer Truth Serum is strictly Bayes-Nash incentive compatible, with strictly positive expected payoffs in the truthful reporting equilibrium.*

PROOF. Consider a sensor s and assume other sensors are honest. Reports from local histogram $\mathbf{x}_{\text{local}}$ are conditionally independent given Λ_s and Γ , so we can apply the law of large numbers to obtain:

$$\lim_{|P| \rightarrow \infty} \mathbf{x}_{\text{local}}(x) = Pr(X_s = x|\Lambda_s, \Gamma)$$

Next, consider reference sensors $\sigma \subset S$ whose measurements are conditionally independent of each other given Γ . Since we assumed $|S| \gg |P| \gg 1$, it follows that $|\sigma|$ can also be large, so we choose one such σ . The reports of sensors in σ are conditionally independent given Γ , hence we have that:

$$\lim_{|\sigma| \rightarrow \infty} \mathbf{x}_{\text{global}}(x) = Pr(X_s = x|\Gamma)$$

Therefore, the expected score of sensor s , who measured $X_s = x$, for reporting $Y_s = y$ is equal to:

$$\begin{aligned} \lim_{|\sigma|, |P| \rightarrow \infty} \int_{\Lambda_s, \Gamma} Pr(\Lambda_s, \Gamma|X_s = x) \log \frac{\mathbf{x}_{\text{local}}(y)}{\mathbf{x}_{\text{global}}(y)} d\Lambda_s d\Gamma &= \\ \int_{\Lambda_s, \Gamma} Pr(\Lambda_s, \Gamma|X_s = x) \log \frac{Pr(X_s = y|\Lambda_s, \Gamma)}{Pr(X_s = y|\Gamma)} d\Lambda_s d\Gamma &= \\ \int_{\Lambda_s, \Gamma} Pr(\Lambda_s|X_s = x, \Gamma) \cdot Pr(\Gamma|X_s = x) \cdot \\ \cdot \log \frac{Pr(\Lambda_s|X_s = y, \Gamma)}{Pr(\Lambda_s|\Gamma)} d\Lambda_s d\Gamma \end{aligned}$$

where the last equality is due the chain rule for conditional probabilities (the product of the probabilities) and the Bayes' law (the term inside the logarithm). The equation can be further reduced to:

$$\begin{aligned} & \int_{\Lambda_s, \Gamma} Pr(\Lambda_s | X_s = x, \Gamma) Pr(\Gamma | X_s = x) \cdot \\ & \cdot \log \frac{Pr(\Lambda_s | X_s = y, \Gamma) Pr(\Lambda_s | X_s = x, \Gamma)}{Pr(\Lambda_s | \Gamma) Pr(\Lambda_s | X_s = x, \Gamma)} d\Lambda_s d\Gamma \\ & \int_{\Gamma} Pr(\Gamma | X_s = x) [\\ & \int_{\Lambda_s} Pr(\Lambda_s | X_s = x, \Gamma) \log \frac{Pr(\Lambda_s | X_s = y, \Gamma)}{Pr(\Lambda_s | X_s = x, \Gamma)} d\Lambda_s \\ & + \int_{\Lambda_s} Pr(\Lambda_s | X_s = x, \Gamma) \log \frac{Pr(\Lambda_s | X_s = x, \Gamma)}{Pr(\Lambda_s | \Gamma)} d\Lambda_s] d\Gamma \end{aligned}$$

which can be represented as:

$$\begin{aligned} & \int_{\Gamma} f_1(\Gamma) [-KL(Pr(\Lambda_s | X_s = x, \Gamma) || Pr(\Lambda_s | X_s = y, \Gamma))] \\ & + f_2(\Gamma)] d\Gamma \end{aligned}$$

Since the only part that depends on the sensor's report is KL divergence $KL(Pr(\Lambda_s | X_s = x, \Gamma) || Pr(\Lambda_s | X_s = y, \Gamma))$, and the KL divergence is strictly positive unless $y = x$ due to the stochastic relevance of X_s , we conclude that sensor s 's score is in expectation maximized for $y = x$ (honest reporting). Moreover, when sensor s is honest $KL(Pr(\Lambda_s | X_s = x, \Gamma) || Pr(\Lambda_s | X_s = y, \Gamma)) = 0$. Notice that function f_2 is also a KL divergence: $f_2(\Gamma) = KL(Pr(\Lambda_s | X_s = x, \Gamma) || Pr(\Lambda_s | \Gamma))$, and, thus, is strictly positive. Hence, the expected payoff of a sensor in honest reporting equilibrium is strictly positive. \square

5.2 Heuristic Reporting

Now we turn to the next important property of the logarithmic PTS mechanism. Suppose that sensors agree to use heuristic reporting strategy independent of their locations, i.e. without performing measurements they report according to a certain policy described by a distribution π_{heur} . This distribution can also include sensors colluding on a particular value. For example, if sensors collude and decide to report x , the reporting distribution is equal to $\pi_{heur}(x) = 1$ and $\pi_{heur}(y) = 0, \forall y \neq x$.

In this case, the ratio $\frac{\mathbf{x}_{local}(x)}{\mathbf{x}_{global}(x)}$ converges to 1. The reason is that sensors no longer provide reports that are informative about variables Λ_s , i.e. they are uniform on a global scale. The direct consequence is that the logarithmic PTS score is equal to 0. One should note two important properties. First, uninformed equilibria result in the expected payoff equal to 0, which is less than what sensors expect to obtain if they are honest. This means that if sensing requires a certain amount of effort, it is enough to appropriately scale the logarithmic PTS score in order to incentivize sensors' owners to make measurements and report honestly. Second, simple collusive strategies, such as all sensors reporting a certain value x , are less profitable than honest reporting.

THEOREM 2. *In the Logarithmic Peer Truth Serum, heuristic reporting equilibria result in zero expected payoff.*

PROOF. Consider a sensor s and assume other sensors report according to a distribution π_{heur} . From the law of

large numbers we have that:

$$\begin{aligned} \lim_{|P| \rightarrow \infty} \mathbf{x}_{local}(x) &= \pi_{heur}(x) \\ \lim_{|\sigma| \rightarrow \infty} \mathbf{x}_{global}(x) &= \pi_{heur}(x) \end{aligned}$$

Therefore, the expected payoff for reporting $y \in \{z | \pi_{heur}(z) > 0\}$ is:

$$\lim_{|\sigma|, |P| \rightarrow \infty} \log \frac{\mathbf{x}_{local}(y)}{\mathbf{x}_{global}(y)} = \log \frac{\pi_{heur}(y)}{\pi_{heur}(y)} = \log 1 = 0$$

\square

One should also note that if a sensor s reports heuristically, its expected payoff when other sensors are honest is negative. Namely, from the previous section we know that in that case the score converges to:

$$\begin{aligned} \lim_{|P|, |\sigma| \rightarrow \infty} \log \frac{\mathbf{x}_{local}(\tilde{x})}{\mathbf{x}_{global}(\tilde{x})} \\ = \log Pr(\Lambda_s | X_s = \tilde{x}, \Gamma) - \log Pr(\Lambda_s | \Gamma) \end{aligned}$$

However, sensor s 's belief regarding Λ_s is now $Pr(\Lambda_s | \Gamma)$ because it does not make any measurement. Due to the properties of the logarithmic score, $\log Pr(\Lambda_s | \Gamma)$ is in expectation greater than $\log Pr(\Lambda_s | X_s = \tilde{x}, \Gamma)$, which means that the score of sensor s is in expectation negative.

PROPOSITION 1. *Consider a sensor s that uses heuristic reporting strategy and suppose other sensors are honest. Then, in the Logarithmic Peer Truth Serum, sensor s has a negative expected payoff.*

PROOF. (Sketch) Consider a sensor s and assume all other sensors are honest. Following the steps of Theorem 1 and noting that sensor s 's belief regarding Λ_s given Γ is equal to $Pr(\Lambda_s | \Gamma)$, we obtain that sensor s 's expected score for reporting $Y_s = y$ is equal to:

$$\begin{aligned} & \int_{\Gamma} Pr(\Gamma) \int_{\Lambda_s} Pr(\Lambda_s | \Gamma) \cdot \log \frac{Pr(\Lambda_s | X_s = y, \Gamma)}{Pr(\Lambda_s | \Gamma)} d\Lambda_s d\Gamma \\ & = \int_{\Gamma} Pr(\Gamma) \cdot [-KL(Pr(\Lambda_s | \Gamma) || Pr(\Lambda_s | X_s = y, \Gamma))] d\Gamma \end{aligned}$$

Since the KL divergence is positive, the expected score is negative. \square

5.3 Misreporting

In the previous sections we have seen that truthful reporting achieves higher payoff than heuristic reporting, which means that the logarithmic PTS can be scaled to cover the costs of sensing. In this section, we further investigate whether sensors can achieve a greater payoff if their collusion strategies are based on their observations. We restrict our attention to misreporting that can be defined via function $\rho : \{x, y, z, \dots\} \rightarrow \{x, y, z, \dots\}$ that maps measurements to reports. Notice that function ρ is not necessarily a bijective function, i.e. for two different measurements x and y , reporting strategy ρ can output the same report $\rho(x) = \rho(y)$.¹ We first show that truthful reporting is at least as good as any other misreporting strategy defined by function ρ .

¹Since ρ is a mapping from finite set of values $\{x, y, z, \dots\}$ to the same set of values $\{x, y, z, \dots\}$, ρ is a bijection if and only if it is an injection.

THEOREM 3. *In the Logarithmic Peer Truth Serum, a misreporting strategy profile defined by a function $\rho : \{x, y, z, \dots\} \rightarrow \{x, y, z, \dots\}$ is not in expectation more profitable than truthful reporting.*

PROOF. (Sketch) Consider a sensor s and assume other sensors are reporting according to a function ρ . The expected frequency of reports equal to z that come from the sensors located in the vicinity of sensor s is:

$$\sum_{y|\rho(y)=z} Pr(X_s = y|\Lambda_s, \Gamma) = Pr(X_s \in \{y|\rho(y) = z\}|\Lambda_s, \Gamma)$$

Similarly, we obtain for the expected frequency of reports equal to z on the global scale:

$$\sum_{y|\rho(y)=z} Pr(X_s = y|\Gamma) = Pr(X_s \in \{y|\rho(y) = z\}|\Gamma)$$

From the law of large numbers we have that:

$$\lim_{|P| \rightarrow \infty} \mathbf{x}_{local}(z) = Pr(X_s \in \{y|\rho(y) = z\}|\Lambda_s, \Gamma)$$

and

$$\lim_{|\sigma| \rightarrow \infty} \mathbf{x}_{global}(z) = Pr(X_s \in \{y|\rho(y) = z\}|\Gamma)$$

Furthermore, using Bayes' rule, we obtain:

$$\lim_{|\sigma|, |P| \rightarrow \infty} \frac{\mathbf{x}_{local}(z)}{\mathbf{x}_{global}(z)} = \frac{Pr(\Lambda_s|X_s \in \{y|\rho(y) = z\}, \Gamma)}{Pr(\Lambda_s|\Gamma)} \quad (1)$$

An analysis equivalent to the one in Theorem 1, gives us that the expected score of sensor s , who observed x , for reporting $\rho(x)$ is equal to:

$$\int_{\Gamma} f_1(\Gamma)[-KL(Pr(\Lambda_s|X_s = x, \Gamma)||Pr(\Lambda_s|X_s \in \{y|\rho(y) = \rho(x)\}, \Gamma))] + f_2(\Gamma)d\Gamma \quad (2)$$

Since KL divergence $KL(Pr(\Lambda_s|X_s = x, \Gamma)||Pr(\Lambda_s|X_s \in \{y|\rho(y) = \rho(x)\}, \Gamma))$ is non-negative, expression (2) cannot be larger than $\int_{\Gamma} f_1(\Gamma) \cdot f_2(\Gamma)d\Gamma$, which is equal to the expected payoff for truthful reporting (see the proof of Theorem 1). \square

The intuition behind this result can be seen from expression (1), which implies that sensor s 's score converges to:

$$\lim_{|\sigma|, |P| \rightarrow \infty} \log \frac{\mathbf{x}_{local}(z)}{\mathbf{x}_{global}(z)} = \log Pr(\Lambda_s|X_s \in \{y|\rho(y) = z\}, \Gamma) + c$$

where $z = \rho(x)$, x is sensor s 's measurement, and c does not depend on report z . Due to the properties of the logarithmic scoring rule, the expected value of $\log Pr(\Lambda_s|X_s \in \{y|\rho(y) = z\}, \Gamma)$ cannot be greater than the expected value of $\log Pr(\Lambda_s|X_s = x, \Gamma)$, which is achieved when $\rho(x) = x$. This means that truthful reporting is at least as good as any other misreporting strategy defined by ρ . Notice that truthful reporting does not require any special coordination among sensors - the coordination is directly provided in a form of their measurements. That is, while sensors can collude by misreporting certain values, such collusive behaviour is hard to coordinate and it does not lead to higher payoff.

Theorem 3 does not state if misreporting strategies defined by a function ρ can achieve payoff equal to the one obtained for truthful reporting. The next result shows that this happens when ρ is a bijection, and only in that case.

PROPOSITION 2. *In the Logarithmic Peer Truth Serum, a misreporting strategy profile $\rho : \{x, y, z, \dots\} \rightarrow \{x, y, z, \dots\}$ achieves the same expected payoff as truthful reporting if and only if ρ is a bijective function.*

PROOF. (Sketch) When ρ is a bijection, $\{y|\rho(y) = \rho(x)\} = \{x\}$, so distributions $Pr(\Lambda_s|X_s \in \{y|\rho(y) = \rho(x)\}, \Gamma)$ and $Pr(\Lambda_s|X_s = x, \Gamma)$ are equal, and, hence, the divergence $KL(Pr(\Lambda_s|X_s = x, \Gamma)||Pr(\Lambda_s|X_s \in \{y|\rho(y) = \rho(x)\}, \Gamma))$ in expression (2) is equal to 0. This is true for any bijective ρ , which means that the payoff is in expectation equal to what sensors expect to obtain in truthful reporting equilibrium. When ρ is not a bijection, there exists a measurement x for which $\{y|\rho(y) = \rho(x)\} \supset \{x\}$, and, by stochastic relevance, $KL(Pr(\Lambda_s|X_s = x, \Gamma)||Pr(\Lambda_s|X_s \in \{y|\rho(y) = \rho(x)\}, \Gamma)) > 0$, implying the strictly lower expected payoff than for honest reporting. \square

As an example consider a sensor s and a ternary measurement space $\{x, y, z\}$ representing low, medium and high levels of pollution, respectively. Suppose that sensor s observes x and believes that the frequency of reports on a global scale is $\mathbf{x}_{global}(x) = 0.2$, $\mathbf{x}_{global}(y) = 0.5$ and $\mathbf{x}_{global}(z) = 0.3$, while on a local scale it is $\mathbf{x}_{local}(x) = 0.3$, $\mathbf{x}_{local}(y) = 0.6$ and $\mathbf{x}_{local}(z) = 0.1$. If sensors decide to report honestly, sensor s expects to obtain $\log(\frac{0.3}{0.2}) \approx 0.4$. Now, suppose sensors collude and they decide to report as follows:

- When they measure x or y , they report x .
- When they measure z , they report z .

In other words, the reporting function is defined as: $\rho(x) = x$, $\rho(y) = x$, $\rho(z) = z$. In this case, sensor s expects to obtain $\log(\frac{0.3+0.6}{0.2+0.5}) = \log(\frac{0.9}{0.7}) \approx 0.25$, which is less than in the honest reporting equilibrium.

6. SIMULATION

We examine the characteristics of the logarithmic PTS using realistic data of Nitrogen Dioxide (NO_2) concentrations over the city of Strasbourg. The data consists of both real measurements collected by ASPA² and estimations of pollution from the physical model ADMS Urban V2.3 [5]. In total, the data set contains concentrations of NO_2 for each hour, expressed in parts per billion (ppb), at 116 different locations over a period of four weeks. The locations are placed as shown in Figure 1 and in the following text are referred to as sensors.

Although the initial measurements take values in continuous domain, we discretize it using four levels of pollution defined as:

- *low*: concentrations 0 – 20 ppb;
- *medium*: concentrations 20 – 40 ppb;
- *high*: concentrations 40 – 60 ppb;
- *extra-high*: concentrations 60 – ∞ ppb.

Each hour, sensors report the measured level of pollution to the center and are rewarded with the logarithmic PTS mechanism. As a criterion for peer selection, we consider distance and define peers of a certain sensor as 15 closest

²www.atmo-alsace.net

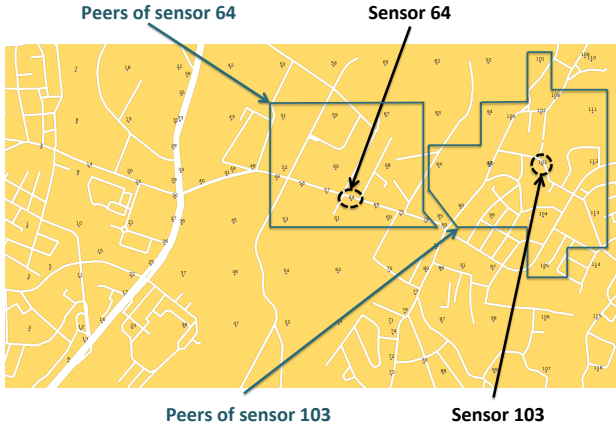


Figure 1: Sensor placement in Strasbourg urban area

sensors. Namely, with 15 sensors, one can obtain fairly good insight in localized aspects of pollution, while satisfying the condition that the number of peers is significantly smaller than the total number of sensors. Local histogram \mathbf{x}_{local} is for a sensor s calculated based on the reports of its 15 peers. Global histogram \mathbf{x}_{global} includes reports of all sensors, except for the sensor s 's report. Moreover, both histograms are smoothed with the Laplace (additive) smoothing operator using parameters $\alpha_{local} = 10^{-4}$ and $\alpha_{global} = 10^{-3}$ (parameters reflect that \mathbf{x}_{global} is calculated based on approximately 8 times more reports than \mathbf{x}_{local}).

To demonstrate the correctness of our results, we examine four different reporting strategies and evaluate their performance by analyzing the average scores of sensors. The four strategies are defined as follows:

- *truthful*: All sensors are honest.
- *collude*: Sensors collude so that those who observe *low* or *medium* report *low*, while those who observe *high* or *extra-high* report *high*.
- *colludeLow*: All sensors collude and report *low*.
- *random*: A sensor whose score is being calculated reports randomly with probabilities $Pr(low) = \dots = Pr(extra-high) = 0.25$; while others sensors are honest.
- *randomAll*: All sensors report randomly with probabilities $Pr(low) = \dots = Pr(extra-high) = 0.25$.

The statistic of the average payoffs is shown in Table 1. These payoffs can be scaled so that the incentives take positive values and cover the cost of sensing.

As expected, random reporting strategies on average lead to non-positive scores. When a single sensor reports randomly, while others are honest, its expected payoff is strictly negative, confirming Proposition 1. When all sensors report randomly, the average payoffs are more concentrated around 0, which is in accordance with Theorem 2. Notice that by Theorem 2, *randomAll* strategy should lead to 0 expected payoff. In our experimental setup, however, we are

Table 1: Average payoffs

Strategy	mean	min	max
<i>truthful</i>	0.037	-1.153	0.291
<i>collude</i>	0.014	-0.27	0.106
<i>colludeLow</i>	0	0	0
<i>random</i>	-0.876	-1.631	-0.36
<i>randomAll</i>	-0.228	-0.362	-0.123

Strategy	median	1st quartile	3rd quartile
<i>truthful</i>	0.047	-0.017	0.102
<i>collude</i>	0.019	-0.009	0.039
<i>colludeLow</i>	0	0	0
<i>random</i>	-0.823	-1.075	-0.673
<i>randomAll</i>	-0.228	-0.258	-0.19

dealing with finite number of sensors, so histograms \mathbf{x}_{global} and \mathbf{x}_{local} are not necessarily equal to each other. This imbalance produces negative average scores because of non-linearity of log function.

Colluding on a single value results in payoff equal to 0, and this trivially follows from the structure of the score. Collusion strategy *collude* has lower mean of the average payoffs than truthful reporting. Moreover, a careful inspection of medians and quartiles shows that such collusion is worse than truthful reporting for the majority of sensors. Namely, median, third quartile and maximum are greater for truthful reporting than for *collude* strategy.

To inspect why minimal and maximal average of sensors' payoffs significantly differ for *truthful* strategy, we show the reports of the corresponding sensors: sensor 103, that has the worst average payoff (-1.153), and sensor 64, that has the best average payoff (0.291). The locations of the sensors are shown in Figure 1, along with the locations of their peers. We give for each sensor: the histogram of the sensor's reports, the average local histogram $\bar{\mathbf{x}}_{local}$, the average global histogram $\bar{\mathbf{x}}_{global}$.




Figure 2: Histograms of the worst sensor

Figures 2 and 3 clearly indicates the main differences between sensors 103 and 64. Sensor 103 often reports a value that is not as common in the local histogram as expected by the global histogram, while sensor 64 does the opposite - it




Figure 3: Histograms of the best sensor

reports a value that is surprisingly common. Since the logarithmic PTS assigns a higher scores to reports that are more common than expected, it is clear that sensor 103 should obtain much lower payoff than sensor 64.

The setting considered in this paper assumes that the network of sensors is dense. We can further investigate how robust the logarithmic PTS is when the density of the sensor network decreases. To do so, we randomly sample subsets of sensors of different sizes (100, 80, 60 and 40 sensors) on daily basis (i.e. each day a different subset is chosen), and we calculate the median of average payoffs. As already noted, the median of average payoffs reflects how good a reporting strategy is for the majority of the sensors. In addition to reducing the number of sensors, we also reduce the number of peers for each sensor. For example, in a random subset of 80 sensors, the set of peers of a certain sensor contains 11 closest sensors.

The detailed results are shown in Table 2. We can see that truthful reporting remains the optimal strategy until the number of sensors decreases to 40, which represents a critical value where collusive strategies *colludeLow* and *collude* become more profitable than truthfulness. This can be explained by the low amount of information used to generate histograms (\mathbf{x}_{local} is constructed from only 7 reports). In *colludeLow* and *collude* strategies, sensors report one and two levels of pollution respectively, so these strategies are less susceptible to random variations in measurements than truthful reporting, where all four levels of pollution are reported. We can make a similar observation for *randomAll* strategy, for which the median of average payoffs diverges from 0 (the expected payoff) as the number of sensors decreases.

Table 2: Median of average payoffs

Strategy	Num. sensors / Num. peers			
	100/13	80/11	60/9	40/7
<i>truthful</i>	0.04	0.03	0.01	-0.007
<i>collude</i>	0.02	0.01	0.003	-0.0003
<i>colludeLow</i>	0	0	0	0
<i>random</i>	-0.81	-0.86	-0.84	-0.99
<i>randomAll</i>	-0.33	-0.47	-0.74	-1.19

As we can see from tables 1 and 2, provided that other sensors are honest, a sensor who reports randomly receives relatively large negative payoff, while for truthful reporting it receives relatively small positive payoff. However, a simple scaling of the the logarithmic PTS can reverse the situation, making revenue for random reporting slightly negative, while revenue for truthful reporting highly positive. The scaling is done by multiplying the logarithmic PTS score with a constant $a > 0$ and adding to it a constant payment b . For example, by additionally rewarding all the sensors with $b = 0.8$, the median of average payoffs in Table 1 for truthfulness becomes 0.847, while for *random* strategy it becomes -0.023 . This way of scaling the payments does not only incentivize honest reporting, but it also discourages random reporters to participate. Notice that parameters $a > 0$ and b should not depend on sensors' reports. In practice, however, one can easily infer their appropriate values using simulations, such as the one presented here.

7. CONCLUSION

In this paper, we have constructed an incentive mechanism that can be applied in a participatory sensing scenario where a large group of sensors make measurements of a spatially distributed phenomenon. The basic idea of our mechanism is to reward sensors based on the statistical significance of their reports. Because of the way we measure the statistical significance, the mechanism has several desirable properties important for its use in participatory sensing settings. The mechanism is minimal in a sense that it elicits sensors' measurements without requiring them to report additional information. Moreover, it does not require sensors to have a common prior belief in order to achieve incentive compatibility. Furthermore, the mechanism is collusion resistant, meaning that collusive strategies, such as heuristic (uninformed) reporting or misreporting, result in worse expected payoffs than truthful reporting. Finally, we have tested our mechanism in a realistic simulator and confirmed that its robust game-theoretic properties hold.

The main direction of our future work would be to analyze if the logarithmic PTS score can be used for filtering low quality sensors without having the access to the ground truth, similar to how [17] applied the BTS to detect the best experts. Namely, Proposition 1 tells us that a sensor who reports randomly can expect to obtain a negative score, provided that the majority of participants is honest. By keeping a track of sensors' scores, we could determine which sensors provide informative data and which sensors should be excluded when calculating the pollution map of a monitored urban area.

Acknowledgments

The work reported in this paper was supported by NanoTera.ch as part of the OpenSense2 project. We thank Jason Jingshi Li for providing a testbed and the anonymous reviewers for useful comments and feedback.

REFERENCES

- [1] K. Aberer, S. Sathe, D. Chakraborty, A. Martinoli, G. Barrenetxea, B. Faltings, and L. Thiele. Opensense: Open community driven sensing of environment. In *ACM SIGSPATIAL International Workshop on GeoStreaming (IWGS)*, pages 39–42, 2010.

- [2] J. Burke, D. Estrin, M. Hansen, A. Parker, and N. Ramanathan. Participatory sensing. In *World-Sensor-Web*, 2006.
- [3] Y. Chen and D. M. Pennock. A utility framework for bounded-loss market makers. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence (UAI2007)*, pages 49–56, 2007.
- [4] D. Christin, A. Reinhardt, S. Kanhere, and M. Hollick. Combinatorial information market design. *The Journal of Systems and Software*, 84(11):1928–1946, 2011.
- [5] R. N. Colvile, N. K. Woodfield, D. J. Carruthers, B. E. A. Fisher, A. Rickard, S. Neville, and A. Hughes. Uncertainty in dispersion modeling and urban air quality mapping. *Environmental Science and Policy*, 5:207–220, 2002.
- [6] A. Dasgupta and A. Ghosh. Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd ACM International World Wide Web Conference (WWW'13)*, 2013.
- [7] A. Dua, N. Bulusu, and W. Feng. Towards trustworthy participatory sensing. In *Proceedings of the 4th USENIX conference on Hot topics in security.*, 2009.
- [8] B. Faltings, J. J. Li, and R. Jurca. Incentive mechanisms for community sensing. *IEEE Transaction on Computers*, 63:115–128, 2014.
- [9] T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102:359–378, 2007.
- [10] R. D. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.
- [11] S.-W. Huang and W.-T. Fu. Enhancing reliability using peer consistency evaluation in human computation. In *Proceedings of the 2013 conference on Computer supported cooperative work*, 2013.
- [12] R. Jurca and B. Faltings. Minimum payments that reward honest reputation feedback. In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC'06)*, pages 190–199, 2006.
- [13] R. Jurca and B. Faltings. Robust incentive-compatible feedback payments. In *Agent-Mediated Electronic Commerce*, volume LNAI 4452, pages 204–218. Springer-Verlag, 2007.
- [14] A. Krause, A. Singh, and C. Guestrin. Combinatorial information market design. *The Journal of Machine Learning Research*, 9:235–284, 2008.
- [15] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51:1359–1373, 2005.
- [16] D. Prelec. A bayesian truth serum for subjective data. *Science*, 34(5695):462–466, 2004.
- [17] D. Prelec and S. Seung. An algorithm that finds truth even if most people are wrong. Working paper, 2006.
- [18] G. Radanovic and B. Faltings. A robust bayesian truth serum for non-binary signals. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI'13)*, 2013.
- [19] G. Radanovic and B. Faltings. Incentives for truthful information elicitation of continuous signals. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI'14)*, 2014.
- [20] B. Riley. Minimum truth serums with optional predictions. In *Proceedings of the 4th Workshop on Social Computing and User Generated Content (SC14)*, 2014.
- [21] L. J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- [22] A. Singla and A. Krause. Incentives for privacy tradeoff in community sensing. In *Proceedings of HCOMP*, 2013.
- [23] B. Waggoner and Y. Chen. Information elicitation sans verification. In *Proceedings of the 3rd Workshop on Social Computing and User Generated Content (SC'13)*, 2013.
- [24] B. Waggoner and Y. Chen. Output agreement mechanisms and common knowledge. In *Proceedings of the 2nd AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*, 2014.
- [25] J. Witkowski and D. C. Parkes. Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC' 12)*, pages 964–981, 2012.
- [26] J. Witkowski and D. C. Parkes. A robust bayesian truth serum for small populations. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI'12)*, 2012.
- [27] P. Zhang and Y. Chen. Elicitability and knowledge-free elicitation with peer prediction. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems (AAMAS '14)*, 2014.