

Incentives for Answering Hypothetical Questions

Radu Jurca
Google Inc.
Switzerland
radu.jurca@gmail.com

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

ABSTRACT

Prediction markets and other reward mechanisms based on proper scoring rules can elicit accurate predictions about future events with public outcomes. However, many questions of public interest do not always have a clear answer. For example, facts such as the effects of raising or lowering interest rates can never be publicly verified, since only one option will be implemented. In this paper we address reporting incentives for opinion polls and questionnaires about hypothetical questions, where the honesty of one answer can only be assessed in the context of the other answers elicited through the poll. We extend our previous work on this problem by four main results. First, we prove that no reward mechanism can be strictly incentive compatible when the mechanism designer does not know the prior information of the participants. Second, we formalize the notion of *helpful* reporting which prescribes that rational agents move the public result of the poll towards what they believe to be the true distribution (even when that involves reporting an answer that is not the agent's first preference). Third, we show that helpful reporting converges the final result of the poll to the true distribution of opinions. Finally, we present a reward scheme that makes helpful reporting an equilibrium for polls with an arbitrary number of answers. Our mechanism is simple, and does not require information about the prior beliefs of the agents.

Keywords

opinion polls, mechanism design, incentive compatibility

1. INTRODUCTION

Soliciting opinions or predictions through electronic channels has become an important component of our digital society. Companies, governments and private individuals commonly run opinion polls, questionnaires or prediction markets to solicit feedback or advice from the users. While the types of information elicited varies widely, we assume in this paper that the task is to report one of a finite number k of

outcomes, for example an observation, or the probability of one of these outcomes.

For the validity of such mechanisms, it is important that the contributions be truthful. For settings that involve predictions of an outcome that will eventually be known, there are several schemes that reward truthful reporting based on comparison with this outcome. Most of these are based on proper scoring rules [18] that compute rewards to be paid based on comparing the prediction with the true outcome once it has become known. Lambert, Pennock and Shoham [13] [14] [12] provide a complete characterization of what information can be truthfully elicited using proper scoring rules. An alternative are prediction markets [3], where participants express their opinion by buying *securities* that pay off if a certain outcome actually occurs, and expire without payoff otherwise. Such markets have achieved remarkable accuracy for problems such as a predicting the outcome of presidential elections or the completion of a software development projects [1], [5], [2], [16]. To solve the problem of lack of liquidity, prediction markets are usually complemented with a market maker that provides unbounded liquidity and thus ensures an efficient market. Such market makers can also be constructed on the basis of proper scoring rules [7] to ensure that truthful reporting is the dominant strategy.

For a large class of questions, however, the true outcome will never be known. These include predictions of hypothetical events, such as

- *What economic growth will result if we raise the interest rates by 0.5 percent?*
- *Will there be a famine in country X if the harvest fails two times in a row?*

where the outcome can never be verified if the action is not taken, but also many questions where verification would be too costly, for example

- *What is the failure rate of vacuum cleaner X ?*
- *What is the success rate of plumber X ?*
- *What is the average connection speed of internet provider X ?*

In such cases, neither scoring rules nor prediction markets can be used to elicit truthful information.

However, it is possible to use the coherence of a report with others to incentivize truthfulness. In contrast to the case of verifiable outcomes, where truthful reporting is the *dominant* strategy, the goal is now to make truthful reporting an equilibrium strategy by comparing the reports provided by different agents and rewarding them based on how well they agree. Provided that most reports are truthful, it is then a best response to also report truthfully, and truthful reporting becomes a Nash equilibrium. Miller, Resnick and Zeckhauser [15] give a full analysis of how proper scoring rules can be adapted to this setting and called it the *peer prediction method*. In essence, the peer prediction method assumes that a report of a value x amounts to reporting a certain probability that another report, the *reference report*, also reports the value x . The mechanism then applies a proper scoring rule to this predicted probability, taking the reference report as the true outcome. Miller et al. [15] show that for any proper scoring rule, such a mechanism has a Nash equilibrium where all agents report their opinion x truthfully.

Jurca and Faltings [8] [9] extend the results of Miller et al. by applying the principal of Automated Mechanism Design [4] to design a mechanism that is optimal under some cost objective. Goel Reeves and Pennock [6] describe a mechanism that is also able to collect the confidence of the reporters.

However, peer prediction methods have an important drawback: in order to apply the scoring rule mechanism, it is necessary to know the probability distribution that the rater assigns to the reference report. This assumption is not very realistic in practice. While it is possible to design scoring rules that are robust against variations in this probability, [20], [10] has shown that only small variations can be tolerated and that even these can require vastly higher rewards.

One relaxation of this requirement comes in a mechanism called the *Bayesian Truth Serum* [17] which works by asking reporters to also estimate the final distribution of reports. The agent's report and the estimate of the final distribution can be used as a basis for rewards that enforce a truthful equilibrium.

In our previous work [11] we presented an opinion poll mechanism for settings with 2 possible answers that does not require knowledge of agent beliefs. When the current poll result is far from a rater's private beliefs, the rater will file reports that may not always be truthful but are certain to drive the poll result closer to its own beliefs. For a population of raters with uniform beliefs (for example, rater that all report on the same observation), the poll converges to an equilibrium where it indicates this belief.

This paper brings four main additions to our previous results. First, we prove that no reward mechanism can be strictly incentive compatible (in the sense that every participant always reports her honest opinion) if the designer does not know the agents' prior information (Section 3). Second, we formalize the notion of *helpful* reporting (Section 4) which prescribes that rational agents move the public result

of the poll towards what they believe to be the true distribution, even when that involves reporting an answer that is not the agent's first preference. Third, we show that helpful reporting converges the final result of the poll to the true distribution of opinions. Finally, we present a reward scheme that makes helpful reporting an equilibrium for polls with an arbitrary number of answers (Section 4.2). We believe that this contribution is an important step forward for demonstrating the practicality of incentive-based mechanisms for real life information elicitation scenarios.

2. THE SETTING

We assume a setting where a principal (he) asks a question to a large set of agents. The question has N possible answers, denoted as $A = \{a_1, a_2, \dots, a_N\}$, and every agent (she) is allowed to report exactly one answer.

Every agent i has a private belief that the answer to the question should be $o_i \in A$, and no agent knows exactly what other agents believe. Agents are assumed rational and may choose to report a false answer. However, agents are not allowed to collude by using side-communication channels to coordinate their reports.

Let $\Delta(A)$ be the set of probability distributions over the set of answers. The population profile ω is described by the probability vector $(\omega_1, \omega_2, \dots, \omega_N) \in \Delta(A)$ where the fraction of agents believing in answer a_i is $\omega_i \in (0, 1)$. The agent's opinions (i.e., the answer they believe in) are drawn independently according to a true (unknown) population profile.

Let Ω be the random variable representing the true population profile. The prior $p(\omega) = Pr[\Omega = \omega]$ is assumed to be common knowledge among the agents, but is not known to the principal. For any distribution $p(\omega)$, the prior probability that an agent endorses the answer a is:

$$Pr[a] = \int_{\omega \in \Delta(A)} Pr[a|\omega]p(\omega)d\omega$$

All answers are assumed possible, such that the prior probabilities $Pr[a]$ are bounded away from 0 by a finite amount.

If $o_i \in A$ is the private opinion of agent i , the private posterior of that agent regarding the distribution of population profiles is computed according to the Bayes' Law:

$$p(\omega|o_i) = Pr[\Omega = \omega|o_i] = \frac{Pr(o_i|\omega)p(\omega)}{Pr[o_i]}$$

and the posterior probability that some other agent believes the answer $a \in A$ can be computed as:

$$Pr[a|o_i] = \int_{\omega \in \Delta(A)} Pr[a|\omega]p(\omega|o_i)d\omega$$

We impose one more restriction on the information structure, namely that every answer a is the best predictor for itself:

$$Pr[a|a] > Pr[a|b] \forall a, b \in A \quad (1)$$

namely that the posterior belief on the probability of answer a is highest when the agent actually believes in answer a . Note, however, that Equation 1 does *not* imply

$Pr[a|a] > Pr[b|a]$; agents who endorse the answer a accept that some other answer b may be more popular among the other poll participants. This intuitive constraint is satisfied by many natural probability distributions (e.g., the Dirichlet distribution).

The principal keeps a running statistics about the answers submitted by the agents. The statistic is updated in rounds, indexed by the variable $t = 1, 2, \dots$. The agents privately report their answers to the principal (no other agents can intercept the communication between the agent and the principal), however, the principal publishes the running statistic at the end of every round. For simplicity, we assume the principal reveals the histogram of the received answers, denoted as $R^t = (R_1^t, R_2^t, \dots, R_N^t)$ where R_i^t is the number of a_i answers submitted by the previous agents in the rounds $(1, 2, \dots, t)$. Let \bar{R}^t be the normalized statistics:

$$\bar{R}_i^t = \frac{R_i^t}{\sum_{i=1}^N R_i^t}$$

We will also refer to R^t and \bar{R}^t as the *public* information or statistics available at time t . Let \mathcal{R} be the set of all possible histograms R^t . Naturally, $\Delta(A)$ is the set of all possible normalized statistics \bar{R}^t .

As compensation for the information, the principal offers payment to the agents. The payment received by an agent depends on her report and on any other information known to the principal. We consider the simplest family of payment functions, where the payment to the agent is a function $\tau(a, a^*, \bar{R}^t)$ depending on:

- the answer $a \in A$ submitted by the agent,
- the *reference* answer a^* reported by some other agent,
- the public statistics R^t of answers submitted in the previous rounds.

Note that the principal cannot condition the payments on the prior $p(\omega)$ which he does not know. We also assume that the reference report a^* is always drawn from the same round as the report a such that the agent and the *reference* agent had the same public information when they answered the question. Naturally, this implies that there are at least two agents reporting in any round t .

The problem of the principal is to design a payment scheme $\tau(\cdot, \cdot, R^t)$ that encourages the agent to reveal their private information. A reporting strategy of an agent is a mapping $s : A \times \mathcal{R} \rightarrow A$, where $s(a, R^t) \in A$ is the answer communicated to the principal when the agent believes in the answer a and the public histogram of reports is R^t . For example the honest reporting strategy is defined as:

$$\hat{s}(a, R) = a; \quad \forall a \in A, \forall R \in \mathcal{R}$$

For a simpler notation, we will drop whenever possible the dependence on R^t when expressing payments and strategies. Thus, $\tau(a, a^*, R^t) = \tau(a, a^*)$ and $s(a, R^t) = s(a)$, however, the reader should keep in mind that payments and strategies may be different for different histograms R^t .

The reporting strategy s is a weak Bayesian Nash Equilibrium (BNE) if no unilateral deviation from s increases the expected payment received by the agent. The payment expected by the agent depends on: (i) the private belief of the agent o_i , (ii) the true private belief of the agent providing the reference report, and (iii) the reporting strategy s used by the reference reporter. Formally, the expected payment V is:

$$E_{a \in A} [V] = \sum_{a \in A} Pr[a|o_i] \tau(o_i, s(a))$$

and the equilibrium condition for the strategy s becomes:

$$\sum_{a \in A} Pr[a|o_i] (\tau(o_i, s(a)) - \tau(a', s(a))) \geq 0$$

for all opinions $o_i \in A$ and deviating reports $a' \neq o_i$. If the inequality is strict, the strategy s is a strict equilibrium.

The assumptions we make in this section are standard for the literature addressing *peer prediction* mechanisms where there is no future event with a public outcome that unambiguously determines what answers are correct. An important difference, however, is that we don't allow the principal to condition the incentives mechanisms on the common knowledge prior beliefs the agents are assumed to have. We emphasize this as an important practical advantage, since the resulting mechanism is simpler, and easier to justify to the users.

3. TRUTHFUL REPORTING

A number of mechanisms [15, 9, 6], show how to design incentive compatible payment mechanisms that depend on the prior beliefs regarding the distribution of population profiles. The basic idea is that different private opinions trigger different posterior beliefs; an agent believing a will compute her expected payment according to the distribution $Pr[x|a]$ while an agent who believes in b will compute her expected payment according to the distribution $Pr[x|b]$. Since the two distributions are different, any payment derived from a scoring rule will make honest reporting a Nash equilibrium.

The problem, however, is that payments need to depend on the beliefs $Pr[x|a]$ and $Pr[x|b]$. When this dependence is not possible, our first result shows that no payment scheme can make the truthful reporting strategy a strict¹ equilibrium.

The intuition behind the proof is that the same payment function that enforces honesty under one belief $p_1(\omega)$ will also encourage a profitable deviation under another belief $p_2(\omega)$. The distributions $p_1(\omega)$ and $p_2(\omega)$ are constructed such that the posteriors $p_1(\omega|a)$ and $p_2(\omega|b)$ are equal, i.e. and agents believing the report a in the first case and b in the second case have the same posterior beliefs regarding the distribution of the reference report. Since the payment is only conditioned on the posterior belief, it may only elicit one of the answers a or b , not both; therefore, either under p_1 or under p_2 the agents have incentives to mis-report. The formal proof of the following theorem shows how to construct the two priors p_1 and p_2

¹Truthful reporting can, however, be a weak equilibrium if the principal pays a constant amount for every report.

THEOREM 1. *There is no payment function $\tau(a, a^*, R^t)$ such that the truthful reporting strategy $s^*(a, R^t) = a$, $\forall a \in A$ is a strict BNE.*

PROOF. We are looking for two probability distributions $p_1(\omega)$ and $p_2(\omega)$ over the space of possible population profiles $\omega \in \Delta(A)$ such that any payment that encourages truthful reporting under p_1 encourages lying under p_2 . We will build p_1 and p_2 from the family of Dirichlet distributions. Let p_1 be defined by:

$$p_1(\omega) = \frac{1}{B(\alpha)} \prod_{i=1}^N \omega_i^{\alpha_i - 1}$$

where $\omega = (\omega_1, \omega_2, \dots, \omega_N) \in \Delta(A)$ is the vector of frequencies for each answer, and $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_N) \in \mathbb{N}^N$, $\alpha_i > 1$ are the Dirichlet coefficients. $B(\alpha)$ is the multinomial beta function.

If $\alpha_0 = \sum_i \alpha_i$, the prior probability of answer a_i is:

$$Pr_1[a_i] = \int_{\omega \in \Delta(A)} Pr[a_i|\omega] p_1(\omega) d\omega = \frac{\alpha_i}{\alpha_0}$$

and the posterior probability of a_i given that the agent truly believes in the answer a_i is:

$$Pr_1[a_i|a_i] = \int_{\omega \in \Delta(A)} Pr[a_i|\omega] p_1(\omega|a_i) d\omega = \frac{\alpha_i + 1}{\alpha_0 + 1}$$

and the posterior probability of a_i given that the agent believes in some other answer a_j is:

$$Pr_1[a_i|a_j] = \int_{\omega \in \Delta(A)} Pr[a_i|\omega] p_1(\omega|a_j) d\omega = \frac{\alpha_i}{\alpha_0 + 1}$$

Let $\tau(\cdot, \cdot)$ be an incentive compatible payment function such that the truthful reporting strategy \hat{s} is an equilibrium. The equilibrium condition states that an agent decreases her expected pay by falsely reporting the answer b instead of a :

$$\sum_{i=1}^N Pr_1[a_i|a] (\tau(a, a_i) - \tau(b, a_i)) > 0; \quad (2)$$

Let us now consider $p_2(\omega)$ as a Dirichlet distribution with the parameters α' defined as follows:

$$\alpha'_a = \alpha_a + 1; \quad \alpha'_b = \alpha_b - 1; \quad \alpha'_c = \alpha_c \quad \forall c \in A \setminus \{a, b\}$$

Note that $Pr_1[a_i|a] = Pr_2[a_i|b]$. Therefore, Equation (2) implies that:

$$\sum_{i=1}^N Pr_2[a_i|b] (\tau(a, a_i) - \tau(b, a_i)) > 0;$$

so that an agent who believes the true answer is b will have the incentive to falsely report a under the prior belief p_2 . As the distribution p_2 was constructed independently of the payment scheme τ , it follows that no payment function can always enforce a truthful equilibrium for all prior beliefs. \square

4. HELPFUL REPORTING

Since truthful reporting cannot be guaranteed as an equilibrium, let us turn our attention to other strategies that are not *always* truthful, but are still *helpful*, in the sense that the public statistics of answers submitted on the equilibrium path converges to the true population profile. Jurca and Faltings [11] describe such a helpful equilibrium for a binary setting, where the question can have only two answers. In this section we extend our previous results to questions with an arbitrary number of answers.

The basic idea of a *helpful* strategy is that agents report truthfully only when the public statistics of the questionnaire is *close enough* to their private beliefs. However, when the public statistics is *not* close enough, the (possibly) false reports push the public statistics towards the true private beliefs. From this perspective, lying reports can still be useful as they help the public information converge towards the private beliefs.

To visualize how lying strategies can be helpful, consider a question with three possible answers, $A = \{a, b, c\}$ and a common knowledge prior over answer distributions that makes the answers a and b twice more likely than answer c : e.g., $Pr[a] = Pr[b] = 0.4, Pr[c] = 0.2$. Also consider a questionnaire that starts in round $t = 1$ with a public statistics $R^0 = (1, 1, 1)$, normalized to $\bar{R}^0 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

However unlikely, the answer c is still possible. An agent who reports in round $t = 1$ and believes in answer c has two choices:

- to report c truthfully,
- to falsely report a or b .

Although the latter alternative involves lying, the false report actually helps the public statistic get closer to the distribution the agent believes to be true. Obviously, when the public information gets close enough to the private beliefs, agents can be incentivised to report truthfully, which not only makes the public statistic accurate, but also allows the future agents to learn and refine their priors regarding the true population profile.

We therefore define a helpful reporting strategy by two constraints:

DEFINITION 1. *A helpful reporting strategy \bar{s} is defined by:*

1. *no agent reports a as long as the public probability of a is much greater than the prior probability of a , i.e.,*

$$\bar{R}_a > Pr[a] + \varepsilon_1 \Rightarrow \bar{s}(\cdot) \neq a;$$

2. *\bar{s} is truthful whenever \bar{R} is close enough to the prior probability $Pr[\cdot]$:*

$$Pr[a] + \varepsilon_1 > R^t[a] > Pr[a] - \varepsilon_2, \quad \forall a \Rightarrow \bar{s}(b, R^t) = b, \quad \forall b;$$

Definition (1) explicitly does not specify what happens when the public frequency of a certain answer is well below the private probability, i.e., $\bar{R}_a < Pr[a] - \varepsilon_2$. As we will show in the rest of this section, the two conditions above are enough to prove that (i) the public information converges to the true population profile and (ii) there are payment functions that support helpful equilibria.

First, let us show that helpful strategies exist.

THEOREM 2. *For any prior distribution $p(\omega)$ and any public statistic \bar{R} there exists a consistent strategy s that satisfies the constraints defined by Definition (1).*

PROOF. If the public statistics \bar{R} happens to be close enough to the priors $Pr[\cdot]$ such that $Pr[a] + \varepsilon_1 > \bar{R}_a^t > Pr[a] - \varepsilon_2$, the truthful strategy is trivially consistent. If, on the other hand, \bar{R} is far from $Pr[\cdot]$ we have to prove that the first constraint of Definition (1) is not over-constraining: i.e., it does not exclude all alternatives. This, again, is trivial, since \bar{R}_a cannot simultaneously be greater than $Pr[a] + \varepsilon_1$ for all a . \square

Note that helpful strategies are parametrized by the two constants ε_1 and ε_2 , but the proof of Theorem 2 does not depend on these parameters.

4.1 Convergence

Second, we will show that helpful strategies converge the public statistics of reported answers to the true population profile.

THEOREM 3. *For any population profile $\omega \in \Delta(a)$ and any prior belief $p(\omega)$ it is possible to define a helpful strategy that converges the public statistic of reports to the true population profile.*

PROOF. To prove this theorem we need to show that one can always find the parameters ε_1 and ε_2 such that the helpful strategy defined by ε_1 , ε_2 and the initial prior $p(\omega)$ makes the public statistics of answers converge to the true population profile for all ω .

We will show that a helpful strategy follows the following play pattern: As long as the public statistics \bar{R}^t is far from the prior $p(\omega)$, the agents will push the public statistics towards the prior, and within a *finite* number of rounds, the public statistic will get close enough to the prior. When this happens, the agents start reporting honestly, and the honest reports allow the other agents to refine their priors by Bayesian updating. It may happen that honest reports push again the public statistic outside the bounds of the new prior. However, a new finite sequence of helpful reports will resolve this divergence. Given enough answers, the principal obtains a statistics which converges to the true population profile.

The first step of the proof is to show that there generally exist the parameters ε_1 and ε_2 such that helpful reporting

pushes the public statistics \bar{R}^t within:

$$Pr[a] + \varepsilon_1 > \bar{R}_a^t > Pr[a] - \varepsilon_2, \forall a$$

within a finite number of rounds.

Take some small (but finite) parameter ε_1 and let $\varepsilon_2 = (N - 1) \cdot \varepsilon_1$. Every time $\bar{R}_a^t > Pr[a] + \varepsilon_1$ a helpful strategy prescribes a report other than a which triggers an update of the public statistics that decreases the frequency of a . In fact, every report different than a decreases the public frequency of a by at least:

$$\bar{R}_a^t - \bar{R}_a^{t+1} = \frac{R_a}{\sum_a R_a} - \frac{R_a}{1 + \sum_a R_a}$$

Since this decrease is finite for any finite round t , it follows that the public frequency of the answer a will decrease below $Pr[a] + \varepsilon_1$ within a finite number of rounds.

The same argument holds for all other answers in A , so that after a finite number of rounds, $\bar{R}_a^t < Pr[a] + \varepsilon_1$ for all $a \in A$. However,

$$\sum \bar{R}_a = \sum P[a] = 1$$

$$1 - \bar{R}_a = \sum_{b \neq a} \bar{R}_b < \sum_{b \neq a} P[b] + (N - 1)\varepsilon_1$$

and therefore:

$$\bar{R}_a > 1 - \sum_{b \neq a} P[b] - (N - 1)\varepsilon_1 = P[a] - (N - 1)\varepsilon_1$$

which completes the first step of our proof.

The second step of the proof is to show that the priors of the agents converge to the true population profile. The argument above shows that there is a potentially infinite sequence of rounds where the helpful strategy will dictate honest reporting. In each of these rounds, the future agents learn from the submitted reports, and update their priors with the reports published by the principal. Given a long enough sequence of Bayesian updates, the priors of the agents converge to the true population profile. \square

4.2 Helpful payment mechanisms

Theorem 3 shows that helpful strategies converge the public statistics of reports to the true population profile, and therefore, the focus of the next result is to show that the principal can design a payment mechanism that makes helpful reporting an equilibrium.

By analogy, we define the *helpful* payment mechanisms as the family of payment functions τ that support a helpful reporting Bayesian Nash Equilibrium. The key property of a helpful payment mechanism is to encourage truthful reporting when the agents' prior beliefs are close enough to the public statistics \bar{R} .

Restating the equilibrium conditions from Equation (2) the payment function τ must satisfy the following set of constraints:

$$\sum_{i=1}^N Pr[a_i|a] (\tau(a, a_i, \bar{R}^t) - \tau(b, a_i, \bar{R}^t)) > 0, \forall a \neq b; \quad (3)$$

THEOREM 4. *There always exist ε_1 and ε_2 such that helpful reporting is a BNE under the following payment function:*

$$\tau(a, b, \bar{R}) = \begin{cases} 0 & \text{if } a \neq b \\ \frac{1}{\bar{R}_a} & \text{if } a = b \end{cases} \quad (4)$$

PROOF. First, let us prove that helpful reporting is a BNE when the the public information \bar{R}^t is far from the prior $Pr[\cdot]$. We must show that no agent has the incentive to report a when $\bar{R}_a > Pr[a] + \varepsilon_1$, for some value of ε_1 that will be derived below. This follows trivially: since the payment (4) only rewards matching reports, any deviation to reporting a will have zero chance of being matched by the reference report (assumed helpful as well) and generates an expected payment of zero.

Next, let us prove that helpful reporting is an equilibrium strategy when:

$$Pr[a] + \varepsilon_1 > \bar{R}_a^t > Pr[a] - \varepsilon_2$$

Let us now consider the expected payment of an agent who believes the answer a , truthfully reports a and get's paid according to the payment function (4):

$$E[V(a, a)] = \sum_{i=1}^N Pr[a_i|a] \tau(a, a_i, \bar{R}) = \frac{Pr[a|a]}{\bar{R}_a}$$

and

$$\frac{Pr[a|a]}{\bar{R}_a} > \frac{Pr[a|a]}{Pr[a] + \varepsilon_1}$$

On the other hand, and agent who believes a but falsely reports b expects a payment:

$$E[V(a, b)] = \sum_{i=1}^N Pr[a_i|a] \tau(b, a_i, \bar{R}) = \frac{Pr[b|a]}{\bar{R}_b} < \frac{Pr[b|a]}{Pr[b] - \varepsilon_2}$$

We will show that there always exist the parameters ε_1 and $\varepsilon_2 = (N - 1)\varepsilon_1$ such that:

$$\frac{Pr[a|a]}{Pr[a] + \varepsilon_1} > \frac{Pr[b|a]}{Pr[b] - \varepsilon_2}$$

The inequality above can be transformed as:

$$Pr[a|a]Pr[b] - Pr[b|a]Pr[a] > Pr[b|a]\varepsilon_1 + Pr[a|a]\varepsilon_2 \quad (5)$$

and by replacing ε_2 we obtain:

$$\varepsilon_1 (Pr[b|a] + (N - 1)Pr[a|a]) < Pr[a|a]Pr[b] - Pr[b|a]Pr[a]$$

and

$$\varepsilon_1 < \frac{Pr[a|a]Pr[b] - Pr[b|a]Pr[a]}{Pr[b|a] + (N - 1)Pr[a|a]} \quad (6)$$

However, $Pr[b|a]Pr[a] = Pr[a|b]Pr[b]$ by Bayes' Law and since our assumptions require $Pr[a|b] < Pr[a|a]$, we have

$$Pr[a|a]Pr[b] - Pr[b|a]Pr[a] > 0, \quad \forall a, b \in A$$

which means that one can always find ε_1 and ε_2 to satisfy the inequality (5). Hence the payment mechanism described by this theorem supports a helpful reporting equilibrium. \square

Note that Theorem 4, just like Theorem 2, does not specify the equilibrium behavior when the public information is far from the prior belief. The only constraint on the reporting strategy is not to report an answer that is already over-represented in the public statistics.

While many strategies will be in equilibrium, two obvious choice are:

- The report that generates the highest expected payment:

$$s(a, \bar{R}) = \arg \max_b \left(\frac{1}{\bar{R}_b} \right)$$

- The report which bridges the biggest relative gap between the private posterior belief of the agent and the public statistics:

$$s(a, \bar{R}) = b^* = \arg \max_b \left(\frac{Pr[b|a]}{\bar{R}_b} \right)$$

This latter strategy is particularly interesting because it is also a best response when a certain subset of agents are altruistic and always report the truth:

- if the reference report comes from a rational agents following the helpful strategy (i.e., also reporting b^*) the expected payment is:

$$\frac{1}{\bar{R}_a}$$

- if the reference report is honest the expected payment is:

$$\frac{Pr[b^*|a]}{\bar{R}_{b^*}}$$

and neither can be improved by reporting something else than b^* .

5. NUMERICAL EXAMPLE

Let us illustrate the payment function and the helpful equilibrium for the numerical example presented at the beginning of Section 4. Consider a question with three possible answers, $A = \{a, b, c\}$ and a common prior belief about population profiles characterized by a Dirichlet distribution with the parameters $\alpha = (2, 2, 1)$. Therefore, the prior probabilities of the answers a , b and c are:

$$Pr[a] = \frac{2}{5} = 0.4; \quad Pr[b] = 0.4; \quad Pr[c] = \frac{1}{5} = 0.2$$

the posterior probabilities are:

$$Pr[a|a] = 3/6, \quad Pr[b|a] = 2/6, \quad Pr[c|a] = 1/6$$

$$Pr[a|b] = 2/6, \quad Pr[b|b] = 3/6, \quad Pr[c|b] = 1/6$$

$$Pr[a|c] = 2/6, \quad Pr[b|c] = 2/6, \quad Pr[c|c] = 2/6$$

and the public statistics before the beginning of round 1 is:

$$\bar{R}_a^0 = 1/3, \quad \bar{R}_b^0 = 1/3, \quad \bar{R}_c^0 = 1/3$$

In round 1 the principal defines the payment mechanism as in Theorem 4:

$$\tau(a, a) = 3, \tau(b, b) = 3, \tau(c, c) = 3, \tau(a_i, a_j) = 0 \forall a_i \neq a_j$$

And the helpful equilibrium requires truthful reporting if and only if:

$$Pr[x] + \varepsilon_1 > \bar{R}_x > Pr[x] - \varepsilon_2 \forall x \in A$$

The parameters ε_1 and ε_2 must satisfy the constraints (6) which become:

$$\begin{aligned} \varepsilon_1 &< \frac{Pr[a|a]Pr[b] - Pr[b|a]Pr[a]}{Pr[b|a] + (N-1)Pr[a|a]} = 0.05 \\ \varepsilon_2 &= (N-1)\varepsilon_1 < 0.1 \end{aligned}$$

The helpful reporting strategy requires the agent to report either a or b regardless of their opinion.

6. DISCUSSION

One limitation of the mechanism we present in this paper is that it is vulnerable to collusion. Since payments are awarded for matching reports, the agents have incentives to coordinate their reports on the same value. As long as the public statistics is far from the prior beliefs, collusive behavior is actually consistent with helpful reporting: the most profitable collusion strategy is:

$$s(a, \bar{R}) = b' = \arg \max_b \frac{1}{\bar{R}_b}$$

which does satisfy the constraints of Definition 1 since $\bar{R}_{b'} > Pr[b']$ for the report b' which corresponds to the smallest value in the probability distribution \bar{R} .

Collusion, however, is an important problem once the public information is sufficiently close to the prior beliefs and helpful reporting prescribes honest reporting. A complete solution to collusive behavior will be addressed in our future work; several directions we are currently investigating are:

- randomized mechanisms, that make it hard for an agent to predict which reference report will be used to compute her payment
- statistical methods that detect collusive patterns based on the distribution of reports submitted in previous rounds
- use the fact that in many practical systems many users will be altruistic and will report honestly regardless of the incentive mechanism
- use more than one reference report and leverage the positive results presented in [9] to automatically compute a payment mechanism that is resistant to collusion.

One practical property of prediction markets is that the subsidies to the market are bounded, such that the principal knows exactly what is the worst case price he has to pay for the information elicited through the market. The Bayesian Truth Serum (BTS) [17] method provides even

stronger guarantees by being budget balanced. We believe, however, that budget-balance is not a desirable property for online polls and questionnaires. The total payments break even only when some agents lose money. The same would be true for prediction markets, despite the subsidy of the market owner. While both prediction markets and the BTS mechanism are ex-ante individually rational (i.e., agents do not expect to lose money by entering the mechanism), they are not ex-post individually rational (IR).

We see two problems with reward mechanisms that are not ex-post IR. First, the negative payoffs require some mechanism to collect payments from the participants. The collection should happen before participation, because the nature of the internet makes it impossible to later track down agents and extract payments. While entry fees might seem reasonable for a market, we cannot imagine an opinion poll where agents pay to submit their opinion.

Second, the possibility of losing money may also deter participation. Opinion polls depend on a large number of agents expressing their opinions, so we strongly believe that every participant should leave the poll with some reward. Under these conditions, budget balance is not a feasible property.

Our mechanism can, however, guarantee upper bounds on the budget spent by the principal. The worst case payment made for one report is $\max_b \frac{1}{\bar{R}_b}$ which although big, is still finite for any probability distribution \bar{R} with a full support (e.g., every answer has at least one count in the public statistic). Therefore, by limiting the number of participants, the principal automatically sets limits on the maximum payment he has to make to the agents.

A second alternative is to scale down the rewards as more information becomes available. The marginal information contributed by every new opinion decreases, hence it is also natural to decrease the reward for later opinions. The rewards defined by Theorem 4 can be scaled by a constant δ that decreases exponentially with the number of already submitted opinions. Even when the poll runs indefinitely, the payments made by the principal will converge to zero, and the total budget will be bounded.

Yet a third alternative to limit the budget is to forbid reports on the answers that are very unlikely. The proof of Theorem 3 shows that agents must at some point report the answer a if they believe the probability $Pr[a]$ is greater than the public statistics \bar{R}_a . If the public probability of a stays low for a long sequence of reports, it must be that the agents do not believe in the answer a , and hence the principal can take the answer out of the poll. The description of the exact mechanism which allows the principal to modify the set of possible answers remains for future work.

Another disadvantage of our mechanism is the speed of convergence. When compared to a simple poll mechanism (i.e., no payment mechanisms, agents are assumed to report truthfully) the helpful reporting equilibrium converges slower: the agents update their private belief only with the reports submitted in the rounds where the public information approximates well enough the prior beliefs. This potentially happens only once every k rounds, hence the slower progress.

When compared to a prediction market, the convergence of our mechanism is even slower. The participants in a prediction market can trade multiple shares and thus instantly set the market price of an alternative to the probability they privately believe to be true.

Peer prediction methods can be modified to also elicit probability distributions over the answers. For example, the BTS mechanism asks all agents to report their expectation of the final result of the poll, and uses a divergence-based payment to incentivize honest reporting. While the elicitation of the full probability distribution over the answers may be practically prohibitive, we are planning to investigate simpler alternatives, where, for example, the agent also indicates by how much the public information is wrong. A concrete mechanism with this property will be addressed in future work.

One last point we would like to stress is that the rewards need not be monetary payments. Status or bonus points, preferential QoS, lottery tickets or any other assets that users value can also be used. Effective micro-payment systems are still hard to implement, fortunately however, users seem to care enough about virtual currencies [19] to make the implementation of such reward mechanisms feasible.

7. CONCLUSION

Obtaining and aggregating the private information of individual agents has the potential to greatly improve the decision process across a wide range of domains. Markets proved very efficient for extracting predictions about claims and facts that can be precisely verified in the near future. All other non-verifiable information, such as implications of alternative policies, long term effects, subjective or artistic impressions can only be elicited through opinion polls and questionnaires.

This paper addresses the reporting incentives for rational agents in such scenarios. We impose a practical constraint on the mechanism by requiring the rewards be independent of the prior beliefs of the agents. In a first result we show that prior-independent payments are not able to support a strict truthful equilibrium in all settings. The main contribution of the paper, however, is to describe a reward scheme that supports a *helpful* equilibrium, where occasional misreports still let the principal derive accurate information from the poll.

Despite the limitations mentioned in Section 6 we believe our work makes an important step towards the practical acceptance of reward mechanisms when eliciting information from the *crowds*. First, the reward mechanism is very simple, and does not require the users to understand how a market or a scoring rule functions. Second, its correctness does not depend on a principal making correct guesses about the beliefs of the agents. Third, helpful strategies are intuitive to explain and coexist well with altruistic reporting. Finally, the mechanism can be readily integrated into existing frameworks, and does not require significant changes to the current processes used to elicit feedback online.

8. REFERENCES

- [1] K. Chen, L. Fine, and B. Huberman. Forecasting Uncertain Events with Small Groups. In *Proceedings of ACM Conference of Electronic Commerce (EC'01)*, 2001.
- [2] K. Chen and C. Plott. Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem. California Institute of Technology Social Science Working Paper No. 1131, 2002.
- [3] Y. Chen and D. Pennock. Designing Markets for Prediction. *AI Magazine*, 31(4):42, 2010.
- [4] V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Proceedings of the Uncertainty in Artificial Intelligence Conference (UAI)*, 2002.
- [5] R. Forsythe, T. Rietz, and T. Ross. Wishes, expectations and actions: Price formation in election stock markets. *Journal of Economic Behavior and Organization*, 39(1):83–110, 1999.
- [6] S. Goel, D. Reeves, and D. Pennock. Collective Revelation: A Mechanism for Self-Verified, Weighted, and Truthful Predictions. In *Proceedings of the ACM Conference on E-Commerce*, pages 265–273, Stanford, USA, 2009.
- [7] R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *The Journal of Prediction Markets*, 1(1):3–15, 2007.
- [8] R. Jurca and B. Faltings. Minimum Payments that Reward Honest Reputation Feedback. In *Proceedings of the ACM Conference on Electronic Commerce (EC'06)*, pages 190–199, Ann Arbor, Michigan, USA, June 11–15 2006.
- [9] R. Jurca and B. Faltings. Collusion Resistant, Incentive Compatible Feedback Payments. In *Proceedings of the ACM Conference on Electronic Commerce (EC'07)*, pages 200–209, San Diego, USA, June 11–15 2007.
- [10] R. Jurca and B. Faltings. Robust incentive-compatible feedback payments. In *Agent-mediated electronic commerce: automated negotiation and strategy design for electronic markets; AAMAS 2006 workshop, TADA/AMEC 2006, Hakodate, Japan, May 9, 2006; selected and revised papers*, pages 204 – 218. Springer-Verlag New York Inc, 2007.
- [11] R. Jurca and B. Faltings. Incentives For Expressing Opinions in Online Polls. In *Proceedings of the ACM Conference on E-Commerce*, pages 119–128, Chicago, USA, 2008.
- [12] N. Lambert, D. Pennock, and Y. Shoham. Eliciting properties of probability distributions. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 129–138. ACM, 2008.
- [13] N. Lambert and Y. Shoham. Truthful Surveys. *Lecture Notes in Computer Science*, 5385:154–165, 2008.
- [14] N. Lambert and Y. Shoham. Eliciting Truthful Answers to Multiple-Choice Questions. In *Proceedings of the tenth ACM conference on electronic commerce*, pages 109–118, 2009.
- [15] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science*, 51:1359 –1373, 2005.
- [16] D. Pennock, S. Lawrence, F. Nielsen, and C. Lee Giles. Extracting collective probabilistic forecasts from web

- games. In *Proceedings of Conference on Knowledge Discovery and Data Mining*, 2001.
- [17] D. Prelec. A bayesian truth serum for subjective data. *Science*, 306(5695):462–466, 2004.
- [18] L. J. Savage. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- [19] E. Servan-Schreiber, J. Wolfers, D. Pennock, and B. Galebach. Prediction markets: Does money matter. *Electronic Markets*, 14(3), 2004.
- [20] A. Zohar and J. Rosenschein. Robust Mechanisms for Information Elicitation. In *Proceedings of the AAAI*, 2006.