

# Incentives for Expressing Opinions in Online Polls

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 efficiently extract and aggregate the private information held by individuals about events and facts that can be publicly verified. However, facts such as the effects of raising or lowering interest rates can never be publicly verified, since only one option will be implemented. Online opinion polls can still be used to extract and aggregate private information about such questions. This paper addresses incentives for truthful reporting in online opinion polls. The challenge lies in designing reward schemes that do not require a-priori knowledge of the participants' beliefs. We survey existing solutions, analyze their practicality and propose a new mechanism that extracts accurate information from rational participants.

## Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence

## General Terms

Algorithms, Design, Economics

## Keywords

opinion polls, mechanism design, incentive compatibility

## 1. INTRODUCTION

Prediction markets have become efficient tools for extracting and aggregating the local private information detained by different individual agents. They function like normal markets where the security that is traded depends on the realization of a specific future event. Typical prediction markets trade securities related to the result of presidential elections, to the earnings of movies, or to the outcome of sport competitions.

\*This work was done in the Artificial Intelligence Lab, EPFL

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'08, July 8–12, 2008, Chicago, Illinois, USA.

Copyright 2008 ACM 978-1-60558-169-9/08/07 ...\$5.00.

The basic idea behind such information markets is that the pricing scheme of the market encourages the participants to buy or sell shares in accordance to their private information. Consider, for example, the security which pays \$1 if the winner of the 2008 US presidential elections is democrat. A rational agent will find it profitable to buy such securities as long as the price she pays for one share is less than  $p$ , the value of her private belief that the United States will be lead by a democrat president after the 2008 election. When several agents participate in the market, the final price of the security will reflect an aggregate *public* value of the probability of a democrat winner, despite the fact that information has only been *privately* available before the trading.

Numerous experiments have shown that the price in prediction markets converges to reflect the true information available to the agents [2], [10], [3], [19], despite game theoretic negative results (e.g., the *no trade* theorem) [16] which predict an equilibrium behavior where rational agents do not trade. Moreover, prediction markets have been known to consistently outperform other traditional prediction tools [9], [23], [18], and are empirically hard to manipulate [25], [22], [1], [20].

One important requirement to run a prediction market is to be able to define a security related to information that can be unambiguously and publicly verified on a precise moment in the future. Even for the information that most trivially satisfies this condition, the need for accuracy can generate significant complications. For example, Foresight Exchange<sup>1</sup> defined the security BUSH04 in the following way:

*G. W. Bush, the president of the United States at the time this claim started trading, will still be president on 2005-02-01 (after the inauguration after the election is usually scheduled). This claim will be TRUE even if elections are postponed or G.W. Bush remains in power by staging a coup. If there are events which make it confusing who the US president is, as of 2005-02-01, this claim is true if G.W. Bush is leading a sovereign government in at least part of the territory of the United States of America (as of 2001-01-01) that has recognition of at least one of the UN Security Council permanent members (Britain, France, China and Russia) other than the United States.*

However, there are many other cases where it would be invaluable to obtain private information about claims that can not be verified. An example of such a question is the effect of raising or lowering central bank interest rates on economic growth and inflation. Since only one policy can be implemented, a prediction market for different alterna-

<sup>1</sup><http://www.ideosphere.com/fx/index.html>

tives will include many securities whose underlying event cannot be publicly verified. An accurate picture of how the economy, or even just the financial markets, would react to different options would be extremely helpful to policymakers. Other examples of such problems are predicting the success of hypothetical products or fashion elements.

Private information about non-verifiable claims can be obtained through opinion polls. Traditional polls conducted by specialized interviewers are usually accurate, but very expensive. On the other hand, online opinion polls are trivial to set up, but lack both verification methods and honest participation incentives for the reporters. Without such incentives, (a) the altruistic reports will get diluted between an ever increasing number of polls making most of the results statistically insignificant, and (b) the few results that are significant can not be guaranteed to be accurate. Furthermore, for questions such as interest rates there are incentives for respondents to manipulate the outcome and get a certain policy implemented, and these must also be overcome.

The goal of this paper is to address reporting incentives for such online opinion polls. The basic idea is to encourage honest reporting by explicit rewards that depend on the opinions of other peers. Intuitively, every report is *benchmarked* against the opinion of the group, and rewarded such that honesty, but not conformity, is optimal.

First we will briefly survey existing incentive schemes that were initially designed for reputation mechanisms, but can be easily adapted to online opinion polls. The main drawback, however, of these mechanisms is that they either cost too much, or require unpractical prior knowledge in the design process. Second, we will investigate the mechanism based on *Bayesian Truth Serum* [21] which rewards more the opinions that turn out to be surprisingly common. Third, we present a novel mechanism that (a) is more suitable for an online process where present reporters can see the opinions expressed by previous users, and (b) requires participants to submit less information and, therefore, is easier to use.

## 2. THE SETTING

We assume a setting where a decision maker is organizing an online opinion poll on a multiple choice question. For the moment we restrict our attention to binary outcomes where every participating agent can answer the poll question with *yes* or *no*. These two answers will also be called the *positive*, respectively the *negative* reports. For example, the poll can be around questions like:

- “Are you afraid of global warming?”, or
- “Would you prefer white over red wine?”, or
- “Is the world economy going through a recession?”.

Whenever possible, we will give directions for extending our results to the general case where the poll question can have several possible answers.

Participating agents are rational, they have a private opinion about the poll question, and are assumed not to collude: i.e., every agent answers the poll at most once, and different agents do not communicate among themselves before reporting. The result of the poll is the (running) average of all votes, computed as the fraction of the participants who endorsed the positive alternative. Note that all opinions are equally weighted in the final outcome of the poll.

We model the private opinion of an agent as a private binary signal the agent receives from nature regarding the poll question. Let  $s_i \in \{0, 1\}$  be the signal received by agent  $i$ , where  $s_i = 0$  indicates that agent  $i$  endorses the negative answer, while  $s_i = 1$  indicates that the agent endorses the positive answer. Agents do not know the private signal of other agents, however, they all share a common belief regarding the prior distribution of preferences in the population. Let  $\Omega$  be a random variable expressing the true distribution of preferences of the agents. Let  $p(\omega) = Pr[\Omega = \omega]$  describe the common knowledge a priori belief that a fraction  $\omega \in [0, 1]$  of the population will endorse the positive answer. All agents are rational and Bayesian, and therefore their private signal influences their expectation regarding the posterior distribution of preferences. For example,

- an agent endorsing the positive answer would believe that the preferences of other agents are distributed according to:

$$\begin{aligned} p(\omega|1) &= \frac{p(1|\omega)p(\omega)}{\int_{\omega'=0}^1 p(1|\omega')p(\omega')d\omega'} \\ &= \frac{\omega p(\omega)}{\int_{\omega'=0}^1 \omega' p(\omega')d\omega'}; \end{aligned} \quad (1)$$

- an agent endorsing the negative answer would believe that the preferences of other agents are distributed according to:

$$\begin{aligned} p(\omega|0) &= \frac{p(0|\omega)p(\omega)}{\int_{\omega'=0}^1 p(0|\omega')p(\omega')d\omega'} \\ &= \frac{(1-\omega)p(\omega)}{\int_{\omega'=0}^1 (1-\omega')p(\omega')d\omega'}; \end{aligned} \quad (2)$$

The assumption that agents with different opinions assess differently the frequency of their own answer is not only a consequence of Bayesian theory. Numerous lab experiments confirmed that people expect to be “typical” and therefore overestimate the popularity of their own choices [21].

The correlation between an agent’s private information and private beliefs allows the design of incentive compatible schemes that reward truth-telling. By conditioning the reward of an agent on the reports of other peers, agents with different private information will compute differently their expected outcome, as the probability distribution for the reports of the peers depends on the private information. This difference can be just enough such that every agent believes that she is maximizing her expected reward by declaring the truth. An example of how this mechanism works will follow in the next section.

We define a reward mechanism by a payment function  $\tau(\cdot, \cdot) \in \mathbb{R}^+$ . Let  $\tau(r_i, r_{-i})$  be the reward to agent  $i$  when her information submitted to the poll is  $r_i$  and the information submitted by all other participants is  $r_{-i}$  (according to standard notation,  $\{-i\}$  denotes the set of all other agents except  $i$ ). For now we are using a generic notation for the information reported by an agent. In the simplest case,  $r_i \in \{0, 1\}$  is the binary answer to the poll question. Nevertheless,  $r_i$  might include additional information like the user’s estimate for the outcome of the poll (e.g., used by the Bayesian Truth Serum [21]). The rewards “paid” to the agents need not be monetary, and can represent virtual

points, in kind services, or any other benefits online users put value to.

The goal of the designer is to define the function  $\tau$  such that all agents find it rational to report truthfully. Formally, given that  $r_i(1)$  and  $r_i(0)$  are the information reported by agent  $i$  to reveal the positive, respectively the negative answer to the poll question, the incentive compatibility constraints on the reward scheme  $\tau$  are the following:

$$\begin{aligned} \sum_{r_{-i}} Pr[r_{-i}|0] \left( \tau(r_i(0), r_{-i}) - \tau(r_i(1), r_{-i}) \right) &> 0; \\ \sum_{r_{-i}} Pr[r_{-i}|1] \left( \tau(r_i(1), r_{-i}) - \tau(r_i(0), r_{-i}) \right) &> 0; \end{aligned} \quad (3)$$

$Pr[r_{-i}|0]$  and  $Pr[r_{-i}|1]$  are the probabilities that the information reported by other agents is exactly  $r_{-i}$  when agent  $i$  privately endorses the positive, respectively the negative answer to the poll question. These probabilities depend on the conditional posterior beliefs  $p(\omega|0)$  and  $p(\omega|1)$  detailed in Eq. (1) and (2) respectively. The two inequalities can be satisfied simultaneously because the probability distributions  $Pr[r_{-i}|0]$  and  $Pr[r_{-i}|1]$  are different. Unfortunately, the mechanism designer cannot exactly compute  $Pr[r_{-i}|0]$  and  $Pr[r_{-i}|1]$  as the prior distribution  $p(\omega)$  is not known.

One last clarification regards to use of the term *online*. So far we have used it to indicate that the opinion poll is hosted on the internet where agents can answer it from the comfort of their home. From this point onwards, the term online will be used exclusively to characterize the process governing the poll. An *online poll* is updated in (almost) real-time, and publicly displays some summary (typically the fraction of positive answers) of the information reported by the previous users. By contrast, an *offline poll* will display the results only at the end, after the submission process has been ended.

### 3. INCENTIVE-COMPATIBLE REWARDS FOR OPINION POLLS

Fundamental results in the mechanism design literature [8, 7] show that side payments can be designed to create the incentive for agents to reveal their private opinions truthfully. Such payment schemes have been constructed based on *proper scoring rules* [15, 11, 4], and exploit the correlation between the private signals and the private beliefs of an agent. The first adaptation of these results to general feedback reporting mechanisms is due to [17], and functions as follows.

A central authority, the poll sponsor in our case, *scores* every answer by comparing it with another report (called the *reference* report) submitted by a different user. The score reflects the value of a proper scoring rule, computed for the values of the two reports. For example, if  $r_{ref(i)}$  denotes the reference report of  $r_i$ , the score of  $r_i$  can be directly proportional to  $\log(Pr[r_{ref(i)}|r_i])$ , the value of the logarithmic scoring rule for the point represented by the conditional posterior probability of the reference report, given  $r_i$ . All proper scoring rules have the property that they maximize the expected score when the report  $r_i$  honestly reflects the agent's private information.

Jurca and Faltings modified the mechanism proposed by Miller et al. in several important aspects. First, they used the technique of *automated mechanism design* [6] to develop

novel *adaptive scoring functions* that have the property that they minimize the budget required to pay the reports [12]. Second, they extended the scoring functions to include several reference reports and a filtering mechanism that probabilistically eliminates some of the false reports. Third, they show that yet another family of scoring functions can generate robust mechanisms where honest reporting is the only equilibrium [13].

To get a practical feel of how these different mechanisms work, let us consider a simple numerical example. The poll asks the participants to assess whether or not without cuts in interest rates, the US economy would enter a recession this year. Such a question could be answered in a prediction market by trading securities conditional on the interests rates being lowered. Nevertheless, the security must define precisely the meaning of *recession*, a task which we believe to be complicated enough to justify the search for alternative information elicitation mechanisms.

This poll closely follows the Davos World Economic Forum, where world leaders have already expressed their views on this issue. These public declarations create a prior belief regarding the frequency of the positive answer, which, for simplicity, is assumed to follow a beta distribution:  $p(\omega) = \frac{\omega^\alpha(1-\omega)^\beta}{B(\alpha,\beta)}$ . To reflect the divergence of opinions slightly skewed towards the positive answer, let  $\alpha = 6$  and  $\beta = 3$ .

A priori, the expected outcome of the poll is  $E[\omega] = \frac{\alpha}{\alpha+\beta} = 2/3$ . An agent who is inclined to believe in an upcoming economic recession will update accordingly her expectation of the final outcome of the poll. From her point of view the frequency of the positive answer follows a beta distribution with parameters  $\alpha = 7$  and  $\beta = 3$ , so the pessimistic agent expects the poll to close with 70% of the participants endorsing the positive answer.

The optimistic agent, on the other hand, who believes the economy suffers only from a temporary slowdown, will expect the frequency of the positive answer to follow a beta distribution with  $\alpha = 6$  and  $\beta = 4$ . Consequently, she expects the output of the poll to be around 40%. Notice that both agents expect their own opinion to be more popular, although they both accept that the majority will endorse the positive answer.

The following reward mechanisms can create honest reporting incentives:

- by using the logarithmic scoring rule as described by Miller et al.,  $\tau(r_i, r_{ref(i)}) = a \cdot \log(Pr[r_{ref(i)}|r_i]) + b$  where  $a$  and  $b$  are two appropriate constants. Given that  $Pr[1|1] = 0.7$  and  $Pr[1|0] = 0.6$  the payments could be:

$\tau(\cdot, \cdot)$	$r_{ref} = 0$	$r_{ref} = 1$
$r = 0$	1.3	3.2
$r = 1$	0	3.9

and the incentive compatibility constraints are satisfied because:

$$\begin{aligned} 0.4\tau(0,0) + 0.6\tau(0,1) &= 2.44 > 0.4\tau(1,0) + 0.6\tau(1,1) = 2.34; \\ 0.3\tau(1,0) + 0.7\tau(1,1) &= 2.73 > 0.3\tau(0,0) + 0.7\tau(0,1) = 2.63; \end{aligned}$$

- using adaptive scoring rules computed through automated mechanism design, much cheaper payments can guarantee the same margins for reporting honestly. For example, if:

$\tau(\cdot, \cdot)$	$r_{ref} = 0$	$r_{ref} = 1$
$r = 0$	1.3	0
$r = 1$	0	0.7

reporting the truth is better than lying by 0.1 units, while the expected payment to a reporter has gone down to roughly 0.5:

$$0.4\tau(0, 0) = 0.52 > 0.6\tau(1, 1) = 0.42;$$

$$0.7\tau(1, 1) = 0.49 > 0.3\tau(0, 0) = 0.39;$$

Considering several reference reports can further decrease the cost of the rewards.

The mechanisms based on scoring rules work for both on-line and offline process. In an offline setting, the poll operator computes the reward scheme once, keeps all submitted reports secret, and distributes the payments after the close of the poll, by choosing randomly the reference report(s). In an online setting, the poll operator has two choices:

1. publish every report as soon as it gets submitted. The reward mechanism has to be recomputed every time, to reflect the new information gained through the last submission. The immediately following report(s) are used as reference to compute the payments.
2. publish batches of reports, and update the reward mechanism offered to the next batch based on the results of the previous batch. Reference reports are always chosen from the same batch;

The second alternative establishes a compromise between updating speed and computational complexity. With more reports in a batch, the delay for releasing the information to the public increases, but the frequency of recomputing the reward mechanism decreases.

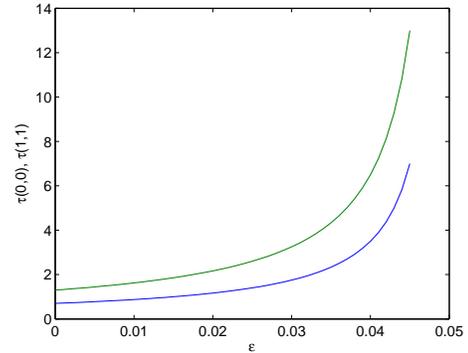
One big disadvantage of the mechanisms presented above is that they require precise knowledge about the prior beliefs of the agents. If the mechanism designer does not know the prior distribution  $p(\omega)$ , the payments cannot be computed. One possible workaround is to construct payment systems that are incentive-compatible for a range of prior beliefs. For example, assume that the mechanism designer does not have precise estimates for  $Pr[1|1]$  and  $Pr[1|0]$ , but believes they must be in the ranges:  $(0.7 \pm \epsilon)$  and  $(0.6 \pm \epsilon)$  respectively. The following payments guarantee a truthtelling margin of 0.1 even when the probabilities  $Pr[1|1]$  and  $Pr[1|0]$  take the worst possible values for an uncertainty level characterized by  $\epsilon = 0.02$ :

$\tau(\cdot, \cdot)$	$r_{ref} = 0$	$r_{ref} = 1$
$r = 0$	2.1	0
$r = 1$	0	1.1

For increasing levels of uncertainty, the rewards become larger and larger, as shown in Figure 1. The general algorithm for designing incentive compatible payments under uncertainty is described in [14], and confirms the same drastic increase in cost as the designer is less certain about the private beliefs of the reporters.

### 3.1 The Bayesian Truth Serum

The Bayesian Truth Serum (BTS) [21] establishes an incentive compatible reward mechanism that does not depend on knowledge about the prior beliefs. The agents are still



**Figure 1: Incentive compatible rewards for increased uncertainty.**

assumed to have common prior beliefs, but these beliefs are not reflected in the design of the mechanism.

BTS works by asking reporters supplementary information. Every agent submits her subjective answer to the poll question, and also reports an estimate of the final distribution over the possible answers. The reward received by each agent takes into account the given answer, the estimated outcome of the poll, the final outcome of the poll, and an average of the predicted outcomes of the poll.

For the example in the previous section, BTS would generate the following rewards to agent  $i$ :

- $\log \frac{\omega^*}{\bar{e}} + KL(\omega^*, e_i)$  if the agent endorses the positive answer, and
- $\log \frac{1-\omega^*}{1-\bar{e}} + KL(\omega^*, e_i)$  if the agent endorses the negative answer;

$\omega^*$  is the actual percentage of respondents endorsing the positive answer,  $e_i$  is the estimate of  $\omega^*$  reported by agent  $i$ ,  $\bar{e}$  is the geometric average of all  $e_i$ 's, and  $KL(\omega^*, e_i) = \omega^* \log \frac{e_i}{\omega^*} + (1-\omega^*) \log \frac{1-e_i}{1-\omega^*}$  is the Kullback-Leibler distance between the actual distribution of reports, and the distribution predicted by agent  $i$ . [21] proves that honest answering is a pareto-optimal Nash Equilibrium of the mechanism.

### 3.2 A framework for improving BTS

Original incentive schemes based on proper scoring rules were significantly improved by automated mechanism design, as explained in the beginning of this section. The basic idea is to define adaptive scoring functions that are optimal for a given context (i.e., given information structure). The same idea can be applied to improve BTS.

Formally, BTS elicits from a poll participant (i.e., agent  $i$ ) the tuple  $r_i = (o_i, e_i)$  where  $o_i \in \{0, 1\}$  is the binary answer to the question, and  $e_i \in [0, 1]$  is an estimation of the final outcome of the poll (i.e., fraction of the respondents endorsing the positive answer). An agent reports truthfully when  $o_i$  equals her private signal  $s_i$ , and when the expected outcome  $e_i$  equals the posterior expected probability of the positive answer:

$$e_i = Pr[1|s_i] = \int_{\omega'=0}^1 \omega' p(\omega'|s_i) d\omega';$$

Not all agents are expected to accurately compute and report  $e_i$ . However, given that sufficient agents participate, individual errors made in the reports  $e_i$  will cancel out, and some aggregate of the reports  $e_i$  will be equal to the same aggregate computed over the errorless reports. Let  $\bar{e}$  denote this aggregate which, for example, could be the arithmetic or the geometric average. When  $\bar{e}$  is the arithmetic average  $\frac{\sum_{i=1}^N e_i}{N}$ , the assumption of canceling errors implies that:

$$\begin{aligned}\bar{e} &= \lim_{N \rightarrow \infty} \frac{\sum_{i=1}^N e_i}{N} = \omega^* Pr[1|1] + (1 - \omega^*) Pr[1|0] \\ &= \omega^* \int_{\omega'=0}^1 \omega' p(\omega'|1) d\omega' + (1 - \omega^*) \int_{\omega'=0}^1 \omega' p(\omega'|0) d\omega'\end{aligned}$$

where  $\omega^*$  is the true frequency of the positive answer in the population.

For large enough number of reporters, the information  $r_{-i}$  truthfully reported by all other agents except  $i$  can be summarized by  $\omega^* \in [0, 1]$  and  $\bar{e} \in [0, 1]$ , where again,  $\omega^*$  is the true frequency of the positive answer, and  $\bar{e}$  is the average of the expected frequencies for the positive answer. The payment mechanism can now be specified by the function

$$\tau : \{0, 1\} \times [0, 1] \times [0, 1] \times [0, 1] \rightarrow \mathbb{R}^+$$

where  $\tau(o_i, e_i, \omega^*, \bar{e})$  is the payment received by the agent  $i$  given that:

- she answers the question with  $o_i \in \{0, 1\}$
- she predicts the outcome of the poll (i.e., fraction of positive answers) will be  $e_i$ ;
- the actual outcome of the poll is  $\omega^*$
- the average expected outcome of the poll is  $\bar{e}$

A rational, risk-neutral agent has the incentive to report honestly (given that all other agents report honestly) if and only if her expected payment for reporting the truth is higher than the expected payment for lying by some margin  $\Delta$ :

$$\int_{\omega=0}^1 p(\omega|s_i) (\tau(s_i, Pr[1|s_i], \omega, \bar{e}) - \tau(\neg s_i, e_i, \omega, \bar{e})) > \Delta; \quad (4)$$

where  $\neg s_i$  is the binary opposite of  $s_i$  and  $e_i \neq Pr[1|s_i]$  is a false report about the expected outcome of the poll. The inequality must hold for all  $s_i \in \{0, 1\}$ , for all  $e_i \neq Pr[1|s_i]$ , and for all prior probability distributions  $p(\omega)$ .

A slight relaxation of the constraints in (4) is to assume that agents truthfully report their expected outcome of the poll,  $e_i^2$ . The payment function becomes  $\tau(s_i, \omega^*, \bar{e})$ , and without loss of generality we can make the notation:

$$g(\omega^*, \bar{e}) = \tau(1, \omega^*, \bar{e}) - \tau(0, \omega^*, \bar{e})$$

The constraints on honest reporting incentives thus become:

$$\begin{aligned}\int_{\omega=0}^1 p(\omega|1) g(\omega, \bar{e}) &> \Delta; \\ \int_{\omega=0}^1 p(\omega|0) g(\omega, \bar{e}) &< -\Delta;\end{aligned} \quad (5)$$

Hence the design problem can be summarized by the following:

<sup>2</sup>This assumption can be easily replaced by an independent payment that penalizes the distance between  $e_i$  and the actual outcome of the poll. This is the approach taken in BTS

PROBLEM 1. Find a function  $g : [0, 1] \times [0, 1] \rightarrow \mathbb{R}$  such that:

$$\int_{\omega=0}^1 p(\omega|1) g(\omega, \bar{e}(\omega)) d\omega > \Delta$$

$$\int_{\omega=0}^1 p(\omega|0) g(\omega, \bar{e}(\omega)) d\omega < -\Delta$$

for all probability distributions  $p : [0, 1] \rightarrow \mathbb{R}^+$ .

$\Delta$  is known and positive.  $p(\omega|0)$  and  $p(\omega|1)$  can be computed by Eq. (1) and (2). The notation  $\bar{e}(\omega)$  is meant to show that the aggregate function on expected outcomes also depends on the actual frequency of the positive answer.

The answer to Problem 1, while a valuable theoretical exercise, will inevitably suffer from the following two practical shortcomings:

- first, the mechanism places a significant burden on the poll participants who must compute and report the estimates of the poll outcome;
- second, the mechanism may only work for an offline process, where the poll owner stores all reports until the end of the poll and does not disclose any partial information to future participants. This can be seen against the philosophy of internet systems and could deter participation.

For these reasons, the next section pursues an alternative mechanism that requires less information from the participants and can be adapted to an online opinion reporting process.

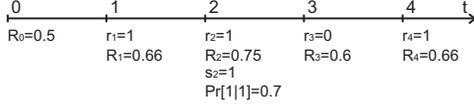
## 4. ONLINE INCENTIVE-COMPATIBLE REWARDS

In an online mechanism the opinions submitted by the poll respondents are immediately reflected in the (partial) result of the poll. The main differences from the offline process is that:

- agents can see the opinions submitted by the previous users
- the users (and the poll operator) do not know how many more opinions will be submitted in the future
- rewards have to be conditioned on a finite number of future reports

Let  $R_t$  be the partial result of the poll, as known after the answer of the  $t^{th}$  participant. If  $n$  is the number of positive answers among the  $t$  reports submitted so far, the partial result  $R_t$  is defined as the fraction  $\frac{n}{t}$ . We will use the notation  $R_{t+1} = R_t \oplus r_{t+1}$  to denote the partial result updated with the binary report  $r_{t+1}$ : i.e.,  $R_{t+1} = \frac{n+r_{t+1}}{t+1}$ . The reward mechanism we propose is very simple:

1. Agents are paid only if their answer agrees with the answer of another agent (the reference report).
2. By default, the report  $r_t$  at time  $t$  is compared to the report  $r_{t+1}$ .



**Figure 2: Example of reports in an online opinion poll.**

3. The agents may choose to be scored against a future reference report. If  $r_t = 0$ , the agent may specify a threshold  $\theta < R_{t-1}$ , and she will be matched against the first report submitted after the result of the poll belongs to the interval  $[\theta, R_{t-1}]$ . Likewise, if  $r_t = 1$  the agent may specify a threshold  $\theta > R_{t-1}$ , and she will be matched against the first report submitted after the result of the poll belongs to the interval  $[R_{t-1}, \theta]$ .
4. The payments for matching positive and negative reports are  $\tau_t(1)$  and  $\tau_t(0)$  respectively. They are computed based on the partial result of the poll available so far:

$$\begin{aligned}\tau_t(0) &= c \cdot R_t \\ \tau_t(1) &= c \cdot (1 - R_t)\end{aligned}\quad (6)$$

where  $c$  is a positive constant.

Consider the example from Section 3, and a sequence of reports as in Figure 2. Let us analyze the second participant. The first reporter submitted a positive answer, and since the starting value of the poll was set to  $R_0 = 0.5$ , the partial outcome of the poll before the report of the second agent is  $R_1 = 0.66$ . The second participant endorses the positive answer and she reports truthfully  $r_2 = s_2 = 1$ . This report takes the partial result of the poll up to 75%. The second agent has two options. First, she can accept to be scored against the default reference report,  $r_3$ , understating that she will get paid  $\tau_1(1)$  if  $r_3$  is positive, or 0 if  $r_4$  is negative. In our example  $r_3$  is negative, so had the agent chosen the default reference report, she wouldn't have received any reward. The second option of the agent is to specify a threshold (e.g.,  $\theta = 0.7$ ) such that she will be scored against the first report submitted after the partial result was in  $[0.66, 0.7]$ . If she took this option, the second report would be scored against the 5th report, as  $R_4 = 0.66 \in [0.66, 0.7]$ .

The next result describes the equilibrium strategy of this mechanism. The equilibrium is characterized by a more general notion of incentive compatibility where the reports are not necessarily truthful, but always act to decrease the distance between the updated partial result of the poll and the subjective belief of the reporter regarding the outcome of the poll. The typical case where the equilibrium strategy will prescribe non-truthful reporting is when the prior information of the agents conflicts strongly with the publicly available result of the poll. Coming back to the example from Figure 2, the first agent reports positively regardless of her private opinion. She lies about her private information because by doing so she helps correct the inaccurate public information available in the poll (i.e., the starting value in this case).

**THEOREM 1.** *The following strategy is a Nash Equilib-*

*rium:*

$$r_t = \arg \min_r |(R_{t-1} \oplus r) - Pr[1|s_t]|$$

**PROOF.** Recall that  $s_t$  is the private information of agent  $t$ , and  $Pr[1|s_t]$  is the posterior private expectation of the agent  $t$  regarding the outcome of the poll. The report  $r_t$  submitted to the poll attempts to approach as much as possible the updated result to the private opinion.

Depending on the value of  $R_{t-1}$  and on the possible private beliefs of an agent, we might be in one of the following three cases:

- $R_{t-1} < Pr[1|0] < Pr[1|1]$  when both agent types (the type endorsing the positive answer and the type endorsing the negative answer) regard the partial result of the poll as “lower” than what it should be;
- $Pr[1|0] \leq R_{t-1} \leq Pr[1|1]$  when the agent who endorses the negative answer believes the partial result is overestimating the final outcome, while the agent who endorses the positive answer believes the partial result is underestimating the final outcome.
- $Pr[1|0] < Pr[1|1] < R_{t-1}$  when again, regardless of their opinion, all agents believe the partial result is overestimating the final outcome.

In the first case, the strategy requires both agent types (agents who believe in the positive alternative and agents who believe in the negative alternative) to report positively. In the last case, both agent types will report negatively, while in the middle case the agents truthfully report their private opinion.

It is simple to prove that all agents who are in the first situation do not have the incentive to report negatively. If they were to report negatively the value of  $R_t$  will be even lower than  $R_{t-1}$  and therefore agent  $t+1$  will report positively. According to the definition of our reward scheme, the reward to agent  $t$  is 0 in this case. Moreover, any future report submitted when the partial result is smaller than  $R_t$  will also be positive, so it does not help to choose a future reference report instead of the next one. A positive report, on the other hand, generates a positive reward with positive probability.

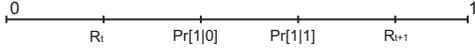
Similarly, in the third case, all agents have the incentive to report negatively as any deviation to a positive report does not bring any payoffs.

For the second case when  $Pr[1|0] \leq R_{t-1} \leq Pr[1|1]$ , take an agent  $t$  who endorses the negative answer, and consequently believes the final outcome of the poll should be  $Pr[1|0]$ . If agent  $t$  truthfully reports her opinion, she may choose to be scored against a future reference report,  $r_j$  that has been submitted in the same conditions: i.e.,  $Pr[1|0] < R_{j-1} < Pr[1|1]$ . The report  $r_j$  will therefore be truthful, and agent  $t$  believes the probability of  $r_j$  being negative is  $Pr[0|0] = 1 - Pr[1|0]$ . Her expected payoff in this case is:

$$Pr[0|0] \cdot \tau_{t-1}(0) = (1 - Pr[1|0]) \cdot \tau_{t-1}(0) \geq (1 - R_{t-1}) \cdot c \cdot R_{t-1};$$

On the other hand, if agent  $t$  were to lie and report positively, her best choice would be to choose a future reference report  $r_j$ , that is also truthful. Since the probability of  $r_j = 1$  is  $Pr[1|0]$  agent  $t$  expects the payoff:

$$Pr[1|0] \cdot \tau_{t-1}(1) = Pr[1|0] \cdot \tau_{t-1}(0) \leq R_{t-1} \cdot c \cdot (1 - R_{t-1});$$



**Figure 3: Example of the overshooting effect**

As a consequence, if agent  $t$  endorses the negative answer, she is better off reporting honestly. Similarly, if agent  $t$  endorsed the positive answer, which makes the strategy enounced in the theorem a Nash equilibrium. ■

Note that the equilibrium strategy in Theorem 1 does not specify when an agent would choose to be scored against a report different from the next one. But the option has to exist in order to ensure that there is always a profitable reporting option. Consider the case from Figure 3 where although the current result of the poll is below both  $Pr[1|0]$  and  $Pr[1|1]$ , the updated result  $R_{t+1} = R_t \oplus 1$  overshoots the private beliefs of the agents, and the next participant, will find an overestimated result.  $r_{t+1}$  would report negatively in this case, and  $r_t$  would not get paid. In our mechanism, however, the agent can first choose the reference report to optimize the rewards, and then follows the equilibrium strategy to determine what to report. On the other hand, the agent cannot choose *any* reference report, since there must be some correlation between the binary answer to the poll question and the choice of reference report. This correlation is enforced by making sure the reference report is chosen from a round where the public information was also lower (respectively higher) than the private information.

An obvious choice for the value of the threshold  $\theta$  is exactly the private belief  $Pr[1|s_t]$ . Nevertheless, the proof in Theorem 1 does not depend on this assumption.

The next important question is whether the final outcome of the poll will converge to the true frequency of the positive answer. In equilibrium, some agents will misreport their opinion, however, the following theorem proves that after sufficient reports, the outcome of the poll converges to the true fraction of agents that endorse the positive answer.

**THEOREM 2.** *The poll converges to the correct outcome.*

**PROOF.** The proof of this result is based on two observations. First, according to the equilibrium, whenever  $Pr[1|0] < R_t < Pr[1|1]$  agent  $t + 1$  reports honestly her opinion. This honest report allows all future agents to learn something new (the private signal of agent  $t + 1$ ) and can use this information to update their own beliefs regarding the distribution of opinions in the population. Given that the partial outcome of the poll will be a sufficient number of times between the spread of private beliefs, the poll will receive a sufficient number of honest opinions, which by Bayesian updating, will converge all beliefs to the true distribution of opinions.

What remains to be proven is that the equilibrium strategy will push often enough the partial outcome of the poll within the interval  $[Pr[1|0], Pr[1|1]]$ . This result is straight forward: for any values  $Pr[1|0] < Pr[1|1] \in [0, 1]$ , and for all  $R_t = \frac{n}{t}$  not belonging to the interval  $(Pr[1|0] < Pr[1|1])$  the strategy of Theorem 1 pushes the partial outcome between  $Pr[1|0]$  and  $Pr[1|1]$  within a finite number of steps. Because the update to  $R_t$  is always smaller than  $\frac{1}{t}$ , once  $t$  becomes larger than  $\lceil \frac{1}{Pr[1|1] - Pr[1|0]} \rceil$  the partial result can-

not oscillate anymore from the left to the right of the interval  $(Pr[1|0] < Pr[1|1])$  without falling inside. ■

The proof of Theorem 2 also gives an insight to the dynamic behavior of the equilibrium. As long as the public outcome of the poll is far from the private belief of the agents, users will submit opinions that decrease the gap between the private and the public information. As soon as the public information becomes close enough to the private information, the agents honestly report their opinions. The partial poll result can be initialized with any value. The first reports will make sure that within a finite number of rounds, the publicly available information accurately reflects the private (unknown to the designer) information of the agents.

Another consequence of Theorem 1 is that the agents who report when the public information is significantly away from the private information will generally expect higher payments. When the partial result  $R_t$  is either below  $Pr[1|0]$  or above  $Pr[1|1]$ , the equilibrium entitles the reporters to an expected payment of  $\tau_t(1)$  and  $\tau_t(0)$  respectively. On the other hand, the agents reporting when  $Pr[1|0] < R_t < Pr[1|1]$  only expect a fraction on the published payments. This automatically creates strong participation incentives, as the participants find it more profitable to submit their opinion as early as possible.

## 5. COLLUSION

The equilibrium described by Theorem 1 is unfortunately not unique. Agents get rewarded for matching reports, therefore one simple strategy that is also a Nash equilibrium is to always report the same opinion (i.e., 0 or 1). The existence of different equilibria brings forth the problem of collusion, as agents may synchronize their reports to game the reward mechanism.

Although a more general solution to address collusion is part of our ongoing work, we believe the current mechanism is resistant against some basic collusion strategies. For example, always reporting the same opinion, although an equilibrium, generates less revenues for the participants than honest reporting. When every agent endorses the negative answer, the partial result of the poll converges fast to 0. The payments for matching negative reports are directly proportional to the running average of the opinions submitted so far, and therefore will also converge to 0. For all possible distributions of opinions, the expected revenue of the honest reporting strategy will at some point become higher than the expected revenue of the lying strategy. The analysis for the symmetric strategy of always reporting 1 gives the same result.

Random reporting, if an equilibrium at all, is also less profitable than honest reporting. A random strategy is characterized by the fact that agents answer the poll using a randomization device whose output does not depend on the agent's private information. First, fixed random reporting strategies (i.e., the probability of an agent reporting 1 is constant over time) are not in equilibrium. The payments change depending on the partial outcome of the poll, therefore two agents reporting in different contexts will expect different payoffs for a negative, respectively a positive report. There is, however, an equilibrium random strategy that depends on  $R_t$ , the partial outcome of the poll.

Let  $\sigma_{R_t}$  be a random strategy depending on the partial outcome  $R_t$ . If  $R_t = \frac{n}{t}$ , where  $n$  out of the  $t$  previous

respondents endorsed the positive answer, let  $\sigma(n, t)$  be the probability that the next agent reports 1:

$$\sigma(n, t) = Pr[r_{t+1} = 1 | n, t]$$

Assume that  $\sigma(n, t) = \frac{n-1}{t-2}$  for all  $n, t \in \mathbb{N}$ ,  $\sigma(0, t) = 0$  and  $\sigma(t, t) = 1$  for all  $t \in \mathbb{N}$ . It is easy to verify that  $\sigma_{R_t}$  is a Nash equilibrium:

- if agent  $t+1$  reports 0, the agent  $t+2$  will also report 0 with probability  $1 - \sigma(n, t+1) = \frac{t-n}{t-1}$  and her expected payoff is:

$$\tau_t(0) \cdot \frac{t-n}{t-1} = c \frac{n(t-n)}{t(t-1)};$$

Moreover, agent  $t+1$  cannot increase her expected payment by choosing to be scored against another reference report: the probability of a negative report following a partial outcome that is higher than  $\theta > \frac{n}{t+1}$  is always lower than  $1 - \sigma(n, t+1)$ .

- if agent  $t+1$  reports 1, the agent  $t+2$  will also report 1 with probability  $\sigma(n+1, t+1) = \frac{n}{t-1}$  and her expected payoff is:

$$\tau_t(1) \cdot \frac{n}{t-1} = c \frac{n(t-n)}{t(t-1)};$$

Moreover, agent  $t+1$  cannot increase her expected payment by choosing to be scored against another reference report: the probability of a positive report following a partial outcome that is lower than  $\theta < \frac{n+1}{t+1}$  is always lower than  $\sigma(n+1, t+1)$ .

Besides the complexity involved in taking random decisions that depend on the current outcome of the poll, the equilibrium brought by  $\sigma_{R_t}$  is also unattractive. In the honest equilibrium the expected payment is at least  $\tau_t(0)\tau_t(1)/c$  (see the proof of Theorem 1, which is greater than the outcome of  $\sigma_{R_t}$ ).

In practical settings it is quite reasonable to assume that the majority of the agents will not attempt to become part of a lying coalition. If the fraction of colluders is small enough, we claim that no symmetric lying strategy (i.e., all colluders use the same strategy) can be profitable for the colluders. The restriction to symmetric strategies is justified by practical considerations: asymmetrical collusion strategies require the synchronization of reports and side communication channels among colluders.

Assume there is a collusion strategy  $\sigma$  such that for some prior beliefs and some value of the partial outcome  $R_t$ , the colluder would report the value that increases the distance between  $R_{t+1}$  and the private information. Assume that up to  $\gamma$  percent of the agents can become part of the coalition on  $\sigma$ , and that the remaining  $1 - \gamma$  fraction reports honestly. As in the proof of Theorem 1, we can consider the following three cases:

- $R_t < Pr[1|0] < Pr[1|1]$ ,
- $Pr[1|0] \leq R_t \leq Pr[1|1]$
- $Pr[1|0] < Pr[1|1] < R_t$

In the first case, the colluder (agent  $t+1$ ) is assumed to report 0 according to  $\sigma$ . Assuming that colluders cannot control the sequence of reports (a reasonable assumption in

most real applications) the reference report used to score  $r_{t+1}$  is honest with probability  $1 - \gamma$ , and part of the coalition with the remaining probability  $\gamma$ . Any honest report submitted when the partial outcome of the poll is lower than  $R_t$  will be positive, hence the best case payment expected by the colluder is:

$$\gamma \cdot \tau_t(0) = \gamma \cdot c \cdot R_t;$$

On the other hand, the payment expected by an honest reporter is  $\tau_t(1) = c \cdot (1 - R_t)$ . Therefore the colluder is worse off than the honest reporter as long as:

$$\gamma < \frac{1 - R_t}{R_t}; \quad (7)$$

The same analysis for the third case sets an upper bound for the maximum coalition fraction to:

$$\gamma < \frac{R_t}{1 - R_t}; \quad (8)$$

For the second case, assume the private signal of the colluder is negative: i.e.,  $s_{t+1} = 0$ . According to  $\sigma$  the colluder should report 1. In the best case, the colluder can expect a matching report from a fallow colluder with probability 1, and a matching report from an honest reporter with probability  $Pr[1|0]$ . The colluder's expected payoff is therefore:

$$\gamma \tau_t(1) + (1 - \gamma) \left( Pr[1|0] \tau_t(1) \right)$$

while an honest reporter in the same circumstances expects:  $(1 - Pr[1|0])\tau_t(0)$ . The honest reporter is better off as long as:

$$\gamma < \frac{R_t - Pr[1|0]}{(1 - R_t)(1 - Pr[1|0])} \quad (9)$$

Similarly, if the private signal of the colluder were positive, the upper bound on the colluding fraction is:

$$\gamma < \frac{Pr[1|1] - R_t}{R_t Pr[1|1]} \quad (10)$$

Conditions (7) and (8) show that it is tempting to collude when the partial result is very close to 0 or very close to 1. For example, when the participants have a clear preference for the positive answer, the reward for a negative report becomes so great that even a small probability of being matched against a fellow colluder makes lying worth while. One possible solution is to force some minimum payments for a report: e.g.,  $\tau_t(0) = c \cdot \max(R_t, b)$  and  $\tau_t(1) = c \cdot \max(1 - R_t, b)$ , where  $b$  is some constant set by the designer. This guarantees that the mechanism is resistant to a collusion fraction of at least  $b$ . Nevertheless, the minimum payments also create supplementary inconveniences:

- the strategies of always reporting 0 or always reporting 1 become more interesting equilibria, and
- in some extreme cases where the private beliefs of the agents are either well below  $b$  or well above  $1 - b$ , honest reporting is not always the best strategy.

Conditions (9) and (10) emphasize another context where collusion becomes attractive, mainly when the public outcome of the poll is very close to the agent's private beliefs (i.e.,  $R_t \simeq Pr[1|0]$  or  $R_t \simeq Pr[1|1]$ ). Intuitively, a report submitted when  $R_t$  matches the private information does not bring an information gain, and therefore the reward of

the honest equilibrium is not substantial enough to prevent collusion. This case, however, does not pose a real threat to the mechanism, since only a limited number of colluders will actually be in a situation to have their private beliefs aligned to the public information.

There seems to be an inherent tradeoff between the resistance to collusion and the range of settings where this mechanism can be used. Whether this tradeoff is a shortcoming of our present reward scheme, or a more general trait of online information elicitation mechanisms remains to be addressed in our future work.

Another frequent type of collusion is when the same agent controls several online identities (or *sybils*) and coordinates the reports submitted under different identities to manipulate the mechanism. For some domains (e.g., some combinatorial auctions), it is possible to construct *false-name-proof* mechanisms [26] where self-interested agents cannot profit from creating multiple identities. For many other domains, on the other hand, false-name-proof mechanisms are known to be impossible (e.g., simple voting mechanisms). We conjecture a similar impossibility result for the setting described in this paper. Intuitively, false-name-proofness implies that an agent who elicits a second report from a different identity should not expect any payment. However, when two separate agents submit each one report, the incentive-compatibility constraints require that they both expect positive rewards. Since the mechanism has no external source of information to distinguish between the two settings, we believe that no incentive-compatible reward mechanism can also be false-name-proof in this domain.

An interesting alternative is proposed by Conitzer [5], who shows that many domains accept false-name-proof mechanisms with limited verification of online identity. As future work, we plan to find the minimum assumptions on an identity verification device that could help deter sybil attacks.

## 6. DISCUSSION

Prediction markets guarantee that the subsidies to the traders are bounded by an upper limit. This is very convenient for market owners as they know the worst case price they have to pay for the information provided through the market. The BTS method provides even stronger guarantees by ensuring that the payments to the poll participants break even. However, we believe that budget-balance is not a desirable property for opinion polls. 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 (or even finite bounds on the budget) are not feasible. Instead, the mechanism can ensure that there is an upper bound on the expected payment per participant. It is only fair that the increased information accuracy obtained through a supplementary report be reflected in the total cost. The maximum subsidy per participant is automatically obtained through the definition of the reward mechanism from Section 4.

We have treated here binary opinion polls, but the results can be easily extended to n-ary poll questions. The private beliefs of the agents will reflect a probability distribution over a vector of possible answer frequencies. Conceptually, the agents would still report to decrease the gap between the public and the private information, however, different metrics used to assess the distance between two probability distributions might affect the final results. A complete characterization of n-ary polls remains for future work.

Another extension is to consider a group of peer reports when paying for an opinion. This allows for more gradual payments where partial agreement is also rewarded. In general, the use of several reference reports is expected to decrease the overall payments that provide honest reporting incentives [12].

The mechanism can be improved by using smarter updating functions for the partial result of the poll. For example, the desire of an agent to be scored against a future reference report may be seen as an indication of overshooting, and should trigger a finer grained update of the  $R_t$ . Moreover, for polls that run over long periods of time one might wish to weigh less the reports from the distant past. However, changes to the updating process, (especially due to information reported by agents) also affect the behavior of the participants, and will be analyzed in more depth in future work.

One last point we would like to stress is that the rewards need not be monetary payments. These rewards can be converted into bonus points, preferential QoS, lottery tickets or any other assets that users value. Effective micro-payment systems are still hard to implement, but fortunately, users care enough about virtual points and currencies [24] 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. This paper addresses the design of incentives for online opinion polls. We survey existing solutions and propose a new mechanism that is simple and effective.

## 8. REFERENCES

- [1] C. Camerer. Can Asset Markets be Manipulated? A Field Experiment with Racetrack Betting. *Journal of Political Economy*, 106(3):457–482, 1998.
- [2] K. Chen, L. Fine, and B. Huberman. Forecasting Uncertain Events with Small Groups. In *Proceedings of ACM Conference of Electronic Commerce (EC'01)*, 2001.

- [3] 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.
- [4] R. T. Clemen. Incentive contracts and strictly proper scoring rules. *Test*, 11:167–189, 2002.
- [5] V. Conitzer. Limited verification of identities to induce false-name-proofness. In *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-07)*, pages 1002–111, Brussels, Belgium, 2007.
- [6] V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Proceedings of the Uncertainty in Artificial Intelligence Conference (UAI)*, 2002.
- [7] J. Crémer and R. P. McLean. Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent. *Econometrica*, 53(2):345–61, 1985.
- [8] C. d’Aspremont and L.-A. Gérard-Varet. Incentives and Incomplete Information. *Journal of Public Economics*, 11:25–45, 1979.
- [9] S. Figlewski. Subjective information and market efficiency in a betting market. *The Journal of Political Economy*, 87(1):75–88, 1979.
- [10] 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.
- [11] S. Johnson, J. Pratt, and R. Zeckhauser. Efficiency Despite Mutually Payoff-Relevant Private Information: The Finite Case. *Econometrica*, 58:873–900, 1990.
- [12] 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.
- [13] 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.
- [14] R. Jurca and B. Faltings. Robust Incentive-Compatible Feedback Payments. In M. Fasli and O. Shehory, editors, *Trust, Reputation and Security: Theories and Practice*, volume LNAI 4452, pages 204–218. Springer-Verlag, Berlin Heidelberg, 2007.
- [15] M. Kandori and H. Matsushima. Private observation, communication and collusion. *Econometrica*, 66(3):627–652, 1998.
- [16] P. Milgrom and N. Stokey. Information, Trade and Common Knowledge. *Journal of Economic Theory*, 26:17–27, 1982.
- [17] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science*, 51:1359–1373, 2005.
- [18] D. Pennock, S. Debnath, E. Glover, and C. Giles. Modeling information incorporation in markets with application to detecting and explaining events. In *Proc. of the 18th Conf. on Uncertainty in Artificial Intelligence*, pages 405–411, 2002.
- [19] 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*, pages 174–183, 2001.
- [20] A. Poteshman. Unusual Option Market Activity and the Terrorist Attacks of September 11, 2001. *Journal of Business*, 79:1703–1726, 2006.
- [21] D. Prelec. A bayesian truth serum for subjective data. *Science*, 306(5695):462–466, 2004.
- [22] P. Rhode and K. Strumpf. Manipulating Political Stock Markets: A Field Experiment and a Century of Observational Data. Working paper, 2007.
- [23] R. Roll. Orange juice and weather. *The American Economic Review*, 74(5):861–880, 1984.
- [24] E. Servan-Schreiber, J. Wolfers, D. Pennock, and B. Galebach. Prediction markets: Does money matter. *Electronic Markets*, 14(3), 2004.
- [25] J. Wolfers and A. Leigh. Three Tools for Forecasting Federal Elections: Lessons from 2001. *Australian Journal of Politics Science*, 37(2):223–240, 2002.
- [26] M. Yokoo, Y. Sakurai, and S. Matsubara. The effect of false-name bids in combinatorial auctions: New fraud in internet auctions. *Games and Economic Behavior*, 46(1):174–188, 2004.