Incentive-compatible Online Opinion Polls

Radu Jurca and Boi Faltings
Ecole Polytechnique Fédérale de Lausanne (EPFL)
Artificial Intelligence Lab
CH-1015 Lausanne

December 12, 2007

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 tied to:

- the result of presidential elections,
- the earnings of movies,
- the outcome of sport competitions,
- the completion of internal projects.

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 that pays $1 if the winner of the 2008 US presidential elections is democrat. An agent who privately believes with probability p that US will have a democrat president in 2008, will find it rational to buy more shares until the price reaches p. Numerous experiments and studies have shown that the price in such prediction markets converges to reflect the true information available in the market, despite game theoretic negative results (e.g., the “no trade” theorem). Moreover, prediction markets have been known to consistently outperforms other traditional prediction tools.

However, one important shortcoming of prediction markets is that they must be linked to a publicly observable event that is precisely defined. This strong requirement can create problems even for the simplest prediction markets: for example, had Al Gore won the recount of votes following the 2000 US presidential election, the security associated with G.W. Bush would have still paid

\$1 since at the time specified in the security, Bush was still the active president of the US. Moreover, there are many other statements that simply cannot be verified within a decent period of time. Several such examples are: subjective or artistic opinions, historical interpretations, distant forecasts, or climate hypothesis.

For all non-verifiable claims, agents can be encouraged to express their honest opinion by conditioning their reward on the opinion of other peers. One such mechanism is the Bayesian Truth Serum [Prelec, 2004] that encourages honest responses to multiple choice subjective questions like:

- Q: "What wine do you prefer?", A: a)Red b)White;
- Q: "What is the average temperature increase by 2050?", A: a) smaller than 2°C, b) smaller than 5°C, c) smaller than 10°C

. Every agent submits her subjective answer to the question, but also provides an estimate of the final distribution over the possible answers. The Bayesian Truth Serum (BTS) mechanism rewards each agent based on the given answer, the estimated probability distribution, and the final distribution of answers given by the respondents. [Prelec, 2004] proves that honest answering is a pareto-optimal Nash Equilibrium of the mechanism.

The objective of this paper is two-fold. First we want to improve the BTS mechanism by finding other reward schemes that make honest reporting a Nash Equilibrium. In particular, we will use Automated Mechanism Design [Conitzer and Sandholm, 2002] to search for the best possible reward mechanism given the available information. Among the improvements we make to BTS are:

- enforce ex-ante participation incentives (ex-ante individual rationality) by ensuring that the reward obtained by every participant is greater or equal to 0;
- ex-ante individual rationality destroys the budget balance property of the BTS mechanism. Since the poll operator has to make positive payments to the participants, our goal is to reduce to a minimum the necessary payments, while still guaranteeing some margins for truth-telling
- reduce the amount of information requested from the participants. In particularly, we will look at ways to relax the BTS requirement to estimate an entire distribution over the answers of the other agents.

One major drawback of the BTS method is that it is an offline process. Agents submit their opinions offline, and once the entire population has answered the questions, the center can compute the payments due to every agent. Moreover, the properties of the BTS method (e.g., incentive compatibility) are guaranteed only for large number of participants.

On the internet, however, opinion polls are mostly online processes, where every new submission immediately reflects into the outcome of the poll. The information from the poll thus becomes instantly available, and the poll operator does not have to count on a predefined number of participants. The second objective of this paper is to present an incentive-compatible reward mechanism that works for an online opinion poll.

Finally we will characterize different properties of the reward mechanisms.

2 Model

N agents participate in a binary opinion poll. The question can be something like: *"Are you afraid of global worming?"* or *"Would you prefer white over red wine?"* or any other question accepting a binary answer. The case of n-ary opinion polls can be generalized using the same tools presented below.

We model the opinion of an agent by some private information the agent has regarding the question at hand. To make things simpler, let us assume that every agent i has a private binary signal $s_i \in \{0, 1\}$, where $s_i = 0$, respectively $s_i = 1$, means that the agent would answer the question negatively, respectively positively. Agents do not know the private signal of other agents, however, they all share a common belief regarding the a priory distribution of preferences in the population. Let Ω be a random variable expressing the true distribution of preferences of the agents. Let $p(\omega) = Pr[\Omega = \omega]$ describe the common knowledge a priory belief that a fraction $\omega \in [0, 1]$ of the population will endorse the positive answer. All agents are rational and bayesian, and therefore expect different distribution of preferences depending on their private opinion: e.g.,

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

$$p(\omega|1) = \frac{p(1|\omega)p(\omega)}{\int_{\omega'=0}^{1} p(1|\omega')p(\omega')d\omega'} = \frac{\omega p(\omega)}{\int_{\omega'=0}^{1} \omega' p(\omega')d\omega'};$$

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

$$p(\omega|0) = \frac{p(0|\omega)p(\omega)}{\int_{\omega'=0}^{1} p(0|\omega')p(\omega')d\omega'} = \frac{(1-\omega)p(\omega)}{\int_{\omega'=0}^{1} (1-\omega')p(\omega')d\omega'};$$

The goal is to encourage agents to truthfully declare their opinions. The poll designer **does not** know the prior distribution $p(\omega)$, and must come up with a payment function $\tau(r_i, r_{-i}) \in \mathbb{R}^+$, where every agent is rewarded based on the information r_i she reports, and on the information r_{-i} that all other agents report.

3 An offline IC reward mechanism

The information reported by agent i is a 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 estimation of the final outcome of the poll. An agent reports truthfully when $o_i = s_i$, and when:

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

Unfortunately, not all agents are expected to accurately compute and report e_i. However, when N is large enough, we assume that individual errors made in the reports e_i will cancel out, so that the average $\bar{e} = \frac{\sum_{i=1}^{N} e_i}{N}$ will accurately reflect the value:

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

where ω^* is the true frequency of the positive answer in the population.

Therefore when N is large enough, 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, ω^* 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 $\tau : \{0,1\} \times [0,1] \times [0,1] \times [0,1] \to \mathbb{R}^+$ 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 ω^*
- 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 Δ:

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

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 Eq. (1) is to assume that the agent truthfully reports $Pr[1|s_i]$ and to design payment mechanisms that only encourage the honest submission of the binary opinion. In this case 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:

$$\begin{aligned}\int_{\omega=0}^{1} p(\omega|1)g(\omega,\bar{e}) > \Delta;\\ \int_{\omega=0}^{1} p(\omega|0)g(\omega,\bar{e}) < -\Delta;\end{aligned} \quad (2)$$

Hence the design problem can be summarized by the following:

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

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

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

for any probability distribution $p : [0,1] \rightarrow \mathbb{R}^+$*, given that:*

- $\Delta \in \mathbb{R}^+$ *is a known, positive constant*
- $$p(\omega|1) = \frac{p(1|\omega)p(\omega)}{\int_{\omega'=0}^{1} p(1|\omega')p(\omega')d\omega'} = \frac{\omega p(\omega)}{\int_{\omega'=0}^{1} \omega' p(\omega')d\omega'};$$
- $$p(\omega|0) = \frac{p(0|\omega)p(\omega)}{\int_{\omega'=0}^{1} p(0|\omega')p(\omega')d\omega'} = \frac{(1-\omega)p(\omega)}{\int_{\omega'=0}^{1} (1-\omega')p(\omega')d\omega'};$$
- $$f(\omega) = \omega p(1|1) + (1-\omega)p(1|0) = \omega \int_{\omega'=0}^{1} \omega' p(\omega'|1)d\omega' + (1-\omega)\int_{\omega'=0}^{1} \omega' p(\omega'|0)d\omega'$$

3.1 Minimum Rewards

The constraint that rewards $\tau(\cdot,\cdot,\cdot)$ are only positive makes it ex-ante individually rational for every agent to participate in the opinion poll (unlike in the BTS method our system does not impose penalties on the participants). Nevertheless, ex-ante individual rationality makes it impossible to balance the budget of the poll operator. Consequently, it is important to design the incentive compatible rewards such that the overall payment of the poll operator is minimized.

We adopt the methodology of *automated mechanism design* [Conitzer and Sandholm, 2002] and specify the reward mechanism through an optimization problem. The objective is to minimize the budget required by the operator, such that every participant finds it optimal to answer the poll truthfully.

3.2 Decrease the information elicitation requirements on the participants

One criticism of the above mechanism is that participants are required to estimate the final outcome of the poll, i.e., submit a probability distribution over the binary answers of the poll. Humans will typically find it very difficult to estimate the exact percentage of participants that endorse the positive answer. Moreover, the estimates they will give will probably contain large errors and might affect the incentives of purely rational respondents.

Human users, however, are very good at estimating whether the final outcome of the poll will be above or below a certain bound. This new information elicitation process is much easier for the user, but also contains less information. The idea is to design a mechanism that is incentive-compatible even for this reduced information set. Particularly, the mechanism designer has to:

- choose the bounds that will be presented to the agents,
- design the reward scheme that is based on the binary answer to the chosen bound being above or below the final outcome.

4 An Online IC Reward Mechanism

In an online mechanisms opinions are submitted sequentially and immediately update the 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

In designing our mechanism, we make the following supplementary assumption: all agents believe that at least one of the future users will honestly submit her opinion. Let us index opinions by t, and let without any loss of generality, let N_1 be the index of the first opinion which is assumed truthful, N_2 be the index of the second opinion which is going to be truthful, and so on. Note that N_2, N_3, etc. might not exist, if, for example, N_1 is the last agent participating in the poll. The (potentially infinite) sequence of opinions can be split into a (potentially infinite) set of partial finite sequences, each terminating with an honest opinion. We will analyze such partial finite sequences of reports, where users $0, 1, \ldots N-1$ are rational, and user N reports honestly. Because we will not make any assumptions on N the results will apply to all finite sequences of reports that end with an honest report, and therefore will generalize for any sequence of opinion reports that satisfies the assumption enounced in the beginning of this section.

As opposed to the offline poll, the online opinion poll also contains a running estimator of the poll's result. Let R_t be the partial outcome of the poll as available at time t (i.e., after the t^th opinion has been submitted). When user $t+1$ submits opinion o_{t+1} the outcome of the poll is updated according to the function:

$$R_{t+1} = f(t, R_t, o_{t+1});$$

4.1 A simple mechanism

The simplest possible mechanism requires users to submit only their binary opinion. The outcome of the poll is the running average of the already submitted reports, and the reward scheme is the following: *The opinion $o_t \in \{0,1\}$ submitted by user t is paid $\tau(o_t)$ if and only if the next participant has the same opinion: i.e., $o_{t+1} = o_t$.* The payments $\tau(0)$ and $\tau(1)$ are different and are scaled such that honest reporting is a Nash Equilibrium [Jurca and Faltings, 2006]:

$$Pr[1|1]\tau(1) - Pr[0|1]\tau(0) > \Delta;$$

$$Pr[0|0]\tau(0) - Pr[1|0]\tau(1) > \Delta;$$

Since the conditional probabilities $Pr[1|1]$ and $Pr[1|0]$ are not known to the center, they will be estimated based on the current outcome R_t of the poll. One important remark is the following: when the outcome of the poll accurately reflects the private beliefs of the agents, the payments $\tau(0)$ and $\tau(1)$ encourage an honest reporting Nash equilibrium (by the definition of the payments). If, however, the public outcome of the poll is far from the private beliefs of the agents, the payments encourage the users to report such that the updated public outcome becomes closer to their private beliefs.

The equilibrium analysis of this mechanism can be done by backward induction. The user $N-1$ knows that the N^{th} opinion will be truthful. Regardless of the current public outcome of the poll, user $N-1$ reports such that R_{N-1} will be closer than R_{N-2} to her private belief. It can be shown that in most circumstances, if user t's strategy is to report such that the public belief becomes closer to the private belief, user $t-1$ also finds it in her best interest to adopt a reporting strategy that pushes the public information closer to the private information.

The dynamic behavior of the equilibrium is therefore the following: as long as the public outcome of the poll is far from the private belief the agents have, 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. Therefore the public outcome of the poll convergence to the true frequency of opinions.

4.2 Allowing neutral reports

One problem of the simple mechanism presented above is that in the initial phase of the equilibrium, the reports that try to decrease the gap between the public and the private information can cause an *overshooting* effect: for example, the current outcome is lower than dictated by the users private beliefs, nevertheless, a positive report will update the current outcome above the value privately believed by the user. In such circumstances the users do not have the incentive to push the current public information towards the true value.

To correct the incentives, we allow the users to submit a *neutral* report. This report does not change the current value of the public outcome, but increases the resolution of the updating process: next time a positive or a negative opinion is submitted, the change to the public outcome of the poll is finer grained.

This new report modifies the equilibrium of the mechanism in the following way:

- as long as the public outcome of the poll is far from the private beliefs, users report to decrease this gap;
- if a positive or a negative report may cause the public estimate to overshoot the private one, users report neutrally
- if the public outcome is close enough to the private beliefs, users submit their opinions truthfully.

As with the previous mechanism, the analysis of the equilibrium can be done by backward induction. Dynamically, the system behaves as explained in the previous subsection, with initial users reporting to push the public estimate closer to the private beliefs, and subsequent users expressing their honest opinions. Occasionally, users may submit neutral ratings to increase the resolution of the updating process. As expected, the outcome of the poll converges to the correct value.

4.3 Asking for additional information

The particularities of the previous mechanisms is that initial opinions are used to align the public information available to the poll operator to the private information available to the participants. Depending on the gap between the two, many initial reports may be wasted.

A more complex mechanism asks the participants to also estimate how far is the current public outcome from what is privately believed to be the accurate result. For example, if a user believes the current estimator is smaller that it should be, she can report that to the center either as a continuous value specifying the distance between the two, or as one out of a finite number of possible reports.

To encourage honest reporting, the center will keep the same reward principle enounced before (opinions are paid if they agree with the next opinion submitted by another user), but will modulate the payments proportional to the accuracy of the reported offset of the public information.

5 Collusion in IC Reward Schemes

The offline IC rewards are vulnerable to collusion. The reward mechanism typically has several equilibria, and agents can coordinate their reports on a false equilibrium. To prevent this, additional constraints can be placed on the design problem, such that honest reporting becomes the unique, or the pareto-optimal Nash Equilibrium. This technique is similar to [Jurca and Faltings, 2007].

A second problem related to collusion is the sybil attack, where the same strategic agent creates multiple online identities and submits multiple reports to the opinion poll. This is a much harder problem that cannot be generally solved by economic incentives. Possible solutions come from partially verifying reports [Conitzer, 2007].

References

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.

V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Proceedings of the Uncertainty in Artificial Intelligence Conference (UAI)*, 2002.

Radu Jurca and Boi 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.

Radu Jurca and Boi 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.

D. Prelec. A bayesian truth serum for subjective data. *Science*, 306(5695): 462–466, 2004.