Truthful Surveys

Nicolas Lambert* and Yoav Shoham

Department of Computer Science
Stanford University
Stanford, CA 94305

Abstract. We consider the problem of truthfully sampling opinions of a population for statistical analysis purposes, such as estimating the population distribution of opinions. To obtain accurate results, the surveyor must incentivize individuals to report unbiased opinions. We present a rewarding scheme to elicit opinions that are representative of the population. In contrast to the related literature, we do not assume a particular knowledge structure. In particular, our method does not rely on a common prior assumption.

1 Introduction

Online surveys, opinion polls and questionnaires are primary tools to gather information on a population and have been growing at fast pace over the past few years. There already exists an extensive literature on the construction of questionnaires and their statistical processing and analysis (see, for example, Montgomery [1] or Kish [2]). However, to derive meaningful results, it is also imperative to get accurate samples. To obtain honest feedback, the surveyor should reward participants appropriately: constant rewards or no reward at all give no incentives to report truthful beliefs. In this paper, we focus on the design of survey mechanisms that incentivize participants to provide truthful samples.

If we are interested in objective information about a future event, such as the probability of a political candidate winning an election, we may use scoring rules and score functions [3, 4]. They induce honest participation by defining payment functions contingent on the outcome of the event. However, they fail to elicit subjective information.

Researchers have addressed the related problem of eliciting subjective beliefs about the quality of a product or service. It is assumed that products have a true quality, distributed according to a prior $P(\omega)$, and that each individual experiencing a particular product of quality ω^* gets a noisy signal t of the quality, distributed according to $P(t|\omega^*)$. Therefore each individual with private signal t forms a posterior belief about the true quality $P(\omega|t)$. Miller et al. [5] show how to design rewards to obtain truthful opinions based on probability scoring rules. Jurca and Faltings [6] minimize the payments needed to offset the potential

* Part of this work was done while the author was visiting Yahoo! Research, New York.

gain from lying. In both cases, agents are only asked for their true signal of the quality, while the payment schemes depend explicitly on the prior distributions $P(\omega)$ and $P(t|\omega)$.

Using the true prior is critical to enforce truthful reports and can be difficult to retrieve. The surveyor may not have enough knowledge on the product being rated. The group of surveyed individuals may be too small to infer a prior from their reports. If a sequential mechanism is being used, the surveyor may not be able to provide incentive-compatible payments to the first raters, hence creating a chain of incentive-biases for the next reports. Prelec [7] suggests an alternative approach by delegating to participants the estimation of the prior. Individuals are asked for *both* their opinion t *and* their estimate of the posterior $P(\omega|t)$. The rewards do not depend on the prior directly, but indirectly through the distributions reported by the participants. Prelec provides a reward function which makes truthful reporting a Nash equilibrium.

However, estimating and communicating a prior for each possible opinion may be unnatural and difficult, especially when there are many possible opinions. Jurca and Faltings [8] create an online mechanism to obtain the true frequency of opinion values as the number of participants increases to infinity. They still assume a common prior, but rewards do not depend explicitly on the prior, and participants are not asked for a posterior. However, their mechanism cannot be used in our setting. Indeed, the method does not provide truthful samples, and opinions are restricted to take binary values.

More importantly, the literature described so far relies on the common prior assumption to guarantee honest reports of opinions. However, in many situations of interest, information is asymmetric and no general assumption can be made about the knowledge of individuals. For example, when rating an hotel, people who often travel in rural areas will form a different belief about the distribution of hotel quality than those who frequently visit large cities. Some individuals may be more informed than others, for example a frequent business traveler staying over an extended period can hold more accurate beliefs than occasional travelers with short stays. In general, when the beliefs over the true quality depends on other information that is not part of the private signal being reported, mechanisms that assume a common prior are no longer incentive-compatible or even individually rational.

To construct robust mechanisms that do not rely on any particular assumption about the knowledge of the agents, it is common to look for dominant-strategy implementations. While a dominant-strategy implementation cannot be achieved in the present setting, we propose mechanisms to obtain independent samples of opinions representative of the population, based on a Nash implementation that does not rely on any particular knowledge structure. In particular, no common prior is needed, and there may be asymmetric information. Our mechanisms provide incentives through a payment scheme that depends only on the reports of opinions of the individuals being surveyed. When at least one participant may be trusted, our mechanisms ensure that at all Nash equilibria correspond to true samples of opinions.

The paper is organized as follows. We present the problem and the model in Section 2. In Section 3, we propose an intermediary mechanism to elicit random values from given distributions. Those results are used in Section 4, which presents our main survey mechanisms. We conclude in Section 5.

2 Model

We consider a large population of individuals, each of whom owns an opinion regarding a given question (e.g., what is the quality of this hotel? what will be the price of a barrel of oil in 10 years?). We assume opinions can be expressed as real values in some closed interval $\mathcal{I}$, for example a scale between 0 (worst hotel quality) and 10 (best hotel quality). The quantity $F(x)$ denotes the proportion of the population with an opinion less than or equal to x. The function F represents the distribution of opinions across the population. Formally, we may consider that the population forms a continuum of individuals in the interval $\mathcal{I}$ distributed according to F (hereafter referred to as the population distribution). We assume that F is absolutely continuous (i.e., admits a density function), and that the density is positive on $\mathcal{I}$. The objective of the surveyor is to obtain n independent samples of opinions, which may be used for example to estimate the population distribution or to perform statistical analysis, such as hypothesis testing, goodness-of-fit, etc.

The process of surveying the population is accomplished by a survey mechanism. Formally, a survey mechanism is a tuple $(\mathcal{I}, n, \Pi)$. $\mathcal{I}$ is the interval of possible values of opinions, n is the number of agents being surveyed, and $\Pi : \mathcal{I}^n \mapsto \mathbb{R}^n$ is the vector of payments. The mechanism is interpreted as follows:

Step 1. The surveyor selects n individuals from the population, uniformly at random, referred to as "agent 1, ..., agent n".
Step 2. Each agent i reports an opinion $r_i \in \mathcal{I}$.
Step 3. Each agent i gets a payment $\Pi_i(r_1, \ldots, r_n)$.

Each individual knows whether she is being surveyed, but does not know the identity of the other agents being surveyed. As agents are selected at random, their opinions are (ex-ante) identically and independently drawn from F. We assume that the agents are rational and seek to maximize their expected payment. The population distribution is, a priori, not known.

We will be interested in mechanisms that satisfy certain properties described below. For a mechanism $(\mathcal{I}, n, \Pi)$:

BUDGET-BALANCE The mechanism is *budget-balanced* when it generates no profit nor loss: for all possible reports $r_1, \ldots, r_n \in \mathcal{I}$,

$$\sum_{1 \leq i \leq n} \Pi_i(r_1, \ldots, r_n) = 0 \ .$$

Anonymity The mechanism is *anonymous* when payments do not depend on the ordering of the agents: for all possible reports $r_1, \ldots, r_n, \in \mathcal{I}$, all agent i, and all permutations σ of $\{1, \ldots, n\}$,

$$\Pi_i(r_1, \ldots, r_n) = \Pi_{\sigma(i)}(r_{\sigma^{-1}(n)}, \ldots, r_{\sigma^{-1}(n)}) \ .$$

The surveyor's objective is to obtain samples of opinions that are representative of the population, and is captured by the following two properties:

Accuracy The mechanism is *accurate* when each agent reporting an opinion drawn (ex-ante) according to the population distribution is a Nash equilibrium.

Strong Accuracy The mechanism is *strongly accurate* when each agent reporting an opinion drawn (ex-ante) according to the population distribution constitute the *only* Nash equilibria.

Note that reporting one's true opinion is an accurate strategy, because the opinion of an agent selected uniformly at random from the population is ex-ante distributed according to F. However, depending on the information available of each individual, there are accurate strategies that are not truthful: for example, each agent reporting the opinion of her neighbor would still lead to accuracy. This is not limiting, as the surveyor is not interested in the opinion of a particular individual, but only in reports of opinions representative of the population.

Although we do not consider an implementation in dominant strategies, our results hold independently of the knowledge structure of the population. In addition to her own opinion, each individual may have some knowledge about the population, about the knowledge of the population, about the knowledge of the knowledge of the population, etc. For example, individuals may be ignorant and know nothing about the population distribution. Or individuals may be omniscient and know the opinion of each individual in the population. Alternatively, there may be asymmetric information: some individuals may be ignorant and others may know exactly the population distribution. There may be a common prior, or different priors conditional on the history of each member of the population. There may be publicly available information, such as the mean opinion, etc. For simplicity, the reader may consider a complete information setting in which opinions all individuals of the population are common knowledge, however our results are much more general.

3 Generating random values

In this section, we present a mechanism to elicit random values drawn from any given distributions. The results of this section will be used for our main survey mechanisms.

3.1 Mechanism description

We consider a group of n agents $1, \ldots, n$. For all agent i, let F_i be the a cumulative distribution on a closed interval $\mathcal{I}$, absolutely continuous with positive density. The distributions $F_1, \ldots, F_n$ are common knowledge. We define the following *random generator mechanism*:

Step 1. Each agent i is asked to report a value randomly drawn from F_i.
Step 2. Each agent i is rewarded a payment given by

$$\begin{aligned}\Pi_i(r_1, \ldots, r_n) = \\ &\frac{1}{n-1}(|\{j \mid F_i(r_i) < F_j(r_j)\}| - |\{j \mid F_i(r_i) > F_j(r_j)\}|) \\ &\qquad + 2F_i(r_i) - \frac{2}{n-1}\sum_{j \neq i} F_j(r_j)\ .\end{aligned}$$

This mechanism creates incentives for each agent i to report a random value drawn from F_i, as shown in the next theorem.

Theorem 1. *The random generator mechanism satisfies the following properties:*

1. *The mechanism is budget-balanced.*
2. *If $F_1 = \cdots = F_n$, the mechanism is anonymous.*
3. *The payments take values in the range $[-1, 1]$.*
4. *There exists a unique Nash equilibrium, corresponding to each agent i reporting a random number drawn according to F_i.*

Proof. Let $\mathbb{1}_B$ be the function that equals 1 if the boolean statement B is true, and 0 otherwise.

Items 1., 2., and 3. are easily shown. The proof appears in appendix.

Item 4. We proceed in two steps. We begin by showing that each agent i choosing a value r_i at random from F_i is a Nash equilibrium, then we show the equilibrium is unique.

Let $\mathcal{I} = [a, b]$, and consider any particular agent i. Assume that all other agent $j \neq i$ chooses to report a value r_j distributed according to F_j. Let r_k be the report of any agent k.

Writing the payment of i as

$$\begin{aligned}\Pi_i(r_1, \ldots, r_n) = \\ &\frac{2}{n-1}\sum_{j \neq i}\left[\frac{1}{2}\mathbb{1}_{F_i(r_i) < F_j(r_j)} - \frac{1}{2}\mathbb{1}_{F_i(r_i) > F_j(r_j)} + F_i(r_i) - F_j(r_j)\right]\ ,\end{aligned}$$

we get the expected payment for agent i, given her report r_i:

$$\begin{aligned}
\mathop{\mathbb{E}}_{r_j \sim F_j, j\neq i}[\Pi_i(r_1,\ldots,r_n)] &= \frac{2}{n-1}\sum_{j\neq i}\Bigg[\int_a^{F_j^{-1}(F_i(r_i))}\left(-\frac{1}{2}\right) f_j(r_j)\,\mathrm{d}r_j \\
&\quad + \int_{F_j^{-1}(F_i(r_i))}^b \left(\frac{1}{2}\right) f_j(r_j)\,\mathrm{d}r_j \\
&\quad + \int_a^b (F_i(r_i) - F_j(r_j)) f_j(r_j)\,\mathrm{d}r_j\Bigg] \\
&= \frac{1}{n-1}\sum_{j\neq i}\Bigg[-\frac{F_i(r_i)}{2} + \frac{1-F_i(r_i)}{2} + F_i(r_i) \\
&\quad - \int_a^b F_j(r_j) f_j(r_j)\,\mathrm{d}r_j\Bigg] \\
&= 0\ .
\end{aligned}$$

Therefore the expected payment of agent i is null for any report r_i. In particular, a randomized value distributed according to F_i is a best response. By symmetry, each agent i choosing a value at random drawn from F_i is a Nash equilibrium.

We now prove that the Nash equilibrium is unique. For all $1 \le i \le n$, let G_i be (cumulative) distributions such that each agent i choosing to report a random value distributed according to G_i is a Nash equilibrium (with the convention that pure strategies correspond to point mass distributions).

Agent i's expected payment is

$$\mathop{\mathbb{E}}_{r_j \sim G_j, j\neq i}[\Pi_i(r_1,\ldots,r_n)] = \frac{1}{n-1}\sum_{j\neq i}\int_a^b H_j(F_i(r_i))\,\mathrm{d}G_i(r_i) \tag{1}$$

under the Riemann-Stieltjes integral, with

$$\begin{aligned}
H_j(\alpha) = \int_a^{F_j^{-1}(\alpha)}\left(-\frac{1}{2}\right)\mathrm{d}G_j(r_j) + \int_{F_j^{-1}(\alpha)}^b \frac{1}{2}\,\mathrm{d}G_j(r_j) \\
+ \int_a^b [\alpha - F_j(r_j)]\,\mathrm{d}G_j(r_j)\ .
\end{aligned}$$

After simplification and rearranging the terms,

$$H_j(\alpha) = \left[\alpha - G_j(F_j^{-1}(\alpha))\right] - \left[-\frac{1}{2} + \int_a^b F_j(r_j)\,\mathrm{d}G_j(r_j)\right]\ . \tag{2}$$

By Hewitt's theorem [9], we may integrate by parts the second term and apply the change of variable $y = F_j(x_j)$:

$$\begin{aligned}
-\frac{1}{2} + \int_a^b F_j(r_j)\,\mathrm{d}G_j(r_j) &= -\frac{1}{2} + [F_j(r_j)G_j(r_j)]_a^b + \int_a^b G_j(r_j)f_j(r_j)\,\mathrm{d}r_j \\
&= \frac{1}{2} - \int_a^b G_j(r_j)f_j(r_j)\,\mathrm{d}r_j \\
&= \frac{1}{2} - \int_0^1 G_j(F_j^{-1}(x))\,\mathrm{d}x \\
&= \int_0^1 \left[x - G_j(F_j^{-1}(x))\right]\,\mathrm{d}x \ ,
\end{aligned}$$

where we observed that $1/2 = \int_0^1 x\,\mathrm{d}x$. We replace the last term of (2) and get

$$H_j(\alpha) = \Gamma_j(\alpha) - \int_0^1 \Gamma_j(y)\,\mathrm{d}y \ , \tag{3}$$

where we defined $\Gamma_j(y) = y - G_j(F_j^{-1}(y))$. Let

$$\zeta_i = \frac{1}{n-1}\sum_{j\neq i} \Gamma_j \ . \tag{4}$$

By putting together (1), (3) and (4), we get

$$\mathop{\mathbb{E}}_{r_j\sim G_j, j\neq i}[\Pi_i(r_1,\ldots,r_n)] = \int_a^b \left[\zeta_i(F_i(r_i)) - \int_0^1 \zeta_i(y)\,\mathrm{d}y\right]\mathrm{d}G_i(r_i) \ . \tag{5}$$

Suppose by contradiction that there exists i such that $\zeta_i \neq 0$. Then we show that there exists some possible report r_i^* such that agent i choosing r_i^* makes a positive expected payment.

We first prove that $\int_0^1 \zeta_i(y)\,\mathrm{d}y < \sup\zeta_i$. If the inequality is false, then $\zeta_i = \sup\zeta_i$ almost everywhere, however since $\zeta_i \neq 0$ and $\zeta_i(0) = \zeta_i(1) = 0$, we can choose $y < 1$ such that $\zeta_i(y) < \sup\zeta_i$. As G_j and F_j^{-1} are nondecreasing, $\Gamma_j(y+\epsilon) < \sup\zeta_i$ for $\epsilon > 0$ small enough, so that ζ_i does not almost everywhere equal $\sup\zeta_i$.

Since $\int_0^1 \zeta_i(y)\,\mathrm{d}y < \sup\zeta_i$, there exists y^* such that $\zeta_i(y^*) - \int_0^1 \zeta_i(y)\,\mathrm{d}y > 0$, and so by taking $r_i^* = F_i^{-1}(y^*)$, we find that agent i choosing the pure strategy r_i^* would make a positive expected payment according to (5).

Since i plays a Nash equilibrium, i's strategy is a best response and her expected payment is at least that obtained by choosing the pure strategy r_i^* and so is strictly positive. Therefore, if the Nash equilibrium is such that $\zeta_i \neq 0$ for some i, then i's payment is strictly positive, otherwise $\zeta_i = 0$ and i's expected profit is null. So if there exists at least on agent i such that $\zeta_i \neq 0$, we have

$$\sum_j \mathop{\mathbb{E}}_{r_k\sim G_k, k\geq 1}[\Pi_j(r_1,\ldots,r_n)] > 0$$

which is impossible as the mechanism is budget-balanced. Hence for all i, $\zeta_i = 0$, which implies $y = G_j(F_j^{-1}(y))$: the only possible Nash equilibrium corresponds to $G_j = F_j$, for all j.

3.2 Graphical interpretation

The *lazy hiker race* gives an intuitive interpretation of our mechanism, and may be described as follows. A group of n hikers starts a march on a mountain with 1 mile high. Each hiker has a designated trail, which is common knowledge. All the trails share a common starting/ending point. Hikers are able to keep track of the distance they cover along their own trail, but cannot observe the progression of others. After 10 hours, the march stops and hikers are ranked in decreasing order of altitude.

Hikers want to win the race, and are strong enough to climb to the top within the time limit. However they are lazy and prefer to win by making as little effort as possible. When there are two hikers, the winner gets the maximum satisfaction (+1) when he wins by being just above the other hiker. He gets the worse satisfaction (−1) when he wins by being at the top, while the other hiker remains at the bottom. The loser's satisfaction is the opposite of the winner's. Satisfaction is linear in the difference of altitude, so that the winner is indifferent between winning and losing when the difference of altitude between the hikers is 0.5 miles. Figure 1 illustrates the case of two hikers. If there are more than two hikers, the satisfaction of a hiker equals the average satisfaction when he compares himself to each other hiker.

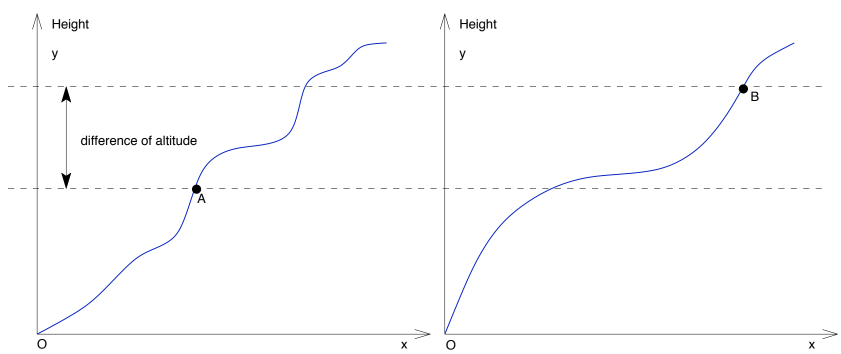

Fig. 1: Profile of trails for two hikers A and B.

Hikers strategize about the distance they should cover so as to maximize their average satisfaction. We observe that, if $F_i(d_i)$ denotes the altitude reached by hiker i after walking a distance[1] d_i, the payment $\Pi_i(d_1, \ldots, d_n)$ of the random

[1] For simplification d_i is normalized and equals to the proportion of the total distance covered.

generator mechanism described previously corresponds exactly to the satisfaction of hiker i when hikers $1, \ldots, n$ cover the respective distances $d_1, \ldots, d_n$ (where we take $\mathcal{I} = [0, 1]$).

To simplify our argument, let's take the case of two hikers. We first note that there is no pure Nash equilibrium: if the loser knows where the winner is, he will change his strategy to place himself slightly above the winner. Therefore hikers should cover a random distance. When one hiker chooses to cover a distance so that his altitude is uniformly random, each hiker gets a null satisfaction on average, no matter what the other hiker decides to do. If, however, one hiker will likely stop at a low altitude, the other hiker would get a likely positive satisfaction by stopping at a medium altitude. More generally, when a hiker makes frequent stops at some altitudes, the other can choose a location so as to get a positive expected satisfaction. Therefore any choices of random distances that result in nonuniform distribution of altitudes cannot lead to a Nash equilibrium. A similar argument applies to groups of any size. Note that uniform distributions of altitudes are obtained only when each hiker i covers a distance d_i chosen at random according to the distribution F_i. For a given distribution F with density $f = F'$, one can verify that the trail with profile given by

$$x(y) = \int_0^y \sqrt{\frac{L}{f(F^{-1}(h))^2} - 1} \, \mathrm{d}h$$

will generate the mixed-Nash equilibrium strategy with distribution F, where L is the desired length of the trail, with $L > \max f^2$. Figure 2 shows some density functions and their associated trails.

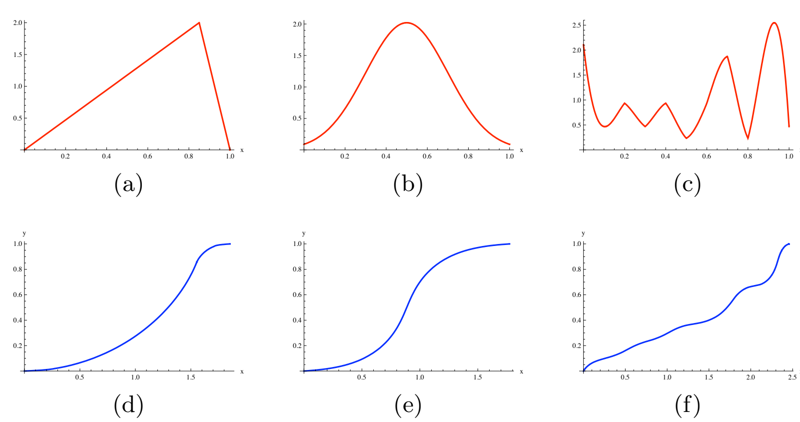

Fig. 2: The equilibrium strategy for trails with profiles (d), (e) and (f) is to choose a random distance with respective densities (a), (b) and (c).

4 Mechanisms for truthful surveys

We now describe our survey mechanisms. Given a random sample of k opinions $x_1, \ldots, x_k$, let $\tilde{F}^{x_1,\ldots,x_k}(x)$ be a statistical estimate of the proportion of the population having an opinion less than or equal to x. In practice, it is common to use a probabilistic model with parameterized densities. The maximum-likelihood parameters may for example be obtained through the Expectation-Maximization algorithm [10]. The statistical estimate is said to be *unbiased* when, for all x,

$$\underset{X_1,\ldots,X_k \sim F}{\mathbb{E}} \left[\tilde{F}^{X_1,\ldots,X_k}(x)\right] = F(x)$$

for all population distribution F. For example, the following empirical distribution is unbiased:

$$\tilde{F}^{x_1,\ldots,x_k}(x) = \frac{1}{k} \sum_{1 \leq i \leq k} \mathbb{1}_{x_i < x} \ .$$

with $\mathbb{1}_{x_i<x} = 1$ if $x_i < x$ and is null otherwise.

Let $\mathcal{G}_1, \ldots, \mathcal{G}_k$ be a partition of the n agents into k groups, $k \geq 2$. Let $\mathcal{G}(i)$ be the group that includes agent i, and let $\mathcal{S}_i = \{1, \ldots, n\} \backslash \mathcal{G}(i)$ be the set of agents that doesn't include the group containing i. Our *basic survey mechanism* uses the random generator mechanism of the previous section to incentivize each agent i to reveal an opinion that corresponds to an empirical distribution of the distribution population. For a given interval of possible opinions $\mathcal{I}$ and a number of agents n, the payments of our survey mechanism are defined by

$$\begin{aligned} \Pi_i(r_1, \ldots, r_n) = \frac{1}{|\mathcal{G}(i)| - 1} &\left[|\{j \mid r_i < r_j\}| - |\{j \mid r_i > r_j\}| \right] \\ &+ 2\tilde{F}_i(r_i) - \frac{2}{|\mathcal{G}(i)| - 1} \sum_{j \in \mathcal{G}(i), j \neq i} \tilde{F}_j(r_j) \end{aligned}$$

where $\tilde{F}_i(x) = \tilde{F}^{\{r_j\}_{j \in \mathcal{S}_i}}(x)$ is an unbiased statistical estimate of $F(x)$ given by the reports of agents in $\mathcal{S}_i$. We prove the following:

Theorem 2. *The basic survey mechanism is budget-balanced, anonymous, and accurate.*

Proof. Budget-balance and anonymity are easily obtained following the steps of Theorem 1. To show accuracy, let i be any agent, and assume all agents $j \neq i$ report an opinion representative of the population distribution. Let further assume w.l.o.g. that $\mathcal{G}(i) = \{1, \ldots, \ell\}$. The expected profit of agent i given the reports of other agents in i's group is

$$\underset{r_{\ell+1},\ldots,r_n \sim F}{\mathbb{E}} [\Pi_i(r_1, \ldots, r_n) | r_1, \ldots, r_\ell] = \tilde{\Pi}_i(r_1, \ldots, r_\ell) \ ,$$

with

$$\tilde{\Pi}_i(r_1, \ldots, r_\ell) = \frac{1}{\ell - 1}(|\{j \mid F(r_i) < F(r_j)\}| - |\{j \mid F(r_i) > F(r_j)\}|) + 2F(r_i) - \frac{2}{\ell - 1} \sum_{1 \leq j \leq l, j \neq i} F(r_j)$$

since the estimates are unbiased, where we observed that $r_i < r_j$ iff $F(r_i) < F(r_j)$ as F is strictly increasing.

By Theorem 1, for all r_i,

$$\mathop{\mathbb{E}}_{r_j \sim F, j \neq i}[\tilde{\Pi}_i(r_1, \ldots, r_\ell)] = 0 \ .$$

Hence,

$$\mathop{\mathbb{E}}_{r_j \sim F, j \neq i}[\Pi_i(r_1, \ldots, r_n)] = \mathop{\mathbb{E}}_{r_j \sim F, j \neq i}\left[\mathop{\mathbb{E}}_{r_j \sim F, j \neq i, j \leq \ell}[\tilde{\Pi}_i(r_1, \ldots, r_\ell)]\right] = 0 \ ,$$

so agent i's expected profit is null for all reports. In particular, reporting a value ex-ante distributed according to F is a possible best-response.

Note that payments take values in the interval $[-3, 3]$. Payments near the interval bounds occur only with low-quality empirical distributions. When the estimates of the population distribution converge to the true population distribution, payments are restricted to the interval $[-1, 1]$. One may also offset/rescale the payments, to get for example payments in the interval $[0, 1]$ so as to provide participation incentives and strict individual rationality.

The mechanism easily adapts to the case of sequential elicitation—as opposed to simultaneous elicitation—often desired with online surveys. The surveyor should form groups of 2 or 3 people, progressively as new reports come in, and reward individuals of a group as soon as the group is finalized. Distribution estimates should be computed from reports of the previous groups only, with the exception of the distribution estimates used for rewarding the first group, which could take as input reports of the second group.

As opposed to the work of Miller et al. [5] and Jurca and Faltings [6], truthful revelation is a non-strict equilibrium of our mechanism. This limitation is due to the lack of common prior: with general knowledge structures, no survey mechanism may implement truthful reporting as a strict Nash equilibrium. Indeed, the expected payment for an agent who reports her true opinion must be maximized under all possible distributions of opinions, and therefore must be constant for all possible reports. We observe by the same argument than any survey mechanism that implements reporting one's true opinion as a Nash equilibrium is also accurate, in the sense of Section 2.

All methods for eliciting subjective information suffer from the multiplicity of Nash equilibria, since rewards can only be a function of information submitted by the agents. However, unlike other methods such as Miller et al. [5] in which non-truthful Nash-equilibria may lead to higher revenue for all agents, in our mechanism all Nash equilibria lead to a null expected payment for all agents.

If there are trusted individuals, we can ensure uniqueness of the Nash equilibria that correspond to accurate samples, so that the surveyor is guaranteed to obtain true random samples. Let $\mathcal{T}$ be a group of trusted individuals who provide their true opinion. $\mathcal{T}$ may not be empty but can be of any positive size, larger groups are generally preferred as they reduce the variance of individual payments. For simplicity we assume that trusted agents form a separate group from the n surveyed agents. The payments of our *trusted-survey mechanism* are defined as follows:

$$\Pi_i(r_1, \ldots, r_n) = \frac{1}{n-1}\left[|\{j \mid r_i < r_j\}| - |\{j \mid r_i > r_j\}|\right] + 2\tilde{F}(r_i) - \frac{2}{n-1}\sum_{j \neq i} \tilde{F}(r_j) \ ,$$

where $\tilde{F}(x) = \tilde{F}^{\{r_j\}_{j \in \mathcal{T}}}(x)$ is an unbiased statistical estimate of $F(x)$ given by the reports of trusted agents in $\mathcal{T}$.

Our next theorem claims that the trusted-survey mechanism is guaranteed to elicit true random samples of opinions. The proof follows from Theorems 1 and 2 and is shown in appendix.

Theorem 3. *The trusted-survey mechanism is budget-balanced, anonymous, and strongly accurate.*

5 Conclusion

We have investigated the problem of incentivizing individuals to elicit accurate samples of opinions that are representative of a population. We have proposed a budget-balanced, anonymous mechanism for which reporting a true sample of opinion, in particular reporting one's true opinion, is a Nash equilibrium for risk-neutral individuals. When some opinions can be trusted, we propose a variation of our mechanism which guarantees that the only Nash equilibria correspond to providing true samples. Although we use a Nash implementation as opposed to a dominant strategy implementation—impossible to achieve in our setting—our results do not depend on the knowledge structure of the population, in particular we do not make use of a common prior.

We believe an important avenue for future work is that of empirical studies. Our analysis has focused on theoretical considerations. However, it is not clear how individuals would behave in practice. Experiments studying and comparing our approach with those whose payments depend on a common prior, either provided by the surveyor as in Miller et al. [5], or provided by the agents as in Prelec [7], or simply assumed by all agents as in Jurca and Faltings [8], would need to be performed to help assess the validity of each method, and their potential applicability to practical contexts. In particular, more work would be needed to understand the limitations raised by the common prior assumption and those implied by the weakness of the Nash equilibria in our mechanisms.

References

1. Montgomery, D.: Design and Analysis of Experiments. Wiley (1984)
2. Kish, L.: Survey sampling. Wiley (1995)
3. Winkler, R., Muñoz, J., Cervera, J., Bernardo, J., Blattenberger, G., Kadane, J., Lindley, D., Murphy, A., Oliver, R., Ríos-Insua, D.: Scoring rules and the evaluation of probabilities. TEST **5**(1) (1996) 1–60
4. Lambert, N., Pennock, D., Shoham, Y.: Eliciting Properties of Probability Distributions. Proceedings of the 9th ACM Conference on Electronic Commerce (2008) 129–138
5. Miller, N., Resnick, P., Zeckhauser, R.: Eliciting Informative Feedback: The Peer-Prediction Method. Management Science **51**(9) (2005) 1359–1373
6. Jurca, R., Faltings, B.: Minimum payments that reward honest reputation feedback. Proceedings of the 7th ACM Conference on Electronic Commerce (2006) 190–199
7. Prelec, D.: A Bayesian Truth Serum for Subjective Data. Science **306**(5695) (2004) 462–466
8. Jurca, R., Faltings, B.: Incentives for Expressing Opinions in Online Polls. Proceedings of the 9th ACM Conference on Electronic Commerce (2008) 119–128
9. Hewitt, E.: Integration by Parts for Stieltjes Integrals. The American Mathematical Monthly **67**(5) (1960) 419–423
10. Dempster, A., Laird, N., Rubin, D., et al.: Maximum likelihood from incomplete data via the EM algorithm. Journal of the Royal Statistical Society **39**(1) (1977) 1–38

Proof of items 1,2 and 3 of Theorem 1

Proof. 1. For all $r_1, \dots, r_n$,

$$\begin{aligned}
\sum_i \Pi_i(r_1, \dots, r_n) &= \frac{1}{n-1}\left[\sum_{i \neq j} \mathbb{1}_{F_i(r_i) < F_j(r_j)} - \sum_{i \neq j} \mathbb{1}_{F_j(r_j) < F_i(r_i)}\right] \\
&\quad + 2\sum_i F_i(r_i) - \frac{2}{n-1}\sum_{i \neq j} F_j(r_j) \\
&= \frac{1}{n-1}\sum_{i \neq j}(\mathbb{1}_{F_i(r_i) < F_j(r_j)} - \mathbb{1}_{F_i(r_i) < F_j(r_j)}) \\
&\quad + 2\sum_i F_i(r_i) - 2\sum_i F_i(r_i) \\
&= 0\ .
\end{aligned}$$

2. Let $F_1 = \cdots = F_n$, and $1 \leq i \leq n$. Then,

$$\begin{aligned}\Pi_{\sigma(i)}(r_{\sigma^{-1}(1)}, \ldots, r_{\sigma^{-1}(n)}) = \frac{2}{n-1} \sum_{j \neq \sigma(i)} \Bigg[& \frac{1}{2} \mathbb{1}_{F_{\sigma(i)}(r_{\sigma^{-1}(\sigma(i))}) < F_j(r_j)} \\ & - \frac{1}{2} \mathbb{1}_{F_{\sigma(i)}(r_{\sigma^{-1}(\sigma(i))}) > F_j(r_j)} \\ & + F_{\sigma(i)}(r_{\sigma^{-1}(\sigma(i))}) - F_j(r_j) \Bigg] \\ = \Pi_i(r_1, \ldots, r_n) \ . \end{aligned}$$

3. For all i, and all $r_1, \ldots, r_n$,

$$\Pi_i(r_1, \ldots, r_n) = \frac{2}{n-1} \sum_{j \neq i} \chi_{i,j}(r_i, r_j)$$

with

$$\chi_{i,j}(x_i, x_j) = \frac{1}{2} \mathbb{1}_{F_i(r_i) < F_j(r_j)} - \frac{1}{2} \mathbb{1}_{F_i(r_i) > F_j(r_j)} + F_i(r_i) - F_j(r_j) \ .$$

As $0 \leq F_i, F_j \leq 1$, if $F_i(r_i) < F_j(r_j)$, $-1 < F_i(r_i) - F_j(r_j) < 0$, and if $F_i(r_i) > F_j(r_j)$, $0 < F_i(r_i) - F_j(r_j) < 1$, hence $-1/2 \leq \chi_{i,j}(r_i, r_j) \leq 1/2$, and $-1 \leq \Pi_i(r_1, \ldots, r_n) \leq 1$.

Proof of Theorem 3

Proof. The proof is a direct consequence of Theorem 1 and is similar to Theorem 2, observing that, for all i,

$$\mathop{\mathbb{E}}_{r_j \sim F, j \in \mathcal{T}} [\Pi_i(r_1, \ldots, r_n)] = \tilde{\Pi}_i(r_1, \ldots, r_n) \ ,$$

with

$$\begin{aligned}\tilde{\Pi}_i(r_1, \ldots, r_n) = \frac{1}{n-1} (|\{j \mid F(r_i) < F(r_j)\}| - |\{j \mid F(r_i) > F(r_j)\}|) \\ + 2F(r_i) - \frac{2}{n-1} \sum_{j \neq i} F(r_j) \ .\end{aligned}$$