Mechanisms for Partial Information Elicitation: The Truth, but Not the Whole Truth

Aviv Zohar and **Jeffrey S. Rosenschein**
School of Engineering and Computer Science
The Hebrew University of Jerusalem
Jerusalem, Israel
{avivz, jeff}@cs.huji.ac.il

Abstract

We examine a setting in which a buyer wishes to purchase probabilistic information from some agent. The seller must invest effort in order to gain access to the information, and must therefore be compensated appropriately. However, the information being sold is hard to verify and the seller may be tempted to lie in order to collect a higher payment.

While it is generally easy to design information elicitation mechanisms that motivate the seller to be truthful, we show that if the seller has additional relevant information it does not want to reveal, the buyer must resort to elicitation mechanisms that work only some of the time. The optimal design of such mechanisms is shown to be computationally hard.

We show two different algorithms to solve the mechanism design problem, each appropriate (from a complexity point of view) in different scenarios.

Introduction

The old aphorism "Knowledge is power", stated by Sir Francis Bacon some four centuries ago, is more relevant now than ever. The need to make informed choices causes correct and accurate information to be a desired and highly-valued commodity. As intelligent artificial agents take on more tasks, and need to act independently within large systems, their need to buy and sell information increases.

The problem with information in stochastic environments is that it is hard to evaluate, and may be easily faked. Any novice can give a prediction regarding the behavior of tomorrow's stock market; by pure chance, those predictions may outperform those of even the most informed financial wizard.

The question that naturally arises is how to pay for information that can only be verified with some probability. This is especially important in cases where in order to obtain the information, the seller itself has to invest some effort. The payments made by the buyer must be carefully set so as to induce the seller to invest the effort into acquiring the true information. Otherwise, the seller might be tempted to avoid the cost of obtaining the information, and simply make something up.

Our Setting

We study a simple information transaction between a pair of agents — a seller and a buyer. The buyer is assumed to be interested in some of the information the seller possesses, while the seller only wishes to gain a payment from the transaction, and is not interested in any other aspect of the information it is passing. The buyer and seller are assumed to be in a one-shot interaction. The buyer therefore has no means of punishing nor of rewarding the seller through future interactions.

Having less knowledge about the exact state of the world, the buyer is naturally at a disadvantage. It must operate under these conditions and decide how to pay for knowledge it does not yet have. In fact, we shall look into situations where the asymmetry in knowledge is even worse — besides the information being sold, the seller has some other relevant information that it will never divulge. This is even worse for the buyer, as it does not possess all the information needed to decide on a payment, even after the fact.

To be able to motivate the seller to give correct information, the buyer must have some way of verifying the information it buys. We assume payment is set not only according to the information that was sold, but also according to some probabilistically linked outcome that the buyer observes after the transaction.

The case of partial revelation of information is very natural to the information elicitation setting. A buyer is often interested in some very specific information, but would not want to pay for anything extra, while the seller would not want to give out any unnecessary information for reasons of privacy, or simply because it would like to sell it to some other party. The classic *mechanism design* approach of using direct and full revelation is therefore unsuitable in these settings.

Throughout the paper we adopt the point of view of the buyer that needs to design the correct payment scheme for information it is about to purchase, and make the assumption often made in game theory that the seller seeks to maximize its gains and will attempt to manipulate the result of the interaction if it can benefit from doing so.

Contribution of the Paper

This paper examines information elicitation problems with partial revelation. While it is computationally easy to design full revelation mechanisms for the problem, we show that in the case of partial revelation mechanisms, even with full knowledge of the underlying probabilities of events, the designer must resort to using mechanisms that work only part of the time, and are computationally hard to design.

We show two approaches to the mechanism design problem that yield two different algorithms. Each algorithm is appropriate under different circumstances.

Our approach is aimed at applications in large open systems, where agents need to trade information without any reputation or reliability record of the party with which they are interacting. Such scenarios are becoming more common due to the growth of the internet, and the multitude of anonymous interactions that it allows. Concrete examples might include various web services, peer-to-peer networking, and the provision of many other types of data that can be sold online.

The rest of the paper is organized as follows. The next section formally introduces the model for information elicitation, and examines some of its basic properties. In the section on Full Revelation Mechanisms, we review the design of payment schemes in cases where no information is kept private. We then turn to Partial Revelation Mechanisms, examine their limitations, and demonstrate that their design is computationally hard. We continue by presenting two different approaches to finding working mechanisms, and by presenting the algorithms they imply. We briefly review related work, and finally present our own conclusions, with some possible avenues for future research.

The Model

We assume the buyer wishes to purchase information about the value of a discrete random variable X from a seller that can learn the value of that variable at a cost c. The seller is also assumed to possess private information about a secret random variable S which it does not wish to reveal. To verify the quality of the information it purchases, the buyer has access to a random variable Ω. Ω, X, S are presumably not independent variables, and knowledge about the value of one of them gives some information regarding the value of the others. Using the variable Ω, the buyer can get some idea if the information it was sold was correct. Without Ω, it would be impossible to create the necessary incentives for truthfulness on the part of the seller. We shall denote the probability distribution for the triplet Ω, X, S by $p_{\omega,x,s} = Pr(\Omega = \omega, X = x, S = s)$. The values the different variables can take, as well as the probability distribution $p_{\omega,x,s}$, are assumed to be common knowledge.

The buyer can now design a payment scheme that will determine the payment it must give to the seller, based on the information the seller gave and on the value of the verification variable Ω. In a full revelation scheme, the agent would be asked to reveal its secret as well as the value of X. It would then be paid an amount $u_{\omega,x,s}$ that depended on the values of X, S it reported, and on the observed value of Ω which the seller used to verify the information. In a partial revelation mechanism, the agent will only be asked about X, and will be paid some amount $u_{\omega,x}$.

We assume that agents seek to maximize their expected gains and that they are risk-neutral. The precise requirements of an appropriate payment scheme are listed below.

Proper Payment Schemes

A good payment scheme must motivate the selling agent to first invest the effort into obtaining the value of X, and then to reveal the true value it found. Additionally, a payment scheme needs to be individually rational — the seller must have a positive expected utility from entering the game. These demands then translate to the following constraints we would optimally want to satisfy, given that the selling agent knows of a secret $s \in S$:

1. **Truth Telling.** Once the seller knows its variable is x, it must have an incentive to tell the true value to the buyer, rather than any lie x'.
$$\forall x, x' \quad s.t. \quad x \neq x', \quad \sum_{\omega} p_{\omega,x,s} \cdot (u_{\omega,x} - u_{\omega,x'}) > 0 \tag{1}$$
Here $p_{\omega,x,s}$ is the probability of what actually occurs, while the payment $u_{\omega,x'}$ is based on the reported value.
2. **Individual Rationality.** A seller must have a positive expected utility from participating in the game:
$$\sum_{\omega,x} p_{\omega,x,s} \cdot u_{\omega,x} > c \cdot p_s \tag{2}$$
3. **Investment.** The *value of information* for the seller must be greater than its cost. Any guess the seller makes without actually computing its value must be less profitable (in expectation) than paying to obtain the true value of the variable and revealing it:
$$\forall x' \quad \sum_{\omega,x} p_{\omega,x,s} \cdot u_{\omega,x} - c \cdot p_s > \sum_{\omega,x} p_{\omega,x,s} \cdot u_{\omega,x'} \tag{3}$$

Definition 1. *A payment scheme u shall be considered* proper *for state (s, x) if the seller has incentive to reveal x truthfully in a case where $S = s, X = x$.*

Definition 2. *We shall say that the scheme is proper for a secret s if $\forall x$, it is proper for state (s, x).*

Truth Above All

It is important to note that by rescaling and shifting the payments, we can transform a mechanism that satisfies only the truth-telling constraints described in the previous section, into a mechanism that satisfies the other constraints as well. This is because multiplying the payments by some positive constant α will not affect the truth-telling constraints, and α can be chosen to satisfy the investment constraints. Then shifting all payments uniformly by β_ω will not affect the truth-telling constraints or the investment constraints, but *can* help satisfy the individual rationality constraint.

While the resulting mechanism will not necessarily be an optimal one, its very existence demonstrates that at least some mechanism is possible. It is therefore sufficient, for purposes of feasibility, to examine solutions for the truth-telling constraints alone. We shall therefore focus our efforts on understanding the structure of the truth-telling constraints.

The Geometric Interpretation of the Truth-Telling Constraints We shall write the truth-telling constraints using vector notation in a way that will help demonstrate their geometric and algebraic properties.

Let us denote the vector $\vec{u}_x$ as the vector of payments $\vec{u}_x = (u_{\omega 1,x} \dots u_{\omega n,x})$ and let us denote by the vector $\vec{p}_{x,s}$ the vector of probabilities $\vec{p}_{x,s} = (p_{\omega 1,x,s} \dots p_{\omega n,x,s})$. We shall also define the vectors $\vec{v}_{x,x'} = \vec{u}_x - \vec{u}_{x'}$. The truth-telling constraints for some secret s can now be written as follows:

$$\forall x, x' \quad \vec{p}_{x,s} \cdot (\vec{u}_x - \vec{u}_{x'}) = \vec{p}_{x,s} \cdot \vec{v}_{x,x'} > 0 \qquad (4)$$

Now, if we think of $\vec{v}_{x,x'}$ as a vector that defines a hyperplane through the origin,[1] we can see that the truth-telling constraint for x, x' simply states that the vector $\vec{p}_{x,s}$ is found on the positive side of this hyperplane. On the other hand, the matching constraint in which x' and x are exchanged gives us $\vec{p}_{x',s} \cdot \vec{v}_{x,x'} < 0$ which means that the vector $\vec{p}_{x',s}$ is found on the negative side of the hyperplane defined by $\vec{v}_{x,x'}$. We can thus find a proper mechanism for a secret s if we manage to separate all vectors $\vec{p}_{x,s}$ correctly using linear separators. Note however, that we are not completely free to select $\vec{v}_{x,x'}$ vectors independently, and we are in fact constrained to satisfy, for all x, x', x'':

$$\vec{v}_{x,x'} = -\vec{v}_{x',x} \quad ; \quad \vec{v}_{x,x''} = \vec{v}_{x,x'} + \vec{v}_{x',x''} \qquad (5)$$

This simple geometric interpretation will shed some light on the difficulties of designing elicitation mechanisms.

Full Revelation Mechanisms

When the seller reveals all the relevant information it possesses, it is computationally easy to design a mechanism that will provide motivation for truth-telling:

Proposition 1. *A full revelation mechanism can be designed in polynomial time.*

Proof: We simply take a linear program composed of all the constraints — for all possible secrets, as was shown above. Such a linear program is solvable in polynomial time using currently known optimization techniques (Bertsimas & Tsitsiklis 1997). □

In fact, such a mechanism will exist, except for some singular cases:

Proposition 2. *A fully revealing mechanism that motivates truthfulness exists iff all vectors $\vec{p}_{x,s}$ are pairwise linearly independent.*

Proof: If there is a pair of vectors that is linearly dependent, it cannot be separated by a hyperplane that passes through the origin, and therefore no payment scheme will be able to satisfy the truth-telling constraints.

Otherwise, no pair of vectors is linearly dependent, and we shall simply demonstrate a proper elicitation mechanism. For a given problem $\vec{p}$ we shall set the payments to be:

$$\vec{u}_{x,s} = \frac{\vec{p}_{x,s}}{||\vec{p}_{x,s}||} \qquad (6)$$

Which gives:

$$\vec{p}_{x,s} \cdot (\vec{u}_{x,s} - \vec{u}_{x',s'}) = \vec{p}_{x,s} \cdot (\frac{\vec{p}_{x,s}}{||\vec{p}_{x,s}||} - \frac{\vec{p}_{x',s'}}{||\vec{p}_{x',s'}||}) =$$

$$= ||\vec{p}_{x,s}|| - \frac{\vec{p}_{x,s} \cdot \vec{p}_{x',s'}}{||\vec{p}_{x',s'}||} \geq 0 \qquad (7)$$

where the last inequality is due to the Cauchy-Schwarz inequality and is a strict inequality whenever $\vec{p}_{x,s}$ and $\vec{p}_{x',s'}$ are independent. This is proof that the payment vector we selected satisfies the truth-telling constraints. Now, with scaling and shifting it can be adjusted to satisfy the other constraints as well. □

Partial Revelation Mechanisms

In partial revelation mechanisms, the seller will keep its private variable S secret. The payment it receives only depends on Ω and X. When designing partial revelation mechanisms, there are often probability distributions that do not allow us to construct an effective mechanism for *all* possible secrets the seller may hold. The example in Figure 1 demonstrates such a case.

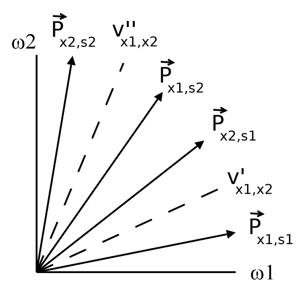

The 2 axes correspond to the probabilities of the two possible results, so all probability vectors are in the 2D plane.

Figure 1: An Elicitation Scenario with 2 Possible Results, and 2 Possible Secrets

It is impossible to find a separating hyperplane that will separate $\vec{p}_{x1,s1}$ from $\vec{p}_{x2,s1}$ and at the same time separate $\vec{p}_{x1,s2}$ from $\vec{p}_{x2,s2}$. The hyperplanes v' and v'' work only for a single secret each. Since the buyer is never told about the actual secret s, he has no way of creating the incentives for truthfulness in both cases.

We must therefore settle on building a mechanism that will work only part of the time. We will naturally aspire to have a good confidence level in our mechanism — to build a mechanism that will work with high probability. There are two possible alternatives we examine here:

1. A single-use, disposable mechanism — where we design the mechanism for only a single transaction. We then want the buyer's confidence in the received answer to be high:

$$\theta_1 = Pr_{s,x}(u \text{ is proper for state } (s,x))$$

[1] The hyperplane is defined as the collection of all vectors perpendicular to $\vec{v}_{x,x'}$.

2. A reusable mechanism — where we design the mechanism for multiple transactions. Here, we want the buyer to have high confidence that once the secret s has been set, he will hear the truth for all possible cases of X.

$$\theta_2 = Pr_s(u \text{ is proper for secret } s)$$

Complexity of Partial Revelation Mechanism Design

Proposition 3. *Deciding if there exists a reusable revelation mechanism with a confidence level over some threshold θ is NP-Complete. Furthermore, the problem of finding the mechanism with the maximal confidence level cannot be approximated within any constant.*

The design problem is in NP. This is because if we are given access to an oracle that tells us which secrets to try and satisfy and which to give up on, we can find a payment scheme that satisfies the right constraints in polynomial time. This is achieved by solving the linear program that consists of the constraints for all the included secrets.

We shall show that constructing fully operational mechanisms is NP-Complete by presenting a reduction from the Independent Set problem. The Independent Set problem, in addition to being NP-Complete, is also hard to approximate. The reduction we shall present is a cost-preserving reduction and therefore demonstrates that our problem is just as hard to approximate as Independent Set.

Given an undirected graph $G(V, E)$ and an integer $k \leq |V|$ the *Independent Set* decision problem is defined as the problem of deciding whether there is a set of vertices $W \subset V$ so that $|W| \geq k$, and such that for every edge $e \in E$, e does not occur on more than one vertex in V.

Proof: [to Proposition 3] Given an Independent Set problem $(G(V, E), k)$ we shall construct a mechanism design problem $(\Omega, X, S, P, \theta)$ in the following manner:

$$\Omega = \bigcup_{e \in E} \{\omega_{e1}, \omega_{e2}\} \quad ; \quad X = \bigcup_{e \in E} \{x_{e1}, x_{e2}\}$$

$$S = V \quad ; \quad \theta = \frac{k}{|V|} \tag{8}$$

We denote by $\vec{\delta}_i$ the vector that is 0 at all coordinates except for coordinate i, where it takes the value of 1, and by α a normalizing constant that equals $\alpha = \frac{1}{2|E||V|}$. P is then defined as follows:

if $v \notin e$ then:

$$\vec{p}_{x_{e1},v} = \alpha \cdot \vec{\delta}_{\omega_{e1}} \quad ; \quad \vec{p}_{x_{e2},v} = \alpha \cdot \vec{\delta}_{\omega_{e2}} \tag{9}$$

otherwise $e = \{v1, v2\}$ and we set:

$$\vec{p}_{x_{e1},v1} = \alpha \cdot \vec{\delta}_{\omega_{e1}} \quad ; \quad \vec{p}_{x_{e2},v1} = \frac{\alpha}{2} \cdot (\vec{\delta}_{\omega_{e1}} + \vec{\delta}_{\omega_{e2}})$$

$$\vec{p}_{x_{e1},v2} = \frac{\alpha}{2} \cdot (\vec{\delta}_{\omega_{e1}} + \vec{\delta}_{\omega_{e2}}) \quad ; \quad \vec{p}_{x_{e2},v2} = \alpha \cdot \vec{\delta}_{\omega_{e2}} \tag{10}$$

With the above construction all secrets have the same probability of occurring:

$$Pr(S = s) = \sum_{\omega,x} p_{\omega,x,s} = 2|E|\alpha = \frac{1}{|V|} \tag{11}$$

Below we sketch the two steps needed to complete the proof:

1. **If the graph G has an independent set of size k then there is a mechanism with a confidence level above the threshold θ.**

   Let us assume that G has an independent set $W \subset V$ of size k. We shall build a payment scheme that will give a proper mechanism for all the secrets matching the vertices in W. For an edge e that has one of its vertices in the independent set,[2] we shall define:

   $$\vec{u}_{x_{e1}} = \frac{\vec{p}_{x_{e1},v}}{||\vec{p}_{x_{e1},v}||} \quad ; \quad \vec{u}_{x_{e2}} = \frac{\vec{p}_{x_{e2},v}}{||\vec{p}_{x_{e2},v}||} \tag{12}$$

   where v is the vertex (from edge e) that was selected for the independent set. If on the other hand e did not have any vertex in the independent set, we simply set

   $$\vec{u}_{x_{e1}} = \vec{\delta}_{\omega_{e1}} \quad ; \quad \vec{u}_{x_{e2}} = \vec{\delta}_{\omega_{e2}} \tag{13}$$

   Due to space limitations, we omit the proof that this selection of payments does indeed satisfy the constraints.

2. **If there is a good mechanism with confidence level above θ then there is an independent set of size k in the graph.**

   Since there is a confidence level of $\frac{k}{|V|}$, there must be at least k satisfied secrets in the mechanism. Each such secret matches a vertex in the original problem. It remains to show that the set W of vertices matching satisfied secrets is independent. Assuming the opposite leads to a contradiction. The secrets matching two vertices that are connected by an edge cannot be satisfied at the same time due to the way the problem was constructed. The probability vectors for each edge ($\vec{p}_{x_{e1},v1}$, $\vec{p}_{x_{e1},v2}$, $\vec{p}_{x_{e2},v1}$, $\vec{p}_{x_{e2},v2}$) were placed in a separate two-dimensional space, and were set similarly to the vectors in Figure 1 — in a way that assures that both pairs cannot be linearly separated at the same time.

□

The high complexity of designing proper mechanisms applies in the single-use, disposable case as well.

Proposition 4. *Deciding if there exists a single-use revelation mechanism with a confidence level over some threshold θ is also NP-Complete.*

Proof: [sketch] A reduction can be shown from the problem Max-Hyperplane-Consistency which is also NP-Complete (Amaldi & Kann 1995). This is the problem of finding a biased hyperplane that will correctly separate a maximal number of points from two sets: a positive set that must be on the positive side of the hyperplane, and a negative set that must be on the opposite side. □

Finding Partial Revelation Mechanisms

We now present two approaches to computing a partial revelation mechanism for a given problem $p_{\omega,x,s}$. As we have already seen, the problem of finding such a mechanism is NP-Complete, and unless P=NP, we cannot hope to locate the

[2] It cannot have both its vertices in the set — only one or none.

optimal mechanism in polynomial time in all cases. However, in some cases, the problem may be simpler than the worst possible case. The two approaches we present differ in the complexity of the algorithm. One algorithm will be better in cases where $|S|$ is small, while the other will be better in cases where $|\Omega| \cdot |X|$ is small.

The algorithms we present are for reusable mechanisms. Similar versions can be constructed for the single-use case.

Considering All Combinations of Secrets

The reductions we used in the proofs of Propositions 3 and 4 both relied on the difficulty of selecting the cases in which we wish the mechanism to work. If we had an oracle that shows us which constraints to try and satisfy, we could easily construct a mechanism. Since we do not possess such an oracle, we can try every possible combination by brute force.

Algorithm 1 [Reusable Mechanism Construction]:

1. For all $W \in 2^S$:
   (a) Locate a mechanism that satisfies all constraints for all secrets in W.
   (b) If such a mechanism exists, compute $\theta_W = \sum_{s \in W} p_s$.
2. Return a mechanism for secrets $\arg\max_W(\theta_W)$.

In the algorithm above, there are $2^{|S|}$ ways to select secrets to satisfy. Each selection then requires $poly(|S||X||\Omega|)$ time to check for feasibility. This therefore gives a running time of $O(2^{|S|} \cdot poly(|S||X||\Omega|))$ which can still be efficient if the number of possible secrets is small.

The Geometric Approach — Partitioning into Cells

The second approach we shall examine is based on a geometric interpretation of the problem. The linear constraint for the mechanism design problem partitions the space of payment vectors into cells. Each cell is a region of the space for which some set of constraints holds, while the rest are violated.

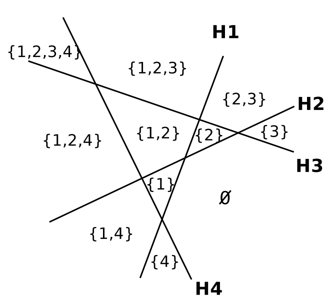

Figure 2: A collection of hyperplanes partitioning the plane into cells.

The mechanism design problem is in fact the problem of locating a non-empty cell that satisfies as many constraints as possible. This naturally leads to an algorithm that builds a list of cells and iterates over them to locate the cell assignment with the highest score.

Algorithm 2 [Geometric]:

1. Construct a list L of cells created by all hyperplanes $\vec{p}_{x,s}$.
2. Select an assignment $\sigma : X \times X \to L$.
3. Try to solve the linear problem that consists of constraints placing $v_{x,x'}$ in the cell $\sigma(x, x')$, and satisfying $\vec{v}_{x,x'} = -\vec{v}_{x',x} \quad ; \quad \vec{v}_{x,x''} = \vec{v}_{x,x'} + \vec{v}_{x',x''}$
4. If a solution is found, compute:
   (a) $W_\sigma \in 2^S$ the list of secrets that assignment σ of vectors $v_{x,x'}$ satisfies.
   (b) $\theta_\sigma = \sum_{s \in W_\sigma} p_s$.
5. Return the payment scheme found for $\arg\max_\sigma(\theta_\sigma)$.

In order to generate the list of cells L needed in the algorithm above, one can simply start from a list containing a single cell that contains the entire vector space and incrementally add hyperplanes. Each hyperplane that is added may partition a cell in the list into two cells, one on either side of the hyperplane, or may leave the cell intact. At every stage one only needs to iterate over the list of existing cells and check if they are split by the new hyperplane.

Complexity of the Algorithm In order to analyze the running time of Algorithm 2, we need to obtain a bound on the number of cells created by the hyperplanes defined by $\vec{p}_{x,s}$. Such a bound is given in (Edelsbrunner 1987). Given m hyperplanes in d-dimensional space, the number of cells is bounded by: $\Phi_d(m) = \sum_{i=0}^{d} \binom{m}{i} = O(m^d)$. The bound is obtained using the VC-Dimension of the concept class implied by cell partitioning and Sauer's lemma (Vidyasagar 1997).

This bound is especially interesting when d is small, since it implies that the number of cells is only polynomial in the number of hyperplanes m.

In our case, we have $|X||S|$ hyperplanes in an $|\Omega|$-dimensional space, which gives a bound of $|L| = O(|X||S|^{|\Omega|})$ cells. Generating the list of cells can be done in $O(|X||S|^{2|\Omega|} \cdot poly(|X||S||\Omega|))$ steps, and going over all possible assignments $\sigma : X \times X \to L$ is then accomplished in $O(|X||S|^{|X|^2|\Omega|} \cdot poly(|X||S||\Omega|))$ time steps.

This algorithm is therefore better in cases where $|S|$ is large, but $|\Omega|$ and $|X|$ are small.

Related Work

Proper Scoring Rules Scoring rules (Savage 1971) are used in order to assess and reward a prediction given in probabilistic form. A score is given to the predicting expert that depends on the probability distribution the expert specifies, and on the actual event that is ultimately observed. For a set Ω of possible events and $\mathcal{P}$, a class of probability measures over them, a scoring rule is then defined as a function of the form: $S : \mathcal{P} \times \Omega \to \mathbb{R}$.

A scoring rule is called *strictly proper* if the predictor maximizes its expected score by saying the true probability of the event, and receives a strictly lower score for any other prediction. That is: $E_{\omega\sim p}[S(p,\omega)] \geq E_{\omega\sim p}[S(q,\omega)]$ where equality is achieved iff $p = q$. (Gneiting & Raftery 2004; Hendrickson & Buehler 1971) show a necessary and sufficient condition for a scoring rule to be strictly proper which allows easy generation of various proper scoring rules by selecting a bounded convex function over $\mathcal{P}$.

An interesting use of scoring rules within the context of a multiagent reputation system was suggested by (Miller, Resnick, & Zeckhauser 2005), who use payments based on scoring rules to create the incentive for agents to honestly report about their experience with some service provider.

Mechanism Design The field of mechanism design deals with the creation of mechanisms that motivate individual agents to act in a way that will be beneficial to the whole of society or to the designer itself. (Maskin & Sjöström 2001) presents a good review of the field. A common approach to the mechanism design problem is to elicit the preferences of participants and then use them in order to decide on the outcome of the mechanism. The direct revelation principle states that whenever a mechanism with the proper incentives exists, there must be a proper mechanism in which the agents reveal everything. This is simply because one can construct the mechanism to receive the preferences from the participants and then act optimally on their behalf.

In settings where information is sold, it is unlikely that the seller would be willing to participate in direct revelation schemes. Since information is the primary commodity, revealing more of it to the mechanism is unwise, and the agent's beliefs about probabilities contain extra information. It remains unwise even if the mechanism is handled by a trusted third party, since revealing extra information would be reflected in payments made by the buyer.

(Conitzer & Sandholm 2002) proposed applying automated mechanism design to specific scenarios as a way of tailoring the mechanism to the exact problem at hand, and thereby developing superior mechanisms.

Other uses for information elicitation exist in multi-party computation (Smorodinsky & Tennenholtz 2005), where some function of the agents' secrets is computed, but agents may have reservations about revealing or computing their own secret. Another area in which information elicitation is implemented is polling. The information market (Bohm & Sonnegard 1999; Wolfers & Zitzewitz 2004) approach has been suggested as a way to get more reliable results than regular polls. There, agents buy and sell options that will pay them an amount that is dependent on the outcome of some event (like some specific candidate winning an election).

Conclusion and Future Research

We have examined information elicitation mechanisms where some relevant information that the seller possesses does not get revealed. We have seen that in these cases, the buyer must resort to mechanisms that work only with some probability. Finding such mechanisms with a high level of confidence is a computationally hard task, and we have demonstrated two approaches for doing so, each appropriate (from a complexity point of view) in different scenarios.

Many issues of information elicitation mechanisms still remain to be addressed. For example, the extension of such mechanisms to scenarios with multiple sellers and buyers would be interesting. In these settings, information held or sold by agents might be obtained through a third party. Issues of cross-validation of information from different sources, and possible collusion among agents that trade in information, could also be interesting. Further exploration can be made of scenarios where the various variables being sold are represented in compact Bayesian-networks. In these cases, the structure of the network and conditional independence between variables may affect mechanisms and prices for information exchange.

Acknowledgment

This work was partially supported by grant #039-7582 from the Israel Science Foundation.

References

Amaldi, E., and Kann, V. 1995. The complexity and approximability of finding maximum feasible subsystems of linear relations. *Theoretical Computer Science* 147(1–2):181–210.

Bertsimas, D., and Tsitsiklis, J. N. 1997. *Introduction to Linear Optimization*, volume 6 of *Optimization and Neural Computation*. Athena Scientific.

Bohm, P., and Sonnegard, J. 1999. Political stock markets and unreliable polls. *Scandinavian Journal of Economics* 101(2):205.

Conitzer, V., and Sandholm, T. 2002. Complexity of mechanism design. In *UAI-2002*, 103–110.

Edelsbrunner, H. 1987. *Algorithms in Combinatorial Geometry*, volume 10 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag.

Gneiting, T., and Raftery, A. E. 2004. Strictly proper scoring rules, prediction, and estimation. Technical Report 463, Department of Statistics, University of Washington.

Hendrickson, A. D., and Buehler, R. J. 1971. Proper scores for probability forecasters. *Annals of Mathematical Statistics* 42:1916–1921.

Maskin, E., and Sjöström, T. 2001. Implementation theory. Working paper, Harvard University and Penn State.

Miller, N.; Resnick, P.; and Zeckhauser, R. 2005. Eliciting honest feedback: The peer prediction method. *Management Science*. Forthcoming.

Savage, L. J. 1971. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* 66(336):783–801.

Smorodinsky, R., and Tennenholtz, M. 2005. Overcoming free-riding in multi-party computation: the anonymous case. *Games and Economic Behavior*. Forthcoming.

Vidyasagar, M. 1997. *A Theory of Learning and Generalization*. Spinger-Verlag.

Wolfers, J., and Zitzewitz, E. 2004. Prediction markets. Working Paper 10504, National Bureau of Econ. Research.