



Eliciting information for decision making

Caspar Oesterheld and Vincent Conitzer

Duke Department of Computer Science



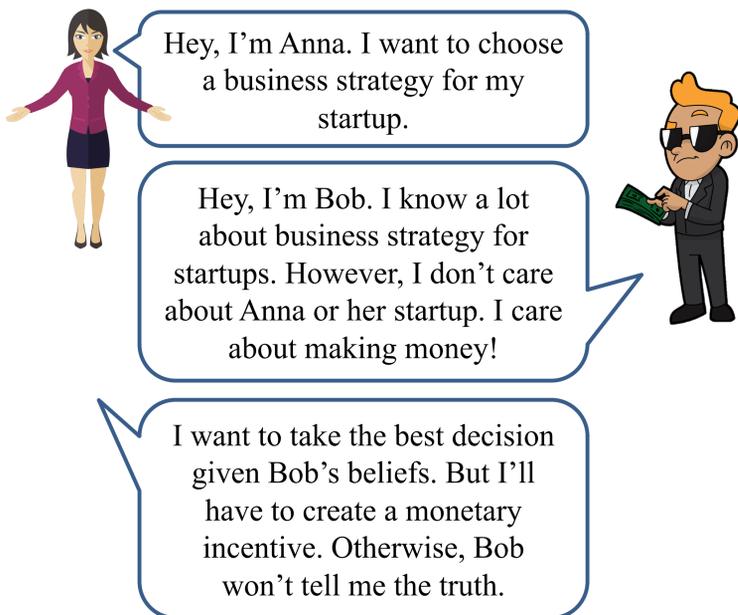
Abstract

We consider a setting in which a principal faces a decision and asks an external expert for a recommendation as well as a probabilistic prediction about what outcomes might occur if the recommendation were implemented. The principal then follows the recommendation and observes an outcome. Finally, the principal pays the expert based on the prediction and the outcome, according to some decision scoring rule. In this paper, we ask the question: What does the class of proper decision scoring rules look like, i.e., what scoring rules incentivize the expert to honestly reveal both the action he believes to be best for the principal and the prediction for that action?

We first show that in addition to an honest recommendation, proper scoring rules can only incentivize the expert to reveal the expected utility of taking the recommended action. The principal cannot strictly incentivize honest reports on other aspects of the conditional distribution over outcomes without setting poor incentives on the recommendation itself. We then characterize proper decision scoring rules as ones which give or sell the expert shares in the principal's project. Each share pays, e.g., \$1 per unit of utility obtained by the principal. Owning these shares makes the expert want to maximize the principal's utility by giving the best-possible recommendation. Furthermore, if shares are offered at a continuum of prices, this makes the expert reveal the value of a share and therefore the expected utility of the principal conditional on following the recommendation.

Keywords: algorithmic game theory, mechanism design, prediction markets, decision markets, principal-agent problems, principal-expert problems, proper scoring rules

The Premise



Hey, I'm Anna. I want to choose a business strategy for my startup.

Hey, I'm Bob. I know a lot about business strategy for startups. However, I don't care about Anna or her startup. I care about making money!

I want to take the best decision given Bob's beliefs. But I'll have to create a monetary incentive. Otherwise, Bob won't tell me the truth.

- A principal (Anna) selects from a set of actions A .
- Actions give rise to outcomes from some set Ω , to which the principal assigns utilities $u: \Omega \rightarrow \mathbb{R}$.
- The expert (Bob) has probabilistic beliefs $P(\cdot | \cdot)$, which specify for each action $a \in A$ and outcome $\omega \in \Omega$ the probability $P(\omega|a)$.
- The principal asks the expert to report P .
- The expert makes a potentially dishonest report \hat{P} .
- The principal takes the best action given \hat{P} .
- Some outcome ω is obtained.
- The principal rewards the expert according to $s(\omega, \hat{P})$.

An Example Mechanism



Hey Bob! I'll give you 3% of my startup.

To maximize my profit, I should give Anna the best possible recommendation.

Hey Anna! I believe business strategy D is best!

I have no reason to give Anna any other information.

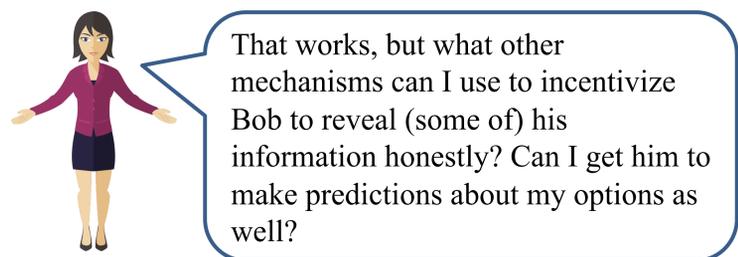
A simple example of a so-called *proper* scoring rule s is:

$$s(\omega, \hat{P}) = c_1 u(\omega) + c_2,$$

where $c_1, c_2 \in \mathbb{R}$ and $c_1 > 0$.

This strictly incentivizes the expert to make the best possible recommendation. But the expert has no reason to reveal any other information honestly.

The Question



That works, but what other mechanisms can I use to incentivize Bob to reveal (some of) his information honestly? Can I get him to make predictions about my options as well?

Impossibilities

- The principal cannot strictly incentivize the expert to reveal anything about non-recommended actions (other than that he believes them to be suboptimal).
 - More surprising: The principal cannot strictly incentivize honest reports on anything other than the recommended action's expected utility without setting poor incentives on the recommendation itself.
- Intuition: If Anna were to reward Bob for making accurate predictions, Bob might recommend an inferior action to make the outcome easier to predict.

Check out the full paper at:
<https://bit.ly/31gYgYt>



Another Example Mechanism



Hey Bob! I'll give you 1% of my startup for free.

Also, do you want to buy another 1% for \$1,000?

Hey Anna! Sure, I want that extra 1%! Do you take cash?

How about another 1% for \$2,000?

Yup!

And another 1% for \$3,000?

No, thanks!

But now that I know how my payoff will be determined, I'll tell you that plan D is best.

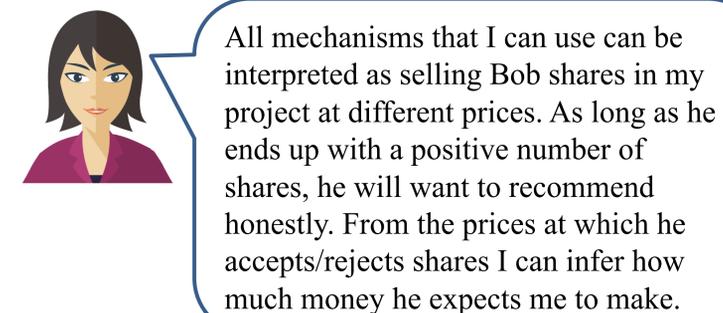
I now know what option Bob believes to be best. I also know that Bob believes the value of my startup (given plan D) to be between \$200k and \$300k.

The Characterization

Main theorem of the paper: A scoring rule s sets good incentives if and only if it can be written as

$$f(\hat{\mu})(y - \hat{\mu}) + \int_0^{\hat{\mu}} f(x) dx + C,$$

for some constant $C \in \mathbb{R}$ and non-decreasing $f: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$, where y is the utility obtained after following the expert's recommendation, $\hat{\mu}$ is the expert's prediction about the principal's expected utility given the recommendation.



All mechanisms that I can use can be interpreted as selling Bob shares in my project at different prices. As long as he ends up with a positive number of shares, he will want to recommend honestly. From the prices at which he accepts/rejects shares I can infer how much money he expects me to make.