

Eliciting beliefs by paying in chance

Alvaro Sandroni · Eran Shmaya

Received: 13 March 2013 / Accepted: 19 March 2013 / Published online: 1 April 2013
© SAET 2013

Abstract This paper describes a general principle that can be used to elicit honest opinions, even if risk attitudes are unknown.

Keywords Scoring rule · Eliciting beliefs

JEL Classification D83

1 Introduction

Several mechanisms, often called proper scoring rules, incentivize an expert to honestly report his subjective probabilities. However, these mechanisms often assume that the expert is risk neutral or that the expert's preferences over risk are known. Recently, [Karni \(2009\)](#) showed a novel mechanism that induces truth reporting even if the expert's preferences over risk are not known. [Lambert \(2011\)](#) characterizes the mechanisms that induce truthful-revelation with unknown preferences. In this paper, we describe a very simple principle that can also be used to dispose the assumption that the expert's preferences are known when beliefs are elicited.

There are two monetary rewards. The expert receives the greater reward with odds proportional to a proper scoring rule. Payments based on this principle delivers incentives for an expert to truthfully reveal what he knows, even if his attitudes toward risk are unknown.

We present two examples to illustrate the use of this principle. In the first example, we consider a traditional proper scoring rules (the Brier Score). In the second example,

A. Sandroni (✉) · E. Shmaya
Department of Managerial Economics and Decision Sciences, Kellogg School of Management,
Northwestern University, 2001 Sheridan Road, Evanston, IL 60208, USA
e-mail: sandroni@kellogg.northwestern.edu

we consider [Prelec \(2004\)](#) “Truth Serum” which elicits honest opinions in the absence of any data. In both cases, the original mechanism is based on known preferences. We show how to transform these mechanisms, with the principle mentioned above, so that the assumption of risk-neutrality (or known attitudes towards risk) are disposed. While we provide no general results, we hope that these two examples illustrate the basic principle clearly.

2 The principle of eliciting beliefs by paying in chance

An event E may or may not occur next period. An expert’s (Bob) subjective probability of E is of interest. The question is how to elicit the Bob’s subjective probability p of event E . Let’s say that the expert is offered the following monetary rewards

$$\begin{aligned} S(\hat{p}, 1) &= 2 - (1 - \hat{p})^2 && \text{if Bob announces probability } \hat{p} \text{ and } E \text{ occurs;} \\ S(\hat{p}, 0) &= 2 - (\hat{p})^2 && \text{if Bob announces } \hat{p} \text{ and } E \text{ does not occur.} \end{aligned}$$

This payment scheme S is a Brier Score, [Brier \(1950\)](#). The critical property of the Brier Score is that it is proper. That is, if the expert maximizes his expected payoff

$$\max_{\hat{p} \in [0,1]} pS(\hat{p}, 1) + (1 - p)S(\hat{p}, 0)$$

then he optimally chooses $\hat{p} = p$. [De Finetti \(1962\)](#) describes proper scoring rules as devices that “oblige each participant to express his true feelings, because any departure from his own personal probability results in a diminution of his own average score as he sees it”. That is, a proper score induces a risk-neutral expert to reveal his subjective beliefs. Several proper scoring rules exist [see [Bernardo \(1979\)](#); [Good \(1997\)](#); [Hendrickson and Buehler \(1971\)](#); [Kiefer \(2010\)](#); [Lindley \(1982\)](#); [Matheson and Winkler \(1976\)](#); [McCarthy \(1956\)](#); [Savage \(1971\)](#); [Schervish \(1989\)](#); [Van Lenthe \(1993\)](#); [Winkler \(1994\)](#)].

The main difficulty with these proper scoring rules is that they assume risk-neutrality or, alternatively, that the risk-attitudes of the expert are known [i.e., the utilities of the expert must be known if the rewards $S(\hat{p}, 1)$ and $S(\hat{p}, 0)$ are interpreted as utils]. If Bob is assumed to be risk-neutral when he is risk-averse, then the Brier Score may induce him to misrepresent his views [see [Winkler and Murphy \(1970\)](#); [Kadane and Winkler \(1988\)](#)]. Hence, it is significant to construct devices that oblige each participant to express his true beliefs, regardless of his attitudes toward risk.

One approach is to learn the experts’ preferences through experiments [see [Offerman et al \(2009\)](#); [Jaffray and Karni \(1999\)](#)]. Another approach was recently proposed by [Karni \(2009\)](#). He constructs a single-stage random scoring rule that induces Bob to honestly reveal his beliefs, regardless of his attitudes toward risk. In fact, [Karni \(2009\)](#) only assumes probabilistic sophistication and dominance [see [Machina and Schmeidler \(1995\)](#)]. So, let x and y be two monetary rewards with $x > y$. [Karni \(2009\)](#) assumes that Bob prefers a lottery that delivers x dollars with probability μ and y dollars with probability $(1 - \mu)$ to a lottery that delivers x dollars with probability μ' and y dollars with probability $(1 - \mu')$ if and only if μ is greater than μ' . That is, Bob prefers a higher chance of a greater monetary reward.

Now consider the following alternative random scoring rule. If Bob announces probability \hat{p} and E occurs then Bob receives the lottery

$$x \text{ with probability } P(\hat{p}, 1) \text{ and } y \text{ with probability } 1 - P(\hat{p}, 1);$$

If Bob announces probability \hat{p} and E does not occur then Bob receives the lottery

$$x \text{ with probability } P(\hat{p}, 0) \text{ and } y \text{ with probability } 1 - P(\hat{p}, 0);$$

where

$$P(\hat{p}, 1) = \frac{S(\hat{p}, 1)}{2} \text{ and } P(\hat{p}, 0) = \frac{S(\hat{p}, 0)}{2}.$$

That is, Bob receives x with odds proportional to his Brier Score (the division by 2 is merely to ensure that $P(\hat{p}, 1) \in [0, 1]$ and $P(\hat{p}, 0) \in [0, 1]$). Thus, assume Bob believes that E occurs with probability p and he announces \hat{p} . Then, he obtains x with probability

$$\frac{pS(\hat{p}, 1) + (1 - p)S(\hat{p}, 0)}{2}.$$

Thus, if S is a proper score (as in the case of the Brier Score) then it is optimal for Bob to truthfully reveals his subjective probability p , regardless of his attitudes toward risk. As in Karni's mechanism, probabilistic sophistication and dominance are the only underlying assumptions.

This example illustrates the principle of eliciting beliefs by paying in chance. The basic procedure is as follows: start with any proper score and normalize it (with some linear, strictly increasing transformation function) so that the score is always between 0 and 1. The higher monetary reward x is delivered with a probability given by the normalized proper score. So, Bob has a better chance of obtaining a greater reward if his score is higher. It then follows that it is optimal for Bob to truthfully reveals his opinions, regardless of his attitudes toward risk.

Lambert (2011) provides a characterization of scoring rules that induce truthful-revelation with unknown preferences. The random scoring rule described above is subsumed in his characterization [see also Allen (1987); Hossain and Okui (2010); Schlag and van der Weele (2010)].

It is easy to see that the principle of eliciting beliefs by paying in chance can be applied to any bounded, proper score. Like Karni (2009) mechanism, this principle can be extended to many events and many experts, by running separate mechanisms to different experts and events. In addition, it also extends in a straightforward way to the case of multiple data points.

Importantly, this principle can be used in multi-agents, game-theoretic settings. We illustrate this point with Prelec (2004) "Truth Serum". Consider several experts who face a subjective question such as "Is Picasso the best painter of the 20th century?" No objective data is observed on this matter, but each expert has an opinion (say yes or no)

on it. The problem is how to induce each expert to announce his honest opinion. Prelec assumes that the opinions of each expert are realizations of a random variable with a common distribution. He designs a Bayesian game in which the action space of expert i , A_i , is to say yes or no and a forecast of fraction of experts who answers yes. A pure strategy, for player i , maps his actual opinions (yes or no) into his action space A_i . Player i 's payoffs is given by a specific utility function

$$U_i : \prod_k A_k \longrightarrow \Re$$

which he constructs. The key point in his paper is that there exists a Nash equilibrium in which all experts announce their opinions honestly.

As in the case of the Brier Score, if the payoffs given U_i are provided in monetary terms then risk-neutrality is assumed. If the payoffs are assumed to be utils then it is assumed that the expert's attitudes toward risk are known. However, let τ be a linear, strictly increasing function such that, for every player i and for every action profile $a \in \prod_k A_k$, $\tau(U_i(a))$ is between 0 and 1.¹ Now consider the game in which given an action profile $a \in \prod_k A_k$, player i receives x with probability $\tau(U_i(a))$ and y with probability $1 - \tau(U_i(a))$, where $x > y$ are monetary rewards. The randomizations of nature determining between x and y is independent for each agent and each action profile. It is straightforward to show that the same truth-revealing Nash equilibrium in Prelec's game remains a Nash equilibrium in this modified game. This follows regardless of the expert's attitudes toward risks, provided that probabilistic sophistication and dominance are satisfied.

3 Conclusion

A simple way to elicit experts' opinions when experts' attitudes toward risk are unknown is to pay the expert in chance: he receives the greater reward with odds based on a proper method of eliciting beliefs that assume risk-neutrality.

Acknowledgments Financial support from the National Science Foundation is gratefully acknowledged.

References

- Allen, F.: Discovering personal probabilities when utility functions are unknown. *Manag. Sci.* **33**, 542–544 (1987)
- Bernardo, J.: Expected information as expected utility. *Ann. Stat.* **7**, 686–690 (1979)

¹ This linear function always exist if the payoffs are uniformly bounded. Let's say that $|U_i(a)| \leq M$ for some $M > 0$. The linear function could then be

$$\tau(z) = \frac{z + M}{2M}.$$

The payoffs in Prelec's paper are based on the Kullback-Leibler measure of relative entropy and, therefore are not bounded. However, similar results can be obtained if payoffs are based on alternative loss functions, such as quadratic loss, that are bounded.

- De Finetti, B.: Does it make sense to speak of good probability appraisals? In: Good, I.J. (ed.), *The scientist speculates: an anthology of partly-baked ideas*, pp. 356–364. Heineman, London (1962)
- Good, I.: *Rational decisions*. Springer, Berlin (1997)
- Hendrickson, A., Buehler, R.: Proper scores for probability forecasters. *Ann. Math. Stat.* **42**, 1916–1921 (1971)
- Hossain, T., Okui, R.: The binarized scoring rule of belief elicitation. Mimeo, UK (2010)
- Jaffray, J., Karni, E.: Elicitation of subjective probabilities when the initial endowment is unobservable. *J. Risk Uncertain.* **18**, 5–20 (1999)
- Kadane, J., Winkler, R.: Separating probability elicitation from utilities. *J. Am. Stat. Assoc.* **83**, 357–363 (1988)
- Karni, E.: A mechanism for eliciting probabilities. *Econometrica* **77**, 603–606 (2009)
- Kiefer, N.: Incentive-compatible elicitation of quantiles. Mimeo, UK (2010)
- Lambert, N.: Probability elicitation for agents with arbitrary risk preferences. Mimeo, UK (2011)
- Lindley, D.: Scoring rules and the inevitability of probability. *Int. Stat. Rev.* **50**, 1–11 (1982)
- Machina, M., Schmeidler, D.: Bayes without Bernoulli: simple conditions for probabilistic sophistication choice. *J. Econ. Theory* **49**, 479–486 (1995)
- Matheson, J., Winkler, R.: Scoring rules for continuous probability distributions. *Manag. Sci.* **22**, 1087–1096 (1976)
- McCarthy, J.: Measures of the value of information. In: *Proceedings of the National Academy of Sciences* **42**(9), pp. 654–655 (1956)
- Offerman, T., Sonnemans, J., van de Kuilen, G., Wakker, P.: A truth serum for non-bayesians: correcting proper scoring rules for risk attitudes. *Rev. Econ. Stud.* **76**, 1461–1489 (2009)
- Prelec, D.: A bayesian truth serum for subjective data. *Science* **306**(5695), 462–466 (2004)
- Savage, L.: Elicitation of personal probabilities and expectations. *J. Am. Stat. Assoc.* **66**, 783–801 (1971)
- Schervish, M.: A general method for comparing probability assessors. *Ann. Stat.* **17**, 1856–1879 (1989)
- Schlag, K., van der Weele, J.: Eliciting probabilities, means, medians, variances and covariances without assuming risk neutrality. Mimeo, UK (2010)
- Van Lenthé, J.: A blueprint of ELI, a new method for eliciting subjective probability distributions. *Behav. Res. Methods Instrum. Comput.* **25**, 425–433 (1993)
- Winkler, R.: Evaluating probabilities: asymmetric scoring rules. *Manag. Sci.* **40**, 1395–1405 (1994)
- Winkler, R., Murphy, A.: Nonlinear utility and the probability score. *J. Appl. Meteorol.* **9**(1), 143–148 (1970)