Econometrica, Vol. 77, No. 2 (March, 2009), 603–606

A MECHANISM FOR ELICITING PROBABILITIES

BY EDI KARNI[1]

This paper describes a direct revelation mechanism for eliciting agents' subjective probabilities. The game induced by the mechanism has a dominant strategy equilibrium in which the players reveal their subjective probabilities.

KEYWORDS: Probability elicitation, direct revelation mechanism.

1. INTRODUCTION

AN INDIVIDUAL'S ASSESSMENT of the likelihoods of events in which he has no stake may be of interest to others. This is the case when the person whose assessment is sought is an expert and the others are decision makers facing a choice between alternative courses of action whose consequences depend on which of the events obtains. Investors, for instance, may be interested in the probability a geologist assigns to discovering mineral deposits beneath a particular tract of land; a patient might want to seek a second opinion on the likelihood of success of a treatment recommended by his physician.

Procedures for eliciting the subjective probabilities of agents whose preferences are represented by subjective expected utility functionals include the proper scoring rule method (Savage (1971)), the promissory notes method (de Finetti (1974)), and the lotteries method (Kadane and Winkler (1988)). The first two procedures entail trade-offs between the incentives and the accuracy of the probability estimate. The third procedure is not incentive compatible.[2]

This paper introduces a new elicitation mechanism that yields accurate elicitation while allowing the incentives to be set at any desirable level.

2. THE ELICITATION MECHANISM

Let S be a set of *states*, one of which is the true state. Subsets of S are *events*. An event is said to obtain if the true state belongs to it. Simple acts are mapping from S to the real numbers, representing monetary payoffs, with finite images. A bet on an event E is a simple act that pays x dollars if E obtains and y dollars otherwise, $x > y$, and is denoted by $x_E y$.

A simple lottery is a finite list of monetary prizes (that is, $(x_1, \dots, x_m) \in \mathbb{R}^m$, $m < \infty$) and a corresponding probability vector $(p_1, \dots, p_m)$, where, for each i, $p_i \geq 0$ is the probability of winning the prize x_i and $\sum_{i=1}^m p_i = 1$.

[1]I am grateful to John Hey for stimulating conversations and to LUISS University for its hospitality. I also benefited greatly from the comments and suggestions of the editor and three anonymous referees.

[2]For a more detailed discussion, see the concluding section.

 DOI: 10.3982/ECTA7833

Denote by D the union of the sets of simple acts and lotteries, let $\succcurlyeq$ be a preference relation on D, and denote by $\succ$ and $\sim$ the asymmetric and symmetric parts of $\succcurlyeq$, respectively. A preference relation $\succcurlyeq$ on D that restricted to the set of finite acts is said to exhibit probabilistic sophistication if it ranks acts or lotteries solely on the basis of their implied probability distributions over outcomes (see Machina and Schmeidler (1995)). In particular, if π is the probability measure implicit in $\succcurlyeq$, then probabilistic sophistication implies that, for all acts f and lotteries $\ell(p, x, y) = [x, p; y, (1-p)]$, $p \in [0, 1]$, $\pi(f^{-1}(x)) = p$ implies $x_{f^{-1}(x)}y \sim \ell(p, x, y)$.

Consider an agent whose assessment of the probability of the event E is of interest. Suppose that the agent's preference relation $\succcurlyeq$ on D displays probabilistic sophistication and dominance in the sense that $\ell(p, x, y) \succcurlyeq \ell(p', x, y)$ for all $x > y$ if and only if $p \geq p'$. Denote by $\pi(E)$ the probability the agent assigns to the event E.

The elicitation mechanism selects a random number r from a uniform distribution on $[0, 1]$ and requires the agent to submit a report, $\mu \in [0, 1]$, of his subjective probability assessment of the event E. The mechanism awards the agent the payoff $\beta := x_E y$ if $\mu \geq r$ and the lottery $\ell(r, x, y)$ if $\mu < r$.

To see that truthful reporting is the agent's unique dominant strategy, suppose that the agent reports $\mu > \pi(E)$. If $r \leq \pi(E)$ or $r \geq \mu$, the agent's payoff is the same regardless of whether he reports μ or $\pi(E)$.[3] If $r \in (\pi(E), \mu)$, the agent's payoff is β; had he reported $\pi(E)$ instead of μ, his payoff would have been $\ell(r, x, y)$. But $r > \pi(E)$, which, by probabilistic sophistication and dominance, implies $\ell(r, x, y) \succ \beta$. Thus the agent is worse off reporting μ instead of $\pi(E)$. A similar argument applies when $\mu < \pi(E)$.[4]

The elicitation method described in the preceding section is quite general. In particular, it may be extended to the case of many agents by running the mechanism separately for each agent.[5] It may also be extended to finitely many events by running the mechanism separately for each of the events. Moreover, if the payoff difference is sufficiently large, the agent is induced to exert the effort necessary to arrive at an accurate assessment of his subjective probability.

An equivalent probability-elicitation auction mechanism is as follows: The mechanism selects r as before and runs a continuous increasing bid auction between the agent and a dummy bidder. The dummy bidder stays in the auction as long as the bid is smaller than r and drops out when the bid equals r. Starting at 0, the bid increases continuously as long as the agent and the dummy bidder are both "in the auction" and stops when one of them drops out or the bid

[3]Regardless of whether the agent reported μ or $\pi(E)$, his payoff is β if $r \leq \pi(E)$ and $\ell(r, x, y)$ if $r \geq \mu$.

[4]In this case, if $r \in (\mu, \pi(E))$, the agent wins $\ell(r, x, y)$; had he reported $\pi(E)$ instead, he would have won β. But $\beta \succ \ell(r, x, y)$.

[5]In the case of many agents, the mechanism may be redesigned so that the random number, r, is generated endogenously. Specifically, the highest reported value is substituted for r and all other reports are treated as μ. For more details see Karni (2008).

reaches 1, whichever is first.[6] The agent is awarded $\ell(r, x, y)$ if he is the first to quit and β otherwise. Clearly, the agent's dominant strategy is to stay in the auction as long as the bid is smaller than $\pi(E)$ and to quit when it is equal to $\pi(E)$.

3. CONCLUDING REMARKS

In situations in which the agent must be induced to take costly measures (e.g., time and effort) in order to arrive at a reliable probability estimate, the mechanisms introduced in this paper perform better than the elicitation procedures discussed in the literature. Consider, for example, the proper scoring rule method. Let E denote the event of interest and denote by E^c its complement. Then, according to this method, the agent's payoff equals the score $-r(\delta_E - \mu)^2$, where r is a positive constant, μ is the agent's reported probability of E, and δ_E is the indicator function of the event E. Let $\tilde{w}$ be the agent's random wealth and denote by F its cumulative distribution function. Consistent with the no-stake requirement, suppose that $\tilde{w}$ is distributed independently of E. If the agent's subjective assessment of the probability of E is $\pi(E)$, then his problem is

$$\text{(1)} \qquad \operatorname*{Max}_{\mu}\Big[\pi(E)\int u(w - r(1-\mu)^2)\,dF(w \mid E) + (1-\pi(E))\int u(w - r\mu^2)\,dF(w \mid E^c)\Big].$$

The necessary condition is

$$\text{(2)} \qquad \frac{\mu^*(r)}{(1-\mu^*(r))} = \frac{K(r)\pi(E)}{1-\pi(E)},$$

where $\mu^*(r)$ denotes the optimal solution and

$$K(r) = \int u'\big(w - r(1-\mu^*(r))^2\big)\,dF(w \mid E) \Big/ \int u'(w - r\mu^*(r)^2)\,dF(w \mid E^c).$$

Thus $\mu^*(r) = \pi(E)$ if and only if $K(r) = 1$. Unless the agent is risk neutral, the elicitation of probabilities by the scoring rule method confounds subjective probabilities and marginal utilities. If, to motivate the agent to assess the probability of the event of interest accurately, it is necessary to expose him to

[6]The agent is "in the auction" until he signals that he quits.

risk by setting a large value of r, then, in general, $\mu^*(r)$ is a biased estimate of $\pi(E)$.[7] To obtain an unbiased assessment of $\pi(E)$, it is necessary to let r tend to zero, but then the agent has no incentive to assess the probability of the event E accurately. The promissory notes method of de Finetti (1974) suffers from the same problem.[8]

The accuracy of the elicitation procedures described here depends critically on the agent having no stake in the event of interest.[9] If he does have a stake in the event, the evaluations of the payoffs of the bet and the lotteries that figure in the mechanism are event dependent, and the preference relation does not exhibit probabilistic sophistication.

REFERENCES

DE FINETTI, B. (1974): *Theory of Probability*, Vol. 1. New York: Wiley. [603,606]

JAFFRAY, J.-Y., AND E. KARNI (1999): "Elicitation of Subjective Probabilities When the Initial Endowment Is Unobserved," *Journal of Risk and Uncertainty*, 18, 5–20. [606]

KADANE, B. J., AND R. L. WINKLER (1988): "Separating Probability Elicitation From Utility," *Journal of the American Statistical Association*, 83, 357–363. [603,606]

KARNI, E. (1999): "Elicitation of Subjective Probabilities When Preferences Are State-Dependent," *International Economic Review*, 40, 479–486. [606]

——— (2008): "A Mechanism Design for Probability Elicitation," Unpublished Manuscript, Johns Hopkins University. [604]

MACHINA, M. J., AND D. SCHMEIDLER (1995): "Bayes Without Bernoulli: Simple Conditions for Probabilistically Sophisticated Choice," *Journal of Economic Theory*, 67, 106–128. [604]

SAVAGE, L. J. (1971): "Elicitation of Personal Probabilities and Expectations," *Journal of the American Statistical Association*, 66, 783–801. [603]

Dept. of Economics, Johns Hopkins University, Baltimore, MD 21218, U.S.A.; karni@jhu.edu.

Manuscript received April, 2008; final revision received November, 2008.

[7]If the agent is risk averse and $\pi(E) \neq 1/2$, then $\mu^*(r)$ is biased toward $1/2$ and the biased increases with r. Similarly, if the agent is risk inclined, then $\mu^*(r)$ is bias toward either 0 (if $\pi(E) < 1/2$) or toward 1 (if $\pi(E) > 1/2$) and the bias increases with r.

[8]The lottery method requires the agent to indicate the probability p that would make him indifferent between $\ell(p, x, y)$ and β. However, because it does not specify the payoff to the agent, it is not incentive compatible.

[9]When the agent has a stake in the events of interest, the other methods also fail (Kadane and Winkler (1988)). Jaffray and Karni (1999) and Karni (1999) developed elicitation methods designed to overcome this difficulty.