Eliciting Properties of Probability Distributions: the Highlights

NICOLAS LAMBERT
Stanford University
and
DAVID M. PENNOCK
Yahoo! Research
and
YOAV SHOHAM
Stanford University

We investigate the problem of incentivizing an expert to truthfully reveal probabilistic information about a random event. Probabilistic information consists of one or more properties, which are any real-valued functions of the distribution, such as the mean and variance. Not all properties can be elicited truthfully. We provide a simple characterization of elicitable properties, and describe the general form of the associated payment functions that induce truthful revelation. We then consider sets of properties, and observe that all properties can be inferred from sets of elicitable properties. This suggests the concept of elicitation complexity for a property, the size of the smallest set implying the property.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics

General Terms: Economics, Theory

Additional Key Words and Phrases: Scoring rule, prediction market, elicitation, forecasting

1. INTRODUCTION

Assume that, given an upcoming uncertain event, we ask an expert about some features of its probability distribution. To incentivize the expert to provide true information, we must reward the expert accordingly. Designing such rewards can be a difficult task. Indeed, to enforce honest reports, the expert's payment should depend, in a nontrivial way, on both the expert's response and the outcome of the event.

In the special case where we are interested in obtaining the full probability distribution, we can use probability scoring rules [Good 1997; Brier 1950; Winkler 1996]. A scoring rule represents a reward function, it is said to be *proper* when the expert gets the maximum expected score by reporting the true distribution of the

Authors' addresses: nlambert@stanford.edu, pennockd@yahoo-inc.com, shoham@stanford.edu

event. Proper scoring rules can be easily described in variety of ways [Savage 1971; Hendrickson and Buehler 1971; Schervish 1989], they have been validated experimentally [O'Carroll 1977] and have appeared in a number of domains [Spiegelhalter 1986; Nelson and Bessler 1989].

However, probability scoring rules don't allow to retrieve partial information on the distribution. For example, we may be interested in a specific parameter, such as the mean, instead of the distribution as a whole. Naturally, such information can always be derived from the full distribution, however this poses practical difficulties. The expert may not be able to give a full distribution, or may wish to keep certain aspects confidential. When the distribution is large or complex, its estimation may be infeasible, and the communication requirements may be high.

To our knowledge, probability scoring rules have been adapted only to two special cases of partial information: eliciting the mean and quantiles of random variables [Savage 1971; Cervera and Munoz 1996; Gneiting and Raftery 2007]. In this paper we determine what partial information can be obtained truthfully, and how to design the corresponding appropriate payment schemes.

2. MODEL

A principal wishes to learn information about a future random event, whose set of outcomes Ω is finite. For simplicity we take a special case of [Lambert et al. 2008] and assume that the probability distribution of the outcomes belong to $\Delta^*(\Omega)$, the set of probability distributions over Ω that assign a positive probability to all outcomes. The information of interest takes the form of parameters of the distribution of outcomes, called *distribution properties*. Formally, they are functions that associate any possible probability distribution with a real value. Common distribution properties include the probability of an event E, $P \mapsto P(E)$, the mean of a random variable X, $P \mapsto \sum_x xP(X = x)$, the variance, moments and centered moments, skewness, kurtosis, or the entropy. We restrict our attention to *nice* properties that are continuous and not locally constant. All the examples above are nice properties, and in the sequel *property* will always refer to *nice property*.

With these conditions, the range of values taken by a property forms an interval, and the expert is allowed to report any value of the interval except its boundary points. Such possible reports are called *admissible*, and similarly a probability distribution is admissible when its corresponding property value is admissible.

To incentivize the expert to provide true information, the principal offers a reward, defined by *score functions* $s(\gamma_1, \ldots, \gamma_k, \omega)$ that may depend on the expert's estimates $\gamma_1, \ldots, \gamma_k$ of the properties of interest and on the outcome of the event ω. For simplicity, we consider a special case of [Lambert et al. 2008] and restrict ourselves to continuously differentiable score functions. In the sequel, score function will always refer to continuously differentiable score function.

A risk-neutral expert reports values $\gamma_1, \ldots, \gamma_k$ that maximize the expected score

$$\mathop{\mathbb{E}}_{\omega \sim P}[s(\gamma_1, \ldots, \gamma_k, \omega)] \ .$$

Using the terminology of the forecasting literature, scores that incentivize the expert to make true predictions are said to be *strictly proper*: with strictly proper scores, the maximum expected score can only be obtained by reporting the true values of the properties. When there exist strictly proper scores for a property or set of properties, it is possible to ask an expert for the true value of the property(ies) without the need to ask for any extra information, and the property or set of properties is *(directly) elicitable.*

If strictly proper scores enforce truthful reports, they don't necessarily reward the accuracy of the prediction: an agent far away from the truth may receive higher expected rewards than an agent whose report is closer, but not equal, to the true values of the properties. To prevent these issues, one can use *accuracy-rewarding* scores, for which the expected score increases with the accuracy of the prediction. More formally, considering properties $\Gamma_1, \ldots, \Gamma_k$, a score function s is accuracy-rewarding if, when an expert A reports $\gamma_1, \ldots, \gamma_k$ while an expert B reports $\gamma'_1, \ldots, \gamma'_k$, given that both experts do not report the same set of values and that, for all i, either $\gamma_i \leq \gamma'_i \leq \Gamma_i(P)$ or $\Gamma_i(P) \leq \gamma'_i \leq \gamma_i$, then B's expected score is strictly higher than A's. Note that accuracy-rewarding scores are always strictly proper, but the opposite is not always true.

3. CHARACTERIZATION OF ELICITABLE PROPERTIES

3.1 Single properties

We start with the problem of eliciting a single property. Even in this simple case, not all properties can be elicited directly, but elicitable properties are simply described: a property is elicitable when its level sets are convex.

THEOREM 3.1. *A distribution property is elicitable if and only if $\Gamma^{-1}(\gamma)$ is convex for all admissible property values γ.*

The theorem can easily be applied to the examples of properties above, to find that an event's probability and the mean are elicitable, while the centered moments, skewness, kurtosis and entropy are not.

Naturally, for properties that are elicitable, it is interesting to know the corresponding score functions that should be used. It can be shown that, for each property, the strictly proper scores form a cone of the space of score functions, and may be simply characterized as weighted integrals of the *signature* of the property.

THEOREM 3.2. *Let Γ be an elicitable property, with (a, b) the range of admissible property values. There exists a unique normalized function $v : (a, b) \times \Omega \mapsto \Re$ such that any score function s is strictly proper for Γ if and only if*

$$s(\gamma, \omega) = s_0(\omega) + \int_{\gamma_0}^{\gamma} \lambda(t) v(t, \omega) dt \ ,$$

for $\gamma_0 \in (a, b)$, a function $s_0 : \Omega \mapsto \Re$ and a continuous weight function $\lambda : (a, b) \mapsto [0, +\infty)$ that is not locally null.

The function v depends only on the property of interest, and so is called the *signature* of the property. Let's consider, for example, the property "probability of an event E". In this particular case, score functions are called probability scoring rules. The signature is given by the function

$$v(p, \omega) = \frac{1}{\sqrt{p^2 + (1-p)^2}} \begin{cases} p & \text{if } \omega \in E \ , \\ 1 - p & \text{if } \omega \notin E \ . \end{cases}$$

We can easily obtain the full set of strictly proper scoring rules by setting arbitrary nonnegative weights that are not locally null. For example, the popular logarithmic score has a weight $\lambda(p) = b\sqrt{1/p^2 + 1/(1-p)^2}$, while that of the quadratic score is $\lambda(p) = b\sqrt{p^2 + (1-p)^2}$. It is interesting to note that, in the case of probability scoring rules, our description of strictly proper scores reduces to the Schervish-Choquet representation [Schervish 1989].

We may also consider accuracy-rewarding scores. It can be shown that, for single properties, strictly proper is equivalent to accuracy-rewarding, and so the previous characterizations remain valid.

3.2 Sets of properties

We are now interested in eliciting multiple properties simultaneously. The process of eliciting sets of properties is different from that of eliciting each property independently. In particular, the individual elements of an elicitable set of properties need not be elicitable, so that eliciting sets of properties may be used to obtain truthful estimates of properties that, alone, cannot be elicited. For example, the variance alone is not elicitable, but the variance and mean together are elicitable. Convexity of the level sets remains a necessary condition for a set of properties to be elicitable, but is no longer sufficient. However, we can characterize the sets of properties that are elicitable with accuracy-rewarding scores.

THEOREM 3.3. *A set of properties $\{\Gamma_1, \ldots, \Gamma_k\}$ is elicitable with an accuracy-rewarding score if and only if each individual property Γ_i is elicitable.*

Therefore, while we can elicit more information using sets of properties with strictly proper scores, we cannot do so with accuracy-rewarding scores. Furthermore, eliciting sets of properties with accuracy-rewarding scores is strictly equivalent to eliciting each property independently, as those scores for the entire set are simply expressed as sums of strictly proper scores for each individual property.

THEOREM 3.4. *A score function s for the properties $\Gamma_1, \ldots, \Gamma_k$ is accuracy-rewarding if and only if there exist strictly proper score functions s_i for Γ_i such that*

$$s(\gamma_1, \ldots, \gamma_k, \omega) = \sum_i s_i(\gamma_i, \omega) \ .$$

3.3 Elicitation complexity

Properties that cannot be elicited directly can always be inferred from some elicitable properties: by asking the probability of each outcome, we can derive the value of any property. Often, we do not need to elicit the full distribution. For example, the variance of a random variable can be inferred from the first and second moment, both elicitable. The size of the smallest set of elicitable properties needed to infer the value of a given property is a measure of the difficulty of obtaining a truthful estimate of that property and is called the *elicitation complexity*: a property Γ has complexity k if it can be inferred from some k elicitable properties, but cannot be inferred from any $k-1$ properties.[1]

[1]A property Γ can be inferred from properties $\Gamma_1, \ldots, \Gamma_k$ when there exists a function f with $\Gamma = f(\Gamma_1, \ldots, \Gamma_k)$.

With n outcomes, all properties have complexity at most $n-1$ since they all can be inferred from the full distribution. Further, it can be shown that for all $1 \leq k \leq n-1$, there exists a property of complexity k. In particular, there exist properties that are as difficult to elicit as the full distribution. For example, the property "maximum likelihood" $P \mapsto \max_{\omega \in \Omega} P(\{\omega\})$ has the maximum complexity $n-1$.

4. CONCLUSION AND FUTURE WORK

This letter deals with the problem of eliciting partial information about the probability distribution of a random event. Given the space constraints, we restricted our presentation to the highlights of the original paper [Lambert et al. 2008], where the reader will find a more complete discussion.

We believe our initial investigation may be further extended in various directions. First, we assume variables take values in a finite set. In many practical situations, the variables are continuous. Is it possible to adapt the characterizations to general probability spaces? Second, our results concerning sets of properties assume accuracy-rewarding scores, but characterizing elicitable properties and valid payments for the case of strictly proper scores remains an open problem. Third, beyond the mean and variance, little is known about the complexity of specific distribution properties. Finally, and most importantly, the results presented here focus on theoretical aspects of elicitation, and still remain to be validated experimentally.

REFERENCES

Brier, G. 1950. Verification of Forecasts Expressed In Terms of Probability. *Monthly Weather Review 78,* 1, 1–3.

Cervera, J. and Munoz, J. 1996. Proper scoring rules for fractiles. *Bayesian Statistics 5*, 513–519.

Gneiting, T. and Raftery, A. 2007. Strictly proper scoring rules, prediction and estimation. *Journal of the American Statistical Association 102,* 477, 359–378.

Good, I. 1997. *Rational Decisions.* Springer-Verlag.

Hendrickson, A. and Buehler, R. 1971. Proper Scores for Probability Forecasters. *The Annals of Mathematical Statistics 42,* 6, 1916–1921.

Lambert, N., Pennock, N., and Shoham, Y. 2008. Eliciting Properties of Probability Distributions. In *Proceedings of the ninth ACM conference on electronic commerce*. 129–138.

Nelson, R. and Bessler, D. 1989. Subjective Probabilities and Scoring Rules: Experimental Evidence. *American Journal of Agricultural Economics 71,* 2, 363–369.

O'Carroll, F. 1977. Subjective Probabilities and Short-Term Economic Forecasts: An Empirical Investigation. *Applied Statistics 26,* 3, 269–278.

Savage, L. 1971. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association 66,* 336, 783–801.

Schervish, M. 1989. A General Method for Comparing Probability Assessors. *The Annals of Statistics 17,* 4, 1856–1879.

Spiegelhalter, D. 1986. Probabilistic prediction in patient management and clinical trials. *Stat Med 5,* 5, 421–33.

Winkler, R. 1996. Scoring rules and the evaluation of probabilities. *TEST 5,* 1, 1–60.