

Introduction to Dynamic Belief Elicitation

CHRISTOPHER P. CHAMBERS

University of California San Diego

and

NICOLAS S. LAMBERT

Stanford University

We present a protocol for eliciting dynamic beliefs from forecasters. At time $t = 0$, forecasters hold beliefs about a random variable of interest that will realize publicly at time $t = 1$. Between $t = 0$ and $t = 1$, forecasters observe private information that impacts their beliefs. We design a class of protocols that, at the outset, elicit forecasters' beliefs about the random variable and elicit beliefs about any private information they expect to receive over time, and then elicit the private information that forecasters receive as they receive it. We show that any alternative elicitation mechanism can be approximated by protocols in our class. The information elicited can be used to solve optimally an arbitrary dynamic decision problem.

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

General Terms: Economics, Theory, Measurement, Design

Additional Key Words and Phrases: Scoring rules, Elicitability, Forecasting, Decision Theory

In this letter we introduce the idea and purpose of dynamic belief elicitation as explored in our working paper [Chambers and Lambert 2014]. In many economic situations, decision makers have less payoff-relevant information than do some other agents. To deal with these situations, the decision maker may delegate the decision to the better informed agents. Alternatively, the decision maker may extract beliefs from the agents and then use these beliefs to decide on the best course of action. Our paper is about the latter case; specifically, it addresses the problem of eliciting beliefs from the agents.

To illustrate, consider a simple stylized situation. The decision maker is an investor who contemplates buying shares in a risky venture. The venture, if successful, yields a payoff $B = 1000$ to the investor. If it fails, it pays nothing. The outcome of the venture will be known at time $t = 1$. The investor must decide at $t = 0$. If she buys the shares, she incurs a cost $C = 450$.

This simple setup is easy to analyze. If the investor is risk neutral without discounting, she is better off buying the shares when the expected benefit exceeds the cost, or $p \times B > C$, where p is the probability that the project is successful. Note that to make the optimal decision, she only needs to assess the probability p . If she is unable to form such assessments, she can elicit probability beliefs from better informed forecasters using *scoring rules* and related methods.

In our context, a scoring rule is a function $s(\hat{p}, x) \in \mathbf{R}$ that takes as input a probability forecast \hat{p} and a realization $x = 1$ if success and $x = 0$ if failure, and assigns a score to the forecaster. It is typically assumed that forecasters want

Authors' addresses: cpchambers@ucsd.edu, nlambert@stanford.edu.

to maximize their average score, for example when the scoring rule is used as a direct incentive scheme with scores mapping to payments or lottery tickets, or when it is used as an indirect incentive scheme that compares average scores of different forecasters over time. (Strictly) Proper scoring rules ensure that it is a strict best response for the forecaster to provide his true belief, meaning that $E[s(p, X)] > E[s(\hat{p}, X)]$ whenever $p = Pr(X = 1)$ and $\hat{p} \neq p$, where X is the random variable associated with the project outcome. For example, the Brier score $s(\hat{p}, x) = 1 - (\hat{p} - x)^2$ is proper [Brier 1950]. More generally, scoring rules elicit information about the uncertainty of random variables. Such information can be used to solve optimally a *static* decision problem, in which the decision maker must choose among a set of alternatives at one particular instant.

However, in many situations of practical interest, decisions can be made at different times. These decision problems are *dynamic*. In such situations, assessing the uncertainty of payoff-relevant random variables at every decision stage is not enough. To decide optimally, the decision maker must know *how and when uncertainty unravels over time*. To illustrate this point, let us revisit our investor story.

The investor can now decide to delay the decision to invest to time $t = \frac{1}{2}$. However, if she invests at that time, she loses the offer made to her at $t = 0$: her cost of participation increases from 450 to 600. By adding the option to invest at a later time, the optimal time-0 decision—invest or delay—depends not only on her current belief about chances of success, but also on *how she expects her belief to change over time*. Consider one instance where the investor believes at $t = 0$ that the project has 50% chances of success, and does not expect to receive any further information by $t = \frac{1}{2}$. With such beliefs, the investor is best off investing at $t = 0$, as she expects to make $1000 \times 0.5 - 450 = 50$ by investing right away (better than delaying the decision to invest or not investing at all). In another instance, suppose that the investor continues to believe in 50% chances of success at $t = 0$, but now expects to conduct further investigation by $t = \frac{1}{2}$. The outcome of the investigation is a signal at $t = \frac{1}{2}$, taking values good news or bad news, each occurring with equal probability. If good news was to arrive, the investor expects to revise her assessment to 80% chances of success, while in case of bad news, she would update her assessment to 20%. With such beliefs, if she delays decision and invest at $t = \frac{1}{2}$ but only upon observing good news, she makes on average $0.5 \times (1000 \times 0.8 - 600) = 100$, which is better than investing right away. So what does this simple analysis tells us? When dealing with dynamic decision problems, beliefs about future beliefs (or equivalently beliefs about future information) matter. This property is common to most dynamic decision problems—unless decisions are reversible, beliefs over future beliefs impact optimal decisions.

If the investor is not well informed, she will want to extract beliefs from well informed forecasters. Scoring rules and related methods can help her extract static beliefs, i.e., beliefs about the chances of success, at both decision stages. However they fail to extract a valuable piece of information: the forecaster's belief at $t = 0$ that concerns his future belief at $t = \frac{1}{2}$. Without such information, the investor cannot make decisions optimally.

Our paper extends the theory of scoring rules to the time dimension, allowing to

extract such information. To fix ideas, consider a forecaster who is about to form a belief at time $t = \frac{1}{2}$ about success of the project ($X = 1$). Denote by P this belief. Note that P is a random variable, because the forecaster generally does not know at the outset what his belief will be later on. At $t = 0$, the forecaster forms a prior about his future belief, captured by a cumulative distribution F about P . We want to design a scoring function $s(\widehat{F}, \widehat{p}, x)$ taking as input a report of belief over beliefs, \widehat{F} , made at time 0, a report of the posterior belief \widehat{p} at time $\frac{1}{2}$, and a realization x , so that to maximize his expected score, the forecaster must provide his true prior belief over beliefs, F , at time 0, and his true posterior belief over X at time $\frac{1}{2}$. Our main results show that such scoring functions exist and show how to construct them; in many cases an analytic expression can be found. Here for example we can use the following scoring function:

$$s(\widehat{F}, \widehat{p}, x) = \int_{\widehat{p}}^1 \left(1 + \int_0^\alpha \widehat{F} \right) (\alpha - x) d\alpha - \frac{1}{2} \int_0^1 \left(\int_0^\alpha \widehat{F} \right)^2 d\alpha.$$

Knowing F and the realization of P turns out to be exactly what is needed to solve optimally any two-period decision problem where the only payoff-relevant variable is X . Solving a three-period decision problems requires beliefs over beliefs over the final posterior belief of X , and so forth.

Our framework includes this instance but is significantly more general. The forecaster may receive high-dimensional private information at multiple times (even continuously or at random times) before the realization of the public random variable of interest, which itself can be high-dimensional. We design a family of protocols that elicit the entire belief structure of the forecaster at the outset, and the updated structure whenever the forecaster observes new private information.

A common approach in the mechanism design and scoring rules literatures relies on convex analysis, making use of the fact that payoffs or scores are the subdifferentials of convex value functions [McCarthy 1956; Savage 1971; Rochet 1985; Gneiting and Raftery 2007]. In our framework, the high dimensionality of the object being communicated and the presence of private information makes this approach less practical. Instead, we apply an idea that originates from Allais and relies on the use of randomization devices [Allais 1953]. Allais applied his technique to elicit preferences of an individual over pairs of choices in the context of revealed preferences. In the context of probability elicitation, similar ideas have been used by [Becker et al. 1964] and [Matheson and Winkler 1976].

Each protocol in our setup is described by a pair (\mathcal{D}, μ) , where \mathcal{D} is a collection of “simple” dynamic decision problems and μ is a probability measure over \mathcal{D} . Every such protocol works as follows: at the outset, the elicitor selects a decision problem $d^* \in \mathcal{D}$ at random according to μ . The forecaster knows \mathcal{D} and μ , but does not know d^* , which the elicitor keeps secret. At time 0, the forecaster is asked to announce his belief structure, and then asked to send updates at subsequent times whenever he receives new information. The elicitor solves the decision problem d^* optimally according to the information communicated by the forecaster over time. Upon realization of the public random variable, the elicitor collects the overall payoff from d^* and transfers it to the forecaster (such a payoff need not be monetary, it can be interpreted as a score). For example, the scoring function $s(\widehat{F}, \widehat{p}, x)$ displayed above

corresponds to the expected payoff for a particular (\mathcal{D}, μ) -protocol. In our more general setup, we use a class \mathcal{D} whose elementary decision problems correspond to making choices when holding a hierarchy of option menus to exercise with random deadlines.

A key tradeoff in the design of these protocols is that the collection \mathcal{D} should be large enough to recover the entire belief structure—a complex, infinite dimensional object—while \mathcal{D} should also be small enough for the existence of an associated μ that induces a strict best response. Our first main result shows that, for a carefully selected class \mathcal{D} , and with appropriate probability measures μ , the protocol is “proper” in that it elicits dynamic beliefs of the forecaster as a strict best response at every instant. Our second main result shows that any protocol that elicits dynamic beliefs delivers scores or payoffs that can be approximated arbitrarily closely using some (\mathcal{D}, μ) -protocol in our class.

REFERENCES

- ALLAIS, M. 1953. Le comportement de l’homme rationnel devant le risque: Critique des postulats et axiomes de l’école américaine. *Econometrica* 21, 4, 503–546.
- BECKER, G., DEGROOT, M., AND MARSCHAK, J. 1964. Measuring utility by a single-response sequential method. *Behavioral Science* 9, 3, 226–232.
- BRIER, G. W. 1950. Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 78, 1, 1–3.
- CHAMBERS, C. P. AND LAMBERT, N. S. 2014. Dynamically eliciting unobservable information. *Working paper*.
- GNEITING, T. AND RAFTERY, A. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102, 477, 359–378.
- MATHESON, J. AND WINKLER, R. 1976. Scoring rules for continuous probability distributions. *Management Science* 22, 10, 1087–1096.
- MCCARTHY, J. 1956. Measures of the value of information. *Proceedings of the National Academy of Sciences of the United States of America* 42, 9, 654.
- ROCHET, J.-C. 1985. The taxation principle and multi-time hamilton-jacobi equations. *Journal of Mathematical Economics* 14, 2, 113–128.
- SAVAGE, L. 1971. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* 66, 336, 783–801.