ISSN 1936-5098

CAE Working Paper #09-13

Incentive-Compatible Elicitation of Quantiles

by

Nicholas M. Kiefer

September 2009

Incentive-Compatible Elicitation of Quantiles

Nicholas M. Kiefer*

September 7, 2009

Abstract

Incorporation of expert information in inference or decision settings is often important, especially in cases where data are unavailable, costly or unreliable. One approach is to elicit prior quantiles from an expert and then to fit these to a statistical distribution and proceed according to Bayes rule. An incentive-compatible elicitation method using an external randomization is available.

MSC 2000 Subject classification, Primary: 62C10; secondary 91B06

Keywords: Bayesian inference, mechanism design, prior assessment

1 Introduction

Incorporation of prior information is important in any decision or inference setting, whether it is done formally or informally. The formal Bayesian approach encourages transparency in assumptions, clear thinking and coherence. One example is

*Cornell University, Departments of Economics and Statistical Science, 490 Uris Hall, Ithaca, NY 14853-7601, US, email:nicholas.kiefer@cornell.edu; US Department of the Treasury, Office of the Comptroller of the Currency, Risk Analysis Division, and CREATES, funded by the Danish Science Foundation, University of Aarhus,Denmark. *Disclaimer: The statements made and views expressed herein are solely those of the author, and do not necessarily represent official policies, statements or views of the Office of the Comptroller of the Currency or its staff.*

risk management in financial institutions in which prudent management requires understanding default probabilities for groups of homogeneous assets. In the case of very safe or new types of assets, there may not be enough data information to support a practical conventional estimator, for example the frequency estimator in the case of binomial defaults. This issue has attracted regulatory and industry as well as academic attention, see Kiefer (2009), and references given there. Kiefer (2010) proposes eliciting prior quantiles for an expert's prior on the value of the default probability for a particular group of assets. For example, the median can be assessed by asking the expert at what value of a default rate θ would he be equally suprised to see a realization above or below θ. These quantiles (perhaps after feedback and revision) are assembled into a distribution, either by fitting a specific functional form or as proposed by Kiefer (2010) fit to a smoothed maximum-entropy distribution. The idea is to impose as little information as possible beyond that elicited from the expert. This distribution is then used to process data information through Bayes rule and the likelihood function. The latter is itself typically a representation of a large number of probabilities in terms of a few parameters, so the statistical treatment and the approximations involved are the same in the prior and the likelihood.

Other examples of elicitation of quantiles and their use to form a prior distribution are cited in O'Hagan, Buck, Daneshkhah, Eiser, Garthwaite, Jenkinson, Oakley, and Rakow (2006) and include applications to drug testing, sales of engines, the effect of nuclear waste on temperature, and future earnings. The incentive compatibility question does not seem to have been addressed.

The difficult part is the elicitation of the quantiles, which requires thought and therefore some effort from the expert. Since the quantiles can never be observed, and therefore the assessment checked, there is an issue of providing an incentive for the expert to provide the required thought. The problem of eliciting probabilities for given sets is well-studied and a widely-used approach is *scoring*. The scoring method does not naturally extend to the quantile assessment problem, as shown in Section 2. Savage (1971) reviews techniques for assessing probabilities and notes an interesting interpretation of probabilities as prices. He notes that the device of outside randomization, used by Marschak (1964) to compel a true valuation for a bid or asked price applies also to probability assessment. An ingenious recent method due to Karni (2009) introduces a second outside source of randomness to eliminate possible effects of risk aversion. This method does extend naturally to eliciting quantiles as shown in Section 2 with a different development than that of Marschak or Karni. Section 3 concludes.

2 Eliciting Quantiles

Assume at the outset that the expert's information about the unknown quantity θ is coherent, that is that it can be described by a probability distribution. Classical discussions of the necessity of describing uncertainty in terms of probability are Savage (1954), De Finetti (1974), and Lindley (1982). We do not review these well-known demonstrations and simply assume that the expert's information is described in a probability distribution with *cdf* $F(\theta)$ and *pdf* $f(\theta)$. We wish to elicit the probability α quantile q_α with $F(q_\alpha) = \alpha$. Assume that $f(q_\alpha) > 0$ so that the quantile is well-defined. Assume for convenience that $Supp(f) \subseteq [0,1]$. This is natural when the uncertain quantity θ is a default probability; in other cases a parameter transformation may be appropriate. The method of scoring for eliciting the probability β associated with a given set, say $[0, q]$, rewards the expert with

$$r(\beta|q) = (I(\theta \in [0, q])g_1(\beta) + (1 - I(\theta \in [0, q])g_0(\beta)$$

where $I()$ is the indicator function $g_1(\beta)$ is a nondecreasing function, $g_0(\beta)$ is a nonincreasing function and β is the elicited probability, after the single realization of θ is seen. The scoring rule is proper if the expectation $Er(\beta|q) = F(q)g_1(\beta) + (1 - F(q))g_0(\beta)$ is maximized at $\beta = q_\alpha$. See Schervish (1989). The optimal choice of β for a risk-neutral expert rewarded by a proper scoring rule is $F(q)$. Fixing β and using this method to elicit the quantile q does not work as the optimal choice of q is 0 for $g_0(\beta) \geq g_1(\beta)$ and 1 for $g_0(\beta) \leq g_1(\beta)$.

The new method based on outside randomness works as follows. Suppose the elicitor wishes q_a. The elicitor has access to a genie which generates a random variable ξ from the uniform distribution on $[0, 1]$ and independently a random variable $d \in \{0, 1\}$ from the Bernoulli with probability α. Nature supplies one realization of θ. The expert supplies a value q for the $\alpha - th$ quantile. The expert receives a reward equal to $rI(\theta \in [0, \xi])$ if $q < \xi$ and rd if $q \geq \xi$. We introduce a utility function explicitly to show how the second random variable d eliminates the effect of risk aversion. The expert's utility of a payoff x is $u(x)$.Then

Theorem 1 *With the reward described above the optimal policy for the expert is to report the true quantile.*

Proof. *Consider the expected utility to the expert of supplying q. First, suppose the genie supplies ξ to the expert. Then the expected utility is*

$$v(q|\xi) = u(r)(F(\xi)I(q < \xi) + F(q_\alpha)I(q \geq \xi))$$

piecewise constant with a break at ξ. Marginalizing with respect to the uniform random variable ξ yields the unconditional expected utility function

$$v(q) = u(r)(\int_q^1 F(t)dt + F(q_\alpha)q)$$

The first-order condition is $v'(q) = u(r)(-F(q) + F(q_\alpha)) = 0$ and the function is concave, so the optimal policy for the expert is to report the true quantile. ■

The binomial random variable is not needed if the expert is risk neutral. In that case the rd payoff can be replaced by $r\alpha$.

An alternative proof from a decision-theoretic point of view and using lotteries can be given. This proof uses preferences over lotteries but does not require the full expected utility framework. Let $(x, p) \in R \times [0, 1]$ denote the lottery that pays $\$x$ with probability p and \$0 with probability $(1 - p)$. The expert payoff is $(r, F(\xi))$ if $\xi > q$ and (r, α) if $\xi \leq q$. Consider the report $q > q_\alpha$. If $\xi > q$ then the expert's payoff is $(r, F(\xi))$ whether he reports q or q_α. If $\xi \leq q_\alpha$ then the expert's payoff is (r, α) whether he reports q or q_α. If $q_\alpha < \xi < q$ then the expert's payoff is (r, α). However, had he reported q_α instead, his reward would have been $(r, F(\xi))$. But $F(\xi) > \alpha$ hence $(r, F(\xi))$ first-order stochastically dominates (r, α) so the expert cannot win and may lose as a result of reporting $q > q_\alpha$.Similarly, reporting $q < q_\alpha$ is dominated.

3 Conclusion

The classical elicitation problem concerns eliciting probabilities for given events. This paper studies the complementary problem of eliciting events for given probabilities. This is the problem involved in obtaining prior quantiles. Although the reward r does not affect the optimality condition, it is clear that the actual effort expended by the expert will depend on the value of the reward (precisely, on its utility). Perhaps some part of a bonus could be tied into the probability assessment. Interesting open questions include: Which quantiles and how many should be assessed? How much accuracy can be expected in a quantile assessment? Can experts be trained to improve their assessments? How can prior quantiles be assessed from a group of experts?

References

De Finetti, B. (1974): *Theory of Probability, Volume 1.* New York: Wiley.

Karni, E. (2009): "A Mechanism for Eliciting Probabilities," *Econometrica*, 77(2), 603–606.

Kiefer, N. M. (2009): "Default Estimation for Low Default Portfolios," *Journal of Empirical Finance*, 16, 164–173.

Kiefer, N. M. (2010): "Default Estimation, Correlated Defaults and Expert Information," *Journal of Applied Econometrics*, forthcoming.

Lindley, D. V. (1982): "Scoring Rules and the Inevitability of Probability," *International Statistical Review / Revue Internationale de Statistique*, 50(1), 1–11.

Marschak, J. (1964): "Actual versus Consistent Decision Behavior," *Behavioral Science*, 9, 103–110.

O'Hagan, A., C. E. Buck, A. Daneshkhah, J. R. Eiser, P. Garthwaite, D. J. Jenkinson, J. E. Oakley, and T. Rakow (2006): *Uncertain Judgements: Eliciting Experts' Probabilities.* Chichester: John Wiley & Sons.

Savage, L. J. (1954): *Foundations of Statistics*. New York: Wiley.

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

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