Bayesian approach works without independence: it has only been assumed here for simplicity and comparison with beliefs. What the Bayesian view does is to force one to consider the subtle nature of the dependence between the witnesses.

III: $p(a_i \mid A) = p(\bar{a}_i \mid \bar{A})$, $(i = 1, 2)$.

This asserts that the witnesses are equally reliable whether A is true or false. Again it is easy to imagine circumstances where this is not true. In some cultures there is a tendency for witnesses to say what they think will please the listener. So if A is the event "the airport is near," veracity is more likely when A is true than when it is false. Consequently one cannot be sure that $p(a_i \mid A)$ and $p(\bar{a}_i \mid \bar{A})$ are both p_i.

The Bishop certainly did not recognize the distinction, as have many writers after him. The Bayesian approach does not demand the equality: it merely forces one to recognize that two types of veracity are possible.

Applied to the Bishop's problem, the rector's approach forces one to consider one's initial belief in the event, the nature of the dependence between the witnesses, and the two forms of reliability that arise. We suggest that, on reflection, it will be admitted that all three features are relevant to the final answer. Even if the independencies and the equalities of the reliabilities are admitted, as the Bishop and the modern equivalent tacitly do, the result is still different from the Bishop's. It is of interest to enquire when they are equal. Equating (2) and $1 - (1 - p_1)(1 - p_2)$ easily gives after a little algebra the condition that

$$(1 - \pi) = p_1 p_2 \pi + (1 - p_1)(1 - p_2)(1 - \pi).$$

The righthand side is $p(a_1, a_2)$, the unconditional probability that both witnesses assert A is true, so that the Bishop and rector only agree (under assumptions II and III) if

$$p(\bar{A}) = p(a_1, a_2).$$

In words, the probability that the event is false has to be equal to the probability that both witnesses assert its truth. This is surely unreasonable.

I put it to the readership: my challenge has survived, probability does do better. Let us support the rector of Tunbridge Wells and not the Bishop of Bath and Wells: let us favor truth and not the establishment. (Bayes was a minister in the unestablished church.)

ACKNOWLEDGMENTS

I am grateful to Richard E. Barlow for useful comments on a first version of this tale, to Sir Richard Eggleston for illuminating discussions on the legal problems with two witnesses, and to Glenn Shafer for drawing my attention to the Bishop's paper.

Comment

David J. Spiegelhalter

It is fairly predictable that I should agree wholeheartedly with Professor Lindley's lucid justification of probability as the correct paradigm for handling uncertainty in expert systems (but how strange it is to see him cast in the role of defender of orthodoxy!). In particular, his emphasis on remembering the background evidence H is crucial to avoid any conception that there is a single "true" probability of an event, and the frequent references to the operational meaning of probability gives a practical as well as a theoretical justification. However, playing the devil's advocate, I see two main reasons why the artificial intelligence community may not be convinced by the argument.

Firstly, he turns all statements expressing uncertainty into expressions of probability concerning (at least theoretically) verifiable events, whereas many constructors of expert systems would prefer to keep their propositions deliberately imprecisely defined in order to look more like human reasoning, and do not provide an operational means of verification. Secondly, even if verifiable events *are* being considered, the scoring rule argument presumes a certain type of evaluation procedure which many might claim was rarely appropriate, since the criteria for the "success" of an expert system may only require a very coarse handling of uncertainty.

Nevertheless, the theoretical arguments concerning optimality and coherence are only one weapon in the armoury. Pearl (1986b), in a recent strong advocacy of probability, uses no normative criteria but concentrates on the power of the theory in adequately modeling complex evidential reasoning, and I feel, in the end, it will be the intuitive appeal and flexibility of probabilistic reasoning that will change the current climate.

Professor Shafer's historical perspective puts the current discussion in an appropriate context, and emphasizes that many of the issues raised in expert system research are by no means novel. The interest in belief function methodology is understandable, as it appears to provide a means of avoiding full subjective assessment of a joint probability distribution, and—by formulating "uncertainty" in terms of reliability of evidence—it seems to attach uncertainty directly to the *rule* rather than the consequences of the rule. All this is very attractive, but users of the methodology also have to take on board a rule of combination that can lead to somewhat unintuitive results (Zadeh, 1986), problems in providing an operational interpretation of the numerical inputs and outputs, and a considerable computational burden.

Shafer does show how computationally efficient schemes are available on simple trees, but this is an extremely restrictive class of model, excluding both multiple causes of the same event, and an element being a member of two taxonomic hierarchies (for example, "gallstones" may also be part of a "dyspepsia" taxonomy). In contrast, efficient probabilistic schemes are now being devised for general graphical structures.

This still leaves the ability of belief functions to deal with "unknown" or "unknowable" probabilities. From a historical point of view, it would be easy to slip into the "likelihood versus Bayesian" debate at this point. But I believe the objective of constructing expert systems enables us to avoid such arguments. In such technological applications, there is real understanding of the problem to be exploited, and from a purely pragmatic point of view, unknown probabilities just do not occur—an assessment can always be obtained by careful questioning. Of course, the subject may not feel too confident in his assessment, and will not be able to list a set of independent sources of evidence for his opinion. But the opinion is there and can be used, although, as Professor Lindley emphasizes, in certain circumstances the imprecision may be relevant. As Professor Shafer points out; explanation of a system's conclusions may be provided at many levels, and probability judgments that have not been "constructed" on specified evidence can, if necessary, be identified. Provided a system's predictive performance is being monitored by scoring rules, it seems quite reasonable in a medical area to exploit "informed guesses" rather than rely on a legalistic paradigm that models unreliable "witnesses."

ADDITIONAL REFERENCE

ZADEH, L. (1986). A simple view of the Dempster-Shafer theory of evidence and its implication for the rule of combinations. *Artificial Intelligence Mag.* 85–90.

Rejoinder

Glenn Shafer

Watson and Dempster and Kong underline the point that belief functions are a form of probability. I can only say that I agree wholeheartedly.

I still have some bones to pick, on the other hand, with Spiegelhalter and Lindley.

Spiegelhalter's comments on the computational situation are misleading. He suggests that computationally efficient schemes for belief functions are available only for a very restrictive class of models, whereas efficient Bayesian schemes "are now being devised" for very general models. In fact, most Bayesian computational schemes have belief-function generalizations. It is true that the Bayesian special cases usually require less computation; Bayesian models require more complicated inputs than belief-function models, and there is less need for computation when you begin with more information. But the trade-off between complexity of input and complexity of computation differs from case to case, and belief-function computations are manageable in a greater variety of situations than Spiegelhalter suggests.

In my article, I discussed Judea Pearl's work on propagating Bayesian belief functions in trees, and I noted that Pearl's Bayesian scheme is a special case of a general scheme for propagating belief functions in trees. This general scheme has now been described in some detail by Shafer, Shenoy, and Mellouli (1986). In recent unpublished work, Pearl and Spiegelhalter have made progress in dealing with Bayesian networks that are not trees. Similar work is also underway for belief functions, with the most important contribution so far being Augustine Kong's dissertation at Harvard (Kong, 1986). In the last chapter of this dissertation, Kong shows how the belief-function scheme of Shafer and Logan (1985) can be adapted to handle multiple diseases with no additional computational cost.