FIESCHI, M., JOUBERT, M., FIESCHI, D., BOTTI, G. and ROUX, M. (1983). A program for expert diagnosis and therapeutic decision. *Med. Informatics* **8** 127–135.

GORDON, J. and SHORTLIFFE, E. H. (1984). The Dempster-Shafer theory of evidence. In Buchanan and Shortliffe (1984), 272–292.

KANAL, L. N. and LEMMER, J., eds. (1986). *Uncertainty in Artificial Intelligence.* North Holland, Amsterdam.

KIM, J. H. and PEARL, J. (1983). A computational model for causal and diagnostic reasoning in inference systems. *Proc. of the Eighth International Joint Conference on Artificial Intelligence.* 190–193. William Kaufmann, Los Altos, Calif.

KULIKOWSKI, C. A. and WEISS, A. M. (1982). Representation of expert knowledge for consultation: the CASNET and EXPERT projects. In *Artificial Intelligence in Medicine* (P. Szolovits, ed.) 21–55. Westview Press, Colorado.

LUCAS, R. W., CARD, W. I., KNILL-JONES, R. P., WATKINSON, G. and CREAN, G. P. (1976). Computer interrogation of patients. *British Med. J.* **2** 623–625.

MILLER, R. A., POPLE, H. E., JR. and MYERS, J. D. (1982). INTERNIST-1: An experimental computer-based diagnostic consultant for general internal medicine. *N. Engl. J. Med.* **307** 468–476.

PAUKER, S. G., GORRY, G. A., KASSIRER, J. P. and SCHWARTZ, W. B. (1976). Towards the simulation of clinical cognition: taking a present illness by computer. *Amer. J. Med.* **60** 981–986.

PEARL, J. (1986a). On evidential reasoning in a hierarchy of hypotheses. *Artificial Intelligence* **28** 9–15.

PEARL, J. (1986b). Fusion, propagation and structuring in belief networks. *Artificial Intelligence* **29** 241–288.

QUINLAN, J. R. (1983). Inferno: a cautious approach to uncertain inference. *Comput. J.* **26** 255–269.

RAUCH, H. E. (1984). Probability concepts for an expert system used for data fusion. *Artificial Intelligence Mag.* 55–60.

SMITH, C. A. B. (1961). Consistency in statistical inference and decision (with discussion). *J. Roy. Statist. Soc. Ser. B* **23** 1–37.

SPIEGELHALTER, D. J. (1986a). A statistical view of uncertainty in expert systems. In *Artificial Intelligence and Statistics* (W. Gale, ed.) 17–56. Addison-Wesley, Reading, Mass.

SPIEGELHALTER, D. J. (1986b). Probabilistic reasoning in predictive expert systems. Cited in Kanal and Lemmer (1986).

SPIEGELHALTER, D. J. and KNILL-JONES, R. P. (1984). Statistical and knowledge-based approaches to clinical decision-support systems, with an application in gastroenterology (with discussion). *J. Roy. Statist. Soc. Ser. A* **147** 35–77.

SZOLOVITS, P. and PAUKER, S. G. (1978). Categorical and probabilistic reasoning in medical diagnosis. *Artificial Intelligence* **11** 115–144.

TUKEY, J. W. (1984). Discussion of Spiegelhalter and Knill-Jones. *J. Roy. Statist. Soc. Ser. A* **147** 62–64.

VAN MELLE, W., SCOTT, A. C., BENNETT, J. S. and PEAIRS, M. A. S. (1981). The EMYCIN Manual. Report HPP-81-16, Computer Science Dept., Stanford Univ.

WERMUTH, N. and LAURITZEN, S. L. (1983). Graphical and recursive models for contingency tables. *Biometrika* **70** 537–552.

Comment

Stephen R. Watson

1. COMMENTS ON SHAFER'S PAPER

One of the things that makes Shafer's theory interesting is that it can be seen as an alternative to the traditional probability theory. Is this really so, however? Firstly, note that one of the strengths of subjective probability theory is the clear cut nature of the axiomatic support for the theory. Indeed, as Lindley's contribution shows, it is possible to claim that probability theory is the only theory one could possibly use to represent uncertainty. Shafer's theory does not as yet have such a clear cut support. For example, although Shafer recognizes the importance of canonical examples, as yet belief function theory is not provided with as strong an axiomatic support as that which is available for probability theory.

It can be claimed, however, that belief functions are indeed rooted in probability theory. It is just that the probability is associated with a power set rather than a simple set. If this interpretation of belief function theory is accepted, then indeed there is no problem, since the philosophical support for probability theory clearly also will support belief function theory. However, Shafer seems in some of his writings not to be very happy with this interpretation of his theory. And if he rejects this interpretation then the problem of a philosophical foundation for belief function theory remains.

The second point I make here concerns concepts of independence. Shafer touches on this point in his paper, but it is worth saying again that concepts of independence in belief function theory are not yet clear. In the application of Dempster's rule to determine the support for a hypothesis on the basis of two pieces of evidence, there is a rather vague notion that the two pieces of evidence should be independent in

---

Stephen R. Watson is Peat Marwick Professor of Management Studies in the University of Cambridge. He is also Head of the Management Studies Group in the Department of Engineering and a Fellow of Emmanuel College. His address is Department of Engineering, University of Cambridge, Mill Lane, Cambridge CB2 1RX, England.

some way. The detailed meaning of this concept of independence is far from clear. Shafer recognizes this difficulty and in his discussion of frames is attempting to overcome it. It is sufficient to say at this point, however, that we do not yet know how to handle dependence concepts in belief function theory in a way which is intuitively understandable.

2. COMMENTS ON LINDLEY'S PAPER

The conviction with which Professor Lindley speaks, and the sheer power of his argument impel users of alternatives to probability theory to respond to his arguments. If we do not accept the inevitability of probability, why not?

Users of Shafer's theory or Zadeh's theory can, and in fact have in the past, respond that they do indeed accept the inevitability of probability. As Dempster has commented, belief function theory is founded on probability, and so there is no contradiction in using belief function theory at the same time as using probability theory. Moreover, one can think of fuzzy set theory as being a heuristic approach in situations where a full probabilistic analysis is far too complicated to be undertaken.

It is, however, also possible to take issue with Lindley's argument. In other words, it is possible to question some of the premises in his argument and thereby avoid the full power of his conclusions. If one investigates the development of subjective probability theory exemplified by Savage's approach, it is possible to ask whether we are prepared to accept the axioms. It is a commonplace now that people do not behave as though they accept Savage's axioms, reasonable as they undoubtedly are. Of course, these axioms are normative and it can be argued, as indeed Lindley does argue, that the fact that we fail to abide by the axioms does not mean that we should not attempt to do so. Indeed he would say that the first act of a rational man is to agree to the axioms, and then attempt to construct his behavior in accordance with these axioms. If, however, we are not prepared to do this, then what happens to us is a matter of practice. It could be argued that if we are consistent in our failure to abide by the axioms, then our opponents can turn us into a money pump or construct a Dutch Book of gambles against us. Of course, we do not do this in practice. We just recognize when we are about to get cornered in this way, and change one of our judgments, possibly in a yet more inconsistent way with our past judgments. There is, therefore, nothing mandatory about accepting Savage's axioms, and we can therefore escape Lindley's conclusions if we wish to.

In his contribution, Lindley gave a very clear account of an alternative way of showing the inevitability of the probability. This was based on the notion of scoring systems. It is indeed quite remarkable that no matter what kind of scoring system one adopts, the numbers that one employs to describe uncertainty must (after an appropriate transformation) satisfy the rules of probability theory. Compelling as this argument is, we have to point out that in practice no Great Scorer exists. There is nobody hovering about us being prepared to dock our pay should we use numbers which fail to conform to the rules of probability theory in our descriptions of uncertainty. Thus while the argument is elegant and powerful, there is nothing inherently irrational in not accepting it, because in practice scoring systems do not exist.

Of course the proof of the pudding is in the eating. If it can be shown that in the long run any person who fails in his assessment of uncertainty to combine his numbers as though they were probabilities will lose out inexorably, then indeed we have a problem in refusing to accept probability theory. But to my understanding practical proofs of this kind are not yet available.

Thus, it is possible to escape the inevitability of probability; it has to be admitted, however, that there is no alternative theory which has the strength of support, and elegant support at that, which is available for probability theory.

The chief drawback with using probability theory is the complexity that sometimes results, and the need to assess an often surprisingly large number of conditional probabilities. In legal work, for example, great difficulty can arise; some interesting work by Schum (1981) shows how problematic probabilistic inference can get. In one simple murder case, with five pieces of evidence, he needed to make 27 probability assessments. Lindley suggests the principle of Occam's razor should be applied to our topic: simplify where possible. Sometimes probabilistic analysis is far from simple.

3. COMMENTS ON SPIEGELHALTER'S PAPER

Spiegelhalter's paper is a most interesting account of the construction of an expert system for medical diagnosis. He gives us some important insights into the practical problems of constructing an expert system, which is both computable and also useful. This raises the general question of how one determines whether a particular expert system, as represented on some computer, is actually a good one or not. The issues involved are very similar to those involved in validating a model. Firstly, one needs the system to be faithful to some normative principle. In my view one should start with probability theory since it has the strongest theoretical base, but be prepared to adopt other approaches as heuristics or as richer representations of the issues involved. It seems that Spiegelhalter's approach has been similar.

Secondly, one could validate an expert system by its comparison with expert performance. One can ask whether the diagnosis achieved by Spiegelhalter's system was better or worse than that achieved by competent diagnosticians. There is of course a debate over whether an expert system should be appraised in this way. Is the goal to reproduce the abilities of an expert, or to improve on the abilities of available human judges? If it is the former, then indeed it is sensible to compare performance with experts, but in this case one wonders why one should not use the experts themselves. This could be answered by observing that very often experts are in short supply. If, on the other hand, our goal is to improve on human inference behavior, then the criterion of conformity with some expert performance is not appropriate. A final measure of the appropriatness of an expert system is user satisfaction. To what extent do the people who interact with the expert system feel that the system is of use to them? In Spiegelhalter's case there are two kinds of people involved, namely the patients and the doctors. As Spiegelhalter observes, it is very important that the doctors are supportive of the endeavor and that they do not feel that their professional competence is in any way being threatened. It is perhaps more important, however, that the patients feel that they are being properly attended to. Spiegelhalter seems to have achieved success on both fronts.

4. SUMMARY

Although the purpose of the conference was to discuss the use of the different theories for the representation of uncertainty in expert systems, the principal speakers, perhaps wisely, devoted their discussion mainly to arguing the cases for the use of their different theories in general. On the basis of the discussions we had at this conference, it seems to me that one can summarize as follows. Probability theory has a strong intellectual support and in principle there is no reason why one should not be satisfied with this theory. Its use does, however, lead to enormous problems of complexity, and as a matter of practice it is necessary to seek for approximations. Fuzzy set theory can be viewed as a heuristic for handling those situations where imprecise inputs and imprecise inferences are required without the need to resort to the greater complexity of probability theory. Belief function theory can be thought of as a way of representing inferences from evidence within the probabilistic framework.

There are yet other alternative approaches to handling uncertain inferences which were not mentioned at the conference, and notable among these is the nonmonotonic logic of Doyle. Recently Cohen (Cohen, Watson and Barrett, 1985) has suggested a combination of Doyle's theory with both Shafer's and Zadeh's which he has referred to as the nonmonotonic probabilist. This seems an exciting possibility for approaching the problem at the heart of this conference.

ADDITIONAL REFERENCES

COHEN, M. S., WATSON, S. R. and BARRETT, E. (1985). Alternative theories of inference in expert systems for image analysis. Technical Report 85-1, Decision Science Consortium, Falls Church, Va.

SCHUM, D. A. (1981). Sorting out the effects of witness sensitivity and response-criterion placement upon the inferential value of testimonial evidence. *Organizational Behavior and Human Performance* **27** 153–196.

Comment

A. P. Dempster and Augustine Kong

The papers by Shafer and Spiegelhalter are valuable summaries by acknowledged leaders in active research fields. There is much food for thought in both papers, and many of the techniques and issues raised by these authors will gradually become better understood as the field of uncertainty assessment in expert systems advances. Our research on models and techniques for belief function analysis (Kong, 1986; Dempster and Kong, 1986) is complementary to that of Shafer and Spiegelhalter. We all seek to provide tools for real applications, based on carefully constructed analyses expressed through mathematically well-articulated principles of uncertain reasoning.

Lindley is on a different track. He rehearses familiar normative arguments for the Bayesian paradigm, evidently seeking to persuade less committed colleagues to abandon their fallacious ways. Unfortunately, he shows no interest in understanding how his

A. P. Dempster is a Professor and Augustine Kong is an Instructor at the Department of Statistics, Harvard University, Cambridge, Massachusetts 02138.