also important differences. It is difficult to use the verb "measure" without pretending that there is a well-defined property to be measured. Talk about canonical examples encourages a more constructive attitude.

One aspect of the constructive nature of Bayesian probability judgment, emphasized by Shafer and Tversky (1985), is the fact that we must construct our starting point. We must construct a probability distribution before we can condition it or multiply it by likelihoods. Bayesian theorists often assert categorically that every new experience must be treated in terms of its likelihood. Lindley, for example, declares that "an AI system faced with uncertainty about A_2 and experiencing A_1 has to update its uncertainty by considering how probable what it has experienced is, both on the supposition that A_2 is true, and that A_2 is false." But since a person may get around to constructing "initial" probabilities only after experiencing A_1, he or she has the option of treating A_1 as part of the evidence for those initial probabilities. Consider Lindley's investigator, who has discovered evidence that a criminal is left-handed. Instead of treating this evidence in terms of its likelihood, the investigator uses it directly in constructing a probability distribution.

There are problems, of course, where the construction can all be done in advance and then applied to many cases. GLADYS deals with this kind of problem; the same framework is applied to one patient after another. If I understand Spiegelhalter correctly, he believes that the bounded nature of expert systems means that this is the only kind of problem with which they can deal.

A finite system that permits construction can, however, deal with an unbounded range of situations. This is one of the fundamental points of the generative theory of grammar. The constructive nature of human reasoning makes us capable of exploring ever new realms of experience, and the ambition of AI is to duplicate this capability. Rule-based expert systems are one attempt to do so. These systems do not handle probabilistic reasoning very well, and many AI theorists would conclude from this that probabilistic reasoning has little role in genuine intelligence. In order to prove them wrong, we must do more than retreat to bounded systems like GLADYS. We must take the problem of automating construction seriously.

ADDITIONAL REFERENCE

SHAFER, G. (1986b). Savage revisited (with discussion). *Statist. Sci.* **1** 463–501.

Comment: A Tale of Two Wells

Dennis V. Lindley

The main issue is whether uncertainty should be described by probability, belief functions, or fuzzy logic; not just in artificial intelligence and expert systems, but generally. Are we to be probabilists, believers, or fuzzifiers? Or do we need some mixture of all three disciplines? To me, the important distinction between the methods rests on the rules of combination of uncertainty statements. Do we operate with the calculus of probability, the rules of belief functions, or with those of fuzzy logic? In my paper the challenge was made "that anything that can be done by these methods (belief functions and fuzzy logic) can better be done with probability." This reply will address one such challenge and I hope to show that Dempster's rule for belief functions does not behave as well as Bayes rule. My discussion is therefore chiefly addressed to Shafer and Zadeh. The omission of any discussion of Spiegelhalter's contribution arises because I agree substantially with it, and highly regard it. I wish that his program for dyspepsia had been more Bayesian and that he had recognized that uncertainty about a probability is usually a reference to the desirability of obtaining more data, so that his conflict ratio should really reflect this. To return to the challenge.

In 1685 the then Bishop of Bath and Wells wrote a paper in which the following problem was discussed. Two witnesses separately report that an event is true. Both are known to be unreliable to the extent that they only tell the truth with probabilities p_1 and p_2 respectively. What reliability can we then place, in the light of the witnesses' testimonies on the truth of the event? The Bishop's answer was $1 - (1 - p_1)(1 - p_2)$. The following is a precis of his argument. If the event is false, both witnesses must have lied, an event of probability $(1 - p_1)(1 - p_2)$. Consequently one minus this is the required reliability.

The result retains its interest today because the