Comment

Glenn Shafer

I found it a pleasure to read the articles by Dennis V. Lindley and David Spiegelhalter. They present an elegant case for the use of Bayesian (i.e., conditional probability) methods in expert systems. Lindley provides a concise summary of arguments he and others have developed over the last several decades in support of the claim that rationality demands the use of conditional probability. Spiegelhalter supplements this with an account of what is actually being accomplished using conditional probability in diagnostic systems, and he also contributes some interesting new arguments.

In response, let me first express my admiration for the practical work Spiegelhalter reports on. The GLADYS system is especially attractive, because it brings close to reality the kind of probability calculation philosophers have always considered ideal—the calculation of conditional probabilities on the basis of massive and directly relevant frequency data. I share Spiegelhalter's excitement with the prospect that widespread use of microcomputers will enable us to attain this ideal much more often than we have in the past.

WHY GENERALIZE PROBABILITY?

Spiegelhalter emphasizes capabilities of the Bayesian language that are often overlooked. Weights of conflict can be used to monitor Bayesian analyses, and weights of evidence can be used to explain the results. There are Bayesian definitions of imprecision and ignorance. We do not need to generalize from Bayes to belief functions in order to formalize these concepts.

The desire to generalize Bayes does not spring, however, from dissatisfaction with the ideal of conditional probability. It springs from the realization that this ideal is sometimes unattainable. Directly relevant frequencies are often unattainable. Sometimes we can make decent conditional probability arguments even without such frequencies, but sometimes we cannot. Sometimes we simply lack evidence for some of the probability judgments that a given conditional probability analysis requires.

The only satisfactory description of uncertainty, Lindley tells us, is probability. He is no less correct than the man who believes that the only satisfactory household is one with a dozen servants. It's wonderful if you can afford it.

STANDARDS OF RATIONALITY

What should we say about the claim that rationality demands we make Bayesian analyses regardless of the availability of the ingredients? For my own part, I find that every argument for this claim boils down to another appreciation of the beauty of the Bayesian ideal.

Lindley believes that Savage's axioms are so self-evident that their violation would look ridiculous. But in fact these axioms derive their appeal from the Bayesian ideal rather than vice versa. If we did not have the picture of conditional probability and expected utility in mind, we would not even be able to understand most of Savage's axioms (Shafer, 1986b).

The idea of a scoring rule also derives from the Bayesian ideal rather than vice versa. It has relatively little force in abstraction from that ideal. If we intend to assign a number to each of two complementary events and accept a penalty for each event based on the number's distance from one if the event happens and its distance from zero if it fails, then we should make the two numbers add to one. But how would we explain this game to a naive listener? We would say that the numbers are supposed to be like probabilities—close to one for events that are expected to happen and close to zero for events that are expected to fail. The game fits the picture of additive, or frequency-like, probability, and it is incomprehensible outside that picture. It does not fit the theory of belief functions, where a degree of belief close to zero indicates inadequate evidence for the event, not assurance that the event will fail.

Another argument for Bayes is based on the relatively sharp preferences given by expected utility calculations. We can calculate upper and lower expectations from belief functions, but these will not give a definite preference between two alternatives as often as the Bayesian calculation will. But would we expect such sharp preferences were it not for our fascination with the Bayesian ideal? Would we really expect an analysis of our evidence and pre-existent preferences to tell us always exactly what to do, leaving no occasion for caprice? In fact, human beings, unlike Buridan's ass, are capable of choosing without sufficient reason, and they often use that capability. Building a similar capability into a computer is one of the easier tasks of artificial intelligence.

CONSTRUCTIVE PROBABILITY

In my contribution to this symposium, I say that Bayesian analyses use games of chance as canonical examples to which to compare actual evidence. Lindley says such games provide a standard by which to measure belief. There are commonalities here, but