WINKLER, R. L. and CUMMINGS, L. L. (1972). On the choice of a consensus distribution in Bayesian analysis. *Org. Behav. Hum. Perform.* **7** 63–76.

Example of an experimental study on the problem of pooling opinions. The results of an experiment are reported in which subjects were asked to subjectively generate a consensus distribution or to choose one from a fixed number of alternatives determined using different linear or logarithmic opinion pools. The weighted average methods are seen to have been most frequently used.

ZECKHAUSER, R. (1968). Group decision and allocation. Discussion paper No. 51, Harvard Institute of Economic Research, Cambridge, MA.

Cited by Raiffa (1968) and Bacharach (1975).

Comment of "Combining probability distributions" by C. Genest and J.V. Zidek.

Glenn Shafer

This is a valuable review article. The annotated bibliography is a useful guide to the literature on the combination of beliefs, and the body of the article puts us in a position to assess the accomplishments and the direction of this literature as a whole, without undue emphasis on the ambitions and limitations of particular contributions.

My own view is that most of this literature is flawed by adherence to one of more of the following fallacies: (1) The Conditional Probability Fallacy: A Bayesian analysis of a problem always takes evidence into account by conditioning a probability distribution on that evidence. (2) The Fallacy of the Coherent Individual: Formation of opinion by a group is fundamentally different from formation of opinion by an individual. (3) The Fallacy of Normalcy: Use of the Bayesian paradigm is normative for an individual. My contribution to the discussion will concentrate on these fallacies.

1. THE CONDITIONAL PROBABILITY FALLACY

My interest was engaged when I read in the introduction that the problem of pooling knowledge would be covered in Section 4. But when I turned to Section 4, I found it entitled "The Supra Bayesian Approach." It appears to be taken for granted in the literature surveyed that the only Bayesian way to pool the knowledge represented by the opinions of several different people is to condition a "supra Bayesian's" probability distribution on these opinions. This is a special case of the conditional probability fallacy.

Like most fallacies, the conditional probability fallacy survives not because of persuasive arguments in its favor but because it so often goes unnoticed. It is a habit of thought resulting from our familiarity with the picture of statistical experimentation associated with parametric inference. The evidence that appears explicitly in this picture is a statistical observation, the result of a statistical experiment. The hypothesis (or parameter) space and the evidence (or observation) space are specified when the experiment is set up, before the observation is made. We are supposed to have a joint probability distribution over the Cartesian product of the hypothesis and evidence spaces, and once we get the evidence (i.e., make the observation) we are supposed to condition this probability distribution on it. Often we are told to do this in a way that uses Bayes's theorem; we must specify the joint distribution by specifying $P(H)$ and $P(E \mid H)$ for each element H of the hypothesis space and each element E of the evidence space, and then we must use Bayes's theorem to calculate $P(H \mid E)$.

In fact, a Bayesian analysis cannot take all the evidence into account by conditioning even in the case of a genuine statistical experiment. If the analysis is to be convincing, there must be some evidence that is used directly as evidence for the numbers $P(H)$ and $P(E \mid H)$. Since this evidence is not part of E, it is not taken into account by conditioning. This point is often overlooked because many advocates of the Bayesian paradigm do not want to acknowledge that the usefulness of Bayes's theorem in a particular problem depends on the existence and quality of the evidence for $P(H)$ and $P(E \mid H)$.

Our habituation to the picture of statistical experimentation also has a more subtle effect. Even after we admit that the evidence E on which we condition is not all our evidence, we tend to assume that E has been singled out from our other evidence by nature, not by ourselves. It is our "new evidence," the evidence that we just got from our experiment, and so it is easily distinguished from our "background information."

It is important to recognize that in the case of everyday evidence, at least, it is usually not true that E is singled out for us. We ourselves, when designing a Bayesian analysis, must decide which part of our evidence we will use to construct a probability distri-

Glenn Shafer is a Professor in the School of Business, University of Kansas, Lawrence, KS 66045.

bution and which part we will take into account by conditioning that distribution. In some cases our design will not involve any conditioning at all. (See the discussion of total evidence designs in Shafer and Tversky, 1985.) Moreover, the evidence we decide to condition on will not necessarily be the evidence we have most recently acquired. We often decide on a design that involves conditioning on certain evidence and then go looking for evidence on which to base the probability distribution to be conditioned.

How do we decide how to partition our evidence? Clearly, we must look for a partition that permits a convincing analysis, an analysis which takes proper account of all our evidence and in which the evidence used to support the construction of the joint probability distribution for E and H is adequate for that purpose. This results in a certain tension. We want to put as much of our evidence as possible into E, because when we do so we feel confident that we have examined it thoroughly and taken it fully into account, but the more we put into E, the more elaborate a probability distribution we need in order to condition on E, and the greater the need for evidence on which to base this probability distribution.

Here is a whimsical way of refuting the conditional probability fallacy. Imagine a person who is told that he should always take new evidence into account by conditioning. He finds some new evidence, and realizes that he must now find evidence on which to base the probabilities to be conditioned. He sets out to find such evidence, but every time he finds some he has made his difficulty worse. The evidence he finds is itself new evidence and so must be taken into account by conditioning, and hence, his need for further evidence is even greater than before.

The relevance of the conditional probability fallacy to the problem of combining the probability judgments of a group of individuals is obvious. If we want to pool the evidence underlying these probability judgments, we must first think about what this evidence is; we must identify the domain and assess the quality of each individual's experience and knowledge. Then we must design a probability analysis that makes good use of all this evidence. Different individuals may supply the different elements needed in this design. One might contribute a judgment of independence, another might contribute a conditional probability, and so on. It is possible that the opinions of some of the individuals might best be taken into account by conditioning a probability distribution constructed from the evidence of the others. It is also possible, although extremely unlikely, that the opinions of all but one might best be taken into account by conditioning a probability distribution constructed from the evidence of that single one, the supra Bayesian. But it is simply incoherent to suggest that the evidence of the whole group might be taken into account by conditioning a probability distribution that they construct together, incoherent because if all their evidence is to be taken into account by conditioning, there is none left for constructing the probability distribution to be conditioned.

In summary, we can say that it is unfair to the Bayesian method to claim that a Bayesian analysis can pool the evidence underlying the probability judgments of a group of individuals only by conditioning a probability distribution on all those judgments.

2. THE FALLACY OF THE COHERENT INDIVIDUAL

The preceding discussion of ways in which a group of individuals might design and carry out a Bayesian analysis suggests that there is no fundamental difference between a group and a single individual in this respect. I believe this is the case.

The literature surveyed by Genest and Zidek shows a deep appreciation of the variety of ways in which the members of a group might differ in their probability judgments. But is not the same capacity for dissonance found within every individual? All of us have had experience in many different domains, and we often find that these different domains of experience provide conflicting arguments about what we should expect in a particular problem. We are now all familiar with psychological studies demonstrating that different ways of expressing a problem can evoke different arguments and hence different probability judgments from a single person. Why, then, should we expect a qualitative difference between the difficulty faced by group and the difficulty faced by an individual when either undertakes to construct consistent probability judgments that take account of all the evidence?

Many groups do not, of course, have the strong sense of identity that most individual human beings have. In many situations we expect individuals to resolve whatever internal conflicts they have and settle on a definite opinion or choice. We do not always expect groups to do this. We know that a group may simply dissolve because of differences of opinion among its members. But a closer look again suggests that this difference between groups and individuals is only a matter of perspective and of degree. As individuals, we see many groups from the inside. We do not see other individuals from the inside, and we often conceal our own inner conflicts even from ourselves. But we know these conflicts do exist. The most difficult problems of decision for an individual are precisely those where there is sharp inner conflict, and resolution of such conflict can be as essential to the survival of the individual as it is to the survival of a group (Janis and Mann, 1977).

3. THE FALLACY OF NORMALCY

Genest and Zidek suggest that the Bayesian paradigm is normative for individuals but not for groups. Why? Let us grant that difficulties arise when a group tries to use the Bayesian paradigm at the same time that the individuals in the group are trying to use it. Does it follow that the group must yield to the individuals? Rather than concluding that the paradigm is normative for individuals but not for groups, why not conclude that it is normative for groups but not for individuals?

We are often told that the normative character of the Bayesian paradigm derives from the appeal of the axioms. Is there any reason to suppose that these axioms are more appealing to individuals than to groups?

For my own part, I do not think that it is reasonable to say that the Bayesian axioms are normative for anyone. Yes, the axioms are appealing. But far more appealing is the single axiom, "Always be right." Yet we do not say that it is normative always to be right.

A norm is a standard which is usually met. If people were almost always right and could be right all the time with just a bit more effort, then we would say that it is normative always to be right. If people nearly always conformed to the Bayesian axioms and could conform completely just by smoothing off some rough edges, then we might say that it is normative to conform to these axioms. When someone asserts that the Bayesian paradigm is normative for individuals but not for groups, he or she is really asserting that individuals, unlike groups, do have well defined opinions that nearly conform to the Bayesian axioms. This assertion can be tested empirically, and it is false.

4. CONSTRUCTIVE PROBABILITY

If we refuse a normative status to the Bayesian axioms, then how do we explain what a Bayesian analysis means? What is an individual or a group doing when it constructs a subjective probability distribution?

I believe that when we construct a subjective probability distribution we are making an argument by analogy. We are comparing our actual evidence in a problem to knowledge of the objective chances in a problem in which the answer to the question that concerns us is determined by chance. Like any argument by analogy, a Bayesian probability argument may or may not be convincing. Our evidence may or may not be analogous in its structure and strength to knowledge of the chances in a chance situation.

As soon as we accept this constructive view of probability judgment (Shafer, 1981), we see that the ambit of subjective probability may be much wider than the Bayesian paradigm. When we make a Bayesian probability argument, we are comparing our evidence to idealized canonical examples where the answers to our questions are determined by chance. But there are other kinds of subjective probability argument, and sometimes these are more convincing. Often it is more convincing to compare our evidence to idealized examples where we know the chances governing an experiment for each possible value of a parameter but have no evidence about the value of the parameter. We are accustomed to thinking of these arguments as objective rather than subjective, but in fact they usually involve probability judgments that are quite subjective. Like Bayesian arguments, they are really subjective arguments that compare actual evidence to idealized canonical examples. They differ from Bayesian arguments only in that they use different canonical examples.

The theory of belief functions (Shafer, 1984) also uses canonical examples that differ from those for the Bayesian paradigm. In these canonical examples, it is the meaning of the evidence, rather than the answer to the question that concerns us, that is determined by known chances. Since it uses probability to model evidence rather than fact, this theory deals with the problem of pooling evidence in a more direct way than the Bayesian theory does.

ACKNOWLEDGMENT

Reseach for this comment was partially supported by Grant MCS-8301282 from the National Science Foundation.

ADDITIONAL REFERENCES

JANIS, I. L. and MANN, L. (1977). *Decision Making: A Psychological Analysis of Conflict, Choice, and Commitment.* The Free Press, New York.

SHAFER, G. (1981). Constructive probability. *Synthese* **48** 1–60.

SHAFER, G. (1984). Belief functions and possibility measures. Working paper no. 163, School of Business, University of Kansas. To appear in Volume 1 of *The Analysis of Fuzzy Information*, edited by James C. Bezdek.

SHAFER, G. and TVERSKY, A. (1985). Languages and designs for probability judgment. *Cognitive Sci.* **9** 309–339.