Comment

Robin M. Hogarth

As the review by Genest and Zidek shows, there is now a burgeoning literature on issues concerning the combining of probability distributions. Moreover, readers will be grateful to them for providing such a comprehensive overview and guide to this literature. In the following comments, I wish to emphasize three points. These relate to (1) whether it is reasonable to expect group opinion to act like the Bayesian model; (2) the importance of determining the commonalities and differences between distributions that are being combined; and (3) how considering the specific decision context can, from a practical viewpoint, often simplify the problem of combining distributions.

(1) On considering whether one should expect group opinion to conform to the Bayesian model, it is first important to consider what one means by group opinion and the purposes for which this has been elicited. I distinguish three forms of "group opinion." The first occurs when an *individual* decision maker assesses a distribution using, as inputs, distributions from other sources, e.g., experts, forecasts from models, and so on. In this case, the combined distribution becomes the opinion of the individual. A second case involves a group of people whose members wish to combine their probability distributions in order to make a particular decision. Examples of this kind of situation could involve business partners or even a married couple or family. A third case concerns a group of people who wish to express an opinion about some issue in the form of a probability distribution in order to communicate this as information for other people. An example would be professional groups (e.g., physicians) providing information to the general public (e.g., about health risks).

From a technical viewpoint, one could treat all these situations identically, e.g., by adopting a "supra Bayesian" model. However, in my view the three situations are conceptually quite different. The first and third situations are extreme cases. In the first, there seems little doubt that one would want the Bayesian model to apply. For example, one would expect the individual to update his or her "consensus" distribution in the light of new evidence according to Bayes' theorem. In the third case, however, all one really wants is a consensus of opinion at a particular point in time. It would seem strange to require the professional association to follow *all* the dictates of the Bayesian model with respect to that opinion. The second case is more complicated. On the one hand, it is possible to treat this like the third case and simply work on assessing a distribution for a particular problem. On the other hand, many people (myself included) feel that one should be able to apply the Bayesian theory to multiple party decision making. What distinguishes these situations? In my view, a critical variable is the extent to which the multiple parties resemble an individual decision maker. However, since resemblance has many dimensions, let me suggest two criteria: (a) whether group members are involved in a stable long-term relationship; and (b) whether they have similar or even identical interests. For example, members of professional associations may have some common interests but their relationships are inherently unstable (e.g., membership is changing constantly). Married couples, or even close business partners, on the other hand, have relationships that can be more easily assimilated to the notion of a single Bayesian decision maker. Without providing an answer, I am suggesting that we think more carefully about what kinds of groups could or should be thought of as Bayesian decision makers. It is not clear to me that an all purpose solution exists for this problem, nor that it would be desirable.

(2) In a recent paper, Clemen and Winkler (1985) show the deleterious effects on the combination process of assuming that distributions are independent when in fact they are not. Conceptually, it is possible to distinguish two types of dependence in aggregation. One is the common notion that two or more distributions may be based on the same data and thus double counting occurs if the data considered in each distribution are treated as independent in the aggregation process. The second is what psychologists term "common method variance." This refers to the possibility that two or more distributions being aggregated result from analyses based on a common method. To illustrate, consider the following thought experiment. Would you have more or less faith in a medical diagnosis if there was agreement between two experts using different methods of diagnosis as opposed to agreement between two experts using the same method?

From a practical viewpoint, what can be done to handle problems of redundancy among distributions? Much prescriptive and descriptive work in decision making attests to the value of decomposing decision

Robin M. Hogarth is Professor of Behavioral Science and Director of the Center for Decision Research at the Graduate School of Business, University of Chicago, 1101 East 58th Street, Chicago, IL 60637.