Aequationes Mathematicae 32 (1987) 74–86
University of Waterloo

0001–9054/87/001074–13$1.50 + 0.20/0

Further evidence against independence preservation in expert judgement synthesis

CHRISTIAN GENEST AND CARL G. WAGNER

Dedicated to Professor Otto Haupt with best wishes on his 100th birthday.

Abstract. When a decision maker chooses to form his/her own probability distribution by combining the opinions of a number of experts, it is sometimes recommended that he/she should do so in such a way as to preserve any form of expert agreement regarding the independence of the events of interest. In this paper, we argue against this recommendation. We show that for those probability spaces which contain at least five points, a large class of seemingly reasonable combination methods excludes all independence preserving formulas except those which pick a single expert. In the case where at most four alternatives are present, the same conditions admit a richer variety of non-dictatorial methods which we also characterize. In the discussion, we give our reasons for rejecting independence preservation in expert judgement synthesis.

1. Introduction

Suppose that you are faced with a decision problem which involves two sources of uncertainty: the future value of the British pound, and whether or not a piece of equipment will fail. In order to help you construct your own probability distribution for these two events, suppose that you consulted two experts who agree with you that the fate of the British pound and that of the piece of equipment are independent events. At the same time, however, your experts disagree on their assessments of probabilities for the different outcomes. How should you reconcile their distributions? Should you (I) pool the experts' opinions for each of the two events separately and then use independence to compute such quantities as the probability that the British pound will rise *and* the equipment will fail? Or should you (II) somehow "average" the experts' distributions at the joint level, even though this might destroy independence?

If confronted with this situation, most decision makers would probably opt for (I),

---

AMS (1980) subject classification: Primary 62A99, 62C05. Secondary 39B40.

Manuscript received December 13, 1985, and in final form, August 12, 1986.

and a large number of those would possibly suggest combining the experts' marginal distributions using a weighted arithmetic mean. This choice is not as innocuous as it may first seem, however, as illustrated in [16, pp. 231–233] with concrete reference to the above example, which was originally formulated by John Pratt. If you compromise at the component level, situations could arise in which maximizing their utility would lead all of your experts to choose act a, say, whereas you would select another act, b! Moreover, this phenomenon may occur even if you and all the members of the panel share the same utilities for consequences. To preserve independence, in other words, you must be prepared to break the consensus which may already exist between the experts at the decisional level. Interestingly, this anomaly cannot occur if option (II) is selected and if you average probabilities at the joint level using a weighted arithmetic mean. Yet, independence preservation is favoured in [16] and recommended explicitly in [9] and [11]. These authors and probably many others would argue that when a group of experts shares a "prior theoretical commitment" to independence, it would be wrong for the decison maker to ignore such an epistemically significant feature of the group's opinions.

In this paper, we argue *against* this position. Namely, we claim that when extracting a consensus from the opinions of a group of experts, there is no reason why a decision maker should want or even expect his/her pooling formula to preserve instances of independence. To substantiate our claim, we present arguments which pertain to both the mathematics and the philosophy of probability amalgamation. Having established some definitions and notation in Section 2, we show in Section 3 that if a group's probabilities are spread on at least five alternatives and if the consensus probability of any state of nature, θ, is assumed to be proportional to some function of the probabilities assigned by group members to θ, then independence preservation can be achieved only at the price of ignoring the opinions of all the experts but one. Curiously, this phenomenon does not occur in the restricted case where there are exactly four alternatives, and in that case a complete characterization of independence preserving aggregation methods is given in Section 4.

These results, which are extensions of those in [5], [10] and [18], do not constitute evidence against independence preservation *per se*. Inasmuch as decision makers agree to the reasonableness of our assumptions on the pooling recipe, however, our impossibility theorem explains why one might not wish to restrict one's attention to independence preserving pooling formulas, unless perhaps a maximum of four alternatives are being considered. In Section 5, we use analogies to attempt to explain why consensual distributions should not be expected to preserve independence, even in these exceptional cases. For a recent review of the multiple facets of "expert resolution", the reader is referred to [7], which contains a lengthy annotated bibliography and accompanying discussion.

2. Mathematical formulation

We will assume throughout that the space Θ of underlying states of nature is countable, and that individuals $i = 1, \ldots, n$ have assigned non-zero subjective probabilities $p_i(\theta_j)$ to each possible point θ_j in $\Theta, j \geq 1$. If $\Delta = \{p: \Theta \to (0,1): \sum_{j \geq 1} p(\theta_j) = 1\}$ represents the collection of all probability distributions whose support is Θ, then $p_1, \ldots, p_n \in \Delta$ and the consensus $T(p_1, \ldots, p_n)$ of these n opinions can itself reasonably be taken to be an element of Δ. The processing rule $T: \Delta^n \to \Delta$ which maps a set of n probability distributions into its consensus is called a *pooling operator*, and such an operator is said to *preserve independence* if, whenever events in the algebra generated by Θ are independent for each individual's probability distribution, then these events are also independent for the consensual distribution. Thus if $p_1, \ldots, p_n \in \Delta$ are such that $p_i(A \cap B) = p_i(A)p_i(B)$ for some subsets A and B of Θ, then $T(p_1, \ldots, p_n)$ is required to satisfy

$$T(p_1, \ldots, p_n)(A \cap B) = T(p_1, \ldots, p_n)(A)T(p_1, \ldots, p_n)(B), \tag{2.1}$$

where, in general, the notation $p(E)$ stands for the sum $\sum_{\theta \in E} p(\theta)$ for arbitrary $p \in \Delta$ and $E \subseteq \Theta$.

A first, useful observation is that unless Θ contains at least four points (as in the British pound example given in Section 1), any pooling formula T preserves independence in a trivial way. This is because in this case, events A and B cannot be independent under a distribution $p \in \Delta$ unless one of A or B is either the whole space, Θ, or the empty set. In those instances, however, equation (2.1) is automatically verified. Henceforth, we will assume that $|\Theta| \geq 4$ and we will say that Θ is *quaternary*. In this case, independence preservation is not a universal property. For example, the work in [5], [10] and [18] implies that if there exist functions $F_j: [0,1]^n \to [0,1]$ such that

$$T(p_1, \ldots, p_n)(\theta_j) = F_j[p_1(\theta_j), \ldots, p_n(\theta_j)] \tag{2.2}$$

for all $\theta_j \in \Theta$ and $p_1, \ldots, p_n \in \Delta$, then T cannot obey (2.1) unless it is dictatorial in nature, e.g. $T(p_1, \ldots, p_n) \equiv p_i$ for some $i = 1, \ldots, n$. Condition (2.2) is quite strong, however, as it implies that T is a linear opinion pool, viz.

$$T(p_1, \ldots, p_n) = \sum_{i=1}^{n} w_i p_i + wq \tag{2.3}$$

where $q \in \Delta$ is arbitrary, $w = 1 - \sum_{i=1}^{n} w_i$ and $w_1, \ldots, w_n \in [-1,1]$ must satisfy certain consistency conditions mentioned in [5]. In particular, (2.2) excludes the geometric aggregation method

$$T(p_1, \ldots, p_n) \propto \prod_{i=1}^{n} p_i^{w_i} \tag{2.4}$$

which is easily seen to preserve independence when Θ contains exactly four points. For this reason, it seems natural to conjecture that if (2.2) were sufficiently relaxed in order to include pooling operators such as (2.4), reasonable pooling procedures might emerge which would be independence preserving. One such generalization was suggested in [4] in a different context and consists of simply replacing the equality sign in (2.2) by a proportionality symbol, with the understanding that the consensus $T(p_1, \ldots, p_n)$ at θ_j would then be equal to the quantity $F_j[p_1(\theta_j), \ldots, p_n(\theta_j)]$ except for a constant which may depend on the p_i's but not on θ_j. Thus a pooling formula in this class would be of the form

$$T(p_1, \ldots, p_n)(\theta_j) = F_j[p_1(\theta_j), \ldots, p_n(\theta_j)] / \sum_{k \geq 1} F_k[p_1(\theta_k), \ldots, p_n(\theta_k)]. \tag{2.5}$$

This class has been used with relative success in [6] in connection with Madansky's axiom of external Bayesianity [13], and so it might also accomodate independence preservation in some interesting ways. As we shall see in Section 4, (2.5) does admit a rich variety of non-dictatorial pooling operators when $|\Theta| = 4$. We present a complete characterization of such methods, subject only to the assumption that the functions F_j are Lebesgue measurable. If $|\Theta| \geq 5$, alas, it turns out that dictatorships are again the only pooling operators of the form (2.5) which preserve independence. This constitutes the main result of the following section.

3. The case of $|\Theta| \geq 5$

First, we show that if $|\Theta| \geq 4$, independence preservation restricts the methods of aggregation in (2.5) to those for which all the functions F_j are identical. In particular, any permutation of the columns of the matrix of individual probabilities $(p_i(\theta_j))$ permutes the corresponding consensual probabilities in the same fashion.

THEOREM 3.1. *Let $T: \Delta^n \rightarrow \Delta$ be an independence preserving pooling operator of the form* (2.5), *and suppose that Θ is quaternary. Then the functions F_j of* (2.5) *are identical for all $j \geq 1$.*

Proof. We show, without loss of generality, that $F_1 = F_2$. Fix $\boldsymbol{x}\in(0,1)^n$, $\boldsymbol{y}\in(0,1)^n$, and let $p_1, \ldots, p_n, q_1, \ldots, q_n \in \Delta$ be such that

$$p_i(\theta_1) = q_i(\theta_2) = x_i(1 - y_i);\ p_i(\theta_2) = q_i(\theta_1) = x_iy_i;$$
$$p_i(\theta_3) = q_i(\theta_3) = (1 - x_i)y_i;\ \text{and } p_i(\theta_j) = q_i(\theta_j), j \geq 4$$

and for each $i = 1, \ldots, n$. That such probability distributions exist is immediate from the fact that Θ is quaternary, since $x_i(1 - y_i) + x_iy_i + (1 - x_i)y_i < 1$ for all $i = 1, \ldots, n$.

Let $A = \{\theta_1,\theta_2\}$, $B = \{\theta_2,\theta_3\}$, and $C = \{\theta_1,\theta_3\}$. Then A and B are independent with respect to each of the individual probability distributions $p_1, \ldots, p_n$, while A and C are independent with respect to each of $q_1, \ldots, q_n$. Using (2.5) and independence preservation for the p_i's and the q_i's respectively, we find that

$$F_1[\boldsymbol{x}(\boldsymbol{1} - \boldsymbol{y})]F_3[\boldsymbol{y}(\boldsymbol{1} - \boldsymbol{x})] = \sigma F_2[\boldsymbol{xy}] \tag{3.1}$$

and

$$F_2[\boldsymbol{x}(\boldsymbol{1} - \boldsymbol{y})]F_3[\boldsymbol{y}(\boldsymbol{1} - \boldsymbol{x})] = \sigma F_1[\boldsymbol{xy}], \tag{3.2}$$

where $\sigma = \sum_{j\geq 4} F_j[p_1(\theta_j), \ldots, p_n(\theta_j)]$, $\boldsymbol{1} = (1, \ldots, 1)$ and vector operations are performed coordinate-wise. By (3.1) and (3.2), it follows that

$$F_1[\boldsymbol{x}(\boldsymbol{1} - \boldsymbol{y})]F_1[\boldsymbol{xy}] = F_2[\boldsymbol{x}(\boldsymbol{1} - \boldsymbol{y})]F_2[\boldsymbol{xy}] \tag{3.3}$$

for all $\boldsymbol{x}, \boldsymbol{y} \in (0,1)^n$. Setting $\boldsymbol{y} = (1/2, \ldots, 1/2)$ in (3.3) yields

$$F_1[\boldsymbol{x}/2] = F_2[\boldsymbol{x}/2],$$

that is,

$$F_1[\boldsymbol{x}] = F_2[\boldsymbol{x}] \tag{3.4}$$

for all $\boldsymbol{x}\in(0,1/2)^n$. Setting $\boldsymbol{x} = \boldsymbol{y} = (\sqrt{z_1}, \ldots, \sqrt{z_n}) \equiv \sqrt{\boldsymbol{z}}$ in (3.3) for arbitrary $\boldsymbol{z}\in(0,1)^n$ yields

$$F_1[\sqrt{\boldsymbol{z}} - \boldsymbol{z}]F_1[\boldsymbol{z}] = F_2[\sqrt{\boldsymbol{z}} - \boldsymbol{z}]F_2[\boldsymbol{z}], \tag{3.5}$$

and since $\sqrt{z} - z \in (0,1/4]^n$, it follows from (3.4) and (3.5) that $F_1[z] = F_2[z]$ for all $z \in (0,1)^n$. □

Let F denote the single function from $(0,1)^n$ to $(0,\infty)$ such that

$$T(p_1, \ldots, p_n)(\theta) = F[p_1(\theta), \ldots, p_n(\theta)] / \sum_{\gamma \in \Theta} F[p_1(\gamma), \ldots, p_n(\gamma)] \tag{3.6}$$

for all $p_1, \ldots, p_n \in \Delta$ and for all $\theta \in \Theta$. An implicit restriction on F is that the (potentially infinite) sum on the right-hand side of (3.6) must converge for all choices of $p_1, \ldots, p_n$. If T preserves independence, however, a lot more can be said about the behaviour of F.

LEMMA 3.2 *If an operator T of the form* (3.6) *preserves independence and if Θ contains at least five points, then there exist constants $w_1, \ldots, w_n$ and c such that*

$$F[x] = \sum_{i=1}^{n} w_i x_i + c \tag{3.7}$$

on the entire domain of F.

Proof. We begin by showing that if $x, x^*, y, y^* \in (0,1)^n$ are such that $x + y = x^* + y^* < 1$, then $F[x] + F[y] = F[x^*] + F[y^*]$. To this end, let $z = x + y$ and choose $2n$ probability distributions $p_1, \ldots, p_n, q_1, \ldots, q_n$ in Δ such that

$$p_i(\theta_1) = q_i(\theta_1) = (1 - z_i)/2,\ p_i(\theta_2) = q_i(\theta_2) = (1 - z_i)^2/2(1 + z_i),$$

$$p_i(\theta_3) = x_i,\ p_i(\theta_4) = y_i,\ q_i(\theta_3) = x_i^*,\ q_i(\theta_4) = y_i^*,$$

and

$$p_i(\theta_j) = q_i(\theta_j),\ j \geq 5$$

for each $i = 1, \ldots, n$. Such p_i's and q_i's must exist because $(1 - z_i)^2/2(1 + z_i) < 1$ and $z_i + (1 - z_i)/2 + (1 - z_i)^2/2(1 + z_i) < 1$ for all $1 \leq i \leq n$.

Now consider events $A = \{\theta_1, \theta_2\}$ and $B = \Theta - \{\theta_1, \theta_3, \theta_4\}$. It is easy to check that A and B are independent with respect to all of $p_1, \ldots, p_n$ and $q_1, \ldots, q_n$. Let $u = (1 - z)/2$ and $w = (1 - z)^2/2(1 + z)$. Combining the fact that T preserves independence with equation (3.6), we find that

$$(F[u] + F[w])(F[w] + \sigma) = F[w](F[x] + F[y] + \lambda) \tag{3.8}$$

when we use the p_i's, while using the q_i's leads to

$$(F[u] + F[w])(F[w] + \sigma) = F[w](F[x^*] + F[y^*] + \lambda), \tag{3.9}$$

where $\sigma = \sum_{j\geq 5} F[p_1(\theta_j), \ldots, p_n(\theta_j)] = \sum_{j\geq 5} F[q_1(\theta_j), \ldots, q_n(\theta_j)]$ and $\lambda = F[u] + F[w] + \sigma$. Equating (3.8) and (3.9) and cancelling the common factor of $F[w]$ on both sides, we conclude that

$$F[x] + F[y] = F[x^*] + F[y^*] \tag{3.10}$$

for all x, x^*, y, y^* in $(0,1)^n$ such that $x + y = x^* + y^* < 1$.

Next, we claim that there exists a constant $c \in I\!R$ such that $c = 2F[x/2] - F[x]$ for all $x \in (0,1)^n$. For, if $x, y \in (0,1)^n$ are abritrary, we can let $z = 1 - \max\{x,y\}$ and use (3.10) to get

$$F[x/2] + F[z] = F[(x + z)/2] + F[z/2]$$

and

$$F[x] + F[z/2] = F[(x + z)/2] + F[x/2].$$

It follows that

$$2F[x/2] - F[x] = 2F[z/2] - F[z],$$

and replacing x by y shows that

$$2F[x/2] - F[x] = 2F[y/2] - F[y] = c \tag{3.11}$$

for all $x, y \in (0,1)^n$.

Finally, consider the function $G:(0,1)^n \to (-c,\infty)$ defined by $G[x] = F[x] - c$ for all $x \in (0,1)^n$. By (3.10), it is clear that $G[x] + G[y] = 2G[(x + y)/2]$ for all x, y in $(0,1)^n$ with $x + y < 1$. Moreover, we have $G[x] = 2G[x/2]$ for all $x \in (0,1)^n$, since

$$G[x] = F[x] - (2F[x/2] - F[x]) = 2(F[x] - F[x/2])$$

and

$$G[x/2] = F[x/2] - (2F[x/2] - F[x]) = F[x] - F[x/2]$$

by (3.11). From this, we conclude easily that G satisfies Cauchy's functional equation $G[x] + G[y] = G[x + y]$ on its entire domain. Since $-G$ is bounded above by c and mid-convex, it is also continuous by Theorem C in [17, p. 215]. That the *only* solution is of the form $\sum_{i=1}^{n} w_i x_i$ on all of $(0,1)^n$ is now a direct consequence of Theorem 4 in [15].

□

If Θ is infinite, it is immediate from Lemma 3.2 and equation (3.6) above that the constant c in (3.7) must be zero. For, otherwise the series in the denominator on the right-hand side of (3.6) would diverge and T would be ill-defined. Moreover, it is clear that the weights $w_1, \ldots, w_n$ must then be non-negative, since $T(p_1, \ldots, p_n)(\theta)$ might otherwise be made strictly negative with an appropriate choice of p_i's. In that case, therefore, T must be a linear opinion pool (2.3) with $w = 0$ and arbitrary $q \in \Delta$. Since T preserves independence, it is easy to see that $T(p_1, \ldots, p_n) \equiv p_i$ for some $1 \leq i \leq n$. The following theorem states that the same conclusion holds true for finite Θ's also, as long as $|\Theta| \geq 5$.

THEOREM 3.3. *Let T be an independence preserving pooling operator of the form* (2.5), *and suppose that $|\Theta| \geq 5$. Then T is a dictatorship.*

Proof. We have just argued that this theorem is true when Θ is infinite, so we can now restrict ourselves to the case where $5 \leq |\Theta| < \infty$. In that case, we know from (3.6) and Lemma 3.2 that there exist real constants $w_1, \ldots, w_n$ and c such that $\sigma + c|\Theta| = 1$ and

$$T(p_1, \ldots, p_n) = \sum_{i=1}^{n} w_i p_i + c \tag{3.12}$$

for all choices of $p_1, \ldots, p_n$ in Δ, where $\sigma = \sum_{n=1}^{n} w_i$. In particular, the operator (3.12) preserves the independence of events $A = \{\theta_1, \theta_2\}$ and $B = \{\theta_2, \theta_3\}$ when $p_1 = \ldots = p_n = p$ are such that $p(\theta_1) = p(\theta_2) = p(\theta_3) = 1/4$. This implies that $\sigma/4 + c = 4[\sigma/4 + c]^2$, an equality which is verified only if $\sigma + 4c = 1 = \sigma + c|\Theta|$ because we must have $T(p, \ldots, p)(\theta_2) > 0$ by definition of T. Since $|\Theta| \geq 5$, this implies that $c = 0$, so that T is a linear opinion pool of the form (2.3) with $w = 0$ and arbitrary $q \in \Delta$. To see that T is a dictatorship, we can now proceed as in the proof of Theorem 2 in [10].

□

4. The case of $|\Theta| = 4$

Suppose that Θ can be written as $\{\theta_1, \theta_2, \theta_3, \theta_4\}$, where the θ_i's are mutually distinct.

In that case, it is easy to check that for non-empty, proper subsets A and B of Θ to be independent, it is necessary for both A and B to be of size 2 and for $A \cap B$ to be a singleton. Moreover, all such instances of independence arise from pairs of real numbers α, $\beta \in (0,1)$, where $p(A - B) = (1 - \alpha)\beta$, $p(A \cap B) = \alpha\beta$, $p(B - A) = \alpha(1 - \beta)$, and $p((A \cup B)^c) = (1 - \alpha)(1 - \beta)$. Hence by Theorem 3.1, a pooling operator of the form (2.5) preserves independence if and only if the functions F_j of (2.5) are identical to some function $F:(0,1)^n \to (0,\infty)$ satisfying

$$F[\boldsymbol{x}\boldsymbol{y}]F[(\boldsymbol{1} - \boldsymbol{x})(\boldsymbol{1} - \boldsymbol{y})] = F[\boldsymbol{x}(\boldsymbol{1} - \boldsymbol{y})]F[\boldsymbol{y}(\boldsymbol{1} - \boldsymbol{x})], \tag{4.1}$$

where $\boldsymbol{x}, \boldsymbol{y} \in (0,1)^n$, $\boldsymbol{1} = (1, \ldots, 1)$ and operations on vectors are performed coordinatewise.

Generalizing results of [8], Abou-Zaid [1] has proved that the Lebesgue measurable solutions of the functional equation

$$G[\boldsymbol{x}\boldsymbol{y}] + G[(\boldsymbol{1} - \boldsymbol{x})(\boldsymbol{1} - \boldsymbol{y})] = G[\boldsymbol{x}(\boldsymbol{1} - \boldsymbol{y})] + G[\boldsymbol{y}(\boldsymbol{1} - \boldsymbol{x})] \tag{4.2}$$

are given by

$$G[\boldsymbol{x}] = \sum_{i=1}^{n} a_i(x_i - x_i^2) + b_i \log(x_i) + d \tag{4.3}$$

with arbitrary reals a_i, b_i and d. A shorter proof could also be based on A. Járai's (private communication) independent observation that every measurable solution of (4.1) must be differentiable infinitely many times. Both derivations ultimately lead to the following theorem.

THEOREM 4.1 *Suppose that* $|\Theta| = 4$, *and let* T *be a pooling operator of the form* (2.5) *such that at least one of the functions* F_j *be Lebesgue measurable. Then* T *preserves independence if and only if there exist arbitrary real constants* $a_1, \ldots, a_n$, $b_1, \ldots, b_n$ *such that*

$$T(p_1, \ldots, p_n)(\theta) \propto \prod_{i=1}^{n} [p_i(\theta)]^{b_i} \exp\{a_i p_i(\theta)[1 - p_i(\theta)]\} \tag{4.4}$$

for all $p_1, \ldots, p_n \in \Delta$ *and all* $\theta \in \Theta$.

Proof. By Theorem 3.1, $F_j = F, j = 1, 2, 3, 4$. By the remarks at the beginning of this section, (2.5) and independence preservation are satisfied if and only if (4.1) holds. But (4.1) is equivalent to (4.2) with $G = \log(F)$, and so (4.4) follows from (4.3). □

The pooling recipes arising from (4.4) include dictatorial aggregation ($a_i \equiv 0, b_i = \delta_{d,i}$, Kronecker's delta for fixed $1 \leq d \leq n$), the so-called logarithmic opinion pool ($a_i \equiv 0$), and the method which disregards individual probability assignments entirely and imposes the uniform distribution in all cases ($a_i \equiv 0$, $b_i \equiv 0$).

5. Heuristics against independence preservation

Theorem 3.3 and similar impossibility theorems such as those in [5], [10] and [18] do not provide direct, irrefutable evidence that instances of independence common to a number of individual distributions are unworthy of preservation in a decision maker's synthetic measure of probability. Indeed, it could be argued, especially in view of Theorem 4.1, that independence preservation is a feasible requirement and that theorems such as that presented in Section 3 merely reflect a conflict between (2.1) and regularity conditions such as (2.2) or (2.5). While the point is well taken, the results of this and other papers on this topic [7] tend to indicate that pooling formulas which obey (2.1) will be so complex that they will otherwise lose any practical appeal. To use (4.4) in the very restricted case of a four-point space, for instance, it would be necessary to determine the values of $2n$ constants a_i and b_i whose interpretation remains unclear. In this final section, we would like to offer arguments of a more heuristic nature which, when we pair them with the mathematical facts, lead us to reject independence preservation as an untenable pooling axiom.

At the outset, a question arises regarding what has been called in [9] the "prior theoretical commitment" that an expert may or may not have towards his/her various statements of independence. It has been argued in [10] that in situations where individuals elicit their subjective distributions by assigning a relative likelihood to each of a finite number of alternatives, "the independence of certain compounds of these propositions is largely fortuitous" [10, p. 343]. Thus, for instance, it is frequent for one to assert that the six faces of a die are equally likely to occur without attributing any epistemic significance to the fact that the events "an even number is obtained" ({2, 4, 6}) and "a multiple of three is obtained" ({3, 6}) are independent. Since the occurrence of these two events depends on the same throw of a die, their independence could not possibly derive from any "physical" consideration of independence, but rather it is a simple consequence of the assessor's current *state of information*. As such, we contend that common attributions of independence need not be preserved, because the consensual distribution of the decision maker is typically founded on more information than that contained in any one of the group members' distributions, and this added information may well destroy independence. As an example of such a loss of independence induced by a data or information gathering process, think of the familiar situation involving a Bayesian analysis of samples from a normal distribution. In this context, it is often remarked that if a joint prior distribution is chosen under which the mean μ and the precision τ parametrizing the normal variates are independent, then μ and τ will have a posterior distribution under which they are dependent after just a single observation has been drawn. This argument is sometimes offered to justify the common use of a normal-gamma distribution in this type of analysis, even though μ and τ are dependent under all members of this conjugate family.

Next, assume along with the decision maker that the experts truly "endorse" every instance of independence, either because they interpret them as empirical of physical laws or perhaps because they view the realization of the state of nature as the outcome of a composite experiment for which they have adopted a product model. To be specific, consider the illustration described in Section 1 where a distribution is needed for the future value of the British pound and the potential failure of a piece of equipment. In such a case, the decision maker may well agree with the experts' model of independence, and this could lead him/her to adopt a formula such as (4.4) for computing joint probabilities. Before committing him/herself to independence preservation, however, the decision maker should consider the implications which this will have on his/her model of the experts. Indeed, when he/she chooses a consensual distribution which preserves the independence of two events A and B, the decision maker implies that his/her judgement about the likelihood of B would not be influenced by the knowledge that A has occurred or not, and vice versa. In particular, this means that the occurrence or non-occurrence of A would not change his/her opinion about the relative expertise of the persons consulted. For, if this information changed his/her evaluation of the experts, it could also change his/her probability of the future event B. Thus independence preservation implies considerably *more* for the decision maker than what had been conveyed by the group of experts. Moreover, the consequences that it has on the model of the experts would appear to be unrealistic in most applications.

The point of the last paragraph can be expressed in a different way as follows. When a number of expert opinions are merged into a consensus using a single formula, implicit as well as explicit assumptions are made about the experts and the way in which their probabilities will influence the decision maker. An example of explicit judgement about the experts is the weight that each one of them will get in the consensus. An implicit assumption at the root of any pooling method is that each member in the group possesses some privileged information about the quantity of interest, θ. If we believe that our experts' opinions are correlated with θ, then we should also be able to learn something about the experts from additional information about θ. This the decision maker cannot always do if he/she subscribes to (2.1).

Another reason for rejecting axiom (2.1) stems from the disagreement which can sometimes occur between the decision maker and the experts about which action is best. As was mentioned in the introduction, this can happen with independence preserving pooling operators even when there is complete agreement on utility assignments. In some situations at least, this would seem to be undesirable.

As a final note, we would like to point out the dubious nature of any type of group agreement preservation requirement, whether it involves independence, consensus of probabilities as in the Zero-Preservation Property of [14], or even mutual accord on the best possible course of action. It is conceivable, after all, that a group of individuals agree and yet be wrong, and it is also possible for individuals to hold common beliefs for different, perhaps even contradictory reasons. It should be borne in mind that pooling operators only provide a decision maker with a computational short-cut to the often long-winded, but fully coherent Bayesian method of updating probability distributions by treating the experts' opinions as data. Such analyses have been carried out with relative success in [3], [12] and [19]. In this framework, it would not be uncommon for a decision maker to destroy independence or to assign to an event a larger probability than that which was agreed upon by the experts, simply because the decision maker believes that each person he/she consulted supplies at least some information which was not available to the other members of the group. The later effect, referred to as a "risky shift" in [2], is not as undesirable as its name implies. It is the logical consequence of one's modelling assumptions regarding the amount and overlap of information that the experts possess. If the decision maker elects to bypass this process by using axioms instead, conditions such as (2.2) or (2.5) might be justified on account of simplicity, but care should otherwise be exerted in considering the modelling consequences that these and other axioms may have.

Acknowledgements

We wish to thank for their valuable assistance Carl Sundberg, who first derived (4.3) from (4.2) under stronger conditions of regularity, and Che Tat Ng, who subsequently brought reference [8] to our attention and posed to S. H. S. Abou-Zaid the problem of generalizing the results of that paper. Our gratitude is also extended to Morris H. DeGroot and Mark J. Schervish for discussions which led to some of the ideas expressed in Section 5.

REFERENCES

[1] ABOU-ZAID, S. H. S., *Functional equations and related measurements.* M. Phil. thesis, University of Waterloo, Waterloo, Ontario, 1984.

[2] BORDLEY, R. F., *A Bayesian model of group polarization.* Org. Behav. Hum. Perform. *32* (1983), 262–274.

[3] FRENCH, S., *Updating of belief in the light of someone else's opinion.* J. R. Statist. Soc. Ser. A *143* (1980), 43–48.

[4] GENEST, C., *A conflict between two axioms for combining subjective distributions.* J. R. Statist. Soc. Ser. B *46* (1984), 403–405.

[5] GENEST, C., *Pooling operators with the marginalization property*. Canad. J. Statist. *12* (1984), 153–163.
[6] GENEST, C., MCCONWAY, K. J. and SCHERVISH, M. J., *Characterization of externally Bayesian pooling operators*. Ann. Statist. *14* (1986), 487–501.
[7] GENEST, C. and ZIDEK, J. V., *Combining probability distributions: a critique and an annotated bibliography*. Statist. Sci. *1* (1986), 114–148.
[8] KANNAPPAN, PL. and NG, C. T., *On functional equations and measures of information, II*. J. Appl. Prob. *17* (1980), 271–277.
[9] LADDAGA, R., *Lehrer and the consensus proposal*. Synthese *36* (1977), 473–477.
[10] LEHRER, K. and WAGNER, C. G., *Probability amalgamation and the independence issue: a reply to Laddaga*. Synthese *55* (1983), 339–346.
[11] LEVI, I., *Consensus as shared agreement and outcome of inquiry*. Synthese *62* (1985), 3–11.
[12] LINDLEY, D. V. and SINGPURWALLA, N. D., *Reliability (and fault tree) analysis using expert opinions*. J. Amer. Statist. Assoc. *81* (1986), 87–90.
[13] MADANSKY, A., *Externally Bayesian groups*. Rand memo RM-4141-PR, The Rand Corporation, Santa Monica, CA, 1964.
[14] MCCONWAY, K. J., *Marginalization and linear opinion pools*. J. Amer. Statist. Assoc. *76* (1981), 410–414.
[15] NG, C. T., *Uniqueness theorems for a general class of functional equations*. J. Austr. Math. Soc. *11* (1970), 362–366.
[16] RAIFFA, H., *Decision analysis: introductory lectures on choices under uncertainty*. Addison-Wesley, Reading, MA, 1968.
[17] ROBERTS, A. W. and VARBERG, D. E., *Convex functions*. Academic Press, New York, 1973.
[18] WAGNER, C. G., *Aggregating subjective probabilities: some limitative theorems*. Notre Dame J. Formal Logic *25* (1984), 233–240.
[19] WINKLER, R. L., *Combining probability distributions from dependent information sources*. Manag. Sci. *27* (1981), 479–488.

Department of Statistics and Actuarial Science,
University of Waterloo,
Waterloo, Ontario, Canada
N2L 3G1

Department of Mathematics,
University of Tennessee,
Knoxville, TN 37996–1300
U.S.A.