Independence Preservation in Expert Judgment Synthesis

Carl Wagner
Department of Mathematics
The University of Tennessee

The Choice Group Seminar

Department of Philosophy, Logic and Scientific Method

The London School of Economics and Political Science

October 20, 2009

THE PEER DISAGREEMENT PROBLEM:

You and I recognize each other as epistemic peers, but disagree in our probability assessments. How, if at all, should we revise our original assessments upon discovering this disagreement ? Should we

1. Agree to disagree, sticking with our original assessments (T.Kelly, *et al*) ; or

2. Revise our original assessments in a way that accords those assessments "equal weight" (A. Elga, D. Christensen, *et al*). This has typically been interpreted to mean that we should revise our priors p_1 and p_2 to their arithmetic mean $p := ½ (p_1 + p_2)$.

● But p, so defined, may fail to preserve instances of independence common to p_1 and p_2.

- And theorems of Wagner and Genest establish that, for a broad class of probability pooling methods, only dictatorial pooling is compatible with universal preservation of independence.

- Shogenji (2007) has cited the above as a possible “conundrum” for the equal weight view.

I don’t agree that these theorems have such dire consequences, and I aim in this talk to show how to preserve *epistemically significant cases of independence* in several ways that are true to the spirit of the equal weight view.

Probability Pooling Theory – A Quick Survey

S = a countable set of possible states of the world.

Δ = the set of all probability distributions on S.

$\Delta^n = \{ (p_1,\ldots, p_n) : \text{each } p_i \,\varepsilon\, \Delta \}$.

Each $(p_1,\ldots, p_n) \,\varepsilon\, \Delta^n$ is a *profile* of the distributions assessed by n individuals.

A *pooling operator* is any function $T: \Delta^n \rightarrow \Delta$. Each pooling operator furnishes a method of reconciling the distributions $p_1,\ldots, p_n$, replacing them by the "consensual" distribution $T(p_1,\ldots, p_n)$.

Probability pooling theory has been modeled on axiomatic social choice theory, as conceived by Arrow and Black. Typical pooling axioms have included:

Irrelevance of Alternatives (IA): For each $s \,\varepsilon\, S$

there exists a function $f_s : [0,1]^n \rightarrow [0,1]$ such

that for all $(p_1,\ldots, p_n) \,\varepsilon\, \Delta^n$,

$$T(p_1,\ldots, p_n)(s) = f_s(p_1(s),\ldots, p_n(s)).$$

Remark. It is implicit in IA that for all $(p_1,\ldots, p_n) \varepsilon \Delta^n$, $\sum_s f_s(p_1(s),\ldots, p_n(s)) = 1$,

with no normalization.

Zero Preservation (ZP): For each s ε S and all

$(p_1,\ldots, p_n) \varepsilon \Delta^n$, if $p_1(s) = \ldots = p_n(s) = 0$, then

$T(p_1,\ldots, p_n)(s) = 0$.

Universal Independence Preservation (UIP):

For all $(p_1,\ldots, p_n) \varepsilon \Delta^n$, and for all subsets E and F of S, if $p_i(E \cap F) = p_i(E)\, p_i(F)$ for i=1,…,n then

$T(p_1,\ldots, p_n)(E \cap F) =$

$$T(p_1,\ldots, p_n)(E)\, T(p_1,\ldots, p_n)(F).$$

Theorem 1 (L&W 1983). If $|S| \geq 3$, a pooling operator satisfies IA, ZP, and UIP if and only if it is dictatorial ,i.e., there exists a $d \in \{1,\ldots,n\}$ such that, for all $(p_1,\ldots, p_n) \in \Delta^n$, $T(p_1,\ldots, p_n) = p_d$.

Theorem 2 (Wagner 1984). If $|S| \geq 3$, a pooling operator satisfies IA and UIP if and only if it is dictatorial or imposed (specifically, there exists an s* in S, such that, for all $(p_1,\ldots, p_n) \in \Delta^n$, $T(p_1,\ldots, p_n)(s) = \delta_{s,s^*}$.

Suppose that ZP is deleted and IA is weakened to

Normalized Pooling (NP): For each $s \in S$ there exists a function $f_s : [0,1]^n \rightarrow [0,1]$ such that for all $(p_1,\ldots, p_n) \in \Delta^n$,

$$T(p_1,\ldots, p_n)(s) =$$

$$f_s(p_1(s),\ldots, p_n(s)) \,/\, \sum_s f_s(p_1(s),\ldots, p_n(s)).$$

In what follows, we restrict consideration to probability distributions P that assign a *strictly positive* probability to each $s \in S$.

Case 1. If $|S| = 3$, independence is trivially preserved, since events E and F cannot be independent with respect to a *strictly positive* distribution P unless one of E or F is equal to S or to the empty set.

Case 2. $|S| = 4$

Theorem 3 (Abou-Zaid 1984; Sundberg and Wagner 1987). Suppose that $|S| = 4$, the pooling operator T is restricted to strictly positive distributions on S, and T satisfies NP, with at least one of the functions f_s being Lebesgue measurable. Then T satisfies UIP if and only if there exist real constants $a_1, \ldots, a_n$ and $b_1, \ldots, b_n$ such that $T(p_1, \ldots, p_n)(s)$ is proportional to

$$\prod_{1 \leq i \leq n} [P_i(s)]^{b_i} \exp \{a_i P_i(s)[1 - P_i(s)]\}.$$

Theorem 4 (Genest and Wagner 1987).
If $|S| \geq 5$, a pooling operator satisfies NP and UIP if and only if it is dictatorial.

Why the G-W theorem is no problem for the equal weight view:

1. Requiring that pooling preserve *every single instance of independence* is unwarranted, since there are clearly cases of independence having no epistemic significance (e.g., the independence of the events "fair die comes up even" and "fair die comes up a multiple of 3.")

2. In the epistemic peer problem we are simply trying to reconcile *two particular probability distributions*, not provide a method of reconciling *every conceivable profile of distributions.* Our task is more properly conceived of as an analogue of the (single profile) social choice theory of Bergson and Samuelson than as an analogue of the (multi-profile) theory of Black and Arrow.

3. Arithmetic (and, more generally, quasi-arithmetic) means hardly exhaust the ways of according distributions p_1 and p_2 equal weight. For example,……..

- Jehle (2007) has suggested that a distribution p that is *equidistant* from p_1 and p_2 in the Euclidean (or some other) metric can reasonably be thought of as giving equal weight to p_1 and p_2.

Question: Suppose that events E and F are independent with respect to p_1 and p_2. Does there always exist a probability distribution p equidistant from p_1 and p_2 that preserves this independence?

Answer (Shattuck and Wagner 2008): Yes, as long as $E \cap F$, $E \cap F^c$, and $E^c \cap F$ are all nonempty. This holds for the Euclidean metric

$$d(p_1, p_2) = (\textstyle\sum_s (p_1(s) - p_2(s))^2)^{1/2},$$

as well as any metric inducing a coarser topology than the Euclidean metric. But, …

● The proof is non-constructive if $|S| \geq 4$.

● The common distance from p to p_1 and p_2 may exceed the distance between p_1 and p_2.

● It is doubtful that the result is true for the independence of more than two events, or for more than two epistemic peers.

Constructive Approaches to the Preservation of Independence:

Example 1. Independent Random Variables.

You and I agree that the outcomes of two tosses of a coin are independent, but you think that the probability of heads is ¼ and I think it is ½. So your distribution over S = {hh, ht, th, tt} is (1/16, 3/16, 3/16, 9/16) and mine is (1/4,1/4,1/4,1/4). The fact that the arithmetic mean of these distributions fails to preserve the independence of "h on the 1st toss" and "h on the 2nd toss" is simply a red herring…..

To reconcile our different assessments under equal weighting, we should clearly first reconcile our different estimates of the probability of heads by taking the arithmetic (or some other) mean of ¼ and ½, and then exploit our agreed-upon independence to construct a distribution over S. Agreed-upon independence of random variables X_1, X_2,… can always be preserved in this way if our disagreement is about the *values of the defining parameters* of these random variables.

Example 2. Independent Partitions of S.

A *partition* of S is a set of nonempty, pairwise disjoint subsets of S, with union equal to S.

Partitions **E** and **F** are *independent with respect to the probability distribution* p if, for all E ε **E** and all F ε **F**, $p(E \cap F) = p(E)\, p(F)$.

This definition extends in the obvious way to any finite family **E**, **F**, **G**, etc. of partitions, and probabilistic independence is best conceptualized in terms of partitions…

For example,

● The so-called *total independence* of a family $\{E_1,\ldots, E_n\}$ of events is equivalent to the independence of the partitions

$$\mathbf{E_1} = \{ E_1, E_1^c\},\ldots, \mathbf{E_n} = \{ E_n, E_n^c\}.$$

● The *events* E and F are independent if and only if the partitions $\mathbf{E} = \{E,E^c\}$ and $\mathbf{F} = \{F,F^c\}$ are independent.

Suppose that $\mathbf{E} = \{E,E^c\}$ and $\mathbf{F} = \{F,F^c\}$ are independent with respect to your distribution p_1 as well as my distribution p_2.

Here is a way to construct a distribution q that preserves this independence, and is true to the spirit of the equal weight view:

1. Let $p := \tfrac{1}{2}(p_1 + p_2)$.

2. Let $\mu_{E\cap F} := p(E)\,p(F)$

$$\mu_{E\cap F^c} := p(E)\,p(F^c)$$

$$\mu_{E^c\cap F} := p(E^c)\,p(F)$$

$$\mu_{E^c\cap F^c} := p(E^c)\,p(F^c)$$

3. Revise p to q by Jeffrey conditionalization on the partition $\{ E\cap F, E\cap F^c, E^c\cap F, E^c\cap F^c \}$, with

(i) $q(E\cap F) = \mu_{E\cap F}$

(ii) $q(E\cap F^c) = \mu_{E\cap F^c}$,

(iii) $q(E^c\cap F) = \mu_{E^c\cap F}$, and

(iv) $q(E^c\cap F^c) = \mu_{E^c\cap F^c}$.

It is easy to check that the partitions $\mathbf{E} = \{E, E^c\}$ and $\mathbf{F} = \{F, F^c\}$ are independent with respect to q.

4. The probability distribution q is the closest distribution to the arithmetic mean p of our priors p_1 and p_2 that satisfies (i) – (iv) above (and thus preserves independence of the partitions **E** = {E,E^c} and **F** = {F,F^c}) with respect to the Kullback-Leibler information number (see, e.g., Diaconis Zabell, Updating subjective probability, *Journal of the American Statistical Association* 77(1982), 822-830).

Concluding Remarks:

1. I am not advocating the equal weight solution to the epistemic peer problem, simply defending it against the charge that equal weighting is incompatible with independence preservation.

2. I am not endorsing arithmetic means as the only way to average. My aim is only to show that there are principled ways to implement the equal weight view that preserve epistemically significant cases of independence, not (at least at this stage) to advocate a uniquely rational way to do this.