CrossMark

Probabilistic opinion pooling generalized.
Part two: the premise-based approach

Franz Dietrich[1] · Christian List[2]

Abstract How can several individuals' probability functions on a given σ-algebra of events be aggregated into a collective probability function? Classic approaches to this problem usually require 'event-wise independence': the collective probability for each event should depend only on the individuals' probabilities for that event. In practice, however, some events may be 'basic' and others 'derivative', so that it makes sense first to aggregate the probabilities for the former and then to let these constrain the probabilities for the latter. We formalize this idea by introducing a 'premise-based' approach to probabilistic opinion pooling, and show that, under a variety of assumptions, it leads to linear or neutral opinion pooling on the 'premises'.

1 Introduction

Suppose each individual member of some group (expert panel, court, jury etc.) assigns probabilities to some events. How can these individual probability assignments be

✉ Christian List
C.List@lse.ac.uk

[1] Paris School of Economics and CNRS, Paris, France

[2] London School of Economics, London, UK

🍀 Springer

aggregated into a collective probability assignment? Classically, this problem has been modelled as the aggregation of probability functions, which are defined on some σ-algebra of events, a set of events that is closed under negation and countable disjunction (and thereby also under countable conjunction). Each individual submits a probability function on the given σ-algebra, and these probability functions are then aggregated into a single collective probability function. One of the best-known solutions to this aggregation problem is linear pooling, where the collective probability function is a linear average of the individual probability functions. Linear pooling has several salient properties. First, if all individuals unanimously assign probability 1 (or probability 0) to some event, this probability assignment is preserved collectively ('consensus preservation'). Second, the collective probability for each event depends only on individual probabilities for that event ('event-wise independence'). Third, all events are treated equally: the pattern of dependence between individual and collective probability assignments is the same for all events ('neutrality').

In many practical applications, however, not all events are equal. In particular, the events in a σ-algebra may fall into two categories (whose boundaries may be drawn in different ways). On the one hand, there are events that correspond to intuitively basic propositions, such as 'it will rain', 'it will be humid', or 'atmospheric CO_2 causes global warming'. On the other hand, there are events that are intuitively non-basic. These can be viewed as combinations of basic events, for instance via disjunction (union) of basic events, conjunction (intersection), or negation (complementation). It is not obvious that when we aggregate probabilities, basic and non-basic events should be treated alike.

For a start, we may conceptualize basic and non-basic events differently, in analogy to the distinction between atomic and composite propositions in logic (the latter being logical combinations of the former). Second, the way we assign probabilities to non-basic events is likely to differ from the way we assign probabilities to basic events. When we assign a probability to a conjunction or disjunction, this typically presupposes the assignment of probabilities to the underlying conjuncts or disjuncts. For example, the obvious way to assign a probability to the event 'rain *or* heat' is to ask what the probability of rain is, what the probability of heat is, and whether the two are correlated.[1] If this is right, the natural method of making probabilistic judgments is to consider basic events first and to consider non-basic events next. Basic events serve as 'premises': we first assign probabilities to them, and then let these probability assignments constrain our probability assignments for other, non-basic events.

In this paper, we propose an approach to probability aggregation that captures this idea: the *premise-based approach*. Under this approach, the group first assigns collective probabilities to all basic events (the 'premises') by aggregating the individiduals' probabilities for them; and then it assigns probabilities to all other events, constrained by the probabilities of the basic events. If the basic events are 'rain' and 'heat', then, in a first step, the collective probabilities for these two events are determined by aggregat-

---

[1] The correlation might be due to causal effects between, or common causes of, rain and heat.

ing the individual probabilities for them. In a second step, the collective probabilities for all other events are assigned. For example, the collective probability of 'rain *and* heat' might be defined as a suitable function of the collective probability of 'rain', the collective probability of 'heat', and an estimated rain/heat-correlation coefficient, which could be the result of aggregating the rain/heat-correlation coefficients encoded in the individual probability functions.

This proposal can be expressed more precisely by a single axiom, which does not require the (inessential) sequential implementation just sketched, but focuses on a core informational restriction: the collective probability of any 'premise' (basic event) should depend solely on the individual probabilities for this premise, not on individual probabilities for other events. We call this axiom *independence on premises*. Our axiomatic analysis of premise-based aggregation is inspired by binary judgment-aggregation theory, where the premise-based approach has also been characterized by a restricted independence axiom, for instance by Dietrich (2006), Mongin (2008), and Dietrich and Mongin (2010). For less formal discussions of premise-based aggregation, see Kornhauser and Sager (1986), Pettit (2001), List and Pettit (2002), and List (2006).

The way in which we have just motivated the premise-based approach and the corresponding axiom is bound to prompt some questions. In particular, although the distinction between 'basic' and 'non-basic' events is arguably not ad hoc, there is no purely formal criterion for drawing that distinction.[2] However, there is another, less controversial motivation for the premise-based approach. Our central axiom—independence on premises—privileges particular events, called the 'premises'. We have so far interpreted these in a very specific way, taking them to correspond to basic events and to constitute the premises in an individual's probability-assignment process. But we can give up this interpretation and define a 'premise' simply as an event for which it is desirable that the collective probability depend solely on the event-specific individual probabilities. If 'premises' are defined like this, then our

---

[2] One could construct basic events from non-basic events, using the operations of negation and disjunction. Formally, while the basic events typically form a generating system of the σ-algebra, there exist many alternative generating systems, and usually none of them is canonical in a technical sense. The task of determining the 'basic' events therefore involves some interpretation and may be context-dependent and open to disagreement. One might, however, employ a syntactic criterion which counts an event as 'basic' if, in a suitable language (perhaps one deemed 'natural'), it can be expressed by an atomic sentence (one that is *not* a combination of other sentences using Boolean connectives). An event expressible by the sentence 'it will rain *or* it will snow' would then count as non-basic. This syntactic criterion relies on our choice of language, which, though not a purely technical matter, is arguably not ad hoc. An n-place connective (e.g., the two-place connective 'or') is called *Boolean* or *truth-functional* if the truth-value of every sentence constructed by applying this connective to n other sentences is determined by the truth values of the latter sentences. For instance, 'or' is Boolean since 'p or q' is true if and only if 'p' is true or 'q' is true. Many languages, especially ones that mimic natural language, contain non-Boolean connectives, for instance non-material conditionals for which the truth-value of 'if p then q' is not always determined by the truth-values of p and q. When the sentence 'if p then q' is not *truth-functionally* decomposable, an event represented by it would count as 'basic' under the present syntactic criterion. The sentence 'CO_2 emissions cause global warming' can be viewed as the non-material (specifically, causal) conditional 'if p then q', hence would describe a basic event. See Priest (2001) for an introduction to non-classical logic.

axiom—independence on premises—is justified by definition (though of course we can no longer offer any guidance as to which events should count as premises).[3]

We show that premise-based opinion pooling imposes significant restrictions on how the collective probabilities of the premises can be determined. At the same time, these restrictions are not undesirable; they do not lead to 'undemocratic' or 'degenerate' forms of opinion pooling. Specifically, given certain logical connections between the premises, independence on premises, together with a unanimity-preservation requirement, implies that the collective probability for each premise is a (possibly weighted) linear average of the individual probabilities for that premise, where the vector of weights across different individuals is the same for each premise. We present several variants of this result, which differ in the nature of the unanimity-preservation requirement and in the kinds of connections that are assumed to hold between premises. In some variants, we do not obtain the 'linearity' conclusion, but only a weaker 'neutrality' conclusion: the collective probability for each premise must be a (possibly non-linear) function of the individual probabilities for that premise, where this function is the same for each premise. These results are structurally similar to, but interpretively different from those in our companion paper (Dietrich and List 2017), to which we shall refer as 'Part I'. Furthermore, our results stand in contrast with existing results on the premise-based approach in binary judgment aggregation. When judgments are binary, independence on premises leads to dictatorial aggregation under analogous conditions (see especially Dietrich and Mongin 2010).

Our results apply regardless of which events are deemed to serve as premises. In the extreme case in which *all* events count as premises, the requirement of independence on premises reduces to the familiar event-wise independence axiom (sometimes called the *strong setwise function property*), and our results reduce to a classic characterization of linear pooling (see Aczél and Wagner 1980; McConway 1981; see also Wagner 1982, 1985; Aczél et al. 1984; Genest 1984a; Mongin 1995; Chambers 2007).[4]

2 The framework

We consider a group of $n \geq 2$ individuals, labelled $i = 1, \ldots, n$, who have to assign collective probabilities to some events.

The agenda: a σ-algebra of events We consider a non-empty set Ω of possible *worlds* (or *states*). An *event* is a subset A of Ω; its complement ('negation') is denoted

---

[3] The terminology 'premise' is still justified, though not in the sense of 'premise of *individual* probability assignment' (since we no longer assume that premises are basic in the individuals' formation of probabilistic beliefs), but in the sense of 'premise of *collective* probability assignment' (because the collective probabilities for these events are determined independently of the probabilities of other events and then constrain other collective probabilities).

[4] Historically, linear pooling goes back at least to Stone (1961). Linear pooling is by no means the only plausible way to aggregate subjective probabilities. Other approaches include *geometric* and, more generally, *externally Bayesian* pooling (e.g., McConway 1978; Genest 1984b; Genest et al. 1986; Russell et al. 2015; Dietrich 2016), *multiplicative* pooling (Dietrich 2010; Dietrich and List 2016), *supra-Bayesian* pooling (e.g., Morris 1974), and pooling of *ordinal* probabilities (Weymark 1997). For literature reviews, see Genest and Zidek (1986), Clemen and Winkler (1999) and Dietrich and List (2016).

$A^c := \Omega \backslash A$. The set of events to which probabilities are assigned is called the *agenda*. We assume that it is a σ-algebra, $\Sigma \subseteq 2^\Omega$, i.e., a set of events that is closed under complementation and countable union (and by implication also countable intersection). The simplest non-trivial example of a σ-algebra is of the form $\Sigma = \{A, A^c, \Omega, \varnothing\}$, where $\varnothing \subsetneq A \subsetneq \Omega$. Another example is the set 2^Ω of *all* events; this is a commonly studied σ-algebra when Ω is finite or countably infinite. A third example is the σ-algebra of Borel-measurable sets when $\Omega = \mathbb{R}$.

An example Let us give an example similar to the lead example in Part I, except that we now take the agenda to be a σ-algebra. Let the set Ω of possible worlds be the set of vectors $\{0, 1\}^3 \backslash \{(1, 1, 0)\}$ with the following interpretation. The first component of each vector indicates whether atmospheric CO_2 is above some threshold (1 = 'yes' and 0 = 'no'), the second component indicates whether there is a mechanism to the effect that *if* atmospheric CO_2 is above that threshold, *then* Arctic summers are ice-free, and the third component indicates whether Arctic summers are ice-free. The triple $(1, 1, 0)$ is excluded from Ω because it would represent an inconsistent combination of characteristics. Now the agenda is $\Sigma = 2^\Omega$.

The opinions: probability functions Opinions are represented by probability functions on Σ. Formally, a *probability function* on Σ is a function $P : \Sigma \rightarrow [0, 1]$ such that $P(\Omega) = 1$ and P is σ-additive (i.e., $P(A_1 \cup A_2 \cup \ldots) = P(A_1) + P(A_2) + \cdots$ for every sequence of pairwise disjoint events $A_1, A_2, \ldots \in \Sigma$). We write $\mathcal{P}_\Sigma$ to denote the set of all probability functions on Σ.

Opinion pooling Given the agenda Σ, a combination of probability functions across the individuals, $(P_1, \ldots, P_n)$, is called a *profile (of probability functions)*. An *(opinion) pooling function* is a function $F : \mathcal{P}_\Sigma^n \rightarrow \mathcal{P}_\Sigma$, which assigns to each profile $(P_1, \ldots, P_n)$ a collective probability function $P = F(P_1, \ldots, P_n)$, also denoted $P_{P_1, \ldots, P_n}$. An example of $P_{P_1, \ldots, P_n}$ is the arithmetic average $\frac{1}{n} P_1 + \cdots + \frac{1}{n} P_n$.

Some logical terminology We conclude this section with some further terminology. Events distinct from $\varnothing$ and Ω are called *contingent*. A set S of events is *consistent* if its intersection $\bigcap_{A \in S} A$ is non-empty, and *inconsistent* otherwise; S *entails* an event B if the intersection of S is included in B (i.e., $\bigcap_{A \in S} A \subseteq B$).

3 Axiomatic requirements on premise-based opinion pooling

We now introduce the axioms that we require a premise-based opinion pooling function to satisfy.

3.1 Independence on premises

Before we introduce our new axiom of *independence on premises*, let us recall the familiar requirement of *(event-wise) independence*. It requires that the collective

probability for any event depend only on the individual probabilities for that event, independently of the probabilities of other events.

Independence For each event $A \in \Sigma$, there exists a function $D_A : [0, 1]^n \to [0, 1]$ (the *local pooling criterion* for A) such that, for all $P_1, \ldots, P_n \in \mathcal{P}_\Sigma$,

$$P_{P_1, \ldots, P_n}(A) = D_A(P_1(A), \ldots, P_n(A)).$$

This requirement can be criticized—in the classical framework where the agenda is a σ-algebra—for being normatively unattractive. Typically only some of the events in the σ-algebra Σ correspond to intuitively basic propositions such as 'the economy will grow' or 'atmospheric CO_2 causes global warming'. Other events in Σ are combinations of basic events, such as 'the economy will grow *or* atmospheric CO_2 causes global warming'. The non-basic events can get enormously complicated: they can be conjunctions of (finitely or countably infinitely many) basic events, or disjunctions, or disjunctions of conjunctions, and so on. It seems natural to privilege the basic events over the other, more 'artificial' events by replacing the independence requirement with a restricted independence requirement that quantifies only over basic events. Indeed, it seems implausible to apply independence to composite events such as 'the economy will grow *or* atmospheric CO_2 causes global warming', since this would prevent us from using the probabilities of each of the constituent events in determining the overall probability.

By restricting the independence requirement to basic events, we treat these as *premises* in the collective probability-assignment process, first aggregating individual probabilities for basic events and then letting the resulting collective probabilities constrain the collective probabilities of all other events. (The probabilities of the premises constrain those other probabilities because the probability assignments in their entirety must be probabilistically coherent.)

Formally, consider a sub-agenda of Σ, denoted X, which we interpret as containing the basic events, called the *premises*. By a *sub-agenda* we mean a subset of Σ which is non-empty and closed under complementation (i.e., it forms an 'agenda' in the generalized sense discussed in Part I). We introduce the following axiom:

Independence on X ('on premises'). For each $A \in X$, there exists a function $D_A : [0, 1]^n \to [0, 1]$ (the *local pooling criterion* for A) such that, for all $P_1, \ldots, P_n \in \mathcal{P}_\Sigma$,

$$P_{P_1, \ldots, P_n}(A) = D_A(P_1(A), \ldots, P_n(A)).$$

In the climate-change example of Sect. 2, the sub-agenda of premises might be defined as $X = \{A_1, A_1^c, A_2, A_2^c, A_3, A_3^c\}$, where A_1 is the event that atmospheric CO_2 is above the critical threshold, A_2 is the event that there is a mechanisms by which CO_2 concentrations above the threshold cause ice-free Arctic summers, and A_3 is the event of ice-free Arctic summers. Conjunctions such as $A_1 \cap A_2$ are not included in the set X of premises here. As a result, independence on X allows the collective probability for any such conjunction to depend not only on the experts' probabilities

for that conjunction, but also, for instance, on their probabilities for the underlying conjuncts (together with auxiliary assumptions about correlations between them).[5]

We have explained why event-wise independence should not be required for non-basic events. But why should we require it for basic events (premises)? We offer three reasons:

- First, if we accept the idea that an individual's probabilistic belief about a given premise is not influenced by, but might influence, his or her beliefs about other events, then we may regard those other beliefs as either by-products of, or unrelated to, the individual's belief about the premise in question. It then seems reasonable to treat those other beliefs as irrelevant to the question of what collective probability to assign to that premise. (More precisely, any beliefs about other events provide no relevant *additional* information once the individual's belief about the premise is given.)
- Second, the premise-based approach can be motivated by appealing to the idea of a 'rational collective agent' that forms its probabilistic beliefs by reasoning from premises to conclusions. This kind of collective reasoning can be implemented by first aggregating the probabilities for the premises and then letting these constrain the probabilities assigned to other events. In the case of binary judgment aggregation, Pettit (2001) has described this process as the 'collectivization of reason'.
- Third, as mentioned in the introduction, one might simply *define* the premises as the events for which it is desirable that the collective probabilities depend solely on the event-specific individual probabilities. This would render the requirement of independence on premises justified by definition.

3.2 Consensus preservation on premises

Informally, our second axiomatic requirement says that whenever there is unanimous agreement among the individuals about the probability of certain events, this agreement should be preserved collectively. We distinguish between different versions of this requirement. The most familiar one is the following:

Consensus preservation For all $A \in \Sigma$ and all $P_1, \ldots, P_n \in \mathcal{P}_\Sigma$, if, for all i, $P_i(A) = 1$, then $P_{P_1,\ldots,P_n}(A) = 1$.[6]

A second, less demanding version of the requirement is restricted to events in the sub-agenda X of premises.

Consensus preservation on X ('on premises') For all $A \in X$ and all $P_1, \ldots, P_n \in \mathcal{P}_\Sigma$, if, for all i, $P_i(A) = 1$, then $P_{P_1,\ldots,P_n}(A) = 1$.

---

[5] These auxiliary assumptions might be given exogenously; or they might be determined endogenously based on the experts' probability functions (e.g., based on how dependent or independent the conjuncts are according to these probability functions).

[6] Equivalently, one can demand the preservation of any unanimously assigned probability 0.

Restricting consensus preservation in this way may be plausible because a consensus on any event outside X may be considered less compelling than a consensus on a premise in X, for reasons similar to those for which we restricted event-wise independence to premises. A consensus on a non-basic event could be 'spurious' in the sense that there might not be any agreement on its basis (see Mongin 2005).[7]

We also consider a third version of consensus preservation, which is still restricted to premises, but refers to conditional probabilities. It says that if all individuals assign a conditional probability of 1 to some premise given another, then this should be preserved collectively.[8]

Conditional consensus preservation on X ('on premises') For all $A, B \in X$ and all $P_1, \ldots, P_n \in \mathcal{P}_\Sigma$, if, for all i, $P_i(A|B) = 1$ (provided $P_i(B) \neq 0$), then $P_{P_1,\ldots,P_n}(A|B) = 1$ (provided $P_{P_1,\ldots,P_n}(B) \neq 0$).

Conditional consensus preservation on X is equivalent to another requirement. This says that if all individuals agree that some premise implies another with probabilistic certainty (i.e., the probability of the first event occurring without the second is zero), then that agreement should be preserved collectively.

Implication preservation on X ('on premises') For all events $A, B \in X$ and all $P_1, \ldots, P_n \in \mathcal{P}_\Sigma$, if, for all i, $P_i(A \backslash B) = 0$, then $P_{P_1,\ldots,P_n}(A \backslash B) = 0$.

The equivalence between conditional consensus preservation on X and implication preservation on X follows from the fact that the clause '$P_i(A|B) = 1$ (provided $P_i(B) \neq 0$)' is equivalent to '$P_i(B \backslash A) = 0$', and the clause '$P_{P_1,\ldots,P_n}(A|B) = 1$ (provided $P_{P_1,\ldots,P_n}(B) \neq 0$)' is equivalent to '$P_{P_1,\ldots,P_n}(B \backslash A) = 0$'. Thus the statement of conditional consensus preservation on X can be reduced to that of implication preservation on X (except that the roles of A and B are swapped).

This equivalence also illuminates the relationship between conditional consensus preservation on X and consensus preservation on X, because the former, re-formulated as implication preservation on X, clearly implies the latter. Simply note that, in the statement of implication preservation on X, taking $B = A^c$ yields $P(A \backslash B) = P(A)$, so that a unanimous *zero* probability of any event A in X must be preserved, which is equivalent to consensus preservation on X.

In fact, conditional consensus preservation on X, when re-formulated as implication preservation on X, is also easily seen to be equivalent to a further unanimity-preservation requirement, which refers to unanimous assignments of probability 1 to a *union of two events* in X (just note that $A \backslash B$ has probability 0 if and only if $A^c \cup B$ has probability 1). This also shows that conditional consensus preservation on X is logically weaker than consensus preservation in its original form (on all of Σ), since it does not require preservation of unanimous assignments of probability 1 to

---

[7] In Part I, we make the opposite move of extending consensus preservation to events outside the agenda, i.e., we extend it to events constructible from events in the agenda using conjunction (intersection), disjunction (union), or negation (complementation). In the present paper, there is no point in extending consensus preservation to other events, since there are no events outside the agenda constructible from events in it (as a σ-algebra, the agenda is closed under the relevant operations).

[8] We are indebted to Richard Bradley for suggesting this formulation of the requirement.

intersections of two events in X, or unions or intersections of *more than two* events in X.

The following proposition summarizes the logical relationships between the different consensus-preservation requirements (in part (a)) and adds another simple but useful observation (in part (b)).

Proposition 1 *(a) For any sub-agenda X of Σ, conditional consensus preservation on X*
- *implies consensus preservation on X;*
- *is implied by (global) consensus preservation;*
- *is equivalent to implication preservation on X, and to each of the following two requirements:*

$$[\forall i \ P_i(A \cup B) = 1] \Rightarrow P_{P_1,\ldots,P_n}(A \cup B) = 1, \text{ for all } \ A, B \in X, \ P_1,\ldots, P_n \in \mathcal{P}_\Sigma;$$
$$[\forall i \ P_i(A \cap B) = 0] \Rightarrow P_{P_1,\ldots,P_n}(A \cap B) = 0, \text{ for all } \ A, B \in X, \ P_1,\ldots, P_n \in \mathcal{P}_\Sigma.$$

(b) For the maximal sub-agenda $X = \Sigma$, all of these requirements are equivalent.

4 A class of applications

So far, all our examples of opinion pooling problems have involved events represented by propositions in natural language, such as 'it will rain'. As argued in Part I, the assumption that the agenda is a σ-algebra is often unnatural in such cases. But there is a second class of applications, in which it is more natural to define the agenda as a σ-algebra (Σ) and to restrict the independence requirement to some sub-agenda X. Suppose we wish to estimate the distribution of a real-valued or vector-valued variable, such as rainfall or the number of insurance claims in some period. Here, the set of worlds Ω could be $\mathbb{R}$, $\mathbb{Z}$, $\mathbb{N}$, or $\{0, 1, \ldots, m\}$, or it could be $\mathbb{R}^k$, $\mathbb{Z}^k$, $\mathbb{N}^k$, or $\{0, 1, \ldots, m\}^k$ (for natural numbers m and k). In such cases, the focus on the σ-algebra of events is more realistic. First, we may *need* a full probability distribution on that σ-algebra. Second, individuals may be *able* to come up with such a probability distribution, because, in practice, they can do the following:

- first choose some parametric class of probability functions (e.g., the class of Gaussian distributions if $\Omega = \mathbb{R}$, Poisson distributions if $\Omega = \mathbb{N}$, or binomial distributions if $\Omega = \{0, 1, \ldots, m\}$);
- then estimate the relevant parameter(s) of the distribution (e.g., the mean and standard deviation in the case of a Gaussian distribution).

Because the agenda in this kind of application (e.g., the σ-algebra of Borel sets over $\mathbb{R}$, or the power set of $\mathbb{N}$) contains very complicated events, it would be implausible to require event-wise independent aggregation for all such events. For instance, suppose $\Omega = \mathbb{R}$, and consider the event that a number's distance to the nearest prime exceeds 37. It would seem artificial to determine the collective probability for that event without paying attention to the probabilities of other events. Here, the sub-agenda X on which event-wise independence is plausible is likely to be much smaller than the full σ-algebra Σ.

Let us give a concrete example. Let Σ consist of the Borel-measurable subsets of $\Omega = \mathbb{R}$. A natural sub-agenda of Σ is $X = \cup_{\omega \in \mathbb{R}} \{(-\infty, \omega], (\omega, \infty)\}$. If we require independence on X with a uniform decision criterion $D = D_A$ $(A \in X)$, where $D(t_1, \ldots, t_n) = \frac{1}{n} t_1 + \cdots + \frac{1}{n} t_n$, we obtain a unique pooling function $F : \mathcal{P}_{\Sigma}^n \to \mathcal{P}_{\Sigma}$, because the collective probabilities for X uniquely extend to a probability function on the entire σ-algebra Σ. Alternatively, one might require independence on the smaller sub-agenda $X = \cup_{\omega \in \{-1, +1\}} \{(-\infty, \omega], (\omega, \infty)\}$, still with the same uniform decision criterion D. This under-determines the pooling function $F : \mathcal{P}_{\Sigma}^n \to \mathcal{P}_{\Sigma}$, because probability assignments for X do not uniquely extend to all of Σ. To fill this gap, one might define the collective probability function as the unique *normal* distribution which assigns the specified probabilities to $(-\infty, -1]$ and $(-\infty, +1]$, as determined by the decision criterion D.[9]

Let us summarize how the present kinds of applications differ from the above-mentioned applications involving events represented by natural-language propositions such as 'it will rain' or 'atmospheric CO_2 causes global warming':

1. Ω is a subset of $\mathbb{R}$ or of a higher-dimensional Euclidean space $\mathbb{R}^k$, rather than a set of 'possible worlds' specified by natural-language descriptions;
2. it is often natural to arrive at a probability function by choosing a parametric family of such functions (such as the family of Gaussian distributions) and then specifying the relevant parameter(s), while this approach would seem ad hoc in the other kind of application;
3. in practice, we may be interested in a probability function on the entire σ-algebra (e.g., in order to compute the mean of the distribution and other moments), rather than just in the probabilities of specific events.

5 When is opinion pooling neutral on premises?

We now show that, if there are certain kinds of interconnections among the premises in X, any pooling function satisfying independence on X and consensus preservation in one of the senses introduced must be *neutral* on X. This means that the pattern of dependence between individual and collective probability assignments is the same for all premises. In the next section, we turn to the question of whether our axioms imply *linear* pooling on premises, over and above neutrality.

Formally, a pooling function for agenda Σ is *neutral on* $X (\subseteq \Sigma)$ if there exists some function $D : [0, 1]^n \to [0, 1]$—the *local pooling criterion* for events in X—such that, for every profile $(P_1, \ldots, P_n) \in \mathcal{P}_{\Sigma}^n$, the collective probability of any event A in X is given by

$$P_{P_1, \ldots, P_n}(A) = D(P_1(A), \ldots, P_n(A)).$$

If $X = \Sigma$, neutrality on X reduces to neutrality in the familiar global sense, briefly mentioned in the introduction.

---

[9] For those special profiles of individual probability functions for which the collective probabilities for $(-\infty, -1]$ and $(-\infty, +1]$ coincide or one of them is zero or one, there is no such normal distribution. A different, non-normal extension must then be used.

Our first result uses the strongest consensus-preservation requirement we have introduced, namely 'global' consensus preservation (on all of Σ). Here, we obtain the neutrality conclusion as soon as the sub-agenda of premises satisfies a very mild condition: it is 'non-nested'. We call a sub-agenda X *nested* if it has the form $X = \{A, A^c : A \in X_+\}$ for some set of events X_+ which is linearly ordered by set-inclusion, and *non-nested* otherwise. For instance, $X = \{A, A^c\}$ is nested (take $X_+ := \{A\}$), as is $X = \{A, A^c, A \cap B, (A \cap B)^c\}$ (take $X_+ = \{A, A \cap B\}$). By contrast, $X = \{A, A^c, B, B^c\}$ is non-nested when the events A and B are logically independent. Also, the above-mentioned sub-agenda $X = \{A_1, A_1^c, A_2, A_2^c, A_3, A_3^c\}$ in our climate-change example is non-nested. Further examples are given in Part I.

Theorem 1 *(a) For any non-nested (finite)[10] sub-agenda X of the σ-algebra Σ, every pooling function $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ satisfying independence on X and (global) consensus preservation is neutral on X.*

(b) For any nested sub-agenda X of the σ-algebra Σ (where X is finite and distinct from $\{\varnothing, \Omega\}$), there exists a pooling function $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ satisfying independence on X and (global) consensus preservation but violating neutrality on X.

The possibilities arising for nested X are illustrated by variants of the two pooling functions constructed in Sect. 4, where Σ is the Borel σ-algebra on $\Omega = \mathbb{R}$ and X is one of the nested sub-agendas $\cup_{\omega \in \mathbb{R}}\{(-\infty, \omega], (\omega, \infty)\}$ and $\cup_{\omega \in \{-1, +1\}}\{(-\infty, \omega], (\omega, \infty)\}$. To obtain pooling functions that are not neutral on X, as described in part (b), we must avoid the use of a uniform decision criterion on all elements of X.[11] Theorem 1 continues to hold if we weaken consensus preservation to conditional consensus preservation on premises, as shown next:

Theorem 2 *(a) For any non-nested (finite) sub-agenda X of the σ-algebra Σ, every pooling function $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ satisfying independence on X and conditional consensus preservation on X is neutral on X.*

(b) For any nested sub-agenda X of the σ-algebra Σ (where X is finite and not $\{\varnothing, \Omega\}$), there exists a pooling function $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ satisfying independence on X and conditional consensus preservation on X but violating neutrality on X.

However, if we weaken the consensus-preservation requirement further—namely to consensus preservation on X—then the neutrality conclusion follows only if the events within the sub-agenda X exhibit stronger interconnections. Specifically, the set X must be 'path-connected', as originally defined in binary judgment-aggregation theory (often under the name 'total blockedness'; see Nehring and Puppe 2010). To define path-connectedness formally, we begin with a preliminary notion. Given

---

[10] The finiteness assumption in Theorems 1(a), 1(b), 2(a), 2(b), 3(a), 4(a), 4(b), 5(a), and 6(a) could be replaced by the assumption that the σ-algebra generated by X is Σ (rather than a sub-σ-algebra of Σ). It might be that some of these finiteness assumptions (or their substitutes)—especially in Theorems 1(b), 2(b), and 4(b)—could be dropped.

[11] For example, for every event of the form $A = (-\infty, \omega]$, we might use the decision criterion defined by $D_A(t_1, \ldots, t_n) = (\frac{1}{n}t_1 + \cdots + \frac{1}{n}t_n)^2$, and for every event of the form $A = (\omega, \infty)$, we might use the decision criterion defined by $D_A(t_1, \ldots, t_n) = 1 - (\frac{1}{n}(1 - t_1) + \cdots + \frac{1}{n}(1 - t_n))^2$.

the sub-agenda X, we say that an event $A \in X$ *conditionally entails* another event $B \in X$—written $A \vdash^* B$—if there is a subset $Y \subseteq X$ (possibly empty, but not uncountably infinite) such that $\{A\} \cup Y$ entails B, where, for non-triviality, $Y \cup \{A\}$ and $Y \cup \{B^c\}$ are each consistent. In our climate-change example with sub-agenda $X = \{A_1, A_1^c, A_2, A_2^c, A_3, A_3^c\}$, A_1 conditionally entails A_3 (take $Y = \{A_2\}$), but none of A_1^c, A_2^c, and A_3 conditionally entails any event in X other than itself.

We call the sub-agenda X *path-connected* if any two events $A, B \in X \backslash \{\varnothing, \Omega\}$ can be connected by a path of conditional entailments, i.e., there exist events $A_1, \ldots, A_k \in X$ ($k \geq 1$) such that $A = A_1 \vdash^* A_2 \vdash^* \cdots \vdash^* A_k = B$, and *non-path-connected* otherwise. For example, suppose $X = \{A, A^c, B, B^c, C, C^c\}$, where $\{A, B, C\}$ is a partition of Ω (and $A, B, C \neq \varnothing$). Then X is path-connected. For instance, to see that there is a path from A to B, note that $A \vdash^* C^c$ (take $Y = \varnothing$) and $C^c \vdash^* B$ (take $Y = \{A^c\}$). Many sub-agendas are not path-connected, including all nested sub-agendas X ($\neq \{\varnothing, \Omega\}$) and the sub-agenda $X = \{A_1, A_1^c, A_2, A_2^c, A_3, A_3^c\}$ in the climate-change example.

Theorem 3 *(a) For any path-connected (finite) sub-agenda X of the σ-algebra Σ, every pooling function $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ satisfying independence on X and consensus preservation on X is neutral on X.*
(b) For any non-path-connected (finite) sub-agenda X of the σ-algebra Σ, there exists a pooling function $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ satisfying independence on X and consensus preservation on X but violating neutrality on X.

6 When is opinion pooling linear on premises?

Our next question is whether, and for which sub-agendas X, our requirements on an opinion pooling function imply *linearity* on premises, over and above neutrality. Formally, a pooling function for agenda Σ is called *linear on X* ($\subseteq \Sigma$) if there exist real-valued weights $w_1, \ldots, w_n \geq 0$ with $w_1 + \cdots + w_n = 1$ such that, for every profile $(P_1, \ldots, P_n) \in \mathcal{P}_\Sigma^n$, the collective probability of any event A in X is given by

$$P_{P_1, \ldots, P_n}(A) = \sum_{i=1}^n w_i P_i(A).$$

If $X = \Sigma$, linearity on X reduces to linearity in the global sense, familiar from the established literature.

As in the case of neutrality, whether our axioms imply linearity on a given sub-agenda X depends on how the events in X are connected and which consensus-preservation requirement we impose on the pooling function. Again, our first result uses the strongest consensus-preservation requirement and applies to a very large class of sub-agendas.

Theorem 4 *(a) For any non-nested (finite) sub-agenda X of the σ-algebra Σ with $|X \backslash \{\Omega, \varnothing\}| > 4$, every pooling function $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ satisfying independence on X and (global) consensus preservation is linear on X.*

(b) *For any other sub-agenda X of the σ-algebra Σ (where X is finite and distinct from* $\{\varnothing, \Omega\}$*), there exists a pooling function* $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ *satisfying independence on X and (global) consensus preservation but violating linearity on X.*

If we weaken consensus preservation to conditional consensus preservation on X, the linearity conclusion still follows, but only if the sub-agenda X is 'non-simple'—a condition stronger than non-nestedness, but still weaker than path-connectedness.[12] The notion of non-simplicity also comes from binary judgment-aggregation theory, where the non-simple agendas are those that are susceptible to majority inconsistencies, the judgment-aggregation analogues of Condorcet's paradox (e.g., Nehring and Puppe 2010; Dietrich and List 2007). Formally, a sub-agenda X is *non-simple* if it has a minimal inconsistent subset $Y \subseteq X$ of more than two (but not uncountably many) events, and *simple* otherwise. (A set Y is *minimal inconsistent* if it is inconsistent but all its proper subsets are consistent.) For example, the sub-agenda $X = \{A_1, A_1^c, A_2, A_2^c, A_3, A_3^c\}$ in our climate-change example is non-simple, since its three-element subset $Y = \{A_1, A_2, A_3^c\}$ is minimal inconsistent. By contrast, a sub-agenda of the form $X = \{A, A^c\}$ is simple.

Theorem 5 *(a) For any non-simple (finite) sub-agenda X of the σ-algebra Σ, every pooling function* $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ *satisfying independence on X and conditional consensus preservation on X is linear on X.*
(b) For any simple sub-agenda X of the σ-algebra Σ (where X is finite and distinct from $\{\varnothing, \Omega\}$*), there exists a pooling function* $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ *satisfying independence on X and conditional consensus preservation on X but violating linearity on X.*

Finally, if we impose only the weakest of our three consensus-preservation requirements—consensus preservation on X—then the linearity conclusion follows only if the sub-agenda X is path-connected and satisfies an additional condition. A sufficient such condition is 'partitionality'. A sub-agenda X is *partitional* if some subset $Y \subseteq X$ partitions Ω into at least three non-empty events (where Y is finite or countably infinite), and *non-partitional* otherwise. As an illustration, recall our earlier example of a sub-agenda given by $X = \{A, A^c, B, B^c, C, C^c\}$, where $\{A, B, C\}$ partitions Ω (with $A, B, C \neq \varnothing$). This sub-agenda is both path-connected (as mentioned above) and partitional.

Theorem 6 *(a) For any path-connected and partitional (finite) sub-agenda X of the σ-algebra Σ, every pooling function* $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ *satisfying independence on X and consensus preservation on X is linear on X.*
(b) For any non-pathconnected (finite) sub-agenda X of the σ-algebra Σ, there exists a pooling function $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ *satisfying independence on X and consensus preservation on X but violating linearity on X.*

It is clear from part (b) that path-connectedness of the premises is necessary for the linearity conclusion to follow. The other condition, partitionality, is not necessary. But it is not redundant:

---

[12] To be precise, path-connectedness implies non-simplicity as long as $X \neq \{\varnothing, \Omega\}$.

Proposition 2 *For some path-connected and non-partitional (finite) sub-agenda X of the σ-algebra Σ, there exists a pooling function $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ satisfying independence on X (even neutrality on X) and consensus preservation on X but violating linearity on X.*[13]

7 Classic results as special cases

As should be evident, if we apply our results to the maximal sub-agenda $X = \Sigma$, we obtain classic results (by Aczél and Wagner 1980; McConway 1981) as special cases. To see why this is the case, note three things. First, when $X = \Sigma$, our various conditions on the sub-agenda X all reduce to a single condition on the size of the σ-algebra Σ.

Lemma 1 *For the maximal sub-agenda $X = \Sigma$ (where $\Sigma \neq \{\Omega, \varnothing\}$), the conditions of non-nestedness, non-simplicity, path-connectedness, and partitionality are all equivalent, and they all hold if and only if $|\Sigma| > 4$, i.e., if and only if Σ is not of the form $\{A, A^c, \Omega, \varnothing\}$.*

Second, when $X = \Sigma$, independence, neutrality, and linearity on X reduce to independence, neutrality, and linearity in the familiar 'global' sense, as already noted. Third, our various consensus-preservation requirements all become equivalent, by Proposition 1.

In consequence, our six theorems reduce to two classic results:[14]

- Theorems 1–3 reduce to the result that all pooling functions satisfying independence and consensus preservation are neutral if $|\Sigma| > 4$, but not if $|\Sigma| = 4$;
- Theorems 4–6 reduce to the result that all pooling functions satisfying independence and consensus preservation are linear if $|\Sigma| > 4$, but not if $|\Sigma| = 4$.

The case $|\Sigma| < 4$ is uninteresting because it means that Σ is the trivial σ-algebra $\{\Omega, \varnothing\}$. Let us slightly re-formulate these two results:

Corollary 1 *For the σ-algebra Σ,*

(a) *if $|\Sigma| > 4$, every pooling function $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ satisfying independence and consensus preservation is linear (and by implication neutral);*
(b) *if $|\Sigma| = 4$, there exists a pooling function $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ satisfying independence and consensus preservation but violating neutrality (and thereby also violating linearity).*

---

[13] In this proposition, we assume that the agenda Σ is not very small, i.e., contains more than $2^3 = 8$ events (e.g., $\Sigma = 2^\Omega$ with $|\Omega| > 3$). Note that, as Σ is a σ-algebra, it has the size 2^k for some $k \in \{1, 2, 3, \ldots\}$ or is infinite.

[14] We require no restriction to a finite Σ, as observed in footnote 10.

Appendix A: Proofs

We now give all proofs. In Sect. A.1, we prove parts (a) of all our theorems by reducing them to results in Part I. In Sects. A.2, A.3, A.4, and A.5, we prove parts (b) of Theorems 1, 3, 4, and 5. Parts (b) of Theorems 2 and 6 require no separate proofs: Theorem 2(b) follows from Theorem 1(b) (since consensus preservation implies conditional consensus preservation on X by Proposition 1) and Theorem 6(b) follows from Theorem 3(b) (since non-neutrality on X implies non-linearity on X). In Sect. A.6, we prove Proposition 2.

A.1 Proof of part (a) of each theorem

We now prove Theorems 1(a) to 6(a). To do so, we first relate premise-based opinion pooling to opinion pooling on a general agenda as introduced in Part I. We begin by generalizing the present paper's framework to agendas that need not be σ-algebras. In general, an *agenda* is a non-empty set X of events (each of which is of the form $A \subseteq \Omega$), where X is closed under complementation (i.e., $A \in X \Leftrightarrow A^c \in X$). It contains the events on which opinions are formed. Given an agenda X, an *opinion function* is a function $P : X \to [0, 1]$ which is coherent, i.e., extendable to a probability function on the σ-algebra $\sigma(X)$ generated by X (i.e., the smallest σ-algebra which includes X, constructible by closing X under countable unions and complements). Let $\mathcal{P}_X$ be the set of all opinion functions for agenda X. If X is a σ-algebra, $\mathcal{P}_X$ consists of all probability functions on X, in line with the notation used above. An *opinion pooling function* for agenda X is a function $\mathcal{P}_X^n \to \mathcal{P}_X$ which assigns to each profile $(P_1, \ldots, P_n)$ of individual opinion functions a collective opinion function, usually denoted $P_{P_1, \ldots, P_n}$. We call the pooling function *linear* and *neutral*, respectively, if it is linear and neutral on X in line with the definition above.

Crucially, a pooling function for a σ-algebra Σ induces new pooling functions for any sub-agendas X on which it is independent. Formally, a pooling function $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ for agenda Σ is said to *induce* the pooling function $F' : \mathcal{P}_X^n \to \mathcal{P}_X$ for (sub-)agenda X if F and F' generate the same collective opinions within X, i.e.,

$$F'(P_1|_X, \ldots, P_n|_X) = F(P_1, \ldots, P_n)|_X \text{ for all } P_1, \ldots, P_n \in \mathcal{P}_\Sigma$$

(and if, in addition, $\mathcal{P}_X = \{P|_X : P \in \mathcal{P}_\Sigma\}$, where this addition holds automatically whenever X is finite or $\sigma(X) = \Sigma$).[15] Our axiomatic requirements on a pooling function for agenda Σ—i.e., independence on a sub-agenda X and various consensus requirements—should be compared with the following requirements on a pooling function for the agenda X (introduced and discussed in Part I). The first two requirements are unrestricted versions of independence and consensus preservation:

---

[15] In this case, each opinion function in $\mathcal{P}_X$ is extendable not just to a probability function on $\sigma(X)$, but also to one on Σ. Probability theorists will be aware that the extendability of a probability function to a larger σ-algebra cannot be taken for granted.

Independence For each event $A \in X$, there exists a function $D_A : [0, 1]^n \rightarrow [0, 1]$ (the *local pooling criterion* for A) such that $P_{P_1,\ldots,P_n}(A) = D_A(P_1(A), \ldots, P_n(A))$ for any $P_1, \ldots, P_n \in \mathcal{P}_X$.

Consensus preservation For all $A \in X$ and $P_1, \ldots, P_n \in \mathcal{P}_X$, if $P_i(A) = 1$ for all individuals i, then $P_{P_1,\ldots,P_n}(A) = 1$.

Note the following criterion for the existence of induced pooling functions:

Lemma 2 (cf. Part I, Lemma 14) *If a pooling function for a σ-algebra Σ is independent on a sub-agenda X (where X is finite or $\sigma(X) = \Sigma$), then it induces a pooling function for agenda X.*

The next two axiomatic requirements are two different extensions of consensus preservation, namely to either *implicitly revealed* or *unrevealed* beliefs. An individual i's *explicitly revealed* beliefs are given by the individual's submitted opinion function P_i. Her *implicitly revealed* beliefs are given by the probabilities of events in $\sigma(X)\backslash X$ which are implied by her explicitly revealed beliefs, i.e., hold under every extension of P_i to a probability function on $\sigma(X)$. If, for instance, P_i assigns probability 1 to $A \in X$, then the agent implicitly reveals certainty of all events $B \supseteq A$ in $\sigma(X)\backslash X$. The following axiom extends consensus preservation to implicitly revealed beliefs:

Implicit consensus preservation   For all $A \in \sigma(X)$ and all $P_1, \ldots, P_n \in \mathcal{P}_X$, if each P_i implies certainty of A (i.e., $\overline{P}_i(A) = 1$ for every extension $\overline{P}_i$ of P_i to a probability function on $\sigma(X)$), then so does $P_{P_1,\ldots,P_n}$.

By contrast, individual i's *unrevealed* beliefs are any probabilistic beliefs which she privately holds relative to events in $\sigma(X)\backslash X$ and which cannot be inferred from the submitted opinion function P_i because different extensions of P_i assign different probabilities to the events in question. The following axiom requires the collective opinion function to be compatible with any unanimously held certainty of an event— including any unrevealed certainty, which is not implied by the submitted opinion functions but is consistent with them. This ensures that no consensus (not even an unrevealed consensus) is ever overruled.

Consensus compatibility   For all $A \in \sigma(X)$ and all $P_1, \ldots, P_n \in \mathcal{P}_X$, if each P_i is consistent with certainty of A (i.e., $\overline{P}_i(A) = 1$ for *some* extension $\overline{P}_i$ of P_i to a probability function on $\sigma(X)$), then so is $P_{P_1,\ldots,P_n}$.

A final requirement pertains to *conditional* beliefs. Note that, based on individual i's opinion function P_i, the conditional belief $P_i(A|B) = P_i(A \cap B)/P_i(B)$ of one agenda event A given another B (where $P_i(B) \neq 0$) may be undefined, since we may have $A \cap B \notin X$ so that $P_i(A \cap B)$ is undefined. Hence, if the agent happens to be privately certain of A given B, then this conditional certainty may be unrevealed. Our axiom of *conditional consensus compatibility* requires that any (possibly unrevealed) unanimous conditional certainty should not be overruled. In fact, we require something subtly stronger: any *set* of (possibly unrevealed) unanimous conditional certainties should not be overruled (see Part I for details).

Conditional consensus compatibility　　For all $P_1, \ldots, P_n \in \mathcal{P}_X$, and all finite sets S of pairs (A, B) of events in X, if every opinion function P_i is consistent with certainty of A given B for all (A, B) in S (i.e., some extension $\overline{P}_i$ of P_i to a probability function on $\sigma(X)$ satisfies $\overline{P}_i(A|B) = 1$ for all pairs $(A, B) \in S$ such that $P_i(B) \neq 0$), then so is the collective opinion function $P_{P_1, \ldots, P_n}$.

The following lemma shows how properties of a pooling function for a σ-algebra translate into corresponding properties of an induced pooling function for a sub-agenda:

Lemma 3 (cf. Part I, Lemma 12) *Suppose pooling function F for σ-algebra Σ induces pooling function F' for sub-agenda X (where X is finite or $\sigma(X) = \Sigma$). Then:*

- *F' is independent (respectively neutral, linear) if and only if F is independent (respectively neutral, linear) on X,*
- *F' is consensus-preserving if and only if F is consensus-preserving on X,*
- *F' is consensus-compatible if F is consensus-preserving,*
- *F' is conditional-consensus-compatible if F is conditional-consensus-preserving on X.*

This lemma follows from a more general result:

Lemma 4 (cf. Part I, Lemma 13) *Consider a σ-algebra Σ and a sub-agenda X (where X is finite or $\sigma(X) = \Sigma$). Any pooling function for X is*

(a) *induced by some pooling function for agenda Σ,*
(b) *independent (respectively neutral, linear) if and only if every inducing pooling function for agenda Σ is independent (respectively neutral, linear) on X, where 'every' can be replaced by 'some',*
(c) *consensus-preserving if and only if every inducing pooling function for agenda Σ is consensus-preserving on X, where 'every' can be replaced by 'some',*
(d) *consensus-compatible if and only if* some *inducing pooling function for agenda Σ is consensus-preserving,*
(e) *conditional-consensus-compatible if and only if* some *inducing pooling function for agenda Σ is conditional-consensus-preserving on X*

　　(where in (d) and (e) the 'only if' claim assumes that X is finite).

Proof of parts (a) of Theorems 1–6　　Using the above translation machinery, one can reduce Theorem 1(a) to Part I's Theorem 1(a), Theorem 2(a) to Part I's Theorem 2(a), and so on until Theorem 6(a). Since the reduction is analogous for each theorem, we only spell it out for Theorem 1. Let X be a non-nested finite sub-agenda of the σ-algebra agenda Σ, and let $F : \mathcal{P}_\Sigma^n \to \mathcal{P}_\Sigma$ be independent on X and (globally) consensus preserving. By Lemma 2, F induces a pooling function F' for agenda X, which is independent and consensus-compatible by Lemma 3, hence neutral by Part I's Theorem 1(a). So F is neutral on X by Lemma 3.　　□

A.2 Proof of Theorem 1(b)

We now write **1** and **0** for the n-dimensional vectors $(1, \ldots, 1)$ and $(0, \ldots, 0)$, respectively. We draw on a measure-theoretic fact:

Lemma 5 (cf. Part I, Lemma 15) *Every probability function on a finite sub-σ-algebra of σ-algebra Σ can be extended to a probability function on Σ.*

Proof of Theorem 1(b) Consider a finite nested sub-agenda $X \neq \{\varnothing, \Omega\}$ of the σ-algebra agenda Σ. (As will become clear, finiteness could be replaced by the assumption that $\sigma(X) = \Sigma$. Under this alternative assumption, the 'Claim' below can be skipped, and the rest of the proof remains almost unaffected.) We construct a pooling function $(P_1, \ldots, P_n) \mapsto P_{P_1, \ldots, P_n}$ for agenda Σ with all relevant properties. Without loss of generality, let $\varnothing, \Omega \in X$.

Claim. If Theorem 1(b) holds in the case that $\sigma(X) = \Sigma$, then it holds in general.

Let Theorem 1(b) hold in the special case. Let $\Sigma' := \sigma(X) (\subseteq \Sigma)$. By assumption, there is a pooling function $F' : \mathcal{P}_{\Sigma'}^n \to \mathcal{P}_{\Sigma'}$ with all relevant properties. Let $\mathcal{A}$ be the set of atoms of the (finite) σ-algebra Σ'. We define $F : \mathcal{P}_{\Sigma}^n \to \mathcal{P}_{\Sigma}$ as follows. Consider $P_1, \ldots, P_n \in \mathcal{P}_{\Sigma}$. Let $P' := F'(P_1|_{\Sigma'}, \ldots, P_n|_{\Sigma'})$. For all $A \in \mathcal{A}$ such that $P'(A) \neq 0$, there is an individual i_A such that $P_{i_A}(A) \neq 0$, since otherwise everyone assigns probability one to $\Omega \backslash A$ while $P'(\Omega \backslash A) \neq 1$, violating consensus-preservation. By Lemma 5, P' can be extended to a probability function P on Σ. As is clear from that lemma's proof (in Part I), we may assume without loss of generality that[16]

$$P(\cdot|A) = P_{i_A}(\cdot|A) \quad \text{for each} \quad A \in \mathcal{A} \quad \text{such that} \quad P(A) \neq 0.$$

Now let $F(P_1, \ldots, P_n)$ be this P. It remains to show that the pooling function F just defined inherits all relevant properties from F'. This is clear for independence on X and non-neutrality on X. To show that F is (globally) consensus-preserving, consider $B \in \Sigma$ and $P_1, \ldots, P_n \in \mathcal{P}_{\Sigma}$ such that $P_1(B) = \cdots = P_n(B) = 1$. To show that $P(B) = 1$, where $P := F(P_1, \ldots, P_n)$, note first that $P(B) = \sum_{A \in \mathcal{A}:P(A) \neq 0} P(B|A)P(A)$. Here (in the notation above) each $P(B|A)$ equals $P_{i_A}(B|A)$, which equals 1 as $P_{i_A}(B) = 1$. So $P(B) = \sum_{A \in \mathcal{A}:P(A) \neq 0} P(A) = 1$. This proves the claim.

Now let $\sigma(X) = \Sigma$, drawing on the above 'Claim'. As X is nested, we may express it as $X = \{A, A^c : A \in X_+\}$ for some subset $X_+ \subseteq X$ which is linearly ordered and contains both $\varnothing$ and Ω.

As an ingredient of our construction, we consider any pooling function for agenda Σ which is neutral (at least) on X and consensus-preserving and whose pooling criterion on X, denoted $D : [0, 1]^n \to [0, 1]$, is at least weakly increasing in each argument. (For instance, we might use dictatorship by individual 1, given by $(P_1, \ldots, P_n) \mapsto P_1$, with pooling criterion given by $D(t_1, \ldots, t_n) = t_1$.) As $X \neq \{\varnothing, \Omega\}$, there is some $A \in X \backslash \{\Omega, \varnothing\}$. As $A \neq \Omega, \varnothing$, there are $P_1, \ldots, P_n \in \mathcal{P}_{\Sigma}$ which all assign probability $1/2$ to A (hence to A^c), so that the collective probabilities of A and of A^c are each given by $D(1/2, \ldots, 1/2)$. As these probabilities sum to 1, it follows that

$$D(1/2, 1/2, \ldots, 1/2) = 1/2. \tag{1}$$

---

[16] In that proof it suffices to choose the Q_As appropriately, since each Q_A equals $P(\cdot|A)$, provided $P(A) \neq 0$.

We now transform this pooling function, which is neutral on X, into a pooling function $(P_1, \ldots, P_n) \mapsto P_{P_1, \ldots, P_n}$ which is non-neutral on X, but still independent on X and consensus-preserving. To do so, we consider a function $T : [0, 1] \to [0, 1]$ such that (i) $T(1/2) \neq 1/2$, (ii) $T(0) = 0$ and $T(1) = 1$, (iii) T is at least weakly increasing, and (iv) T is Lipschitz continuous, i.e., there is a $K > 0$ such that $|T(x) - T(y)| \leq K |x - y|$ for all $x, y \in [0, 1]$. (T could be defined by $T(x) = \min\{2x, 1\}$.) Now consider any $P_1, \ldots, P_n \in \mathcal{P}_\Sigma$. We have to define $P_{P_1, \ldots, P_n}$. We write P for the result of applying the neutral pooling function to $(P_1, \ldots, P_n)$. To anticipate, our definition will imply that

$$P_{P_1, \ldots, P_n}(C) = T(P(C)) \quad \text{whenever} \quad C \in X_+.$$

As a first step towards our definition, we define $P_{P_1, \ldots, P_n}$ on the subdomain

$$\widetilde{X} := \{A \cap B : A, B \in X\} = \{B \backslash A : A, B \in X_+ \quad \text{such that} \quad A \subseteq B\}.$$

The restriction of $P_{P_1, \ldots, P_n}$ to $\widetilde{X}$, to be denoted g, is defined as follows. Each $C \in \widetilde{X}$ is uniquely representable as $C = B \backslash A$ with $A, B \in X_+$ and $A \subseteq B$ (and $A = B = \varnothing$ if $C = \varnothing$), and we let

$$\begin{aligned} g(C) &= T(P(B)) - T(P(A)) \\ &= T(D(P_1(B), \ldots, P_n(B))) - T(D(P_1(A), \ldots, P_n(A))). \end{aligned}$$

It follows that

$$g(C) = \begin{cases} T(P(C)) = T(D(P_1(C), \ldots, P_n(C))) & \text{if } C \in X_+ \\ 1 - T(P(C^c)) = 1 - T(D(P_1(C^c), \ldots, P_n(C^c))) & \text{if } C \in X \backslash X_+, \end{cases} \tag{2}$$

because, firstly, each $C \in X_+$ can be written as $C \backslash \varnothing$ where $C, \varnothing \in X_+$, and, secondly, each $C \in X \backslash X_+$ can be written as $\Omega \backslash C^c$ where $\Omega, C^c \in X_+$ and where $T(P(\Omega)) = T(1) = 1$.

Note that $\widetilde{X}$ is a *semi-ring* on Ω, since (i) $\varnothing \in \widetilde{X}$, (ii) $C, C' \in \widetilde{X} \Rightarrow C \cap C' \in \widetilde{X}$, and (iii) for all $C, C' \in \widetilde{X}$, the difference $C \backslash C'$ is a union of finitely many—in fact, at most *two*—events in $\widetilde{X}$. We next show that the function g on this semi-ring is σ-additive. First, g is finitely additive, i.e., for all disjoint $C_1, C_2 \in \widetilde{X}$, if $C_1 \cup C_2 \in \widetilde{X}$, then $g(C_1 \cup C_2) = g(C_1) + g(C_2)$, by definition of g and additivity of P. To show σ-additivity, consider pairwise disjoint $C_1, C_2, \ldots \in \widetilde{X}$ such that $\cup_{m=1}^\infty C_m \in \widetilde{X}$. We have to show that

$$\delta_M := g(\cup_{m=1}^\infty C_m) - \sum_{m=1}^M g(C_m) \to 0 \quad \text{as} \quad M \to \infty.$$

For all $M \in \{1, 2, \ldots\}$, note that the difference $(\cup_{m=1}^\infty C_m) \backslash (\cup_{m=1}^M C_m) = \cup_{m=M+1}^\infty C_m$ need not belong to $\widetilde{X}$, but can be partitioned into a finite set $\mathcal{C}^M$ of

events in $\widetilde{X}$ (as $\cup_{m=1}^{\infty} C_m$ belongs the the semi-ring $\widetilde{X}$). So, $\mathcal{C}^M \cup \{C_1, \ldots, C_M\}$ partitions $\cup_{m=1}^{\infty} C_m$. Careful inspection of g's definition reveals that $\delta_M = \sum_{C \in \mathcal{C}^M} g(C)$. So, as $g(C) \leq K P(C)$ for each $C \in \widetilde{X}$ (by definition of g and property (iv) of T), $\delta_M \leq K \sum_{C \in \mathcal{C}^M} P(C) = K P(\cup_{m=M+1}^{\infty} C_m)$. As $M \rightarrow \infty$ we have $P(\cup_{m=M+1}^{\infty} C_m) \rightarrow 0$ (by σ-additivity of P), and so $\delta_M \rightarrow 0$, as required.

As g is non-negative, σ-additive, and also σ-finite (i.e., Ω is a union of countably many events in $\widetilde{X}$ of finite g-measure, which trivially holds as $\Omega \in \widetilde{X}$), Caratheodory's Extension Theorem tells us that g extends uniquely to a measure on $\sigma(\widetilde{X}) = \sigma(X) = \Sigma$. Let $P_{P_1, \ldots, P_n}$ be this extension. $P_{P_1, \ldots, P_n}$ is indeed a *probability function* since $P_{P_1, \ldots, P_n}(\Omega) = 1$ as $\Omega \in \widetilde{X}$ and $g(\Omega) = T(1) = 1$.

Finally, we must prove that the pooling function $(P_1, \ldots, P_n) \mapsto P_{P_1, \ldots, P_n}$, as just defined, is independent on X, (globally) consensus-preserving, and non-neutral on X.

Independence on X This holds because, for all $P_1, \ldots, P_n \in \mathcal{P}_\Sigma$, the function $P_{P_1, \ldots, P_n}$ extends g, which satisfies (2). Note that the pooling criterion D_C for $C \in X_+$ is defined as $T \circ D$, while the pooling criterion D_C for $C \in C \backslash X_+$ is defined by $\mathbf{t} \mapsto 1 - T \circ D(\mathbf{1} - \mathbf{t})$.

Non-neutrality on X Here it suffices to show that, for some $C \in X \backslash \{\Omega, \varnothing\}$, the pooling criteria D_C and D_{C^c} differ. This follows from the following argument. First, $X \backslash \{\Omega, \varnothing\} \neq \varnothing$ as $X \neq \{\varnothing, \Omega\}$. So we may pick $C, C^c \in X \backslash \{\Omega, \varnothing\}$; say, assume $C \in X_+$ and $C^c \in X \backslash X_+$. So, as just shown, $D_C = T \circ D$ and $D_{C^c} = 1 - T \circ D(\mathbf{1} - \cdot)$. Hence $D_C \neq D_{C^c}$, since $D_C(1/2, \ldots, 1/2) \neq D_{C^c}(1/2, \ldots, 1/2)$, as is clear from the fact that $T(1/2) \neq 1/2$ and that

$$D_{A_j}(1/2, \ldots, 1/2) = T \circ D(1/2, \ldots, 1/2) = T(1/2),$$
$$D_{A_j^c}(1/2, \ldots, 1/2) = 1 - T \circ D(1 - 1/2, \ldots, 1 - 1/2)$$
$$= 1 - T \circ D(1/2, \ldots, 1/2) = 1 - T(1/2).$$

Consensus preservation Let $P_1, \ldots, P_n \in \mathcal{P}_\Sigma$ and $A \in \Sigma$ such that $P_1(A) = \cdots = P_n(A) = 1$. We show that $P_{P_1, \ldots, P_n}(A) = 1$. Let P be the result of pooling $P_1, \ldots, P_n$ using the (at least on X) neutral pooling function defined above. As that pooling function is consensus-preserving, $P(A) = 1$. It suffices to show that $P_{P_1, \ldots, P_n} \leq K P$, as this implies that $P_{P_1, \ldots, P_n}(A^c) \leq K P(A^c) = K(1 - P(A)) = K(1 - 1) = 0$, so that $P_{P_1, \ldots, P_n}(A) = 1$. Now, to show that $P_{P_1, \ldots, P_n} \leq K P$, note first that, by property (iv) of T, $g \leq K P|_{\widetilde{X}}$, and so $K P|_{\widetilde{X}} - g \geq 0$. Since g and $K P|_{\widetilde{X}}$, and hence also $K P|_{\widetilde{X}} - g$, are σ-additive, σ-finite and non-negative functions on the semi-ring $\widetilde{X}$, each of them extends uniquely to a measure on $\sigma(\widetilde{X}) = \Sigma$ by Caratheodory's Extension Theorem. The first two extensions are $P_{P_1, \ldots, P_n}$ and $K P$, respectively. So the third one must be $K P - P_{P_1, \ldots, P_n}$. Hence $K P - P_{P_1, \ldots, P_n} \geq 0$, and thus $P_{P_1, \ldots, P_n} \leq K P$. □

A.3 Proof of Theorem 3(b)

Let X be a non-path-connected and finite sub-agenda of the σ-algebra Σ. As in the proof of Theorem 1(b), we begin by proving that we may assume without loss of generality that $\sigma(X) = \Sigma$.

Claim 1 If Theorem 3(b) holds when $\sigma(X) = \Sigma$, then it holds in general.

Assume Theorem 3(b) holds if $\sigma(X) = \Sigma$ and let $\Sigma' := \sigma(X) \, (\subseteq \Sigma)$. By assumption, there exists $F' : \mathcal{P}_{\Sigma'}^n \to \mathcal{P}_{\Sigma'}$ which, on X, is independent, consensus-preserving, and non-neutral. Consider some $F : \mathcal{P}_{\Sigma}^n \to \mathcal{P}_{\Sigma}$ which, for any $P_1, \ldots, P_n \in \mathcal{P}_{\Sigma}$, generates a probability function in $\mathcal{P}_{\Sigma}$ extending $F'(P_1|_{\Sigma'}, \ldots, P_n|_{\Sigma'})$ (where such an extension exists by Lemma 5 and finiteness of Σ'). The so-defined F inherits all relevant properties from F': it is, on X, independent, consensus preserving, and non-neutral. This proves the claim.

Now let $\sigma(X) = \Sigma$. Notationally, for any sub-σ-algebra $\bar{\Sigma} \subseteq \Sigma$, let $\mathcal{A}(\bar{\Sigma})$ be its set of atoms (i.e., minimal elements of $\tilde{\Sigma} \setminus \{\varnothing\}$). We now define a pooling function for agenda Σ and show that it has the desired properties. As an ingredient to the definition, let $D' : [0, 1]^n \to [0, 1]$ and $D'' : [0, 1]^n \to [0, 1]$ be the local pooling criteria of two distinct linear pooling functions; and let $\bar{A} \in X \setminus \{\varnothing, \Omega\}$ be a (by assumption existing) event such that not for all $A \in X \setminus \{\varnothing, \Omega\}$ there is $\bar{A} \Vdash^* A$, where $\Vdash^*$ is the transitive closure of $\vdash^*$. Consider any $(P_1, \ldots, P_n) \in \mathcal{P}_{\Sigma}^n$. To define $P_{P_1,\ldots,P_n} \in \mathcal{P}_{\Sigma}$, we start by defining probability functions on two sub-σ-algebras of Σ, denoted Σ' and Σ'' and defined as the σ-algebras generated by the sets

$$X' := \{A \in X : \bar{A} \Vdash^* B \text{ for both } B \in \{A, A^c\}\},$$
$$X'' := \{A \in X : \bar{A} \Vdash^* B \text{ for no } B \in \{A, A^c\}\},$$

respectively. (X' and X'' might be empty, in which case Σ' and Σ'', respectively, are $\{\varnothing, \Omega\}$.) Let $P'_{P_1,\ldots,P_n} \in \mathcal{P}_{\Sigma'}$ and $P''_{P_1,\ldots,P_n} \in \mathcal{P}_{\Sigma''}$ be defined by

$$P'_{P_1,\ldots,P_n}(A) = D'(P_1(A), \ldots, P_n(A)) \text{ for all } A \in \Sigma',$$
$$P''_{P_1,\ldots,P_n}(A) = D''(P_1(A), \ldots, P_n(A)) \text{ for all } A \in \Sigma''.$$

These two functions are indeed probability functions (on Σ' and Σ'', respectively), as they are linear averages of of probability functions.

Claim 2 The σ-algebras Σ' and Σ'' are logically independent, that is: if $A' \in \Sigma'$ and $A'' \in \Sigma''$ are non-empty, so is $A' \cap A''$.

Suppose the contrary. Then, as each non-empty element of Σ' includes an atom of Σ' and hence a non-empty intersection of events in X', and similarly for Σ'', there are consistent sets $Y' \subseteq X'$ and $Y'' \subseteq X''$ such that $Y' \cup Y''$ is inconsistent. Let Y be a minimal inconsistent subset of $Y' \cup Y''$. Then Y is not a subset of Y' or Y'', as Y' and Y'' are consistent. So there are $A \in Y \cap X'$ and $B \in Y \cap X''$. Note that $A \vdash^* B^c$, a contradiction since $A \in X'$ and $B^c \in X''$. This proves Claim 2.

We now extend $P'_{P_1,\ldots,P_n}$ and $P''_{P_1,\ldots,P_n}$ to a probability function on the σ-algebra $\tilde{\Sigma} := \sigma(\Sigma' \cup \Sigma'') = \sigma(X' \cup X'')$, in such a way that the events in Σ' are probabilistically independent of those in Σ''. By Claim 2, the atoms of $\tilde{\Sigma}$ are precisely the intersections of an atom of Σ' and one of Σ'': $\mathcal{A}(\tilde{\Sigma}) = \{A' \cap A'' : A' \in \mathcal{A}(\Sigma'), A'' \in \mathcal{A}(\Sigma'')\}$. Let $\tilde{P}_{P_1,\ldots,P_n}$ be the unique measure on $\tilde{\Sigma}$ that behaves as follows on the atoms:

$$\tilde{P}_{P_1,\ldots,P_n}(A' \cap A'') = P'_{P_1,\ldots,P_n}(A') P''_{P_1,\ldots,P_n}(A''), \tag{3}$$

for all $A' \in \mathcal{A}(\Sigma')$, $A'' \in \mathcal{A}(\Sigma'')$. Now $\tilde{P}_{P_1,\dots,P_n}$ is a *probability function* as

$$\sum_{A \in \mathcal{A}(\tilde{\Sigma})} \tilde{P}_{P_1,\dots,P_n}(A) = \sum_{A' \in \mathcal{A}(\Sigma'), A'' \in \mathcal{A}(\Sigma'')} P'_{P_1,\dots,P_n}(A') P''_{P_1,\dots,P_n}(A'')$$

$$= \sum_{A' \in \mathcal{A}(\Sigma')} P'_{P_1,\dots,P_n}(A') \underbrace{\sum_{A'' \in \mathcal{A}(\Sigma'')} P''_{P_1,\dots,P_n}(A'')}_{=1} = 1.$$

Check that restricting $\tilde{P}_{P_1,\dots,P_n}$ to Σ' and Σ'' yields $P'_{P_1,\dots,P_n}$ and $P''_{P_1,\dots,P_n}$, respectively. So

$$\tilde{P}_{P_1,\dots,P_n}(A) = \begin{cases} D'(P_1(A), \dots, P_n(A)) & \text{for all } A \in \Sigma' \\ D''(P_1(A), \dots, P_n(A)) & \text{for all } A \in \Sigma''. \end{cases} \tag{4}$$

Before we can extend $\tilde{P}_{P_1,\dots,P_n}$ to the full σ-algebra Σ, we prove another claim. For all $A \in X$ such that $\bar{A} \vdash\vdash^* A$ but not $\bar{A} \vdash\vdash^* A^c$, define

$$A_{P_1,\dots,P_n} := \begin{cases} A & \text{if } P_i(A) > 0 \text{ for some } i \\ A^c & \text{if } P_i(A) = 0 \text{ for all } i. \end{cases}$$

Claim 3 For all atoms C of $\tilde{\Sigma}$ $(= \sigma(X' \cup X''))$ with $\tilde{P}_{P_1,\dots,P_n}(C) > 0$, the event $C \cap \left(\cap_{A \in X: \bar{A} \vdash\vdash^* A \text{ and not } \bar{A} \vdash\vdash^* A^c} A_{P_1,\dots,P_n}\right)$ is an atom of Σ.

Let C be as specified, and write $C_{P_1,\dots,P_n}$ for the event in question. As noted above, $C = A' \cap A''$ with $A' \in \mathcal{A}(\Sigma')$ and $A'' \in \mathcal{A}(\Sigma'')$. By $\tilde{P}_{P_1,\dots,P_n}(C) > 0$ and (3), we have $\tilde{P}'_{P_1,\dots,P_n}(A') > 0$ and $\tilde{P}''_{P_1,\dots,P_n}(A'') > 0$. Since $A' \in \mathcal{A}(\Sigma')$, we may write $A' = \cap_{A \in Y'} A$ for some set $Y' \subseteq X'$ containing exactly one member of each pair $A, A^c \in X'$. Similarly, $A'' = \cap_{A \in Y''} A$ for some set $Y'' \subseteq X''$ containing exactly one member of each pair $A, A^c \in X''$. Note also that $\cap_{A \in X: \bar{A} \vdash\vdash^* A \text{ and not } \bar{A} \vdash\vdash^* A^c} A_{P_1,\dots,P_n}$ can be written as $\cap_{A \in Y_{P_1,\dots,P_n}} A$, where the set

$$Y_{P_1,\dots,P_n} = \{A_{P_1,\dots,P_n} : A \in X, \ \bar{A} \vdash\vdash^* A, \ \text{not } \bar{A} \vdash\vdash^* A^c\}$$

consists of exactly one member of each pair $A, A^c \in X \backslash (X' \cup X'')$. So $C_{P_1,\dots,P_n} = \cap_{A \in Y' \cup Y'' \cup Y_{P_1,\dots,P_n}} A$, where the set $Y' \cup Y'' \cup Y_{P_1,\dots,P_n}$ consists of exactly one member of each pair $A, A^c \in X$. So, since $\Sigma = \sigma(X)$, $C_{P_1,\dots,P_n}$ is an atom or is empty. Hence it suffices to show that $C_{P_1,\dots,P_n} \neq \varnothing$. Suppose the contrary. Then $Y' \cup Y'' \cup Y_{P_1,\dots,P_n}$ is inconsistent, hence has a minimal inconsistent subset Y. We distinguish two cases and derive a contradiction in each.

Case 1 There is some $B \in Y \cap Y_{P_1,\dots,P_n}$ with $\bar{A} \vdash\vdash^* B$. Consider some $B' \in Y \backslash \{B\}$. We have (i) not $\bar{A} \vdash\vdash^* B'$ (otherwise by $B' \vdash^* B'^c$ we would have $\bar{A} \vdash\vdash^* B'^c$, hence $B \in X'$, a contradiction as $B \in Y_{P_1,\dots,P_n}$). Further, (ii) $\bar{A} \vdash\vdash^* (B')^c$ (as $\bar{A} \vdash\vdash^* B$ and $B \vdash^* (B')^c$). By (i) and (ii), letting $A := (B')^c$, the event $A_{P_1,\dots,P_n}$ $(\in \{A, A^c\})$ is well-defined. Since $Y_{P_1,\dots,P_n}$ contains $A_{P_1,\dots,P_n}$ $(\in \{A, A^c\})$ and contains $B' = A^c$ but not $(B')^c = A$, we must have $A_{P_1,\dots,P_n} = A^c$. So, for all i, $P_i(A) = 0$ and

hence $P_i(B') = 1$. Note that this holds for all $B' \in Y \backslash \{B\}$. So $P_i(\cap_{B' \in Y} B') = P_i(B)$ for all i. Hence, as Y is inconsistent, $P_i(B) = 0$ for all i. Thus $B_{P_1, \ldots, P_n} = B^c$. So $B^c \in Y_{P_1, \ldots, P_n}$, a contradiction as $B \in Y_{P_1, \ldots, P_n}$.

Case 2 There is no $B \in Y \cap Y_{P_1, \ldots, P_n}$ with $\bar{A} \vdash\vdash^* B$. Then all $B \in Y \cap Y_{P_1, \ldots, P_n}$ take the form $A_{P_1, \ldots, P_n} = A^c$, so that $P_i(A) = 0$ for all i, i.e., $P_i(B) = 1$ for all i. So, (*) $P_i(\cap_{B \in Y} B) = P_i(\cap_{B \in Y \backslash Y_{P_1, \ldots, P_n}} B)$ for all i. Now, either (i) $Y \subseteq Y_{P_1, \ldots, P_n} \cup Y'$, or (ii) $Y \subseteq Y_{P_1, \ldots, P_n} \cup Y''$, because otherwise there are $A' \in Y'$ and $A'' \in Y''$, and $A' \vdash^* (A'')^c$, whence $\bar{A} \vdash\vdash^* (A'')^c$, a contradiction as $(A'')^c \in X''$. First suppose case (i) holds. Then $Y \backslash Y_{P_1, \ldots, P_n} \subseteq Y'$, and so (*) implies that (**) $P_i(\cap_{B \in Y} B) \geq P_i(\cap_{B \in Y'} B) = P_i(A')$ for all i. Since by assumption $\tilde{P}_{P_1, \ldots, P_n}(A') > 0$, there is (by (4)) at least one i with $P_i(A') > 0$, hence by (**) with $P_i(\cap_{B \in Y} B) > 0$. So $\cap_{B \in Y} B \neq \varnothing$, i.e., Y is consistent, a contradiction. Similarly, in case (ii), one can show that Y is consistent, a contradiction. This completes the proof of Claim 3.

Let $P_{P_1, \ldots, P_n}$ be the unique measure on Σ behaving as follows on any atom C of Σ. If C takes the form as in Claim 3, i.e., $B = C \cap (\cap_{A \in X: \bar{A} \vdash\vdash^* A \text{ and not } \bar{A} \vdash\vdash^* A^c} A_{P_1, \ldots, P_n})$ where $C \in \mathcal{A}(\tilde{\Sigma})$ and $\tilde{P}_{P_1, \ldots, P_n}(C) > 0$, then let $P_{P_1, \ldots, P_n}(B) := \tilde{P}_{P_1, \ldots, P_n}(C)$. Otherwise let $P_{P_1, \ldots, P_n}(B) := 0$.

Claim 4 $P_{P_1, \ldots, P_n}$ extends $\tilde{P}_{P_1, \ldots, P_n}$ (in particular, is a *probability function*).

It suffices to show that $P_{P_1, \ldots, P_n}$ coincides with $\tilde{P}_{P_1, \ldots, P_n}$ on $\mathcal{A}(\tilde{\Sigma})$. Consider any $C \in \mathcal{A}(\tilde{\Sigma})$. As Σ is a refinement of $\tilde{\Sigma}$,

$$P_{P_1, \ldots, P_n}(C) = \sum_{B \in \mathcal{A}(\Sigma): B \subseteq C} P_{P_1, \ldots, P_n}(B). \tag{5}$$

There are two cases.

Case 1 $\tilde{P}_{P_1, \ldots, P_n}(C) = 0$. Then, for all $B \in \mathcal{A}(\Sigma)$ with $B \subseteq C$, we have $P_{P_1, \ldots, P_n}(B) = 0$ (by definition of $P_{P_1, \ldots, P_n}$), and so by (5) we have $P_{P_1, \ldots, P_n}(C) = 0 = \tilde{P}_{P_1, \ldots, P_n}(C)$, as desired.

Case 2 $\tilde{P}_{P_1, \ldots, P_n}(C) > 0$. Then, among all atoms $B \in \mathcal{A}(\Sigma)$ with $B \subseteq C$, there is by definition of $P_{P_1, \ldots, P_n}$ exactly one such that $P_{P_1, \ldots, P_n}(B) > 0$ (namely $B = C \cap (\cap_{A \in X: \bar{A} \vdash\vdash^* A \text{ and not } \bar{A} \vdash\vdash^* A^c} A_{P_1, \ldots, P_n})$), and $P_{P_1, \ldots, P_n}(B) = \tilde{P}_{P_1, \ldots, P_n}(C)$. So by (5) $P_{P_1, \ldots, P_n}(C) = \tilde{P}_{P_1, \ldots, P_n}(C)$. This completes the proof of Claim 4.

Claim 5 For all $A \in X$ such that $\bar{A} \vdash\vdash^* A$ and not $\bar{A} \vdash\vdash^* A^c$, $P_{P_1, \ldots, P_n}(A)$ is 1 if, for some individual i, $P_i(A) > 0$, and 0 otherwise.

By definition of $P_{P_1, \ldots, P_n}$, all atoms of Σ with positive probability are subsets of the event $\cap_{A \in X: \bar{A} \vdash\vdash^* A \text{ and not } \bar{A} \vdash\vdash^* A^c} A_{P_1, \ldots, P_n}$. So this event has probability 1. Hence, for all $A \in X$ such that $\bar{A} \vdash\vdash^* A$ and not $\bar{A} \vdash\vdash^* A^c$, we have $P_{P_1, \ldots, P_n}(A_{P_1, \ldots, P_n}) = 1$, and so

$$P_{P_1, \ldots, P_n}(A) = \begin{cases} 1 \text{ if } A_{P_1, \ldots, P_n} = A, \; i.e., & \text{if } P_i(A) > 0 \text{ for some } i \\ 0 \text{ if } A_{P_1, \ldots, P_n} = A^c, \; i.e., & \text{if } P_i(A) = 0 \text{ for all } i. \end{cases}$$

This proves Claim 5.

By Claim 4, we have constructed a well-defined pooling function $(P_1, \ldots, P_n) \mapsto P_{P_1, \ldots, P_n}$ for agenda Σ. By (4) and Claims 4 and 5, we know its behaviour on the entire

sub-agenda X: the pooling function is independent on X and the local pooling criteria D_A of events $A \in X$ are given by

  (i) the linear criterion D' if $A \in X'$,
 (ii) the different linear criterion D'' if $A \in X''$,
(iii) a non-linear criterion $\hat{D}$ (taking the value 0 at $\mathbf{0}$ and the value 1 everywhere else) if $\bar{A} \vdash\vdash^* A$ but not $\bar{A} \vdash\vdash^* A^c$,
(iv) the different non-linear criterion $1 - \hat{D}(\mathbf{1} - \cdot)$ if not $\bar{A} \vdash\vdash^* A$ but $\bar{A} \vdash\vdash^* A^c$.

These pooling criteria also ensure unanimity preservation on X. To check non-neutrality, it suffices to show that at least two of the four different types of events (i)–(iv) do indeed occur. This is so because $\bar{A}$ is of type (i) or (iii) and because by assumption there exists an $A \in X$ such that not $\bar{A} \vdash\vdash^* A$, i.e., such that A has type (ii) or (iv).   □

A.4 Proof of Theorem 4(b)

Consider any finite sub-agenda $X \neq \{\varnothing, \Omega\}$ (of the σ-algebra agenda Σ) which is nested or satisfies $|X \backslash \{\varnothing, \Omega\}| \leq 4$. If X is nested, the claim follows from Theorem 1(b), as non-neutrality on X implies non-linearity on X. Now assume $|X \backslash \{\varnothing, \Omega\}| \leq 4$. We reduce the claim to Part I's Theorem 4(b). By that result, there is a pooling function F' for agenda X which is independent, consensus compatible, and not linear. By Lemma 4, F' is induced by a pooling function for agenda Σ which is independent on X, (globally) consensus-preserving, and not linear on X.   □

A.5 Proof of Theorem 5(b)

Consider a simple sub-agenda X of σ-algebra Σ, where X is finite and not $\{\varnothing, \Omega\}$. We construct a pooling function which, on X, is independent (in fact, neutral), conditional-consensus-preserving, and non-linear. We may assume without loss of generality that $\sigma(X) = \Sigma$, because the 'Claim' in the proof of Theorem 1(b) holds analogously here as well.

As an ingredient of the construction, we use an arbitrary pooling function $(P_1, \ldots, P_n) \mapsto P_{P_1, \ldots, P_n}^{\text{lin}}$ which, at least on X, is linear and conditional-consensus-preserving. The function could be simply given by $(P_1, \ldots, P_n) \mapsto P_1$, which is even *globally* linear and conditional-consensus-preserving. Let D^{lin} be its pooling criterion for all events in X. To anticipate, the pooling function $(P_1, \ldots, P_n) \mapsto P_{P_1, \ldots, P_n}$ to be constructed will have the pooling criterion $D : [0, 1]^n \to [0, 1]$ for each event in X, where

$$D(t_1, \ldots, t_n) := \begin{cases} 0 & \text{if } D^{\text{lin}}(t_1, \ldots, t_n) < 1/2, \\ 1/2 & \text{if } D^{\text{lin}}(t_1, \ldots, t_n) = 1/2, \\ 1 & \text{if } D^{\text{lin}}(t_1, \ldots, t_n) > 1/2. \end{cases} \tag{6}$$

Consider any $P_1, \ldots, P_n \in \mathcal{P}_\Sigma$. We must define $P_{P_1, \ldots, P_n}$. We use the following notation (which suppresses the parameters $P_1, \ldots, P_n$):

$$p(A) := P^{\text{lin}}_{P_1,\ldots,P_n}(A) \text{ for all } A \in \Sigma,$$
$$X_{\geq 1/2} := \{A \in X : p(A) \geq 1/2\},$$
$$X_{>1/2} := \{A \in X : p(A) > 1/2\},$$
$$X_{=1/2} := \{A \in X : p(A) = 1/2\}.$$

Notice that for all $A \in X$ we have $A \in X_{>1/2} \Rightarrow A^c \notin X_{>1/2}$ and $A \in X_{=1/2} \Leftrightarrow A^c \in X_{=1/2}$. We now prove two claims (which use X's simplicity).

Claim 1 $X_{=1/2}$ can be partitioned into two (possibly empty) sets $X^1_{=1/2}$ and $X^2_{=1/2}$ such that (i) each $X^j_{=1/2}$ satisfies $p(A \cap B) > 0$ for all $A, B \in X^j_{=1/2}$ and (ii) each $X^j_{=1/2} \cup X_{>1/2}$ is consistent (whence $X^j_{=1/2}$ contains exactly one member of each pair $A, A^c \in X_{=1/2}$).

To show this, note first that $X_{=1/2}$ has a subset Y such that $p(A \cap B) > 0$ for all $A, B \in Y$ (e.g., $Y = \varnothing$). Among all such subsets $Y \subseteq X_{=1/2}$, let $X^1_{=1/2}$ a *maximal* one, and let $X^2_{=1/2} := X_{=1/2} \backslash X^1_{=1/2}$. By definition, $X^1_{=1/2}$ and $X^2_{=1/2}$ form a partition of $X_{=1/2}$. We show that (i) and (ii) hold.

(i) Property (i) holds for $X^1_{=1/2}$ by definition, and for $X^2_{=1/2}$ by the following argument. Let $A, B \in X^2_{=1/2}$ and for a contradiction let $p(A \cap B) = 0$. By the maximality property of $X^1_{=1/2}$, there are $A', B' \in X^1_{=1/2}$ such that $p(A \cap A') = 0$ and $p(B \cap B') = 0$. Thus, $p(A \cap C) = p(B \cap C) = 0$ where $C := A' \cap B'$. Since the intersection of any two of the sets A, B, C has zero p-probability, we must have $p(A) + p(B) + p(C) = p(A \cup B \cup C) \leq 1$, a contradiction because $p(A) = p(B) = 1/2$ and $p(C) = p(A' \cap B') > 0$ (the latter because $X^1_{=1/2}$ satisfies (i)).

(ii) For a contradiction, let some $X^j_{=1/2} \cup X_{>1/2}$ be inconsistent. Then (since X and hence $X^j_{=1/2} \cup X_{>1/2}$ are finite) there is a minimal inconsistent subset $Y \subseteq X^j_{=1/2} \cup X_{>1/2}$. Since X is simple, we have $|Y| \leq 2$, say $Y = \{A, B\}$. Since $A \cap B = \varnothing$, we have $p(A) + p(B) = p(A \cup B) \leq 1$. And since $p(A), p(B) \geq 1/2$, it follows that $p(A) = p(B) = 1/2$, i.e., $A, B \in X^j_{=1/2}$. Hence, by (i), we have $p(A \cap B) > 0$, a contradiction as $A \cap B = \varnothing$.

Claim 2 $\cap_{C \in X^1_{=1/2} \cup X_{>1/2}} C$ and $\cap_{C \in X^2_{=1/2} \cup X_{>1/2}} C$ are atoms of the σ-algebra Σ, i.e., ($\subseteq$-)minimal elements of $\Sigma \backslash \{\varnothing\}$ (they are the *same* atoms if and only if $X_{=1/2} = \varnothing$, i.e., if and only if $X^1_{=1/2} = X^2_{=1/2} = \varnothing$).

To show this, first write X as $\{C^0_j, C^1_j : j = 1, \ldots, J\}$, where $J = |X|/2$ and each pair C^0_j, C^1_j consists of an event and its complement. We may write Σ as

$$\Sigma = \{\cup_{(k_1,\ldots,k_J) \in K} (C^{k_1}_1 \cap \cdots \cap C^{k_J}_J) : K \subseteq \{0,1\}^J\}. \tag{7}$$

Recall that Σ is the σ-algebra generated by X. The inclusion '$\supseteq$' in (7) is obvious, and the inclusion '$\subseteq$' holds because the right side of (7) includes X (since any $C^k_j \in X$ can be written as the union of all $C^{k_1}_1 \cap \cdots \cap C^{k_J}_J$ for which $k_j = k$) and is a σ-algebra (check closedness under taking unions and complements).

From (7) and the pairwise disjointness of the intersections of the form $C_1^{k_1} \cap \cdots \cap C_J^{k_J}$, it is clear that every consistent such intersection is an atom of Σ. Now $\cap_{C \in X_{=1/2}^j \cup X_{>1/2}} C$ is (for $j \in \{0, 1\}$) precisely such a consistent intersection. Indeed, $\cap_{C \in X_{=1/2}^j \cup X_{>1/2}} C$ is consistent by Claim 1, and contains a member of each pair A, A^c in X. The latter holds by Claim 1 if $p(A) = p(A^c) (= 1/2)$, and otherwise because there is a $B \in \{A, A^c\}$ with $p(B) > 1/2$, i.e., with $B \in X_{>1/2} \subseteq X_{=1/2}^j \cup X_{>1/2}$. This proves Claim 2.

We can now define $P_{P_1,\ldots,P_n}$. By Claim 1, we may pick $\omega^1 \in \cap_{C \in X_{=1/2}^1 \cup X_{>1/2}} C$ and $\omega^2 \in \cap_{C \in X_{=1/2}^2 \cup X_{>1/2}} C$, where we assume that $\omega^1 = \omega^2$ if $X_{=1/2} = \varnothing$, i.e., if $\cap_{C \in X_{=1/2}^1 \cup X_{>1/2}} C = \cap_{C \in X_{=1/2}^2 \cup X_{>1/2}} C = \cap_{C \in X_{>1/2}} C$. Let δ_{ω^1} and δ_{ω^2} be, respectively, the Dirac measures on Σ at ω^1 and ω^2, given for all $A \in \Sigma$ by $\delta_{\omega^j}(A) = 1$ if $\omega^j \in A$ and $\delta_{\omega^j}(A) = 0$ if $\omega^j \notin A$. Let

$$P_{P_1,\ldots,P_n} := \frac{1}{2}\delta_{\omega^1} + \frac{1}{2}\delta_{\omega^2},$$

where ω^1 and ω^2 depend on $P_1, \ldots, P_n$ via $X_{=1/2}^1$, $X_{=1/2}^2$, $X_{>1/2}$. So $P_{P_1,\ldots,P_n}(A)$ is 1 or 1/2 or 0 depending on whether $A (\in \Sigma)$ contains both, exactly one, or none of ω^1 and ω^2; and $P_{P_1,\ldots,P_n} = \delta_\omega$ if $\omega^1 = \omega^2 = \omega$, i.e., if $X_{=1/2} = \varnothing$. We finally show that the so-defined pooling function $(P_1, \ldots, P_n) \mapsto P_{P_1,\ldots,P_n}$ has all desired properties.

Independence on X We in fact show something stronger, i.e., neutrality on X with pooling criterion D given in (6). Let $P_1, \ldots, P_n \in \mathcal{P}_\Sigma$, $A \in X$ and $(t_1, \ldots, t_n) := (P_1(A), \ldots, P_n(A))$. We prove that $P_{P_1,\ldots,P_n}(A) = D(t_1, \ldots, t_n)$ by considering three cases and using the above notation p, $X_{>1/2}$, $X_{=1/2}^1$, $X_{=1/2}^2$, ω^1, ω^2.

Case 1 $p(A) = D^{\text{lin}}(t_1, \ldots, t_n) < 1/2$. Here $D(t_1, \ldots, t_n) = 0$. So we must prove that $P_{P_1,\ldots,P_n}(A) = 0$, i.e., that $\omega^1, \omega^2 \notin A$. Assume for a contradiction that $\omega^1 \in A$ (the proof is analogous if we instead assume $\omega^2 \in A$). Then A includes $\cap_{C \in X_{=1/2}^1 \cup X_{>1/2}} C$, as this set contains ω^1 and is by Claim 2 an atom of Σ. So $A^c \cap [\cap_{C \in X_{=1/2}^1 \cup X_{>1/2}} C] = \varnothing$. Hence the set $\{A^c\} \cup X_{=1/2}^1 \cup X_{>1/2}$ is inconsistent, so has a minimal inconsistent subset Y. As X is simple, $|Y| \le 2$. Now $\varnothing \notin Y$ as $A^c \ne \varnothing$ (by $p(A^c) = 1 - p(A) > 1/2$) and as all $B \in X_{=1/2}^1 \cup X_{>1/2}$ are non-empty (by $p(B) \ge 1/2$). So $|Y| = 2$. Further, Y is not a subset of $X_{=1/2}^1 \cup X_{>1/2}$, as this set is consistent by Claim 1. So $Y = \{A^c, B\}$ for some $B \in X_{=1/2}^1 \cup X_{>1/2}$. As $A^c \cap B = \varnothing$ and as $p(A^c) = 1 - p(A) > 1/2$ and $p(B) \ge 1/2$, we have $p(A^c \cup B) = p(A^c) + p(B) > 1/2 + 1/2 = 1$, a contradiction.

Case 2 $p(A) = D^{\text{lin}}(t_1, \ldots, t_n) > 1/2$. Then $D(t_1, \ldots, t_n) = 1$. Hence we must prove that $P_{P_1,\ldots,P_n}(A) = 1$, i.e., that $P_{P_1,\ldots,P_n}(A^c) = 0$. The latter follows from case 1 as applied to A^c, since $p(A^c) = 1 - p(A) < 1/2$.

Case 3 $p(A) = D^{\text{lin}}(t_1, \ldots, t_n) = 1/2$. Then $D(t_1, \ldots, t_n) = 1/2$. So we must prove that $P_{P_1,\ldots,P_n}(A) = 1/2$, i.e., that A contains exactly one of ω^1 and ω^2. As $p(A) = 1/2$, exactly one of $X_{=1/2}^1$ and $X_{=1/2}^2$ contains A and the other one contains A^c, by Claim 1. Say $A \in X_{=1/2}^1$ and $A^c \in X_{=1/2}^2$ (the proof is analogous if instead

$A \in X^2_{=1/2}$ and $A^c \in X^1_{=1/2}$). So $A \supseteq \cap_{C \in X^1_{=1/2} \cup X_{>1/2}} C$, whence $\omega^1 \in A$. Further, $\omega^2 \notin A$ because A is disjoint from A^c, hence from its subset $\cap_{C \in X^2_{=1/2} \cup X_{>1/2}} C$ which contains ω^2.

Non-linearity on X Pooling cannot be linear, since otherwise for any fixed $A \in X \backslash \{\Omega, \varnothing\}$ ($\neq \varnothing$) the collective probabilities $P_{P_1,\ldots,P_n}(A)$ could take any given values $t \in [0, 1]$ (for instance by letting $P_1(A) = \cdots = P_n(A) = t$), a contradiction, since by definition $P_{P_1,\ldots,P_n}(A) \in \{0, 1/2, 1\}$.

Conditional-consensus-preservation on X Let $A, B \in X$ and $P_1, \ldots, P_n \in \mathcal{P}_\Sigma$ such that $P_i(A \cup B) = 1$ for all i. We show that $P_{P_1,\ldots,P_n}(A \cup B) = 1$, which establishes conditional-consensus-preservation on X by Proposition 1(a). For all i, $P_i(A) + P_i(B) \geq P_i(A \cup B) = 1$, and hence $P_i(A) \geq 1 - P_i(B) = P_i(B^c)$. So, as $D^{\mathrm{lin}} : [0, 1]^n \to [0, 1]$ takes a linear form with non-negative coefficients and hence is weakly increasing in every component,

$$
\begin{aligned}
D^{\mathrm{lin}}(P_1(A), \ldots, P_n(A)) &\geq D^{\mathrm{lin}}\left(P_1\left(B^c\right), \ldots, P_n\left(B^c\right)\right) \\
&= D(\mathbf{1}) - D^{\mathrm{lin}}\left(P_1(B), \ldots, P_n(B)\right) \\
&= 1 - D^{\mathrm{lin}}\left(P_1(B), \ldots, P_n(B)\right).
\end{aligned}
$$

So, with p as defined earlier, $p(A) \geq 1 - p(B)$, i.e., $p(A) + p(B) \geq 1$. We distinguish between three cases.

Case 1 $p(A) > 1/2$. Then, by the above proof of independence on X, $P_{P_1,\ldots,P_n}(A) = 1$. So $P_{P_1,\ldots,P_n}(A \cup B) = 1$, as desired.

Case 2 $p(B) > 1/2$. Then, again by the above proof of independence on X, $P_{P_1,\ldots,P_n}(B) = 1$. Hence, $P_{P_1,\ldots,P_n}(A \cup B) = 1$, as desired.

Case 3 $p(A), p(B) \leq 1/2$. Then, as $p(A) + p(B) \geq 1$, we have $p(A) = p(B) = 1/2$. Let $X_{>1/2}, X^1_{=1/2}, X^2_{=1/2}, \omega^1, \omega^2$ be as defined above. Note that $A, B \in X^1_{=1/2} \cup X^2_{=1/2}$. It cannot be that A and B are both in $X^1_{=1/2}$: otherwise A^c and B^c are both in $X^2_{=1/2}$ by Claim 1, whence $p(A^c \cap B^c) > 0$ (again by Claim 1), a contradiction since

$$
p\left(A^c \cap B^c\right) = p\left((A \cup B)^c\right) = 1 - p(A \cup B) = 1 - 1 = 0
$$

(where $p(A \cup B) = 1$ because $p(A \cup B) = P^{\mathrm{lin}}_{P_1,\ldots,P_n}(A \cup B)$ and $P_i(A \cup B) = 1$ for all i). Analogously, it cannot be that A and B are both in $X^2_{=1/2}$. So one of A and B is in $X^1_{=1/2}$ and the other one in $X^2_{=1/2}$; say $A \in X^1_{=1/2}$ and $B \in X^2_{=1/2}$ (the proof is analogous otherwise). So $A \supseteq \cap_{C \in X^1_{=1/2} \cup X_{>1/2}} C$ and $B \supseteq \cap_{C \in X^2_{=1/2} \cup X_{>1/2}} C$, and hence $\omega^1 \in A$ and $\omega^2 \in B$. Thus $\omega^1, \omega^2 \in A \cup B$, whence $P_{P_1,\ldots,P_n}(A \cup B) = 1$. $\square$

A.6 Proof of Proposition 2

Consider the σ-algebra agenda Σ, and let $|\Sigma| > 2^3 = 8$, i.e., $|\Sigma| \geq 2^4 = 16$. Then Σ includes a partition of Ω into four non-empty events. Let X be the sub-agenda consisting of any union of *two* of these four events. In the proof of Part I's Proposition 2 we construct a pooling function for this agenda X which is neutral,

consensus-preserving, and non-linear.[17] By Lemma 4, this pooling function is induced by a pooling function for agenda Σ which, on X, is neutral, consensus-preserving, and non-linear. □

References

Aczél J, Wagner C (1980) A characterization of weighted arithmetic means. SIAM J Algebra Discr Methods 1(3):259–260

Aczél J, Ng CT, Wagner C (1984) Aggregation theorems for allocation problems. SIAM J Algebra Discr Methods 5(1):1–8

Chambers C (2007) An ordinal characterization of the linear opinion pool. Econ Theor 33(3):457–474

Clemen RT, Winkler RL (1999) Combining probability distributions from experts in risk analysis. Risk Anal 19(2):187–203

Dietrich F (2006) Judgment aggregation: (im)possibility theorems. J Econ Theory 126(1):286–298

Dietrich F (2010) Bayesian group belief. Soc Choice Welf 35(4):595–626

Dietrich F (2016) A theory of Bayesian groups. Working paper

Dietrich F, List C (2007) Arrow's theorem in judgment aggregation. Soc Choice Welf 29(1):19–33

Dietrich F, List C (2016) Probabilistic opinion pooling. In: Hitchcock C, Hajek A (eds), Oxford Handbook of Probability and Philosophy, Oxford University Press, Oxford

Dietrich F, List C (2017) Probabilistic opinion pooling generalized. Part one: general agendas. Soc Choice Welf. doi:10.1007/s00355-017-1034-z (this issue)

Dietrich F, Mongin P (2010) The premise-based approach to judgment aggregation. J Econ Theory 145(2):562–582

Genest C (1984a) Pooling operators with the marginalization property. Can J Stat 12(2):153–163

Genest C (1984b) A characterization theorem for externally Bayesian groups. Ann Stat 12(3):1100–1105

Genest C, McConway KJ, Schervish MJ (1986) Characterization of externally Bayesian pooling operators. Ann Stat 14(2):487–501

Genest C, Zidek JV (1986) Combining probability distributions: a critique and annotated bibliography. Stat Sci 1(1):114–135

Kornhauser LA, Sager LG (1986) Unpacking the Court. Yale Law J 96(1):82–117

List C, Pettit P (2002) Aggregating sets of judgments: an impossibility result. Econ Philos 18(1):89–110

List C (2006) The discursive dilemma and public reason. Ethics 116(2):362–402

McConway KJ (1978) The combination of experts' opinions in probability assessments: some theoretical considerations. Ph.D. thesis, University College London

McConway KJ (1981) Marginalization and linear opinion pools. J Am Stat Assoc 76(374):410–414

Mongin P (1995) Consistent Bayesian aggregation. J Econ Theory 66:313–351

Mongin P (2005) Spurious unanimity and the Pareto principle. LSE Choice Group working paper series, London School of Economics

Mongin P (2008) Factoring out the impossibility of logical aggregation. J Econ Theory 141(1):100–113

Morris PA (1974) Decision analysis expert use. Manag Sci 20:1233–1241

Nehring K, Puppe C (2010) Abstract arrovian aggregation. J Econ Theory 145(2):467–494

Pettit P (2001) Deliberative democracy and the discursive dilemma. Philos Issue 11:268–299

Priest G (2001) An introduction to non-classical logic. Cambridge University Press, Cambridge

Russell JS, Hawthorne J, Buchak L (2015) Groupthink. Philos Stud 172(5):1287–1309

Stone M (1961) The opinion pool. Ann Math Stat 32(4):1339–1342

Wagner C (1982) Allocation, Lehrer models, and the consensus of probabilities. Theor Decis 14(2):207–220

Wagner C (1985) On the formal properties of weighted averaging as a method of aggregation. Synthese 62(1):97–108

Weymark J (1997) Aggregating ordinal probabilities on finite sets. J Econ Theory 75(2):407–432

---

[17] That proof took the four mentioned events to be singleton, but nothing depends on this.